

GFD-I.080
Open Grid Services Architecture
<http://forge.gridforum.org/projects/ogsa-wg>

Authors:

I. Foster, Argonne & U.Chicago (Editor)
H. Kishimoto, Fujitsu (Editor)
A. Savva, Fujitsu (Editor)
D. Berry, NeSC
A. Djaoui, CCLRC-RAL
A. Grimshaw, UVa
B. Horn, IBM
F. Maciel, Hitachi
F. Siebenlist, ANL
R. Subramaniam, Intel
J. Treadwell, HP
J. Von Reich, HP
24 July 2006

The Open Grid Services Architecture, Version 1.5

Status of this Document

This document provides information to the community regarding the specification of the Open Grid Services Architecture (OGSA™). It does not define any standards or technical recommendations. Distribution is unlimited.

Copyright Notice

Copyright © Open Grid Forum (2002-2006). All Rights Reserved.

Trademarks

OGSA is a trademark of the Open Grid Forum.

Abstract

Successful realization of the Open Grid Services Architecture (OGSA) vision of a broadly applicable and adopted framework for distributed system integration, virtualization, and management requires the definition of a core set of interfaces, behaviors, resource models, and bindings. This document, produced by the OGSA working group within the Global Grid Forum (GGF), provides a high-level OGSA definition. The document focuses on requirements and the scope of important capabilities required to support Grid systems and applications in both e-science and e-business. The capabilities described are Execution Management, Data, Resource Management, Security, Self-Management, and Information. The description of each capability includes, to some extent, possible interrelationships with other capabilities. Capabilities are, however, largely independent of each other and there is no requirement that they all be present in an OGSA system. Further, capabilities themselves are not monolithic and it is possible that only a subset of a given OGSA capability may be present in an implementation of an OGSA system.

Contents

Abstract.....	1
1 Introduction	4
2 Requirements.....	5
2.1 Interoperability and Support for Dynamic and Heterogeneous Environments.....	6
2.2 Resource Sharing Across Organizations	7
2.3 Optimization.....	7
2.4 Quality of Service (QoS) Assurance.....	8
2.5 Job Execution.....	8
2.6 Data Services.....	9
2.7 Security	10
2.8 Administrative Cost Reduction.....	10
2.9 Scalability.....	11
2.10 Availability.....	11
2.11 Ease of Use and Extensibility.....	12
3 Capabilities.....	12
3.1 Overview	12
3.2 OGSA Framework	14
3.3 Infrastructure Services	16
3.3.1 Objectives	16
3.3.2 Approach.....	16
3.3.3 Web Services Foundation	16
3.3.4 Naming.....	17
3.3.5 Security Services	18
3.3.6 Representing State.....	18
3.3.7 Notification.....	18
3.3.8 Transactions.....	19
3.3.9 Orchestration.....	19
3.3.10 Ensuring Interoperability: OGSA Profiles	20
3.4 Execution Management Services	20
3.4.1 Objectives	20
3.4.2 Approach.....	21
3.4.3 EMS Services.....	22
3.4.4 Resources	22
3.4.5 Job Management.....	22
3.4.6 Selection Services.....	24
3.4.7 Interactions with Other OGSA Services	25
3.4.8 Example Scenarios	26
3.5 Data Services.....	28
3.5.1 Objectives	28
3.5.2 Approach.....	28
3.5.3 Heterogeneity and Extensibility	31
3.5.4 Functional Capabilities.....	32
3.5.5 Properties	35
3.5.6 Interactions with Other OGSA Services	36
3.6 Resource Management Services	37
3.6.1 Objectives	37
3.6.2 Model	38
3.6.3 Management Capabilities.....	39

3.6.4	Properties	40
3.6.5	Interactions with Other OGSA Services	40
3.7	Security Services	41
3.7.1	Objectives	41
3.7.2	Model	42
3.7.3	Example Scenarios	44
3.7.4	Functional Capabilities.....	46
3.7.5	Properties	48
3.7.6	Interactions with Other OGSA Services	48
3.8	Self-Management Services	49
3.8.1	Objectives	49
3.8.2	Basic Attributes	50
3.8.3	Example Scenarios	50
3.8.4	Functional Capabilities.....	51
3.8.5	Properties	53
3.8.6	Interactions with Other OGSA Services	53
3.9	Information Services.....	54
3.9.1	Objectives	54
3.9.2	Models.....	54
3.9.3	Example Scenarios	55
3.9.4	Functional Capabilities.....	57
3.9.5	Properties	59
3.9.6	Interactions with Other OGSA Services	59
4	Security Considerations.....	59
	Editor Information.....	60
	Contributors.....	60
	Acknowledgements.....	60
	Glossary	60
	Intellectual Property Statement	61
	Disclaimer.....	61
	Full Copyright Notice	61
	References	61

1 Introduction

Grid systems and applications aim to integrate, virtualize, and manage resources and services within distributed, heterogeneous, dynamic “virtual organizations” [Grid Anatomy] [Grid Physiology]. The realization of this goal requires the disintegration of the numerous barriers that normally separate different computing systems within and across organizations, so that computers, application services, data, and other resources can be accessed as and when required, regardless of physical location.

Key to the realization of this Grid vision is standardization, so that the diverse components that make up a modern computing environment can be discovered, accessed, allocated, monitored, accounted for, billed for, etc., and in general managed as a single virtual system—even when provided by different vendors or operated by different organizations. Standardization is critical if we are to create interoperable, portable, and reusable components and systems; it can also contribute to the development of secure, robust, and scalable Grid systems by facilitating the use of good practices.

We present here a service-oriented architecture, the *Open Grid Services Architecture* (OGSA™), that addresses this need for standardization by defining a set of core capabilities and behaviors that address key concerns in Grid systems. These concerns include such issues as: How do I establish identity and negotiate authentication? How is policy expressed and negotiated? How do I discover services? How do I negotiate and monitor service level agreements? How do I manage membership of, and communication within, virtual organizations? How do I organize service collections hierarchically so as to deliver reliable and scalable service semantics? How do I integrate data resources into computations? How do I monitor and manage collections of services?

This document is divided into two main parts, focused on requirements and capabilities, respectively. In §2, we provide an abstract definition of the set of requirements that OGSA is intended to address. This analysis is based on requirements, technical challenges, use cases, previous experience, and the state of the art in related work. The abstract rendering is not constrained by any assumptions about the underlying infrastructure, but rather is intended to frame the “Grid” discussion and define solutions.

In §3, we translate the requirements of §2 into a coherent set of capabilities that collectively define OGSA. We first describe infrastructure services and assumptions that constrain our development of the OGSA design, in particular explaining how OGSA both builds on, and is contributing to the development of, the growing collection of technical specifications that form the emerging Web Services Architecture [WS-Architecture]. Then we present a refinement of the required functionality into capabilities: Execution Management, Data, Resource Management, Security, Self-Management, and Information services.

This document uses a number of terms whose meaning may require more explanation than given in the text. A companion document, the *OGSA Glossary of Terms* [OGSA Glossary], provides unambiguous definitions of such terms.

Other companion documents will provide descriptions of specific services, making clear where existing service specifications can be used unchanged and where modified or new service specifications are needed. They will also describe the current state of any work known to be underway to define such extensions or definitions.

This informational document, a product of the Global Grid Forum’s OGSA working group (OGSA-WG), defines OGSA version 1.5, and is one component of a set of documents that, over time, will fully define OGSA, both informatively and normatively. The full document set, the

status, and the planned evolution of OGSA are described in *Defining the Grid: A Roadmap for OGSA™ Standards* [OGSA Roadmap].

2 Requirements

This definition of OGSA 1.5 is driven by a set of functional and non-functional requirements, which themselves are informed by the use cases listed in Table 1 and described in companion documents [OGSA Use Cases][OGSA Use Cases Tier 2]. These use cases cover infrastructure and application scenarios for both commercial and scientific areas. They do not constitute a formal requirements analysis, but have provided useful input to the architecture definition process.

Table 1: The OGSA use cases

Use case	Summary
Commercial Data Center (CDC)	Data centers will have to manage thousands of IT resources, including servers, storage, and networks, while reducing management costs and increasing resource utilization.
Severe Storm Modeling	Enable accurate prediction of the exact location of severe storms based on a combination of real-time wide area weather instrumentation and large-scale simulation coupled with data modeling.
Online Media and Entertainment	Delivering an entertainment experience, either for consumption or interaction.
National Fusion Collaboratory (NFC)	Defines a virtual organization devoted to fusion research and addresses the needs of software developed and executed by this community based on the application service provider (ASP) model.
Service-Based Distributed Query Processing	A service-based distributed query processor supporting the evaluation of queries expressed in a declarative language over one or more existing services.
Grid Workflow	Workflow is a convenient way of constructing new services by composing existing services. A new service can be created and used by registering a workflow definition to a workflow engine.
Grid Resource Reseller	Inserting a supply chain between the Grid resource owners and end users will allow the resource owners to concentrate on their core competences, while end users can purchase resources bundled into attractive packages by the reseller.
Inter Grid	Extends the CDC use case by emphasizing the plethora of applications that are not Grid-enabled and are difficult to change: e.g. mixed Grid and non-Grid data centers, and Grid across multiple companies. Also brings into view generic concepts of utility computing.
Interactive Grids	Compared to the online media use case, this use case emphasizes a high granularity of distributed execution.
Grid Lite	Extends the use of Grids to small devices—PDAs, cell phones, firewalls, etc.—and identifies a set of essential services that enable

	the device to be part of a Grid environment.
Virtual Organization (VO) Grid Portal	A VO gives its members access to various computational, instrument-based data and other types of resources. A Grid portal provides an end-user view of the collected resources available to the members of the VO.
Persistent Archive	Preservation environments handle technology evolution by providing appropriate abstraction layers to manage mappings between old and new protocols, software and hardware systems, while maintaining authentic records.
Mutual Authorization	Refines the CDC and NFC use cases by introducing the additional requirement of the job submitter authorizing the resource on which the job will eventually execute.
Resource Usage Service	Facilitates the mediation of resource usage metrics produced by applications, middleware, operating systems, and physical (compute and network) resources in a distributed, heterogeneous environment.
IT Infrastructure and Management*	Job execution, cycle sharing and provisioning scenarios.
Application Use Cases*	Peer-to-Peer PC Grid computing, file sharing and content delivery scenarios.
Reality Grid*	Distributed and collaborative exploration of parameter space through computational steering and on-line, high-end visualization.
The Learning GRID*	User-centered, contextualized and experiential-based approaches for ubiquitous learning in the framework of a Virtual Organization.
HLA-based Distributed Simulation*	A distributed collaborative environment for developing and running simulations across administrative domains based on the IEEE high-level architecture (HLA).
GRID based ASP for Business*	An infrastructure for Application Service Provision (ASP) supporting different business models based on Grid technology.
Grid Monitoring Architecture*	Grid monitoring system scalable across wide-area networks and able to encompass a large number of dynamic and heterogeneous resources.

* Use cases appearing in Tier 2 document

Analysis of these use cases and other relevant input led us to identify characteristics of Grid environments and applications, and functional and non-functional requirements that appear both important and broadly relevant. We summarize our findings in the following sections.

2.1 Interoperability and Support for Dynamic and Heterogeneous Environments

Some use cases involve highly constrained or homogeneous environments that may well motivate specialized profiles. However, it is clear that, in general, Grid environments tend to be heterogeneous and distributed, encompassing a variety of hosting environments (e.g., J2EE, .NET), operating systems (e.g., Unix, Linux, Windows, embedded systems), devices (e.g., computers, instruments, sensors, storage systems, databases, networks), and services, provided by

various vendors. In addition, Grid environments are frequently intended to be long-lived and dynamic, and may therefore evolve in ways not initially anticipated.

OGSA must enable interoperability between such diverse, heterogeneous, and distributed resources and services as well as reduce the complexity of administering heterogeneous systems. Moreover, many functions required in distributed environments, such as security and resource management, may already be implemented by stable and reliable legacy systems. It will rarely be feasible to replace such legacy systems; instead, we must be able to integrate them into the Grid.

The need to support heterogeneous systems leads to requirements that include the following:

- *Resource virtualization.* Essential to reduce the complexity of managing heterogeneous systems and to handle diverse resources in a unified way.
- *Common management capabilities.* Simplifying administration of a heterogeneous system requires mechanisms for uniform and consistent management of resources. A minimum set of common manageability capabilities is required.
- *Resource discovery and query.* Mechanisms are required for discovering resources with desired attributes and for retrieving their properties. Discovery and query should handle a highly dynamic and heterogeneous system.
- *Standard protocols and schemas.* Important for interoperability. In addition, standard protocols are also particularly important as their use can simplify the transition to using Grids.

2.2 Resource Sharing Across Organizations

The Grid is not a monolithic system but will often be composed of resources owned and controlled by various organizations. One major purpose of OGSA is to support resource sharing and utilization across administrative domains, whether different work units within an enterprise or even different institutions. Mechanisms are needed to provide a context that can be used to associate users, requests, resources, policies, and agreements across organizational boundaries. Sharing resources across organizations also implies various security requirements, a topic we address in §2.7.

Resource sharing requirements include:

- *Global name space.* To ease data and resource access. OGSA entities should be able to access other OGSA entities transparently, subject to security constraints, without regard to location or replication.
- *Metadata services.* Important for finding, invoking, and tracking entities. It should be possible to allow for access to and propagation, aggregation, and management of entity metadata across administrative domains.
- *Site autonomy.* Mechanisms are required for accessing resources across sites while respecting local control and policy (see Delegation in §2.7).
- *Resource usage data.* Mechanisms and standard schemas for *collecting and exchanging resource usage (i.e., consumption) data across organizations*—for the purpose of accounting, billing, etc.

2.3 Optimization

Optimization refers to techniques used to allocate resources effectively to meet consumer and supplier requirements. Optimization applies to both suppliers (supply-side) and consumers (consume-side) of resources and services. One common case of *supply-side optimization* is

resource optimization. For example, resource allocation often provides for worst case scenarios (e.g., highest expected load, backup against failures) and leads to resource underutilization. Resource utilization can be improved by flexible resource allocation policies such as advance reservation of resources with a bounded time period and the pooling of backup resources.

Demand-side optimization must be able to manage various types of workload, including the demands of aggregate workloads, which can be difficult to predict. An important requirement in this area is the ability to dynamically adjust workload priorities in order to meet the overall service level objectives. Mechanisms for tracking resource utilization, including metering, monitoring and logging; for changing resource allocation; and for provisioning resources on-demand are the required foundation of demand-side optimization.

2.4 Quality of Service (QoS) Assurance

Services such as job execution and data services must provide the agreed-upon QoS. Key QoS dimensions include, but are not limited to, availability, security, and performance. QoS expectations should be expressed using measurable terms.

QoS assurance requirements include:

- *Service level agreement.* QoS should be represented by agreements which are established through negotiation between service requester and provider prior to service execution. Standard mechanisms should be provided to create and manage agreements.
- *Service level attainment.* If the agreement requires attainment of Service Level, the resources used by the service should be adjusted so that the required QoS is maintained. Therefore, mechanisms for monitoring service quality, estimating resource utilization, and planning for and adjusting resource usage are required.
- *Migration.* It should be possible to migrate executing services or applications to adjust workloads for performance or availability.

2.5 Job Execution

OGSA must provide manageability for execution of user-defined work (jobs) throughout their lifetime. Functions such as scheduling, provisioning, job control and exception handling of jobs must be supported, even when the job is distributed over a great number of heterogeneous resources.

Job execution requirements include:

- *Support for various job types.* Execution of various types of jobs must be supported including simple jobs and complex jobs such as workflow and composite services.
- *Job management.* It is essential to be able to manage jobs during their entire lifetimes. Jobs must support manageability interfaces and these interfaces must work with various types of groupings of jobs (e.g., workflows, job arrays). Mechanisms are also required for controlling the execution of individual job steps as well as orchestration or choreography services.
- *Scheduling.* The ability to schedule and execute jobs based on such information as specified priority and current allocation of resources is required. It is also required to realize mechanisms for scheduling across administrative domains, using multiple schedulers.
- *Resource provisioning.* To automate the complicated process of *resource allocation, deployment, and configuration*. It must be possible to deploy the required applications and data to resources and configure them automatically, if necessary deploying and re-

configuring hosting environments such as OS and middleware to prepare the environment needed for job execution. It must be possible to provision any type of resource not just compute resources, for example, network or data resources.

2.6 Data Services

More and more fields of science, technology and commerce require efficient access to and the movement and management of large quantities of data in a distributed environment. Many also require support for the sharing and integration of distributed data, for example enabling access to information stored in databases that are managed and administered independently, with appropriate security assurances. Archiving of data and data management are essential requirements. OGSA must simplify the creation of such data-oriented applications and make them resilient to changes in the environment.

Data services requirements include:

- *Policy specification & management.* The key to developing a self-managed, scalable, efficient data Grid is the ability to specify policies. Policies permeate the architecture. Examples include specification of who can access data, where data will be required, what transformations are permitted on the data, whether use is exclusive, what performance or availability is required, how much resources can be used, what consistency is mandated between replicas, and similar constraints.
- *Data storage.* Disk and tape systems, amongst others, store data. Common interfaces support the provision of storage and management of quotas, lifetime, and properties such as encryption and persistency.
- *Data access.* Clients require easy and efficient access to various types of data (such as databases, files, streams and integrated/federated data) through a uniform set of interfaces is required, independent of its physical location or platform, by abstracting underlying data resources. Mechanisms are also required to interface with existing security infrastructure or, where these do not exist, to control access rights at different levels of granularity.
- *Data transfer.* High-bandwidth transfer of data is required, independent of the physical attributes of the data sources and sinks, which can exploit relevant features of those sources and sinks if required. Point to point transfer of bytes must be supported, along with more complex transfer schemes that can serve multiple endpoints and preserve the semantics of the data.
- *Data location management.* These services manage where data is physically located OGSA should support multiple methods for making data available to a client at a given location, according to the policy requirements of the client and of the data resource. Methods should include transfer, copying, caching, and replication of data.
- *Data update.* Although some data resources are read only, many if not most provide some users with update privileges. OGSA must provide update facilities which ensure that the specified consistency can be maintained when cached or replicated data is modified.
- *Data persistency.* Data should be preserved according to specified policy and its association with its metadata should be maintained in accordance with that policy. It should be possible to use one of many possible persistency models.
- *Data federation.* OGSA should support data integration for heterogeneous and distributed data. Heterogeneous data includes data organized according to different schemas and data stored using different technologies (e.g., relational vs. flat file). It is also required to

be able to search heterogeneous and distributed data resources in a uniform way and to select an appropriate return format.

For one vision of how data services could form part of a comprehensive Grid solution, from a database-centric viewpoint, see [DAIS].

2.7 Security

Safe administration requires controlling access to services through robust security protocols and according to provided security policy. For example, obtaining application programs and deploying them into a Grid system may require authentication and authorization. Also sharing of resources by users requires some kind of isolation mechanism. In addition, standard, secure mechanisms are required which can be deployed to protect Grid systems while supporting safe resource-sharing across administrative domains.

Security requirements include:

- *Authentication and authorization.* Authentication mechanisms are required so that the identity of individuals and services can be established. Service providers must implement authorization mechanisms to enforce policy over how each service can be used. The Grid system should follow each domain's security policies and also may have to identify users' security policies. Authorization should accommodate various access control models and implementations.
- *Multiple security infrastructures.* Distributed operation implies a need to integrate and interoperate with multiple security infrastructures. OGSA needs to integrate and interoperate with existing security architectures and models.
- *Perimeter security solutions.* Resources may have to be accessed across organizational boundaries. OGSA requires standard and secure mechanisms that can be deployed to protect organizations while also enabling cross-domain interaction without compromising local security mechanisms, such as firewall policy and intrusion-detection policy.
- *Isolation.* Various kinds of isolation must be ensured, such as isolation of users, performance isolation, and isolation between content offerings within the same Grid system.
- *Delegation.* Mechanisms that allow for delegation of access rights from service requestors to service providers are required. The risk of misuse of delegated rights must be minimized, for example by restricting the rights transferred through delegation to the intended job and limiting their lifetimes.
- *Security policy exchange.* Service requestors and providers should be able to exchange dynamically security policy information to establish a negotiated security context between them.
- *Intrusion detection, protection, and secure logging.* Strong monitoring is required for intrusion detection and identification of misuses, malicious or otherwise, including virus or worm attacks. It should also be possible to protect critical areas or functions by migrating attacks away from them.

2.8 Administrative Cost Reduction

The complexity of administering large-scale distributed, heterogeneous systems increases administration costs and the risk of human errors. Support for administration tasks, by automating administrative operations and consistent management of virtualized resources, is needed.

Policy-based management is required to automate Grid system control, so that its operations conform to the goals of the organization that operates and utilizes the Grid system. From the low-level policies that govern how the resources are monitored and managed to high-level policies that govern how business processes such as billing are managed, there may be policies at every level of the system. Policies may include availability, performance, security, scheduling, and brokering.

Application contents management mechanisms can facilitate the deployment, configuration, and maintenance of complex systems, by allowing all application-related information to be specified and managed as a single logical unit. This approach allows administrators to maintain application components in a concise and reliable manner, even without expert knowledge about the applications.

Problem determination mechanisms are needed, so that administrators can recognize and cope quickly with emerging problems.

2.9 Scalability

A large-scale Grid system can create added value such as drastically reducing job turn around (or elapsed) time, allowing for utilizing huge number of resources, thereby enabling new services. However, the large scale of the system may present problems, since it places novel demands on the management infrastructure.

The *management architecture* needs to scale to potentially thousands of resources of a widely varied nature. Management needs to be done in a hierarchical or peer-to-peer (federated/collaborative) fashion.

High-throughput computing mechanisms are required for adjusting and optimizing parallel job execution in order to improve throughput of the entire computational process, as well as optimizing a single computation.

2.10 Availability

High availability is often realized by expensive fault-tolerant hardware or complex cluster systems. Because of the widespread use of IT systems to provide essential public infrastructure services, an increasing number of systems are required to operate at a high level of availability. Since Grid technologies enable transparent access to a wider resource pool, *across* organizations as well as *within* organizations, they can be used as one building block to realize stable, highly-reliable execution environments. But due to the heterogeneity of the Grid, components with longer or more unpredictable mean-time-to-repair (MTTR) characteristics than those generally used in existing high-reliability systems have to be used, presenting difficult problems.

In such a complex environment, policy-based autonomous control (see Policy-based Management in §2.8) and dynamic provisioning (see Provisioning in §2.5) are keys to realizing systems of high flexibility and recoverability.

Disaster recovery mechanisms are needed so that the operation of a Grid system can be recovered quickly and efficiently in case of natural or human-caused disaster, avoiding long-term service disruption. Remote backup and simplifying or automating recovery procedures is required.

Fault management mechanisms can be required so that running jobs are not lost because of resource faults. Mechanisms are required for monitoring, fault detection, and diagnosis of causes or impacts on running jobs. In addition, automation of fault-handling, using techniques such as checkpoint recovery, is desirable.

2.11 Ease of Use and Extensibility

The user should be able to use OGSA to mask the complexity of the environment if so required. As much as possible, tools, acting in concert with run-time facilities, must manage the environment for the user and provide useful abstractions at the desired level. Tempering this ease-of-use objective is the knowledge that there are “power users” with demanding applications that will require, and demand, the capability to make low-level decisions and to interface with low-level system mechanisms. Therefore it should be possible for end-users to choose the level at which they wish to interact with the system.

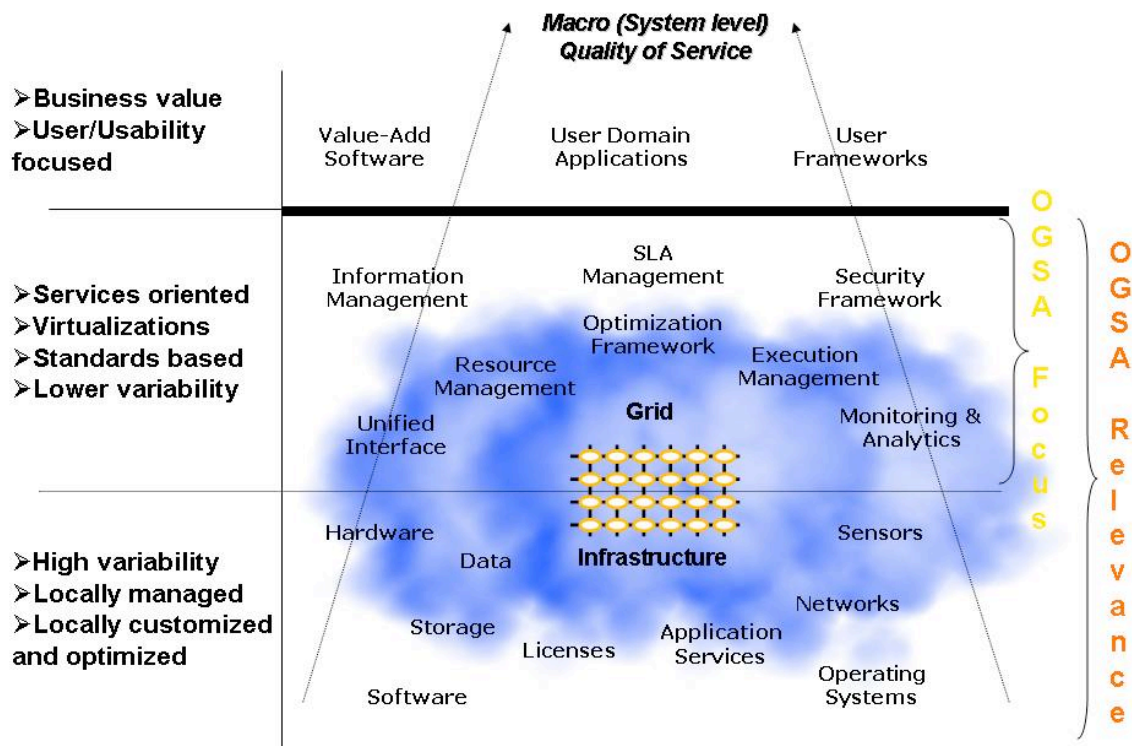
It is not possible to predict all of the many and varied needs that users will have. Therefore, mechanism and policy must be realized via extensible and replaceable components, to permit OGSA to evolve over time and allow users to construct their own mechanisms and policies to meet specific needs. Further, the core system components themselves must be extensible and replaceable. Such extensibility will allow third party (or site-local) implementations which provide value-added services to be developed and used. Extensibility and customization must be provided for in a way that does not compromise interoperability.

3 Capabilities

We now turn to the specific capabilities required to meet the requirements introduced above.

3.1 Overview

OGSA is intended to facilitate the seamless use and management of distributed, heterogeneous resources. In this architecture, the terms “distributed,” “heterogeneous” and “resources” are used in their broad sense. For example: “distributed” could refer to a spectrum from geographically-contiguous resources linked to each other by some connection fabric to global, multi-domain, loosely- and intermittently-connected resources. “Resources” refers to any artifact, entity or knowledge required to complete an operation in or on the system. The utility provided by such an infrastructure is realized as a set of *capabilities*. Figure 1 shows the logical, abstract, semi-layered representation of some of these capabilities. Three major logical and abstract tiers are envisioned in this representation.



The first (bottom) tier of the representation in Figure 1 depicts the *base resources*. Base resources are those resources that are supported by some underlying entities or artifacts that may be physical or logical, and that have relevance outside of the OGSA context. Examples of such physical entities include CPUs, memory, and disks, and examples of such logical artifacts include licenses, contents, and OS processes. The virtualizations are tightly-coupled with the entities that they virtualize, and hence the nomenclature used to name the underlying entities or artifacts is also used to name their virtualizations. These resources are usually locally owned and managed, but may be shared remotely. The configuration and customization is also done locally. Since the actual entities or artifacts can change rapidly, and can be from multiple sources, these resources can vary greatly in their characteristics, quality of service, version, availability etc.

Although in this discussion a distinction of base resources is made to tie OGSA concepts to traditional notions of resources, further discussion in this document makes no specific distinction between base resources and other services as resources, and only the generalized notion of resources is used.

The second (middle) tier represents a higher level of virtualization and logical abstraction. The virtualization and abstraction are directed toward defining a wide variety of *capabilities* that are relevant to OGSA Grids. These capabilities can be utilized individually or composed as appropriate to provide the infrastructure required to support higher-level applications or “user” domain processes. This set of capabilities, as defined in OGSA, is relatively invariant and standard. The manner in which these capabilities are realized or implemented and further composed and extended by user-domain applications determines the macro (system level) QoS of the larger infrastructure, as experienced by an end-user. *It should be noted that the capabilities shown in the diagram only represent a sample of the OGSA capabilities, and that a more complete set is discussed in the rest of this chapter.*

It is worth going into more detail on the relationship between the middle and bottom tiers in Figure 1. The service-oriented nature of OGSA implies that virtualized resources that are represented as services are peers to other services in the architecture (for example services in the middle and top tiers). The peer relationship implies that service interaction can be initiated by any service in the architecture. Furthermore the services in the second tier need to use and manage the virtualizations (resources) in the bottom tier to deliver the capabilities that an individual service (or collection/composition of services) is to provide. This close relationship is indicated by the finer (thin line) demarcation of the bottom and middle layers in Figure 1. Hence the entities in the lower layer may be regarded as relevant to, and part of, the OGSA discussion.

At the third (top) tier in the logical representation are the applications and other entities that use the OGSA capabilities to realize user- and domain-oriented functions and processes, such as business processes. These are, for the most part, outside the purview of OGSA, but they drive the definition of the architecture from the use cases that the infrastructure should support (see Table 1 for a list of use cases that are currently known. It should be noted that this list of use cases is not exhaustive and is expected to grow in the future).

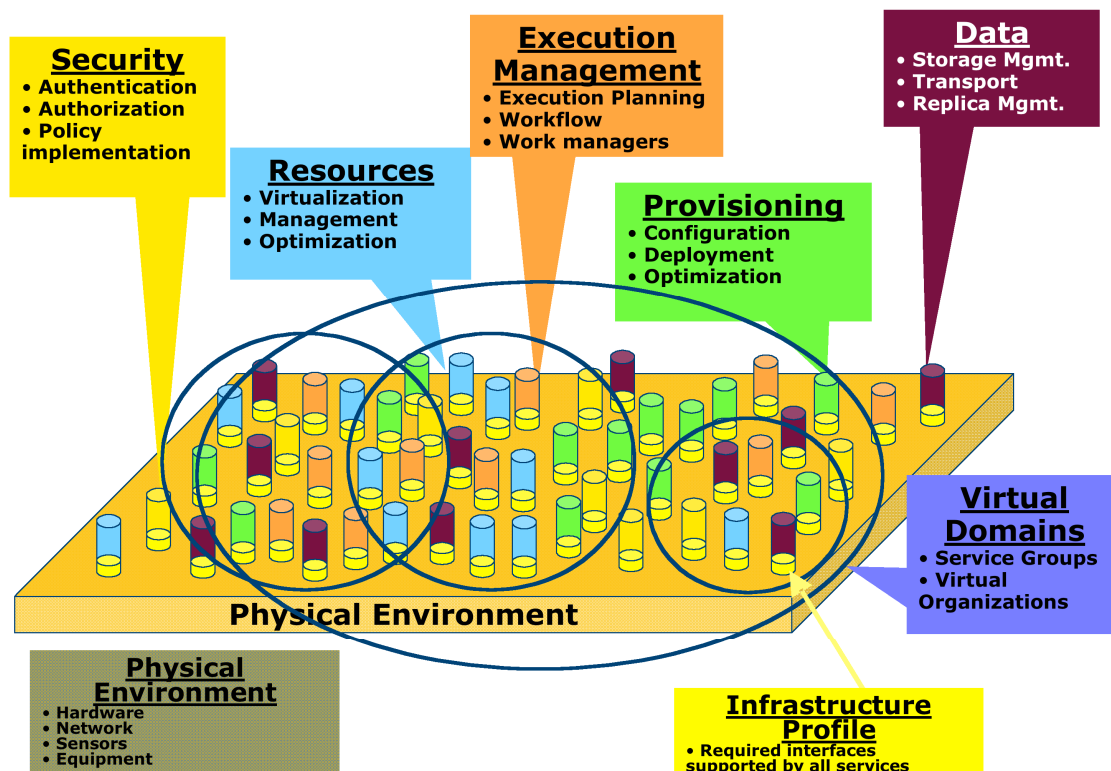
All of these tiers need to interoperate and work synergistically to deliver the required quality of service (QoS). Since this is the QoS of the entire system, including the application tier (or at the very least the services participating in the specific user scenario) that determines the user experience, this is designated as the “Macro Quality of Service.” This is shown by the dotted arrows in Figure 1.

3.2 OGSA Framework

OGSA realizes the logical middle layer in Figure 1 in terms of *services*, the *interfaces* these services expose, the individual and collective *state* of resources belonging to these services, and the *interaction* between these services within a *service-oriented architecture* (SOA). The OGSA services framework is shown in Figure 2 and Figure 3. In the figures cylinders represent individual services. The services are built on Web service standards, with semantics, additions, extensions and modifications that are relevant to Grids.

A few important points to note:

- An important motivation for OGSA is the *composition paradigm* or *building block* approach, where a set of capabilities or functions is built or adapted as required, from a minimalist set of initial capabilities, to meet a need. No prior knowledge of this need is assumed. This provides the adaptability, flexibility and robustness to change that is required in the architecture.
- The entire set of OGSA capabilities does not have to be present in a system. In other words, a system may be built using only a subset of capabilities defined by OGSA. Further, capabilities themselves are *not* monolithic. A system may choose to utilize or provide only a subset of services from any capability. (See also §3.3.10 for more details on interoperability.)
- OGSA represents the services, their interfaces, and the semantics/behavior and interaction of these services. It should be noted that the software architecture driving the implementation of the internals of these services is outside the OGSA working group's scope.
- The architecture is not *layered*, where the implementation of one service is built upon, and can only interact with, the layer upon which it is logically dependent; nor is it *object-oriented*—though many of the concepts may seem to be object-based.



Services are loosely coupled peers that, either singly or as part of an interacting group of services, realize the capabilities of OGSA through *implementation*, *composition*, or *interaction* with other services (see Figure 2 and Figure 3). For example, to realize the “orchestration” capability a group of services might be structured such that one set of services in the group drives the orchestration (i.e., acts as the “orchestrator”), while other services in the group provide the interfaces and mechanisms to be orchestrated (i.e., be the “orchestratees”). A specific service may implement or participate in multiple collections and interactions to realize different capabilities. On the other hand, it is not necessary that *all* services participate to realize a particular capability.

The services may be part of, or participate in, virtual collections called *virtual domains* (see Figure 2) to realize a capability, as in *service groups*, or to share a collective context or manageability framework, as in *virtual organizations*.

OGSA services require and assume a *physical environment* that may include well-known physical components and interconnects such as computing hardware and networks, and perhaps even physical equipment such as telescopes.

It is expected that there may be a core set of non-null interfaces, standards and common knowledge/bootstrap that services must implement to be part of an OGSA Grid. This set of common implementations and manifestations to support OGSA is referred to as the *infrastructure services* or the *Grid fabric*. We describe standards that are expected to be part of the Grid fabric in §3.3.

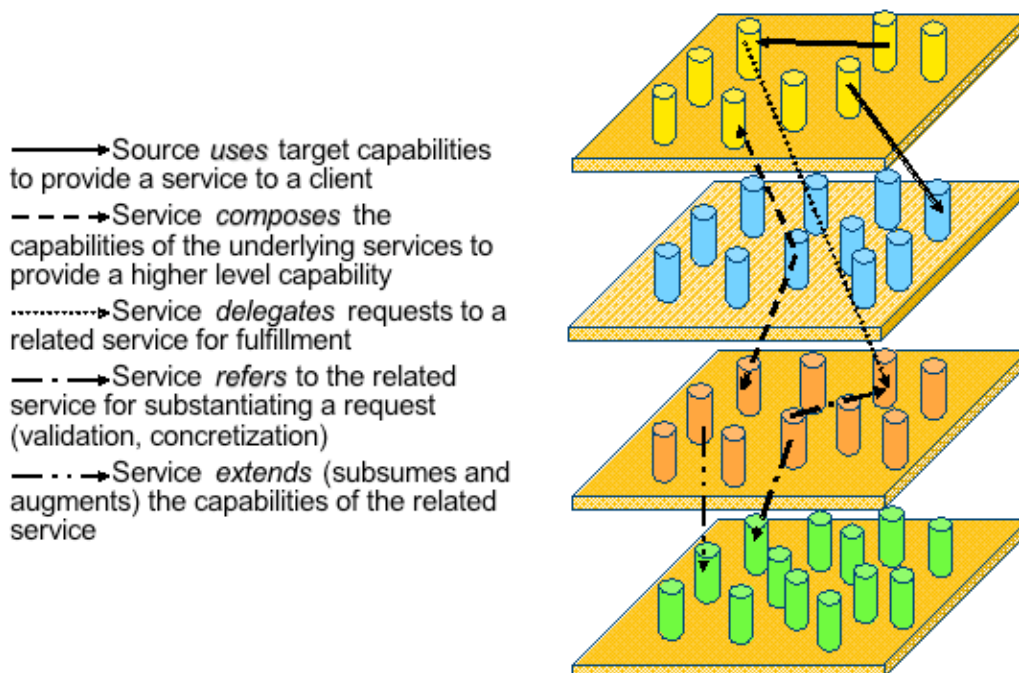


Figure 3: Service Relationships

Figure 3 shows a different view of Figure 2 focusing on possible relationships and interactions between OGSA services. Cylinders represent individual services on a physical infrastructure layer. Each “level” in the figure is a collection representing a single capability—one “level” may realize the Execution Management capability while another may realize the Data management capability. (Note, however, that no hierarchical or stacking relation is intended between the different “levels”.) The relationships shown are between specific services and may span capabilities. Some examples of the types of relationships that services may exhibit are *uses*, *composes*, *delegates*, *refers*, and *extends*. Aspects of service interactions for other purposes such as management or manageability and declarative specification of service profiles can also be modeled with relationships.

3.3 Infrastructure Services

OGSA capabilities share and build on a number of common components. This section describes these infrastructure components.

3.3.1 Objectives

Our goal in defining OGSA is to define a coherent and integrated set of components that collectively address the requirements identified in §2, within the context of a service-oriented architecture. We must necessarily make assumptions about the infrastructure on which we build if we are to make concrete statements about higher-level services. There are many examples of specifications that tried to be too abstract and did not achieve their goals. For example, CORBA¹ 1.1 failed to achieve interoperability because service names depended on implementation. Thus, just as the designs of TCP, DNS, and other higher-level TCP protocols and services are informed by the properties of the IP substrate on which they build, so our design of OGSA is influenced by our choice of underlying mechanisms. Here we list, and to some extent justify, those choices.

3.3.2 Approach

The primary assumption is that work on OGSA both builds on, and is contributing to the development of, the growing collection of technical specifications that form the emerging Web Services Architecture [WS-Architecture]. Indeed, OGSA can be viewed as a particular profile for the application of core WS standards. We made this choice because of a strong belief in the merits of a service-oriented architecture and our belief that the Web Services Architecture is the most effective route to follow to achieve a broadly adopted, industry-standard service-oriented rendering of the functionality required for Grid systems.

3.3.3 Web Services Foundation

This choice of Web services as an infrastructure and framework means that we assume that OGSA systems and applications are structured according to service-oriented architecture principles, and that service interfaces are defined by the Web Services Description Language (WSDL). For now, we assume WSDL 1.1, with a move to WSDL 2.0 planned once that latter specification is finalized and broadly adopted. We also assume XML as the lingua franca for description and representation (although recognizing that other representations may be required in some contexts, for example when performance is critical) and SOAP as the primary message exchange format for OGSA services. In addition, we seek to develop service definitions that are consistent with the interoperability profiles defined through the WS Interoperability (WS-I²) process.

¹ Common Object Request Broker Architecture

² <http://www.ws-i.org/>

While thus working within a Web services framework, it is clear that Web services standards as currently defined do not, and are not designed to, meet all Grid requirements. In some cases, existing specifications may require modification or extension. Thus, OGSA architects have been involved in the definition of WSDL 2.0 and in the review of WS-Security and related specifications, and we identify in this document other areas in which extensions to existing specifications are desirable. In other cases, Grid requirements motivate the introduction of entirely new service definitions.

3.3.4 Naming

Traditional distributed systems usually support a two or three-layer naming scheme. In OGSA, the naming service, *OGSA-naming*, uses a three-level convention. Every named OGSA entity is associated with an (optional) *human-oriented name*, an *abstract name*, and an *address*.

3.3.4.1 Naming Scheme

The *human-oriented name* is usually human-readable and may belong to a name space. Name spaces are usually hierarchic and usually have syntactic restrictions. Hierarchic name spaces permit each part of a name to map to a particular context. For example, in the Unix filename “`node1:/var/log/error`,” the context for `var` is the root directory of a machine named “`node1`,” the context for “`log`” directory is “`var`,” and the context for “`error`” file is “`log`.” OGSA-naming does not require human-oriented names to be unique. Many different naming schemes exist and this document does not attempt to prescribe the set of all supported schemes.

The *abstract name* is a persistent name that does not specify a particular location. It is recommended that *abstract names* are globally unique in space and time. The *AbstractName* element of a WS-Name [WS-Naming] and the Legion Object Identifier [Legion] are examples of abstract names. A mechanism, outside the scope of this document, is required to map *human-oriented names* to *abstract names*.

The *address* specifies the location of an entity. Examples of *addresses* are a WS-Addressing endpoint reference (EPR) [WS-Addressing], a memory address, and an IP address/port pair. A mechanism, outside the scope of this document, is required to bind *abstract names* to *addresses*.

3.3.4.2 Naming Policy

The architectural construct for an address in OGSA is the WS-Addressing Endpoint Reference (EPR). WS-Addressing EPRs support extensibility—additional elements beyond those required by the specification may be added to an EPR. It is therefore possible to define WS-Addressing profiles that support additional functionality. An example of a profile that defines the functionality required for the abstract names, described in §3.3.4.1, is the WS-Naming profile [WS-Naming]. Clients or Web services endpoints that choose to support a profile may be able to take advantage of the extra information to recover from certain types of communication failures or for more efficient communication. Clients or Web services that do not support a profile can still operate as normal.

To ensure consistency in addressing, OGSA *requires* the use of WS-Addressing EPRs; and *encourages* the profiled usage of WS-Addressing EPRs where appropriate and as long as such profiled usage allows services that can only use non-profiled (plain) EPRs to operate. More formally, OGSA services are expected to adhere to the following policy when returning addresses:

- An OGSA-compliant service when returning an address must use a WS-Addressing EPR.
- An OGSA-compliant service may return a profiled usage of a WS-Addressing EPR.

- An OGSA-compliant service should return a WS-Naming profiled WS-Addressing EPR where appropriate.
- An OGSA-compliant service must not require that a specific profiled usage of a WS-Addressing EPR be returned.
- An OGSA-compliant service must not require that a WS-Naming profiled WS-Addressing EPR be returned.

3.3.5 Security Services

One key area in which Grid requirements motivate extensions to existing specifications is security. Security issues arise at various layers of the OGSA infrastructure. We use WS-Security standard protocols to permit OGSA service requests to carry appropriate tokens securely for purposes of authentication, authorization, and message protection. End-to-end message protection is required by some scenarios addressed by the OGSA infrastructure, and thus OGSA must also provide for higher-level protection mechanisms such as XML encryption and digital signatures in addition to, or in place of, point-to-point transport-level security, such as TLS and IPsec. In addition to message-level security, an interoperable and composable infrastructure needs security components to be themselves rendered as services. There are various efforts underway to specify service definitions for these security services. For example, an OGSA authorization service may use the proposed WS-Agreement [WS-Agreement] standard along with OASIS³ standards including SAML and XACML to express security assertions and access control descriptions. When and where appropriate, OGSA will adopt, or define, those security services.

3.3.6 Representing State

A key area in which Grid requirements have motivated new specifications—specifications that have relevance beyond Grid scenarios—is in the area of state representation and manipulation: the ability to model, access and manage state, and related activities. Such capabilities have a fundamental role to play not only in constructing Grid systems, but also in general systems management and other areas.

Two sets of specifications that address this topic are in the process of standardization: the WS-Resource Framework⁴ and WS Management⁵ family of specifications. As each of these seems likely to be adopted by different segments of the Grid community, we have taken the approach of allowing for the definition of multiple distinct infrastructure-level specification stacks, so that any appropriate set of specifications can be used as the foundation for building an OGSA Grid. While initially this approach may lead to the development of Grids that are built on infrastructure stacks that are not interoperable, we expect that over time either the responsible standards bodies will converge the specifications to provide a single industry standard or we will develop a mechanism that will allow them to interoperate despite the differences.

3.3.7 Notification

In a dynamic Grid environment it is critical that components can request and receive timely notification of changes in one another's states.

³ <http://www.oasis-open.org/home/index.php>

⁴ WS-Resource Framework consists of WS-ResourceProperties, WS-ResourceLifetime, WS-ServiceGroup and WS-BaseFaults. See http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=wsrf.

⁵ WS-Management consists of WS-Transfer, WS-Eventing, WS-Enumeration, WS-ManagementCatalog.

We assume notification or “eventing” capabilities such as those defined within the WS-Notification specifications⁶ or WS-Eventing specifications. These specifications define notification mechanisms that support subscription to, and subsequent notification of changes to, state components.

3.3.8 Transactions

This architecture does not define a transaction mechanism per se. We expect this functionality to be provided by other developments in the Web Services community. Currently there are two families of specifications under development. On the one hand, there are the WS-TX specifications⁷, and on the other is the WS Composite Application Framework family⁸.

In both cases, the transaction specification depends upon carrying a transaction context along with every message exchange. It is left up to the implementation to decide whether to honor, or ignore, the transaction context. There is no need to explicitly change the signature of an interface to accommodate transactions. Thus, for the purposes of the OGSA architecture, we assert that transactions layer on top of the basic architecture in a transparent fashion. We strongly suggest that the descriptive information for the implementation include information on how it handles this transaction context.

Finally, it must be observed that it is up to the clients of a particular service to decide if the decision of a particular implementation to honor transactions, or not, is sufficient for the needs of that client.

3.3.9 Orchestration

Many OGSA services are expected to be constructed in part, or entirely, by invoking other services—the EMS Job Manager (see 3.4.5.2) is one such example. There are a variety of mechanisms that can be used for this purpose.

- *Choreography*. Describe required patterns of interaction among services, and templates for sequences (or more structures) of interactions.
- *Orchestration*. Describe the ways in which business processes are constructed from Web services and other business processes, and how these processes interact.
- *Workflow*. Define a pattern of business process interaction, not necessarily corresponding to a fixed set of business processes. All such interactions may be between services residing within a single data center or across a range of different platforms and implementations anywhere.

OGSA, however, will not define new mechanisms in this area. It is expected that developments in the Web services community will provide the necessary functionality⁹. The main role of OGSA is therefore to determine where existing work will not meet Grid architecture needs, rather than to create competing standards.

⁶ WS-Notification consists of WS-Base Notification, WS-Topics, and WS-Brokered Notification. See http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=wsn.

⁷ WS-TX consists of WS-Coordination, WS AtomicTransaction and WS BusinessActivity. See http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=ws-tx.

⁸ WS Composite Application Framework consists of WS-Context, WS-Coordination Framework and WS Transaction Management. See http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=ws-caf.

⁹ A good summary of many relevant standards is at <http://xml.coverpages.org/bpm.html>.

3.3.10 Ensuring Interoperability: OGSA Profiles

The specifications mentioned above are just a few of a rapidly-growing number of Web service specifications, many of which will be used in specifying and developing OGSA components. It is common for specification authors to build in flexibility, for example by making implementation of certain interfaces optional, and while this is helpful for developers in some scenarios it reduces the chances that components based on different implementations will be fully interoperable. Since interoperability is a basic premise of standards-based Grid components, this would present a problem if not addressed.

The Web services interoperability (WS-I)¹⁰ organization addresses the problem of inconsistent implementations by providing a set of normative *interoperability profiles*—guidelines for ensuring consistent and interoperable use of selected specifications. In defining OGSA we are not only developing service definitions that are consistent with the WS-I profiles [WS-I BP 1.1][WS-I BSP 1.0], but we are also building on the profile concept by providing a set of *OGSA Profiles*, each of which defines the use of a set of underlying specifications or other profiles for a specific purpose. By developing a comprehensive and consistent set of Profiles that together address all of the required Grid capabilities we will eventually produce a normative definition of OGSA.

The first Profile is the OGSA WSRF Basic Profile [OGSA WSRF BP], which is based on the WSRF and WS-Notification specification families, and addresses fundamental requirements. We are also defining OGSA Basic Security Profiles [OGSA BSP—Core][OGSA BSP—Secure Channel], and we expect that other Profiles will be developed over time, by various groups, to address capabilities such as job execution and data access. Equivalent profiles based on the WS-Management specifications may be developed as a parallel activity.

The OGSA Profile Definition [OGSA Profile Definition] provides guidelines to be used when developing Profiles.

3.4 Execution Management Services

Execution Management Services (OGSA-EMS) are concerned with the problems of instantiating and managing, to completion, units of work. Examples of units of work may include either OGSA applications or legacy (non-OGSA) applications (a database server, a servlet running in a Java application server container, etc).

3.4.1 Objectives

The following example illustrates some of the issues to be addressed by EMS. An application needs a cache service. Should it use an existing service or create a new one? If it creates a new service, where should it be placed? How will it be configured? How will adequate resources (e.g., memory, disk, CPU) be provided for the cache service? What sort of service agreements can the cache service make? What sort of agreements does it require?

Similarly, suppose a user wants to run a legacy program. In application areas and industries as diverse as bioinformatics, electronic design automation, weather forecasting, aerospace, financial services, and many others there is a class of applications that take some input data (often in files) and parameters and generate outputs. Many of these application instances are run in what are often called “embarrassingly-parallel parameter space studies.” A good example of this in bioinformatics is BLAST—an application that compares a DNA or protein sequence against a target database and generates a list of similarity scores. An example from aerospace is Overflow—a computation fluid dynamics (CFD) code that is used in aircraft simulations. Issues to be addressed when executing such legacy applications include: where will the program run;

¹⁰ <http://www.ws-i.org>

how are the data files and executables staged to the execution location; what happens if execution fails; will execution be restarted, and if so how; and so on. Hereafter we will use BLAST as a canonical example.

More formally, EMS addresses problems with executing units of work, including their placement, provisioning, and lifetime management. These problems include, but are not limited to:

- *Finding execution candidate locations.* What are the locations at which a unit of work can execute because they satisfy *resource* restrictions such as memory, CPU and binary type, available libraries, and available licenses? Given the above, what *policy* restrictions are in place that may further limit the candidate set of execution locations?
- *Selecting execution location.* Once it is known where a unit of work *can* execute, the question is where *should* it execute? Answering this question may involve different selection algorithms that optimize different objective functions or attempt to enforce different policies or service level agreements.
- *Preparing for execution.* Just because a unit of work can execute somewhere does not necessarily mean it can execute there without some setup. Setup could include deployment and configuration of binaries and libraries, staging data, or other operations to prepare the local execution environment.
- *Initiating the execution.* Once everything is ready, actually starting the execution and carrying out other related actions such as registering it in the appropriate places.
- *Managing the execution.* Once the execution is started it must be managed and monitored to completion. What if it fails? Or fails to meet its agreements? Should it be restarted in another location? What about state? Should the state be checkpointed periodically to ensure restartability? Is the execution part of some sort of fault-detection and recovery scheme?

These are the major issues to be addressed by EMS, which clearly covers a wide range of tasks, and will involve interactions with many other OGSA services (e.g., provisioning, logging, registries, and security) that are expected to be defined by other OGSA capabilities. Refer to §3.4.7 for more details.

EMS is important because we cannot just assume a static environment and rely on all the information needed being present in a registry such as UDDI. We expect Grids to be used in a large number of settings where the set of available resources, and the load presented to those resources, are highly variable and require high levels of dependability. For example, in any dynamically provisioned computing environment, the set of resources in use by an application may vary over time, and satisfying the application requirements and service level agreements may require temporarily acquiring the use of remote resources. Similarly, to respond to unexpected failures and meet service-level guarantees may require finding available resources and restarting executions on those resources. The common theme is the requirement to monitor application needs and to respond dynamically to those needs until completion.

3.4.2 Approach

The solution consists of a set of services that decompose the EMS problem into multiple, replaceable components. Different domains may use different subsets of these services to realize their objectives.

EMS services enable applications to have *coordinated* access to underlying resources—CPU, disk, data, memory, and services—regardless of their physical locations or access mechanisms. EMS services are the key to making resources easily accessible to end-users, by automatically matching the requirements of a Grid application with the available resources.

EMS consists of a number of services working together. Below we describe these services. Before we proceed, though, a few caveats and comments are in order.

First, not all services will be used all of the time. Some Grid implementations will not need some of the services—or may encapsulate some services within other services and not make them directly available. In general, we have tried to decompose this capability into services around those functions that we have seen again and again in different Grid implementations—in other words functions that are commonly used, and for which one or more different implementations may be desirable.

Second, this is the first pass at the definitions. It is not our objective in this document to completely define the services. Rather it is our intention to identify key components and their higher-level interactions.

Third, we want to emphasize that these definitions and services will be applicable to general Web service execution—not just to the execution of legacy “jobs.”

3.4.3 EMS Services

There are three broad classes of EMS services:

- *Resources* that model processing, storage, executables, resource management, and provisioning;
- *Job management*; and
- *Resource selection services* that collectively decide where to execute a unit of work.

We also assume the availability of data management services (§3.4), security services (§3.6), and logging services (§3.8.3.3). Interactions with these services will be developed in later versions of this document.

3.4.4 Resources

3.4.4.1 Service Container

A service container, hereafter just a container, “contains” running entities, whether they are “jobs” (described later) or running Web services. A container may encapsulate, for example, a queuing service, a Unix host, a J2EE hosting environment, or a collection of containers (a façade or a VO of job containers). Containers have resources properties that describe static information such as what kind of executables they can take—OS version, libraries installed, policies, and security environment—as well as dynamic information such as current load and QoS information.

A basic container interface may expose only a small set of operations [OGSA BES]. Extended interfaces that provide additional services beyond the basic service container are expected. In particular, a container is expected to implement a manageability interface—for example, assuming the WSRF Basic Profile [OGSA WSRF BP], this interface could be that of a WSDM managed resource.

Finally, we expect containers to use reservation services, logging services, information services, job management services, and provisioning services.

3.4.5 Job Management

3.4.5.1 Job

The OGSA-EMS definition of a *job* incorporates and extends the notion of a traditional job. The *job* encapsulates all there is to know about a particular unit of work (i.e., an instance of a running application (such as BLAST) or a service). A job is the smallest unit that is managed. It represents the manageability aspect of a unit of work: it is not the same as the actual running

application, or the execution aspect of the unit of work. (Note, however, that in this document the term “job” may also be used informally in place of “unit of work,” for example, in statements such as “submitting a job,” or “executing a job.”)

A job implements a manageability interface—for example, assuming the WSRF Basic Profile, this interface could be that of a WSDM managed resource. A job is named by an EPR. It is created at the instant that it is requested, even though at that point no resources may have been committed. The job keeps track of execution state (e.g., started, suspended, restarted, terminated, completed), resource commitments and agreements, job requirements, and so on. Many of these are stored in a *job document*.

A *job document* describes the *state* of the job—e.g., the submission description (JSDL [JSDL]), the agreements that have been acquired, its job status, metadata about the user (credentials etc.), and how many times the job has been started. We do *not* include in “state” application-specific details such as the internal memory of the executing application program.

The job document may be exposed as a resource property of the job. The logical view is of one large document that consists of one or more—possibly many—subdocuments. These subdocuments can be retrieved independently. The organization of the subdocuments will be subject to further specification.

3.4.5.2 *Job Manager*

The Job Manager (JM) is a higher-level service that encapsulates all of the aspects of executing a job, or a set of jobs, from start to finish. A set of jobs may be structured (e.g., a workflow or dependence graph) or unstructured (e.g., an array of non-interacting jobs). The JM may be a portal that interacts with users and manages jobs on their behalf.

The JM will likely interact with an Execution Planning Services (see §3.4.6.1), the deployment and configuration system, containers, and monitoring services. Further, it may deal with failures and restarts, it may schedule jobs to resources, and it may collect agreements and reservations.

The JM implements a manageability interface. For example assuming the WSRF Basic Profile this interface could be that of a WSDM collection—a collection of manageable entities. A WSDM collection can expose as its methods some of the methods exposed by the members of its collection.

The JM is responsible for orchestrating the services used to start a job or set of jobs, by, for example, negotiating agreements, interacting with containers, and configuring monitoring and logging services. It may also aggregate job resource properties from the set of jobs it manages.

Examples of JMs are:

- A “queue” that accepts “jobs,” prioritizes them, and distributes them to different resources for computation. (Similar to JobQueue [JobQueue] or Condor [Condor].) The JM would track jobs, may prioritize jobs, and may have QoS facilities, a maximum number of outstanding jobs, and a set of service containers in which it places jobs.
- A portal that interacts with end-users to collect job data and requirements, schedule those jobs, and return the results.
- A workflow manager that receives a set of job descriptions, QoS requirements, their dependence relationships, and initial data sets (think of it as a data flow graph with an initial marking), and schedules and manages the workflow to completion—perhaps even through a number of failures. (In this case a node could be another workflow job manager). (This is similar in concept to parts of DAGman¹¹.)

¹¹ <http://www.cs.wisc.edu/condor/dagman/>

- An array job manager that takes a set of jobs that have slightly different parameters but are otherwise identical, and manages them through completion—for example, Nimrod [Nimrod].

3.4.6 Selection Services

3.4.6.1 Execution Planning Services (EPS)

An Execution Planning Service (EPS) is a service that builds “schedules,” where a schedule is a mapping (relation) between services and resources, possibly with time constraints. A schedule can be extended with a list of alternative “schedule deltas” that basically say “if this part of the schedule fails, try this one instead.”

An EPS will typically attempt to optimize some objective function such as execution time, cost, reliability, etc. An EPS will not enact the schedule; it will simply generate it. The enactment of a schedule is typically done by the JM. An EPS will likely use information services and Candidate Set Generators (CSG, see below). For example, first call a CSG to get a set of resources, then get more current information on those resources from an information service, then execute the optimization function to build the schedule.

3.4.6.2 Candidate Set Generator (CSG)

The basic idea is quite simple: determine the set of resources on which a unit of work can execute—“where is it *possible* to execute?”, rather than “where *will* it execute?” This may involve issues such as what binaries are available, special application requirements (e.g., 4GB memory and 40GB temporary disk space, *xyz* library installed), and security and trust issues (“I won’t let my job run on a resource unless it is certified Grade A+ by the Pure Computing Association,” or “they won’t let me run there until my binary is certified safe,” or “will they accept my credit card?”).

A Candidate Set Generator (CSG) generates a set of EPRs of containers in which it is possible to run a job described by a JSDL document. The set of container resources to search over may either be a default for the particular service or be passed in as a parameter.

We expect CSGs to be primarily called by EPSs, or by other services such as JMs that are performing EPS-like functions. We expect CSGs to use information services, to access jobs to acquire appropriate pieces of the job document and to interact with provisioning and container services to determine if it is possible to configure a container for a particular execution.

3.4.6.3 Reservation services

Reservation services manage reservations of resources, interact with accounting services (there may be a charge for making a reservation), revoke reservations, etc. This may not be a separate service, rather an interface to get and manage reservations from containers and other resources. The reservation itself is likely to be an agreement document that is signed.

A reservation service presents a common interface to all varieties of reservable resources on the Grid. Reservable resources could include (but are not limited to) computing resources such as CPUs and memory, graphics pipes for visualization, storage space, network bandwidth, special-purpose instruments (e.g., radio telescope), etc.

A reservation could also be an aggregation of a group of lower-level reservations, as might be negotiated and “resold” by a resource broker.

Reservation services will generally be used by many different services: a JM might create reservations for the groups of jobs which are being managed, or an EPS might use reservations in order to guarantee the execution plan for a particular job. It could also be the case that the creation of reservations will be associated with the provisioning step for a job.

3.4.7 Interactions with Other OGSA Services

This section details the interactions between the EMS and other parts of OGSA.

3.4.7.1 *Deployment & Configuration Service*

Often before a service or data container can be used by a unit of work, it must be configured or provisioned with additional resources, see §3.8.4.6. For example, before running BLAST on a host, a user must ensure that the BLAST executable and its configuration files are accessible to the host. A more in-depth example is the configuration of a complex application and installation of appropriate databases, or installing Linux on a host as a first step to using the host as a compute resource.

3.4.7.2 *Naming*

OGSA-EMS uses OGSA-naming, see §3.3.4. For example, in a sophisticated job queuing system which has checkpoint and restart feature for availability or load balancing purpose, an *address* of the job may specify the location of a job on a particular machine. The *abstract name* will identify the job in a location-independent and universal way. For example, the *abstract name* should be the same before and after a job migration. The *human-oriented name* may be a user-friendly short job name that can be disambiguated by referring to the context in which it is used.

3.4.7.3 *Information Service*

In brief, information services are databases of attribute metadata about resources. Within EMS, information services are used by many of the different services: for example, containers need to publish information about their attributes so that CSG services can evaluate the suitability of a container for a job; an EPS might read policy information for a VO from an information service. How the information service gets its information is unspecified, although we expect “freshness” to be an attribute on data. In this sense, OGSA information services (see §3.9) are similar to MDS services in Globus [Globus MDS] and collections in Legion [Legion].

3.4.7.4 *Monitoring*

Simply starting something up is often insufficient. Applications (which may include many different services or components) often need to be continuously monitored, for both fault-tolerance reasons and QoS reasons. For example, the conditions on a given host that originally caused the scheduler to select it may have changed, possibly indicating the need for rescheduling. (See also §3.9.4.4.)

3.4.7.5 *Fault-Detection and Recovery Services*

Fault-detection and recovery services may or may not be a part of monitoring, and may include support for managing simple schemes for stateless functions that allow trading off performance and resource usage; slightly more complex schemes that manage checkpointing and recovery of single-threaded (process) jobs; and still more complex schemes that manage applications with distributed state, such as MPI jobs. (See also §3.8.4.7.)

3.4.7.6 *Logging*

Logging (see §3.9.4.3) is important in the EMS context since a number of other services may be expected to be built on top of it. For example, assuming logging as the basis of the whole chain.

- *Metering* is using the log to keep track of resource usage.
- *Auditing* is using the log to review activity, and may be a component of security functions such as non-repudiation.

- *Billing* is yet another service, not defined by OGSA, which may use auditing, metering logs or other data to generate bills, or chargeback.
- *Provenance* information may be stored using the logging service. Uses for provenance information include regulatory compliance and judging the quality of scientific data, among others.

Services such as billing and accounting may not be defined by OGSA, but they are critical for its success. They can, for example, enable schedulers to interact with resources to establish prices, as well as for resources to interact with accounting and billing services. Like a credit card, some resources may need to see if the user has enough credit to pay. The scheduler may also need to interact with the accounting services.

3.4.8 Example Scenarios

The best way to understand these services is to see how they are used to realize concrete scenarios. We have selected three scenarios to demonstrate: a system patch tool, deploying a data caching service, and a legacy application execution.

3.4.8.1 System Patch Tool

Often operating system patches or library updates need to be applied to a large number of hosts. This can be done in several ways. One commonly-used technique is to run a script on each host in a system that copies the appropriate files. These scripts are often initiated using tools such as “rsh” or “ssh,” and are called from shell scripts that iterate over a set of hosts to be patched. Alternatively, hosts may periodically check if they need an update, and if so, run some script to update the OS.

Suppose the OS version number is a piece of metadata maintained by containers and collected by an information service, and that the objective is to patch all operating systems that don’t have all the patches. One simple way to approach this problem is to first query information services for a list of containers whose OS version number is below some threshold. Then, instruct a Job Manager (JM) to run the patching service on each container in the list. In this case the JM does not need to interact with execution planning services (EPS) because it knows where it wants to run the service. Instead, the JM interacts directly with each container—and possibly a deployment and configuration service—to execute the patching service on the container. (The deployment and configuration service may be needed to install the patch service first.)

3.4.8.2 Data Cache Service

Imagine a data cache service that caches data (files, executables, database views, etc.) on behalf of a number of clients, and maintains some notion of coherence with the primary copies. When a client requests a cache service, one can either deliver a handle to an existing cache service or create a new cache service, depending on the location of the client, the location of the existing caches, the load on existing caches, etc.

Once a decision has been made to instantiate a new cache service, an EPS is invoked to determine where to place the data cache. The EPS uses a CSG to determine where it is possible to run the service—constrained by a notion of locality to the client. Once a location has been selected, the service is instantiated on a container.

3.4.8.3 Legacy Application

Our third example illustrates a rather typical scenario. Suppose a user wants to run a legacy BLAST job. Further, suppose the user is interacting with a portal or queue job manager. There are four basic phases in getting the BLAST job started:

1. Job definition phase. What are the input files? What are the service level requirements? E.g., job must complete by noon tomorrow. What account will be billed for the job? Etc.
2. Discover the resources available and select the resources required to execute the job.
3. Enact the schedule and all that may be involved, e.g., provisioning of resources, accounting, etc.
4. Monitor the job through its lifetime(s). Depending on the service level agreements the job may need to be restarted if it fails to complete for any reason.

To realize this case using EMS the JM creates a new legacy job with the appropriate job description (e.g., written in JSDL [JSDL]). The JM then calls an EPS to get a schedule (step 1). The EPS in turn calls a CSG (step 2), which calls information services to determine where the job can be executed based on binary availability and policy settings (step 3). The EPS selects a service container (step 4), after first checking with the service container that the information is accurate. The EPS returns the schedule to the JM (step 5). The JM then interacts (if necessary) with reservation and deployment and configuration services to set up the job execution environment (step 6). This may involve interaction with the data container as well. The service container is invoked to start the job (step 7). Logging services are used for accounting and audit trail. When the job terminates the job manager is notified by the container. If the job terminates abnormally, the whole cycle may repeat again (see Figure 4).

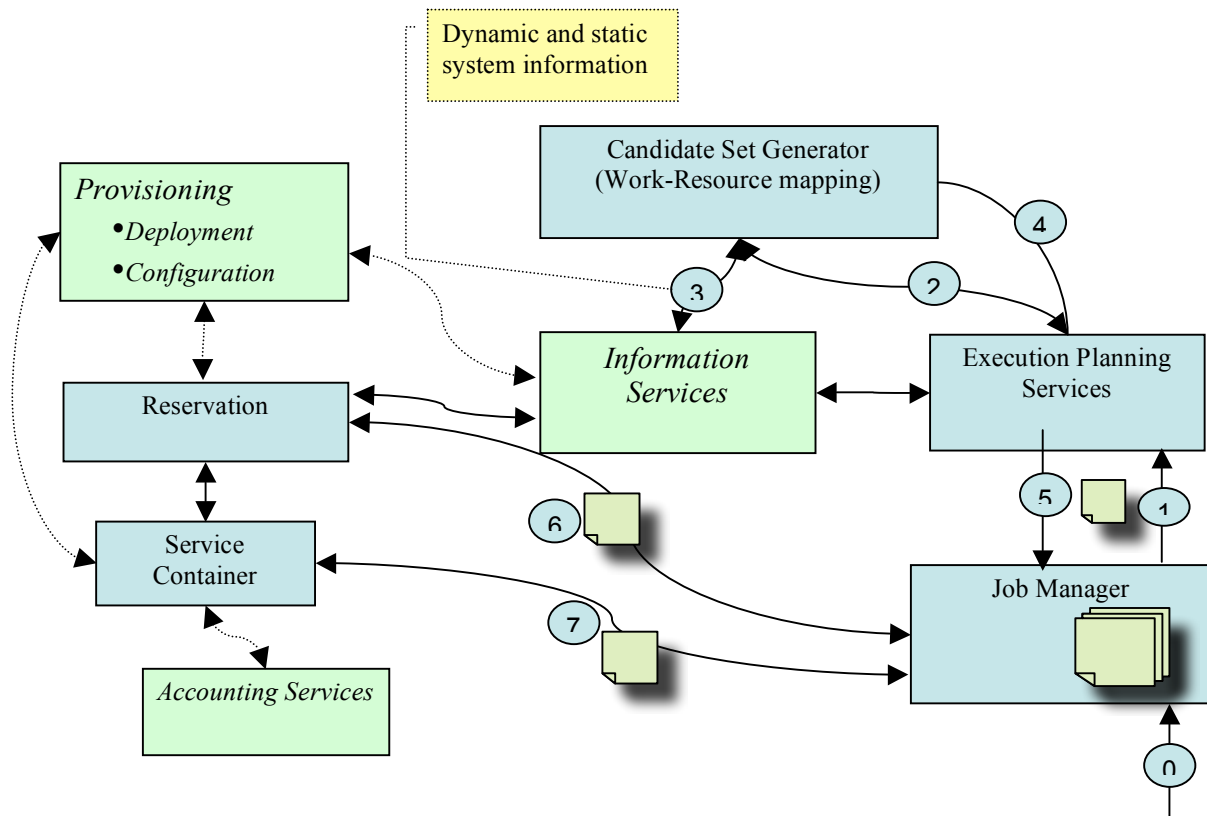


Figure 4: Interactions of EMS services to execute a legacy BLAST job

3.5 Data Services

This section describes those OGSA services concerned with the management of, access to and update of data resources, along with the transfer of data between resources. These are collectively called “data services”.

3.5.1 Objectives

Data services can be used to move data as required; manage replicated copies; run queries and updates; and federate data resources. They also provide the capabilities necessary to manage the metadata that describes this data in particular the provenance of the data itself.

For example, suppose an Execution Management Service (EMS) needs to access data that is stored elsewhere. Data services allow that EMS to access the data remotely or to stage a copy to a local machine. If it accesses the data remotely, a cache service may be used to cache some of the data locally. The data may be available in multiple locations, with the data services ensuring a consistency policy between the replicas. As with all OGSA services, the data services may provide specific policy limitations or service guarantees.

Conversely, suppose that a data federation service wishes to define a schema for data stored in different formats at different locations. It can use OGSA data services to specify how queries against this schema can be mapped to the underlying resources, where joins or data transformations should be executed and where the data should be delivered.

The data services do not rely on or specify the semantics of the data they hold. They operate with generic data.

3.5.2 Approach

3.5.2.1 Types of Data Resource

A data resource is any entity that can act as a source or sink of data. The heterogeneous nature of the Grid means that many different types of data must be supported. These include, but are not limited to:

- *Flat Files*. The simplest form of data is a file with application-specific structure, such as fixed-length records. These files may be accessed using POSIX-like read and write operations. Some file formats support database-like queries: e.g. comma-separated values, which can be queried like relational tables, and XML files, which can be queried using XML Query¹² or other appropriate query languages. Data access services need to support these types of data and also need to be extended to support specialized queries over any new file formats.
- *Streams*. A continuous flow of data is called a stream. The data access services support queries and transformations over streams.
- *DBMS*. Several kinds of database management systems may be part of Grids. These include relational, XML, and object-oriented databases, among others.
- *Catalogues*. A catalogue can structure data and track other data services. A simple example of a catalogue is a directory, which imposes structure by grouping a set of files. Nested directories are equivalent to a hierarchic namespace. Metadata catalogues can provide information about data that is stored in other data services.
- *Derivations*. Some data is the result of asynchronous queries or transformations of other data. These derivations are often managed like finite streams rather than single items.
- *Data Services* themselves can be data resources for other services, as can be sensor devices or programs that generate data.

¹² <http://www.w3.org/TR/xquery/>

3.5.2.2 *Example Scenarios*

Data service capabilities are fundamental to the Grid. This section describes a few examples of how they can be used to support a range of activities.

- *Remote access.* The simplest use of the OGSA data services is to access remote data resources across the Grid. The services hide the communication mechanism from the client; if necessary they can also hide the exact location of the remote data. An optimization of this is to cache some of the data in a local resource.
- *Staging.* When jobs are executed on a remote resource, the data services are often used to copy input data to that resource ready for the job to run, and then to move the result to an appropriate place.
- *Replication.* To improve availability and to reduce access times, the same data can be stored in multiple locations across the Grid. Updates to data in one location may be propagated to the others according to a specified consistency policy.
- *Federation.* The OGSA data services can integrate data from multiple data sources that are created and maintained separately. Clients may access the integrated data via the OGSA data access interface. When a client queries this interface, the query is compiled into sub-queries and operations that extract the appropriate information from the underlying federated resources and return it in the appropriate format. The implementation is also likely to involve data transformation, data combination and caching.
- *Derivation.* The OGSA data services support the automatic generation of one data resource from another.
- *Metadata.* One data service can be used to store descriptions of data held in other data services. As an example, a replicated file system may choose to store descriptions of the files in a central catalogue.

3.5.2.3 *Basic structure*

Figure 5 illustrates the basic entities of the data architecture. Resources are managed by services that may have interfaces for data access, for acting as sources or sinks of data transfer operations, and for describing the resources via properties. Some resources, such as file systems and databases, are storage based, and the architecture also includes interfaces for managing that storage. Data transfer services control transfer of data between resources and provide a transfer control interface. Actual transfer of data between resources uses appropriate protocols and takes as direct a path as possible.

All these interfaces are used by clients, which may be other services, or APIs such as SAGA, POSIX and NFS. Registries, which are described in the Information Services section of this document, play an essential role in the data architecture.

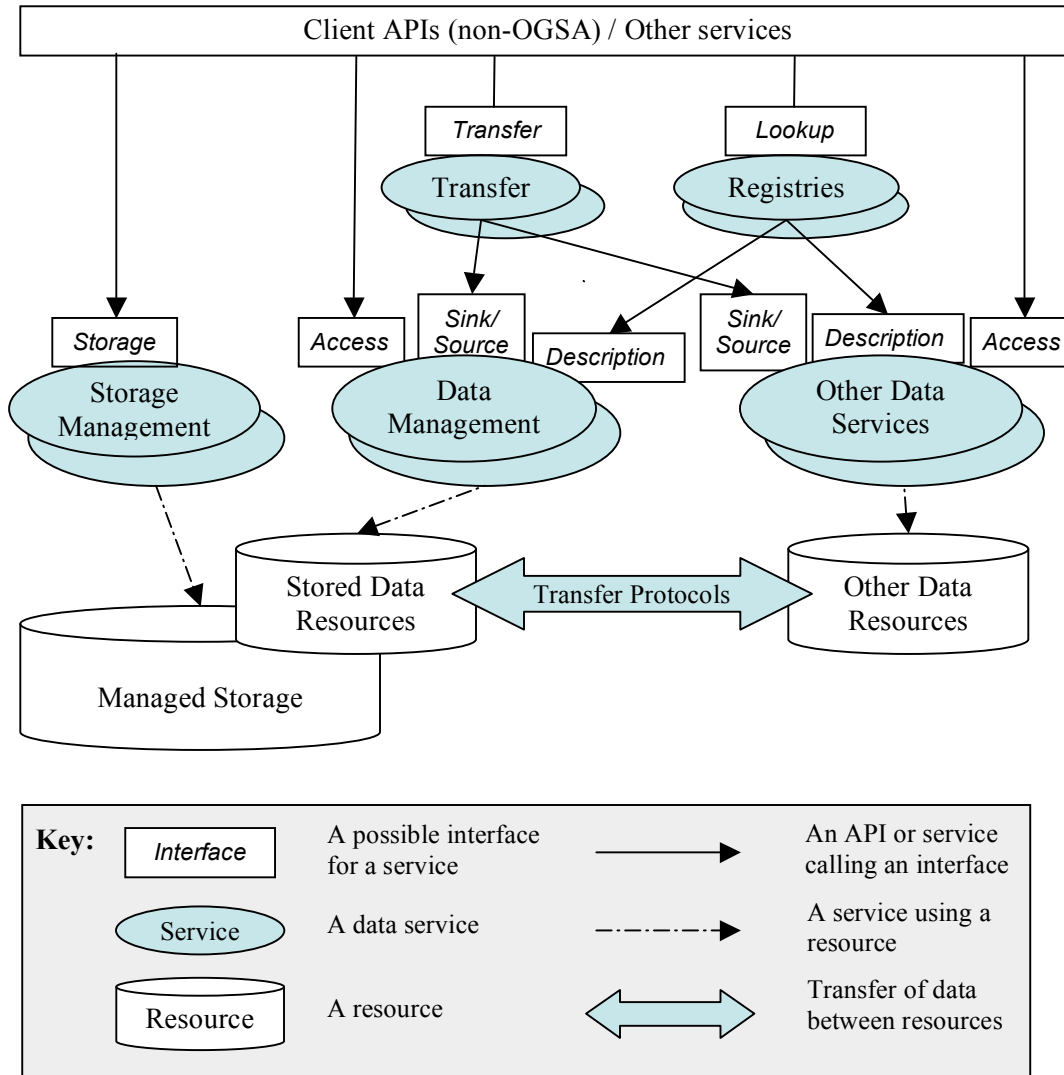


Figure 5: Basic entities in the data architecture

This diagram is necessarily schematic and abstract. In reality, some services may have only an access interface or only a transfer interface. Transfer operations are not restricted to linking stored resources to other resources but may transfer data from any resource with a source interface to any with a sink interface. Some operations may combine functionality from several interfaces: e.g. some access operations might specify that the results should be transferred to another resource; for particular simple transfers such operations might not require the use of a separate data transfer service.

Generic management and provisioning interfaces have been omitted, as have security services and their related access decision points.

3.5.2.4 Composite entities

Figure 6 shows how the basic services can be combined to achieve more complex behaviors. Replication services can use the interfaces provided by the basic service to maintain replicas of data, to improve availability or performance. Cache services can be combined with other services

to produce a new resource. Federation services can combine several data services to form a new service. These composite services can then be treated in the same way as any other service, and provide the same interfaces.




Figure 6: Composite entities in the data architecture

As above, this diagram is necessarily schematic and abstract. For example, it does not show the data transfer operations between the services involved in a data federation or between a cache and its source. As before, generic management and provisioning interfaces have been omitted, as have security services and their related access decision points.

3.5.3 Heterogeneity and Extensibility

This section explains how the architecture fits with non-OGSA data resources and access mechanisms, thus enabling OGSA functionality to augment existing and future data infrastructures.

3.5.3.1 Transparency and Virtualization

A distributed system may contain a variety of data resources. These resources may use different models to structure the data, different physical media to store it, different software systems to manage it, different schema to describe it, and different protocols and interfaces to access it. The data may be stored locally or remotely; may be unique or replicated; may be materialized or derived on demand. OGSA data services provide different levels of *virtualizations* over these data resources. Virtualizations are abstract views that hide these distinctions and allow the data resources to be manipulated without regard to them.

Conversely, although the data services allow the clients to ignore these distinctions, some clients may prefer to exploit them. For example, a client may wish to make use of a particular query language for a given database, or to specify the location of the particular data resource to use. One client may require native access to the data, while another needs to tune the performance parameters of the data resource. To support such clients, the OGSA data interfaces provide extension points so that services can allow clients to bypass the virtualization interfaces and access the resource-specific interfaces directly. These layered interfaces allow clients to choose the combination of power and abstraction that suits them best.

For example, basic byte I/O is provided by an interface that provides read and write operations like those in POSIX and similar systems. The services that provide these operations may perform sophisticated optimizations such as caching, replication or optimized data transfer. The

virtualized interface will hide these details from the client. Those clients that require more detailed control may use the resource-specific interfaces to manipulate the cache services, replica services or data transfer services, with the corresponding loss of transparency.

3.5.3.2 Client APIs

Many users will wish to use the OGSA data services with legacy applications that use existing APIs such as NFS, CIFS, JDBC, ODBC, ADO, POSIX IO, SAX or DOM. Often it will be too expensive or impractical to rewrite the clients to use new interfaces. The OGSA services can easily be wrapped with a veneer that emulates these existing APIs, although OGSA itself will not define such wrappers.

Typically these wrappers will map the operations of the existing APIs to the corresponding messages to send to the OGSA data services. Therefore they may provide access to only a subset of the full OGSA functionality. The extent of the OGSA functionality that they make available will depend on the scope of the API concerned.

3.5.3.3 Extensible Data Type Support and Operation

It is not possible to predict all the data resources with which the OGSA data services will be used. In any case, new data resources will be created in the future. The layered architecture described above allows the addition of new data services that access new types of resources.

Similarly, it is not possible to anticipate all the operations that may be needed on a given resource. The layered interfaces described above support services that provide additional operations beyond those specified in the OGSA document. The virtualization layer gives clients the option of ignoring these extensions, while those clients that require them can bypass the virtualization layer as necessary.

3.5.4 Functional Capabilities

This section describes the functional capabilities provided by the OGSA data services. Different subsets of the services are needed to implement the different capabilities. Implementations are free to provide only some of these capabilities.

3.5.4.1 Data Transfer

The OGSA data services offer data transfer from one location to another. This may be to create a copy of the original data or migrate it completely. QoS can be specified, such as reliable transfer, the maximum bandwidth to use, the time when delivery is required, or delivery guarantees.

The base functionality transfers bytes from one source to a single sink. Thin layers above this supports transfer to multiple sinks, and allow the preservation of semantic information, such as file hierarchies, byte ordering or encodings. This data transfer functionality may be used independently of the underlying protocol, in which case the transfer service takes on the responsibility to meet the QoS requirements, or it may allow the client to configure the underlying protocol if desired.

Most services have a security policy decision point before transferring data. This may be separate from the authorization of the operation itself. For example, an access operation may check authorization to perform the operation at all. It may then check whether the resulting transfer is allowed, which may depend on the contents of the data, and whether any restrictions apply to that transfer (such as required encryption). Security policies may also be attached to the data that is transferred, as well as to data resources themselves.

3.5.4.2 Storage Management

Storage management services control the provision of storage to applications and other OGSA services. They may provide raw storage or space in a file system. They manage quotas, lifetime, and properties such as encryption and persistency.

3.5.4.3 Simple Access

Simple access services provide operations for reading and writing (logically) consecutive bytes from a data resource. These operations may be POSIX-like[ByteIO]. The resource accessed may be local or remote: the virtualization interface hides the details of the data location. This allows the data services to optimize both data location and data access.

3.5.4.4 Queries

The OGSA data access services provide mechanisms for applying queries against data resources. In simple cases these may run an SQL query over a relational database, an XML query over an XML database, or a regular expression over a text file. Other services may implement text mining over a set of documents, or distributed queries over federated databases.

Synchronous queries return the data in the response to a request, while asynchronous queries expose the derived data as new resources. Services may also deliver the results of a query to a specified set of other services.

Query services may optimize a query before sending it to the resource. The resources may further optimize the query and may also handle issues such as concurrent access to the data.

3.5.4.5 Federation

A data federation service will analyze each query that it receives to determine how to best answer the query. This may involve the generation of sub-queries against one or more of the data resources making up the federation, applying transformations to the results of those sub-queries, combining those query results in (arbitrarily) complex ways and then transforming the result into the format requested by the client. It may also determine where intermediate processing is done in order to minimize network traffic. It will allow data sources to be added and removed from the federation, provided they satisfy that service's semantic restrictions.

3.5.4.6 Location Management

The decision of where to store given data may be made for a variety of reasons, such as speed of access for a given client or guarantees of availability. OGSA data services provide several services that manage the location of data.

Data cache services hold recently used data at a given location to avoid unnecessary additional transfers. Data caches can be configured in terms of, for example, lifetime and update consistency.

Replication services manage multiple copies of data either to increase availability through redundancy or to improve performance by reducing access times. They can be configured in terms of update consistency and replication policies.

Services for individual users allow them to upload and manage their own data with the above facilities. Security settings control the access permitted to other users.

3.5.4.7 Update

OGSA data services provide a range of mechanisms for updating data resources, depending on the semantics of the data resource and the nature of the data to be uploaded. Small examples include updating a record in a database or adding a new file to a distributed file system. Other

examples include bulk loading data to such resources. Some services may also support the upload of filters, aggregators or transformers to a given data resource.

Data services may specify some form of transactional behavior for update operations. When a data resource has replicated versions or is the source for derived data services, the updates may be propagated to the replicated or derived versions. In this case, and in the case where several clients are updating the same data resource, the services may implement various forms of consistency maintenance, e.g., to ensure that a client always sees the results of its own updates in any queries that it issues itself. These consistency requirements may be specified as QoS policies.

3.5.4.8 Transformation

Various components of a system may transform data. OGSA does not attempt to distinguish between minor transformations (such as reversing byte order) and more fundamental transformations. In the most general sense, any service that consumes data of one sort and generates other data could be viewed as a transformation.

Data services may themselves transform data. For example, they may provide built-in functionality to convert data from one format to another, or filter it, before moving it or updating it. They may also derive their exported data from an arbitrary computation applied to a set of base data resources. These transformations may be instigated explicitly by certain operations, or they may be programmed to be triggered automatically in response to certain conditions. The interfaces described here do not expose these transformations explicitly. They may include the ability to specify a desired output format from those available, which may implicitly call one or more transformations.

Data services may support stored procedures: application-provided code that the service will execute in client mode. This makes the service a form of container. The usual OGSA-EMS interfaces can be used to control the provisioning of stored procedures. As with other OGSA containers, the procedures themselves will typically only run on certain services, i.e., containers and procedures are not generic.

3.5.4.9 Security Mapping Extensions

Database management systems often implement sophisticated security mechanisms. Some of these provide a large range of possible operations and access control at the level of individual records. Therefore the OGSA data services support extensions to the standard OGSA security infrastructure that allow effective use of fine-grained mechanisms in OGSA context. For example, data services can translate security related information from the OGSA model to that needed by the data resource.

3.5.4.10 Resource and Service Configuration

Data resources often provide sophisticated configuration options. These can be made available to clients via the OGSA data services. In addition, the services may provide additional operations for configuring the virtualization of the resource provided by the service. For example, a service that exports a relational schema and a SQL access interface might allow for specific tables from an underlying database resource to be made part of the service's data virtualization, or for views on the underlying database to be made available as tables within that virtualization.

3.5.4.11 Metadata Catalogues

Metadata catalogues are data services that store descriptions of data held in certain other data services. OGSA provides support for maintaining links and consistency between OGSA services and the metadata that describes them. For resource discovery, this functionality is part of OGSA information services, but there are other limited uses that can be met by the direct application of

data services. A range of different mastering and peering strategies are supported for the deployment of metadata catalogues and the resolution of conflicting updates. These requirements can be specified as QoS policies.

The metadata for OGSA data services may include information about the structure of the data, including references to the schemas that describe the data. For some services this is not practical, as the data resources include many schemas that are modified frequently, and in these cases schema information will be provided by the services themselves.

3.5.4.12 Data Discovery

Data discovery requires languages or ontologies for describing the data and a query language that operates over these descriptions. The infrastructure for data discovery—registries, notification, etc.—is provided by the OGSA information services.

3.5.4.13 Provenance

Users of data services may also wish to see information about the provenance and quality of the data provided by the services. This may be at the level of the whole resource or of its component parts, sometimes to the level of individual elements. This in turn requires the services or other processes that generate the data to also maintain the consistency of the provenance information. It also requires an architecture for maintaining this information. Complete provenance information can allow the data to be reconstructed by following the workflow that originally created it. Provenance information may be provided by the service itself, or it may be maintained in a metadata catalog or a logging service.

3.5.5 Properties

Properties are non-functional capabilities of services. Whereas functional capabilities are defined by entries in the service interfaces, non-functional properties are inherent in the design.

3.5.5.1 Scalability

The OGSA data services handle large scale in several dimensions, including size of data sets, number of data sets, size of data flows, and number of sites.

3.5.5.2 Performance

The OGSA data services are designed to minimize the copying and movement of data. This is a key factor in the overall performance of the Grid.

OGSA data services provide performance information to scheduling services, such as their data access rates and their network location. As a particular example, some mass storage systems have significant latency for data access while they load the necessary tapes.

3.5.5.3 Availability

A data service should make every effort to meet agreed client requirements for availability. However, in a distributed environment, occurrences beyond its control may mean that it is unable to do so. The OGSA data services provide features for graceful degradation in the event of network or other failures. For example, query services may be configured to return partial results when only a subset of sources is available.

3.5.5.4 Legal and Ethical Restrictions

The OGSA data services may have to operate within an environment where a variety of legal and ethical policies affect their operation. For example, some policies may restrict the entities that can access personal data and limit the operations that they can perform (confidentiality). Privacy

concerns may limit the queries that can be made about individuals, although in some cases the policies may permit queries that return results about a group as a whole, such as average income or total salary.

Data are often covered by copyright limitations, which restrict the right to create copies of data. In the European Union, the similar (but distinct) “database right” restrictions apply specifically to databases.

The security mechanisms provided by OGSA allow these restrictions to be specified. When used with data services, these mechanisms must allow the specification of policies that apply at the level of groups (e.g. tables) or elements within a resource. Scenarios involving complex data will be particularly exacting test cases for these mechanisms.

3.5.6 Interactions with Other OGSA Services

This section summarizes the interactions between the data services and the other parts of OGSA.

3.5.6.1 Transactions

Data services are the classic example for transactions, and many data resources provide independent transactional support. There are several ways that transactions can be implemented in distributed systems. In general, transactions should be supported for any interaction of OGSA services, not just data services. Distributed transactions can be provided using a number of proposed specifications (see §3.3.8). Data services may also provide guarantees of their own transactional behavior.

3.5.6.2 Logging

Similarly to other OGSA services, data services use the logging services, for example for audit (see 3.9.4.3). Conversely, the logging services themselves may use data services to store the logs.

3.5.6.3 Deployment and Configuration Services

Often before a data resource must be deployed on the appropriate machine or configured in a particular way. The OGSA deployment and configuration services provide this functionality (see 3.8.4.6)

3.5.6.4 Workflow

Workflows can instruct data access services to send query results to third parties and to apply transformations as the data is moved. This task requires that the workflow services have intimate knowledge of the capabilities of the data access services. Also, a call to a distributed query service may cause a variety of data movements and operations on other machines. For these reasons a data service can provide information to the workflow manager to ensure that the workflow enactment takes full account of the effect of distributed queries on the network and other resources. (See also §3.3.9.)

3.5.6.5 Execution Planning and Job Management

The OGSA-EMS Execution Planning Service (see §3.4.6.1) can use monitoring information such as bandwidth, utilization patterns and packet size to choose the best approach for moving a given data set to suit the agreed quality of service. Providing this information is primarily the responsibility of the OGSA information services, but the OGSA data services will provide such raw information as necessary

The Job Manager service (see §3.4.5.2) may use data services in order to stage data where it is needed for the execution of a particular job.

3.5.6.6 Resource Reservation

Data services may need to reserve certain resources in order to operate (see §3.4.6.3). For example, a data transfer operation will require storage space and network bandwidth, while a distributed query system may require compute power in order to perform join operations. Similar considerations apply when the provision of a data service is scheduled for a certain time. Equally, data services must provide interfaces to allow them to be reserved.

3.5.6.7 Discovery

Discovery services are vital to the data architecture. The data services may use the discovery services not just for registering services themselves, but also for registering the data sets that are stored by those services. This requires languages or ontologies for describing data. Discovery services may also register the locations of schema definitions.

3.5.6.8 Security

Security services provide authentication and access control for the operations of data services and the underlying data. They support the mapping of OGSA identities and roles to resource-specific identities and roles. VO management services similarly provide control of these mappings. Security services also provide support for checking the integrity and encryption of data transfers and provide the means for a data service to “tag” returned data with its security properties.

3.5.6.9 Network Management

Network management can be crucial when planning the transfer of large amounts of data. Data services will provide the necessary information, and possibly the initial trigger, for appropriate OGSA services to configure the network parameters to suit the amount of data to be transferred and the time constraints specified on those transfers. This functionality can be provided by an appropriate resource management service (see §3.6).

3.5.6.10 Naming

The OGSA data services use OGSA-naming (see §3.3.4) for naming data sets. For example, in a replicated file system, an *address* may specify the location of a file on a particular machine. The *abstract name* will identify the file in a location-independent manner, perhaps using a directory service. The *human-oriented name* may be a user-friendly short file name that can be disambiguated by referring to the context in which it is used.

3.5.6.11 Notification

We can use the notification service to notify clients and other services about management events, performance and resource issues, and to implement consistency mechanisms. We can also use it to externalize database triggers. Database management systems often provide a mechanism that specifies actions to be taken when certain conditions (triggers) are met. OGSA can easily cause these triggers to invoke the notification services.

In addition, some implementations of the notification services may use the data services to store the notification messages and support client queries.

3.6 Resource Management Services

3.6.1 Objectives

Resource management performs several forms of management on resources in a Grid. In an OGSA Grid there are three types of management [OGSA RM] that involve resources:

- Management of the physical and logical resources themselves (e.g., rebooting a host, or setting VLANs on a network switch)
- Management of the OGSA Grid resources exposed through service interfaces (e.g., resource reservation, job submission and monitoring)
- Management of the OGSA Grid infrastructure, exposed through its management interfaces (e.g., monitoring a registry service)

3.6.2 Model

Different types of interfaces realize the different forms of management in an OGSA Grid. These interfaces can be categorized into three levels, shown in the middle column of Table 2, and also on the right in Figure 7.

Table 2: Relationships between types of management and interfaces

Type of management	Level of interface	Interface
Management of the physical and logical resources	Resource level	CIM/WBEM, SNMP, etc.
	Infrastructure level	WSRF, WSDM, etc.
Management of the OGSA Grid Resources	OGSA functions level	Functional interfaces
Management of the OGSA Grid Infrastructure		Specific manageability interfaces

We provide below a detailed description of each level and its interfaces. Note that the descriptions focus on the manageability interfaces, *not* on the implementation (e.g., on the services that implement them). Also note that a service may implement multiple interfaces (which are possibly unrelated in terms of functionality), and that a service may be separated from the functionality that it represents (e.g., a manageability provider for a resource that is separate from this resource). Therefore a description based on services would be imprecise, and a description based on interfaces is chosen instead.

In Figure 7, the OGSA capabilities cover all levels, extending to capabilities in the resources that are needed to implement these OGSA capabilities. The interfaces are shown as small circles.

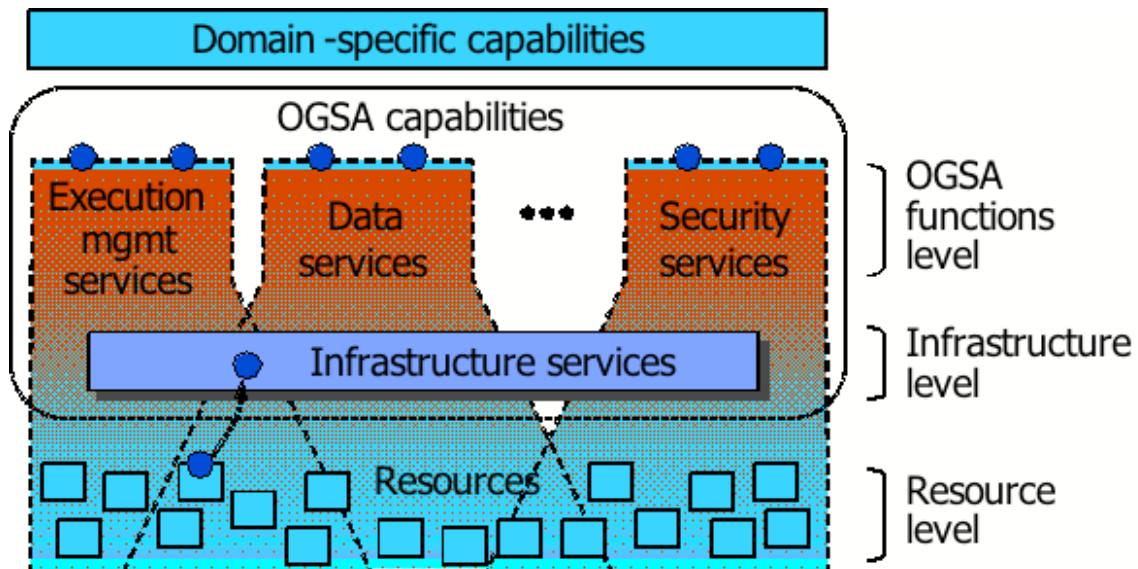


Figure 7: Levels of management in OGSA

At the resource level, resources are managed directly through their native manageability interfaces (for discrete resources, these are usually SNMP, CIM/WBEM, JMX, or proprietary interfaces). Management at this level involves *monitoring* (i.e., obtaining the state of the resource, which includes events), *setup and control* (i.e., setting the state of the resource), and *discovery*. These resources are managed by following the description given by an *information model*, which defines their properties, operations, events, and their relationships with each other.

The infrastructure level provides the base management behavior of resources, forming the basis for both manageability and management in an OGSA environment. Standardization of this base management behavior is required in order to integrate the vast number and types of resources—and the more limited set of resource managers—that are introduced by multiple suppliers. The infrastructure level provides:

- A *base manageability model*, which allows resources in OGSA to be manipulated through the standard Web services means for discovery, access, etc. In the WSRF Basic Profile [OGSA WSRF BP] this corresponds to representing resources as WS-Resources [WS-Resource]. This model allows the resources to become manageable, at least to a minimum degree, by enabling discovery, termination, introspection, monitoring, etc. The information model of the resources is accessed through the base manageability model.

By using a single basic manageability model per OGSA Basic Profile and possibly a single information model, the relationships (dependencies and component definitions), operational status, statistics, etc. become consistent throughout all layers of management.

- A *generic manageability interface* that is common to all services implementing OGSA capabilities. This manageability interface has functionality such as introspection, monitoring, and creation and destruction of services.

At the OGSA functions level, there are two types of management interfaces, denoted by the two circles on the top of each of the capabilities shown in Figure 7:

- *Functional interface*: Some common OGSA capabilities (e.g., OGSA EMS) are a form of resource management. Services that provide these capabilities expose them through functional interfaces (e.g., create and destroy a job).
- *Manageability interface*: Each capability has a specific manageability interface through which the capability is managed (e.g., monitoring of registries or monitoring of a job manager). This interface could extend the generic manageability interface, adding any manageability interfaces that are specific to the management of this capability.

3.6.3 Management Capabilities

On the WSRF Basic Profile, management functionality at the infrastructure level is envisioned to use OASIS WSDM, which is expected to become a key standard for manageability across the IT landscape. WSDM has separate documents to address management *of* Web services (MOWS) [MOWS] and management *using* Web services (MUWS) [MUWS part 1][MUWS part 2]. MUWS provides the base manageability model, i.e., how to represent resources, plus basic manageability functions that are common to the OGSA, such as state and status representation, events, metrics, and relationships among resources. MOWS provides the generic manageability interface to manage the services in an OGSA Grid.

As for information models, modifications and extensions to CIM as well as a set of CIM profiles are under consideration. Due to its breadth and extensibility, CIM can be used in the management

tasks of multiple OGSA capabilities such as security, execution management, self-management, etc. CIM can also be used over different Basic Profiles, improving the unification of OGSA.

At the OGSA functions level, the resource management capability includes (but is not limited to) typical distributed resource management activities and IT systems management activities. These activities may be policy-based, i.e., may enforce policy assertions that are put in place to support requirements such as authentication scheme, transport protocol selection, QoS metrics, privacy policy, etc. Functionalities included in the OGSA resource management capability include resource reservation, monitoring and control, VO management, security management, problem determination and fault management, policy management (i.e., management of the policies themselves), service groups and discovery services, metering, deployment, and discovery.

3.6.4 Properties

3.6.4.1 Scalability

A management architecture needs to scale to potentially thousands of resources. OGSA should therefore allow different forms of management—hierarchical, peer-to-peer (federated/collaborative), etc.—to achieve this scalability.

3.6.4.2 Interoperability

A management architecture must be able to span software, hardware and service boundaries, e.g., across the boundaries between different products, so standardized and broad interoperability is essential to avoid “stovepipes.”

3.6.4.3 Security

There are two security aspects in management. *Management of security* concerns the management of the security infrastructure, including the management of authentication, authorization, access control, VOs and access policies. *Secure management* refers to using the security mechanisms on management tasks. Management should be able to ensure its own integrity and to follow access control policies of the owners of resources and VOs.

3.6.4.4 Reliability

A management architecture should not force a single point of failure. For purposes of reliability, a resource may be virtualized by multiple services each exposing a single URL as the management endpoint. In such situations, the system that provides manageability capabilities must be aware that for certain queries, such as metrics, the manageability provider must aggregate the results from the multiple services that virtualize that single resource.

3.6.5 Interactions with Other OGSA Services

The infrastructure services implement the interfaces at the infrastructure level, such as WSRF and WSDM. Since these interfaces define the base management behavior of resources, all OGSA services will use these interfaces when managing these resources.

The information services provide resource management functionalities, especially for monitoring, that many OGSA services will use. Two examples are registries, used for resource discovery, and logging, which can be used by problem determination.

The execution management services and data services are consumers of resource management functionality, using it for discovery, provisioning, monitoring, etc. of execution and storage-related tasks. They are also providers of resource management functionality at the OGSA functions level.

Self-management services are used on the functional interfaces of resource management services to control resources (e.g., applications, devices, and their allocation) on a Grid. They can also be used on the specific manageability interfaces of these resource management services, to control the Grid infrastructure itself (e.g., monitoring registries and, when they get overloaded, deploying more instances). The same applies also to security services which, for instance, control access for both the functional and specific manageability interfaces.

3.7 Security Services

This section describes objectives, models, functional capabilities, and properties of the security services, and discusses the interactions with the rest of OGSA.

3.7.1 Objectives

OGSA security services facilitate the enforcement of the security-related policy within a (virtual) organization [Grid Anatomy][Grid Physiology][VO Security].

In general, the purpose of the enforcement of security policy is to ensure that the higher-level business objectives can be met. This is often a delicate balance, as there is an increased cost associated with increased policy enforcement, and an increased loss exposure with less stringent enforcement, while potential profits or rewards may be adversely affected by the increased complexity and inflexibility of stricter policy requirements. Note that in some cases, legislative rules and regulations mandate conformance to associated security policy.

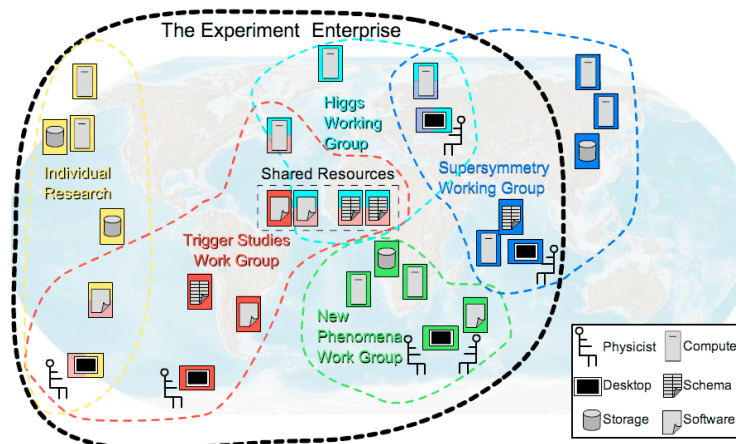


Figure 8: Cross-Organizational Collaborations

The property that Grid-specific applications may span multiple administrative domains (see Figure 8) implies that each of these domains will have its own business objectives to meet, which translates to the individual domains separately establishing and enforcing their own policies, which can differ greatly in complexity and strictness. Note that all interactions associated with a thread of work in a Grid application must adhere to the domain-locally-enforced policies as well as to the policies established for the VO—i.e. the cross-organizational (business) agreement.

To meet the business and associated security objectives, the security policy can, for example, specify the enforcement of message integrity and confidentiality, authentication of interacting entities, minimum authentication strength, secure logging and audit, separation of responsibilities, intrusion and extrusion detection, authorization policy checks, least privilege operations, mandatory access control mechanisms, discretionary access control mechanisms, trust and

assurance level of the environment, application isolation, avoidance of DoS attacks, redundancy, and training.

OGSA security architectural components must support, integrate, and unify popular security models, mechanisms, protocols, platforms, and technologies in a way that enables a variety of systems to interoperate securely. The components must be able to support integrating with existing security architectures and models across platforms and hosting environments. This means that the architecture must be *implementation-agnostic*, so that it can be instantiated in terms of any existing security mechanisms (e.g., Kerberos [Kerberos V5], PKI[RFC3280]); *extensible*, so that it can incorporate new security services as they become available [RFC2903][Liberty]; and *integratable* with existing security services. Also, services that traverse multiple domains and hosting environments need to be able to interact with each other, thus introducing the need for interoperability at multiple levels: protocol, policies and identity. In addition, certain situations can make it impossible to establish trust relationships among sites prior to application execution. Given that the participating domains may have different security infrastructures (e.g. Kerberos or PKI) it is necessary to realize the required trust relationships through some form of federation among the security mechanisms.

3.7.2 Model

This section describes a model to facilitate the reasoning about, and the specification of, the OGSA security services. The presented model is not a formal model in the mathematical sense, but is meant to provide a language to describe, and to obtain common understanding about, the security policies that are to be enforced.

In general, one can make the observation that entities are interacting through mechanisms within a context. Entities are things such as users, subjects or services. Interaction mechanisms span all the different communication methods, such as mail, telephone, HTTP, SOAP, SSL/TLS, etc. A context puts the interaction in perspective, and could localize the interaction to a single machine, or could associate it with a service invocation within a VO; with an established secure association; with a distributed transaction and so on.

All these entities, mechanisms and contexts can be described by sets of attributes or properties. Some of these attribute types and values may be used for unique identification, others are used for classification or grouping. Furthermore, some of these attributes are *inescapable*. An inescapable attribute can identify an entity by itself without any reference to an outside authority. Examples of inescapable attributes are a shared secret, a private/public key pair, fingerprints, etc. All other attribute values are essentially bound to an entity's inescapable attribute by an issuer or attribute authority.

Security policies are statements about these different entities, interaction mechanisms and contexts, and specify restrictions on the associated attribute values, properties and their relationships. The policy statements (or rules) will be able to be expressed in terms of these entities (e.g. identities, user attributes), resources (i.e. end points), and environment characteristics (e.g. time, location, purpose, or the trust level of the requester or request path). These policies will be about various aspects including authorization, authentication, trust, identity mapping, delegation, assurance levels, etc.

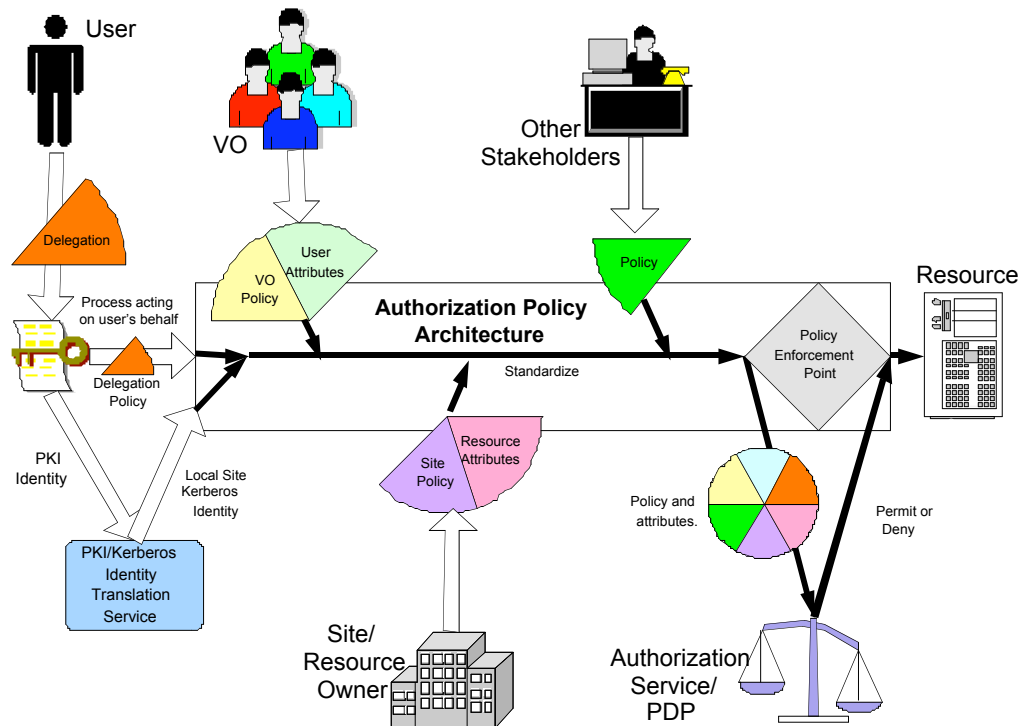


Figure 9: Example input for the policy decision and enforcement

The model defines the security services as entities with interaction patterns that facilitate the administration, expression, publishing, discovery, communication, verification, enforcement and reconciliation of the security policy. In other words, the security policy enforcement is the ultimate goal, and the security services are designed and deployed to support that goal.

To make this model more concrete, we will identify a number of these entities, interaction mechanisms and contexts, and discuss some of their attributes and common relationships. Figure 9 shows an example of the policy enforcement to protect the use of a resource. In order to understand the model, it helps to walk through the example in Figure 9 from the two opposite ends: on one side the initial user authentication, and on the other the enforced authorization, and to realize that a user will only be allowed to access the resource if the authorization policy evaluated at the enforcement point will yield permit. The resource's authorization policy will be expressed as rules for the attribute values of the relevant entities, such as, for example, the user's name, her role, her VO-membership, etc. With the initial authentication of the user, only inescapable attribute values are presented and verified, such as the possession of the private key associated with the presented public key. Normally, there is a mismatch between the fact that the policy is expressed in derived attributes and the initial authentication only yields a key, and this can be resolved by finding the matching attribute assertions that bind the key to the attributes

used in the policy expression. Several security services are depicted in Figure 9; these are used to federate the user's public key credentials to Kerberos ones, and to obtain attributes from different sources, such as the VO or the local site, that assert specific attribute bindings used in the enforced policy statements.

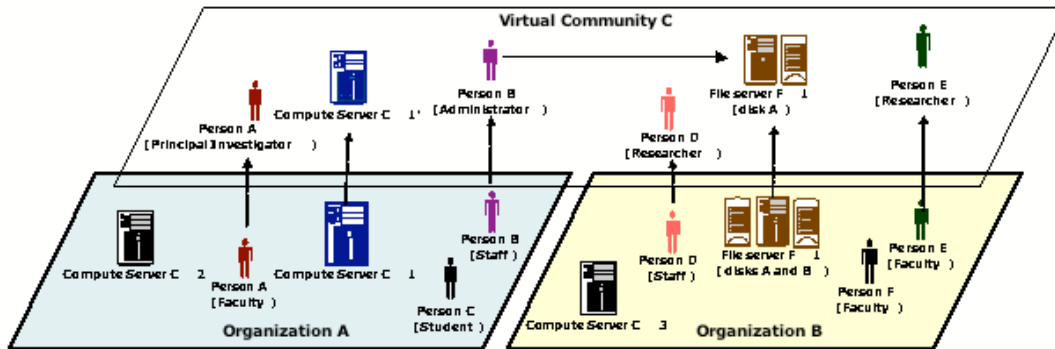



Figure 10: The Virtual Community Concept

Conceptually, the presented model is no different from the general Web services security model and other distributed computing architectures defined by standards bodies[RFC2903], consortia [Liberty], or in company whitepapers¹³. The Grid applications, however, put a distinct focus on the entities and patterns related to cross-organizational settings, and the security services that enable those interactions [VO Security][Role-Based VO][VOMS-1.1][Fine-Grain Auth][Fine-Grain Auth RM][Dynamic Access Control][PRIMA][Grid Auth Framework] [Grid AAA Req][CAS][SAZ][GS Security]. It can be noted that such interactions are also prevalent in business applications where business-to-business interactions (or organizations boundaries within an enterprise) bring about similar requirements and settings. For example, Figure 10 shows two organizations and an overlaid virtual community that is governed by its own policy. Note that the entities within this virtual community context have their own specific attributes and properties, which are different from, and have no relation to, those in their home domain. This highlights the fact that our security services model has to support the concurrent enforcement of multiple policies, which each has to be evaluated within their own proper context.

3.7.3 Example Scenarios

3.7.3.1 Digital Library

A digital library program for educational material is operated by a public organization. A number of schools and public libraries in the nation participate in the library program. The program provides teachers and students of the participating schools or libraries with a means of sharing educational materials such as digital books, videos, photos and all other digital materials for education that would originally have been stored by individual schools or libraries. Each of the participating schools and libraries is responsible for its users and resource (educational materials) management, such as registration or removal of users and resources. Some schools, for example universities, may consist of several sub-organizations. In such a case, each sub-organization in a university could form a VO by itself and the university VO would then become an aggregation of these school VOs. This scenario is depicted in Figure 11.



The following is a list of example policies of how users of the library program have access to the shared materials.

- All the students enrolled in a participant school can have read access to the materials for students.
- All the teachers can have read access to materials both for students and faculties.
- Some teachers who are certified by a participant school can also register new materials.

The school libraries in the university may allow access from inside the university, but will not allow accesses from the Digital Library VO.

3.7.3.2 Least Privilege Delegation

The delegation of rights is a fundamental capability needed to let services work on behalf of other entities. With this rights-delegation comes the associated risk that any of these services may be compromised and use those rights in inappropriate ways. To limit the exposure, one would like to limit the delegated rights to only those rights truly needed by the service. This *least-privilege delegation model* requires that one is able to match the invoked service operations with the exact “amount” of rights, which is a non-trivial requirement. Many Grid applications use the concept of jobs, in which job directives are specified in their own language. The job requirements are then matched with the capabilities and availability of resources by discovery, brokers, and scheduler services. The language used for the expression of these job directives and resource capabilities should be able to match up with the directives used to express the equivalent rights needed. Any mismatch is likely to result in a deployment where essentially too many rights will have to be given to services to ensure that the job directives can be executed. This issue is illustrated in Figure 12.

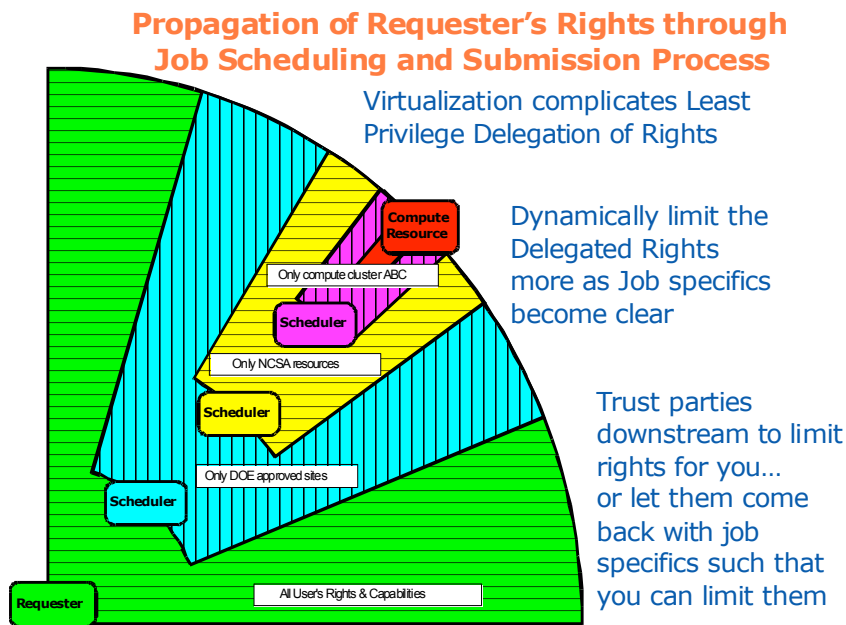


Figure 12: Least-Privilege Delegation and Job Description

3.7.3.3 Secure Logging in a Distributed Environment

Logging services, and secure access to the logs for reconciliation purposes, becomes a much harder problem in a distributed setting, where the services and associated logs may reside in different administrative domains. This scenario is depicted in Figure 13.



Logging services should have security characteristics so that logs are secured, are capable of being tamper-proof, and ensure integrity of the messages. When we get to that level of sophistication, it leads to the notion of auditing—where events are recorded in a secure fashion. These can be any events, including security events, business events, transaction events, etc. Security services and infrastructure will need to deal with generating security events that can be consumed by the event infrastructure so that they can be audited, or acted upon (e.g. an intrusion defense system may react to a set of DoS events).

3.7.4 Functional Capabilities

In this section we describe the security functional capabilities and how they relate to the corresponding security services and usage patterns. More detailed description of services and usage patterns will be provided in later versions of this document.

As an example, Figure 14 illustrates how the requestor and service provider both call-out to different infrastructure security services to ensure policy compliance. Note that in the picture the call-outs are made from within the stubs, and outside of and transparently to the application. The expectation is that most of the policy enforcement could be taken care of this way, which has the desirable benefit of keeping the security-specific code to a minimum for the application developers.

Furthermore, Figure 14 clearly shows that in order for service invocations to comply with the requestor’s, the service provider’s and the VO’s policy, call-outs are made to different security service instances that are managed in those different organizations.

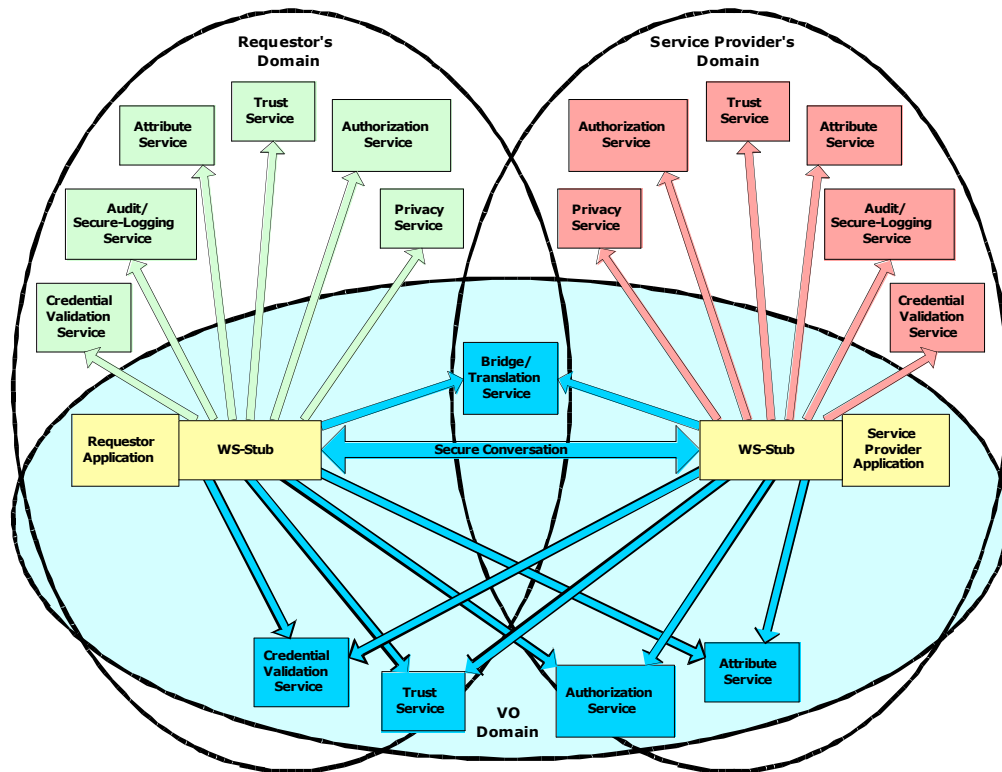


Figure 14: Security services in a virtual organization setting

We enumerate the functional capabilities and corresponding security services in the following:

- *Authentication.* Authentication is concerned with verifying proof of an asserted identity. This functionality is part of the Credential Validation and Trust Services in Figure 14. One example is the evaluation of a user-id and password combination, in which a service requestor supplies the appropriate password for an asserted user-id. Another example involves a service requestor authenticating through a Kerberos mechanism, and a ticket being passed to the service provider's hosting environment, which determines the authenticity of the ticket before the service is instantiated.
- *Identity mapping.* The Trust, Attribute and Bridge/Translation Services in Figure 14 provide the capability of transforming an identity that exists in one identity domain into an identity within another identity domain. As an example, consider an identity in the form of an X.500 Distinguished Name (DN), which is carried within a X.509 V3 digital certificate. The combination of the subject DN, issuer DN and certificate serial number may be considered to carry the subject's or service requestor's identity. The scope of the identity domain in this example is considered to be the set of certificates that are issued by the certificate authority. Assuming that the certificate is used to convey the service requestor's identity, the identity mapping service via policy may map the service requestor's identity to an identity that has meaning (for instance) to the hosting environment's local platform registry. The identity mapping service is not concerned with the authentication of the service requestor; rather it is strictly a policy-driven name-mapping service
- *Authorization.* The authorization service is concerned with resolving a policy-based access-control decision. The authorization service consumes as input a credential that embodies the identity of an authenticated service requestor and, for the resource that the service requestor requests, resolves, based on policy, whether or not the service requestor is authorized to access the resource. It is expected that the hosting environment for OGSA-compliant services will provide access control functions, and it is appropriate to further expose an abstract authorization service depending on the granularity of the access-control policy that is being enforced.
- *Credential conversion.* The Trust, Attribute and Bridge/Translation Services in Figure 14 provide credential conversion from one type of credential to another type or form of credential. This may include such tasks as reconciling group membership, privileges, attributes and assertions associated with entities (service requestors and service providers). For example, the credential conversion service may convert a Kerberos credential to a form that is required by the authorization service. The policy-driven credential conversion service facilitates the interoperability of differing credential types, which may be consumed by services. It is expected that the credential conversion service would use the identity mapping service.
- *Audit and secure logging.* The audit service, similarly to the identity mapping and authorization services, is policy-driven. The audit service is responsible for producing records that track security-relevant events. The resulting audit records may be reduced and examined so as to determine whether the desired security policy is being enforced. Auditing and subsequently reduction tooling are used by the security administrators within a VO to determine the VO's adherence to the stated access-control and authentication policies.

- Privacy.* The privacy service is primarily concerned with the policy-driven classification of personally identifiable information (PII). Service providers and service requestors may store personally identifiable information using the privacy service. Such a service can be used to articulate and enforce a VO's privacy policy.

A different view of the relationships between the relevant security components is shown in Figure 15 as a layered stack of related services.

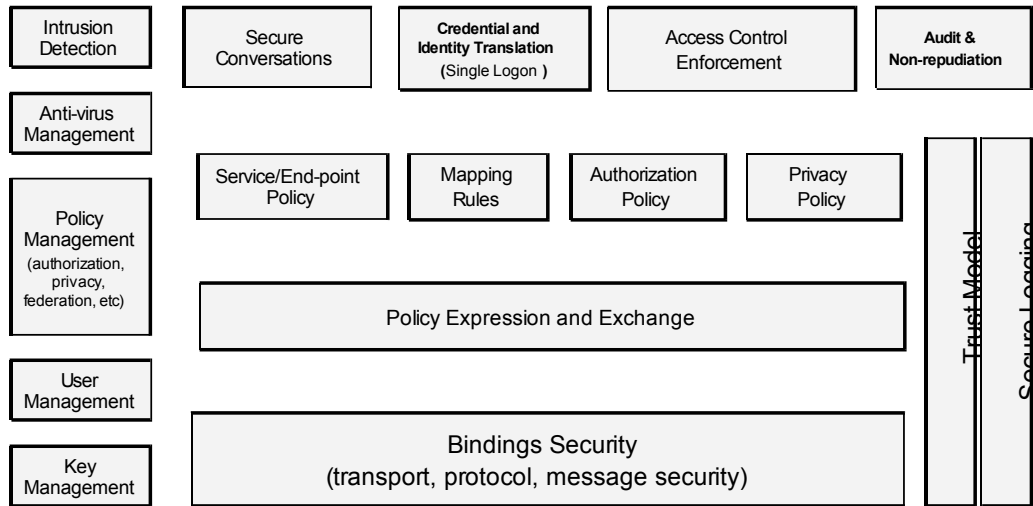


Figure 15: Components of Grid Security Model

All security interfaces used by a service requestor and service provider need to be standardized within OGSA. Compliant implementations will be able to make use of existing services and defined policies through configuration. Compliant implementations of a particular security-related interface would be able to provide the associated and possibly alternative security services.

3.7.5 Properties

In general, the properties of the security services depend on the technological requirements that follow from the policy that has to be enforced. For example, in order to be able to enforce a stated policy, certain service levels have to be met by the security services, which translates in properties such as maximum latency, response-time, availability, and recovery.

In many cases, the implementation of the security services will be able to obtain the desired properties through the use of other services. For example, the attribute information service could use the data services to access the assertion information in a directory (e.g., LDAP) or database (e.g., RDBMS), and it could make use of the data mirroring features of those services to achieve the desired availability property needed to meet the enforcement of the stated security policy.

3.7.6 Interactions with Other OGSA Services

In general, the invocation of any OGSA service is subject to enforcement of all relevant security policies. In some cases, this enforcement is implicit and hard-coded; in other cases the ability to plug-in or call-out to external security infrastructure services is essential for deployment. In this view, all OGSA services depend on and are layered above the security services.

In some cases, the security services and requirements are intimately connected to other OGSA services on a higher level. For example, an attribute service implementation can make use of

OGSA data services to retrieve policy related information from a registry or database. In this view, security services can be consumers of other OGSA services.

3.8 Self-Management Services

3.8.1 Objectives

Self-management was conceived as a way to help reduce the cost and complexity of owning and operating an IT infrastructure. In a self-managing environment, system components—including hardware components such as computers, networks and storage devices, and software components such as operating systems and business applications—are self-configuring, self-healing and self-optimizing.

These self-managing attributes, described further in §3.8.4, suggest that the tasks involved in configuring, healing and optimizing the IT system can be initiated based on situations that the components themselves detect, driven by business needs, and that these tasks are performed by those same technologies. Collectively, these intuitive and collaborative characteristics enable enterprises to operate efficiently with fewer human resources, while decreasing costs and enhancing the organizations' ability to react to change. For instance, in a self-managing system, a new resource is simply deployed and then optimization occurs. This is a significant shift from traditional implementations, in which a significant amount of analysis is required before deployment, to ensure that the resource will run effectively.

Although it is expected that the self-managing attributes will be pervasive in OGSA, it is not the case that every single service in a self-managing system will demonstrate all, or even a subset, of these attributes. Rather, these attributes are part of the autonomous nature of the system as a whole.

One of the main objectives of self-management is to support service-level attainment for a set of services (or resources, depending on the taxonomy)—with as much automation as possible, to reduce the costs and complexity of managing the system. In an operational environment, it is often necessary to control various aspects of the behavior of a solution component in a manner that cannot be determined a priori by the component developer. This is achieved in a self-managing system through the deployment of policies to govern the behavior of system components derived from business objectives; a role called *Service Level Manager*. Service level managers (SLMs) are responsible for setting and adjusting policies, and then changing the behavior of the managed resource or service in response to observed conditions in the system to ensure overall compliance with business objectives. SLMs are themselves managed by policies that are either embedded in their implementation or retrieved from other SLMs. Therefore it is possible that a service may have both self-managing aspects and also be involved in a self-managing activity.

Thus composition and hierarchy are expected between different SLMs, thereby significantly reducing the complexity of operation of the system and initial system design, since complex SLMs can be built by involving simple SLMs in the process.

While the self-management capability is a significant part of the OGSA, this work is still at a preliminary stage and hence only some aspects of self-management are described here. A more detailed analysis will be provided in later versions of this document.

3.8.2 Basic Attributes

The collection of attributes collectively needed for various stages of self-management are given below. Their existence at a conceptual level is discussed; no assumptions are made about their implementation.

3.8.2.1 *Service Level Agreement*

Service level agreements (SLAs) include business or IT agreements between the service provider and the users of the service. SLAs provide guidance, expressed in terms of measurable intent, as to the purpose and delivery objective of the service or resource that is being managed.

3.8.2.2 *Policy*

A policy is used to govern the behavior of an SLM and the manageable resources under its control. Policies governing the behavior of self-managing entities are derived from service level agreements (SLAs), and deployed as part of the management context under which the manageable resources are used. SLMs orchestrate real-time changes in the dynamic management infrastructure, based upon policy which governs their behavior.

3.8.2.3 *Service Level Manager Model*

This model provides the framework to instantiate and work with an SLM such that various SLMs and human operators can interact and understand each other without having knowledge about each other built in at design time. There is an interface component to the SLM model, as well as a representation of an SLM and its component management services (not the managed services) and the control loops that they form, and instantiation and maintenance of SLMs.

SLMs are typically modeled after a *generic control loop pattern* that may be used to control and adjust various service activities. This pattern is a cycle consisting of Monitoring, Analysis and Projection, and Action phases. SLMs carry out *service level attainment* activities. Service Level Management is further described in §3.8.4.1.

3.8.3 Example Scenarios

Self-management capabilities are fundamental to the Grid. This section describes two examples of real-world usage: Job-level management and Grid system-level management. Similar concepts can be derived by analyzing how self-management works for other scenarios, e.g., security and so on. In the following scenarios the terminology used is that of the OGSA Execution Management Services capability and the Commercial Data Center use case.

3.8.3.1 *Job Level management*

The IT business activity manager submits a job and negotiates the job's SLA. The job must be executed so that it satisfies the agreement. At a later stage, the IT business activity manager may wish to renegotiate the agreement to address new business requirements. It does so with the job manager and updates the existing agreement. When the agreement is updated the resource requirements are recalculated in the service level attainment loop (analysis and projection, see Figure 16), and the provisioning steps (including resource allocation and deployment) are triggered (Action) as a result of the changing conditions. After the provisioning steps, the resources are in a ready state for the required components of the job to start, including starting executable resources such as application server or DBMS.

The IT business activity manager can monitor the load and resource utilization of the provisioned resources using the Monitoring Service or the Metering Service. The IT business activity manager also can obtain the state of the running job from the job manager.

3.8.3.2 *Grid System Level management*

At the Grid system level, one aim is to improve resource utilization while maintaining the SLAs of running jobs. The Grid system service level manager may have to add new resources to prepare for expected load increases, and to release surplus resources in order to reduce costs. The resources allocated to a job may also have to be adjusted based on policy—e.g., the priority of a job relative to other jobs in the Grid system.

In the Analysis and Projection phase, information about the available resources and current load, and estimates of the expected future utilization and load are evaluated. An expected increase in utilization may trigger (Action) provisioning steps to add resources to the Grid system's available pool, e.g., by arranging to shift resources from other systems or by releasing resources that are currently being used by lower-priority jobs.

3.8.4 **Functional Capabilities**

Self-management is essentially self-configuration, self-healing, and self-optimization. This shows the essential difference of self-management from other OGSA categories: namely, it is not just about the components that are involved in doing self-management, but the method by which it is done—the *sauce*, or the way in which the components interact, control loops are formed and systems behave intelligently, based on environmental changes. The mechanisms that bring self-management about can be described as follows:

- Self-configuring mechanisms adapt dynamically to changes in the IT system, using policies provided by the IT professional. Such changes could trigger provisioning requests leading to, for example, the deployment of new components or the removal of existing ones, maybe due to a significant increase or decrease in the workload.
- Self-healing mechanisms can detect improper operations of and by the resources and services, and initiate policy-based corrective action without disrupting the IT environment. Self-healing has an element of self-protection included in it as well. This means that components can detect hostile behaviors as they occur, and take corrective actions to make themselves less vulnerable. The hostile behaviors could include unauthorized access and usage, virus infection and proliferation, and denial-of-service attacks.
- Self-optimizing mechanisms are able to tune themselves to the best efficiency to meet end-user or business needs. The tuning actions could mean reallocating resources to improve overall utilization or optimization by enforcing an SLA. Self-optimization makes use of self-configuration in its implementation.

It is important to note that the self-configuring, self-healing, self-optimizing mechanisms are not independent of one another. They work together to allow changes to be made to the configuration of one or more aspects of the IT system.

All of the OGSA service categories are utilized in achieving self-management. In addition, the following sections highlight some special functional requirements that are important for self-management, but that may not be reflected in the same way in other OGSA service categories.

3.8.4.1 *Service Level Management*

Service level management ensures that the desired QoS is maintained. This is done through the activities of the SLMs as described in §3.8.1. These *service level attainment* activities are described in more detail here

- *Monitoring*. The SLM receives and processes resource instrumentation through a monitoring component. Service execution and monitoring of resource utilization—for example,

monitoring the load and utilization of resources and the running states of service components, and detecting faults, may be made possible by using the monitoring services from the general Resource Management capability in OGSA. Such monitoring information is used as input to the Analysis phase.

- *Analysis and projection.* Analysis is performed against the instrumentation to evaluate and determine compliance with the established policy and SLAs. The manager gets to know if resources are meeting QoS objectives and operating within defined policies. Analysis can also predict future resource behavior based on history and projected requirements. When system behavior is not consistent with overall goals, the manager evaluates alternative courses of action to effect changes in the set of configured resources in its sphere of influence, and selects a plan of action in accordance with its configured policies.
- *Action.* The manager then executes the plan, either by interacting with underlying managed resources or by communicating with other managers responsible for other aspects of the system. For example a workload manager may adjust priorities and process shares of tasks executing within a cluster of processors to meet service level objectives and policies. Or, in cases where local action against a pool of resource is either impossible (in violation of policy or constrained) or ineffective, a resource manager may “appeal” to other management functions to remedy the situation. For example, a workload manager might make a provisioning request to add additional processors to a cluster.

This control loop executes continuously, assessing the current state of the system against the expressed service level objectives and making adjustments as necessary to bring the system into compliance. Other services are needed for these activities to be successful—for example *capacity planning, entitlement of resources, problem and root cause analysis, predictive analysis, etc.*

A number of management components may be involved in service level attainment activities. The QoS and SLA requirements addressed by these management activities can be extremely varied. For example, they may include performance attainment objectives such as processor capacity or utilization (workload management), or more qualitative objectives such as security levels. Service level attainment activities affect each other directly or indirectly. Therefore the relationship and order of execution between the different management components in the control loops and service level managers should be fairly loose.

3.8.4.2 Policy and Model based Management

Service level managers orchestrate real-time changes in the dynamic management infrastructure based upon policy that governs their behavior. The intelligence in the self management is basically directed by the policies and models about the SLMs and the services that are being managed.

3.8.4.3 Entitlement

Entitlement is about negotiating with other resource holders (other SLMs, human operators, resource pools etc.) to obtain the right resources. During this process different SLAs are compared to find out which resource need is more important. Entitlement is an earlier phase to resource reservation and provides a looser binding—an option to reserve. In contrast, a resource reservation guarantees access to the resource.

3.8.4.4 Planning

Planning refers to calculating the optimum requirements of a service in terms of the resources it needs. This can be at an initiation stage, due to some problem/root cause being identified, or due to a command from a higher level.

3.8.4.5 Capacity Management

Capacity management includes the actual actions relating to updating the current state and requirement of resources. It includes things such as linking with inventory, asset management, moving the resources from the current location into the service domain, moving the resources from the service domain into a storage area or to another domain, etc.

3.8.4.6 Provisioning

Provisioning, including its sub-activities of deployment and configuration [CDDL], is an important activity that is done in support of self-managing actions as resources are prepared for their expected use. Provisioning is supported by a number of other OGSA services, such as the Application Contents Service [ACS] that maintains the deployment contents and configuration descriptions.

3.8.4.7 Analytics

- *Problem and root cause analysis.* Analyze monitored data to find out the express issue and the root cause of detected problems. Includes filtering capabilities, correlation etc.
- *Predictive analysis.* Analyze monitored data to predict future behavior—to plan for future resource needs, or to predict future problems.

3.8.5 Properties

Service Level Management components implement self-managing, policy-based capabilities across multiple QoS dimensions of the management infrastructure. Key QoS dimensions include, but are not limited to, availability, security, and performance. These cornerstone QoS dimensions are expressed in terms of measurable intent captured in SLAs. Availability might be measured in terms of “minutes of outage.” Desired security capabilities might be described as privileges provided by a class of service for associated users. Performance characteristics could be specified in terms of well-known throughput or response-time objectives.

3.8.5.1 Availability

Specifically denotes the metrics that show the failure rate of the service—of all types, including the macro cases such as system crash, and other, less observable cases, such as the failure of a service to deliver due to throughput drop and buffer flushing.

3.8.5.2 Security

Security is an extremely important aspect of the IT environment. These measures may include security violations, probability of security violations, and identity management metrics such as the time to set up or delete a user, etc.

3.8.5.3 Performance

Gives measurements and metrics—quantifiable data about how the system is performing, such as cumulative CPU load factor, that go into computing the performance metrics of a specified service.

3.8.6 Interactions with Other OGSA Services

Self-management interacts with almost all other aspects of OGSA. As the required interactions are still being worked out, only some services are listed here as an example of what interactions are expected, with no attempt at being exhaustive.

- *Discovery* to find and integrate new resources and services.
- *Logging and monitoring* to provide the information needed to determine the state of the system.
- *Resource reservation* to facilitate more predictable resource usage.
- *Workflow* to automate the actions that the SLMs have to carry out when addressing abnormal conditions.
- *Composition* to construct complex or higher level SLMs from simple ones.
- *Security*, and in particular *Authentication and Authorization*, are essential as different system components may be involved in management actions.
- *Resource management*, and in particular service or resource manageability models, are required to provide the representation of the service or resource that is being managed. The SLMs can read this information, and can act on the intelligence in it, in response to a changing environment or a command from a higher level. Further work to determine how best to model a service or resource needs to be done.

3.9 Information Services

3.9.1 Objectives

The ability to efficiently access and manipulate information about applications, resources and services in the Grid environment is an important OGSA capability. In this section, the term *information* refers to dynamic data or events used for status monitoring; relatively static data used for discovery; and any data that is logged. In practice, an information service needs to support a variety of QoS requirements for reliability, security, and performance. The scope of the OGSA information service covers publication through consumption. Activities prior to publication (e.g. sniffing network packets) or subsequent to consumption are out of scope.

While it may be possible to design one single information service that deals with all information delivery patterns and QoS, OGSA is best served by multiple information services, some general, and some optimized to meet specific use cases. However, caution should be exercised so as not to end up with a number of fractured information services, each capable of answering a limited number of use cases. The challenge is to balance the desire for generality and requirements on domain-specific semantics and QoS. In theory, generality can be achieved by abstracting common (high-level) functionality and features covering the requirements of as many use cases as possible. While common semantics might encourage interoperable implementations they can also be too high-level and not fully expose desired behavior. The question of how far the level of abstraction can be raised, without compromising the usability of the services, is a topic that needs further investigation.

Clients of the OGSA information services include, but are not limited to, execution management services, accounting services, problem determination services, resource reservation services, resource usage services, and application monitoring. To facilitate interoperability and reuse, the information services themselves should be built on top of OGSA infrastructure capabilities such as notification (see §3.3.7). Information services could also make use of other OGSA capabilities such as data access and distributed query processing.

3.9.2 Models

The characterization of an information service depends greatly on factors such as the demand placed on the source of information (e.g., static versus dynamic, publication rate), its purpose (e.g., discovery, logging, monitoring) and QoS requirements. However, we see similar, recurring structures in information services. Information is made available for consumption, either from the originating producer, or through an intermediary (e.g. logging service, notification broker) acting

on behalf of the originating producer. Either one or more consumers wish to obtain information from one or more producers, or one or more producers wish to send information to one or more consumers. Producers and consumers should be decoupled and not be required to have any prior knowledge of each other. Consumers may contact a producer (or intermediary) and pull information in one call or they may use a subscription mechanism to receive information as it becomes available.

OGSA is not prescriptive on the data model used to implement an information service or the language used to query for information. Current systems broadly fall into those that are based on XML and XPath/XQuery query languages (e.g., Globus MDS4¹⁴) and those that use the relational model and the SQL query language (e.g., R-GMA [R-GMA]).

Metadata is associated with information (e.g., events or messages) for describing its structure, properties and usage. For interoperability, a standard event scheme for OGSA information services is desirable. In some cases, such as when performance is paramount and interoperability is not a concern; user-defined, optimized events may be more appropriate.

An information service might allow producers and consumers to discover each other by making detailed descriptions about themselves available for querying. A special distributed registry or point-to-point mechanism could be used for that purpose. The description could include, for example, the type of producer or consumer, what information they produce or consume, and their endpoint URLs.

3.9.3 Example Scenarios

3.9.3.1 Directory Scenario

A user needs to locate a service description that meets some desired functional and other criteria. He queries a central directory service where service descriptions are published, and receives an answer. The directory owner decides who can publish information into the directory and who is authorized to query it. UDDI is an example legacy directory service. OGSA directory services could be based on WSRF service group concepts.

3.9.3.2 Logging Scenario

A number of distributed computing usage scenarios require logging services. These scenarios include those based on problem determination, usage metering, failure recovery, transaction processing, and security.

A problem determination scenario might proceed as follows. A developer is writing code for a Disaster Recovery application. Following corporate programming guidelines for logging, he inserts log statements into his code to provide information regarding any unexpected situations. He is responsible for coding a section of the application that accesses an RDBMS to obtain attributes for candidate fail-over resources. He codes this section of the application such that RDBMS failures are trapped and log records are generated to reference the RDBMS failure. After development has been completed, the application is deployed into an operating environment where log records are processed by a handler chain. In this environment, based on the severity level (e.g., fatal, warning, informational), records exceeding an operator-specified level are stored in a log by a file handler. Due to the deployment of an erroneous new security policy, the code reaches an abnormal state and logs a record indicating that the underlying RDBMS has denied access to the application. The operator had configured the operating environment to store log records of this severity. After the failure, an analyst, in the role of a log consumer, uses a console

¹⁴ See http://www.globus.org/toolkit/mds/#mds_gt4

application to peruse the logs. He finds the log statement indicating the RDBMS problem and directs the database administrator to correct the security policy.

3.9.3.2.1 *Grid Monitoring Architecture (GMA) Scenario*

A scientist wants to run an interactive simulation/rendering job that must be finished within the next half hour. He must find 100 computing elements that are fast enough, have enough memory between them, and are linked by fast network connections. He needs up-to-date information and submits a complex query to the OGSA information service. After a short time he receives the information on the computing resources that satisfy his query. From the set of possible candidate resources, he selects the nearest ones and proceeds with his job.

In answering the query, the information service first discovers the information producers for the computing, storage and network resources that are relevant and then gets the necessary information before performing a “join” to identify the resources that satisfy the query. The scientist then needs to monitor his job and respond to partial failures in the execution environment, possibly using the same GMA-based service (provided that the information has an associated timestamp). GMA is an example of a multi-purpose information services architecture [GMA].

3.9.3.3 *Producer/Consumer Patterns*

Frequently, producers (or publishers) and consumers can directly communicate with one another. For those cases where a direct exchange between a producer and consumer is not appropriate or not possible, the basic pattern of decoupling producers from consumers using an intermediary is widely used. In the Producer-Intermediary-Consumer pattern, producers put data *into* an intermediary, and consumers extract data *from* it. In a general sense, this is the pattern followed by any data store. That is, producers write to and consumers read from a file, RDBMS, ODBMS etc. In addition to supporting producer and consumer interfaces, an intermediary may also support a management interface. The management interface controls those functions that are not directly associated with reading or writing to the data store. Operations controlling retention policy, backup policy, etc. would be accessed through a management interface.

In distributed computing, the above pattern is followed by several infrastructure services whose task it is to provide information to consumers. We broadly refer to these services as information services. They include: logging services, and notification services. We assume that the reader is generally familiar with these concepts and we defer OGSA-specific discussions on each service.

Each of these services has evolved to support a distinct set of domain-specific semantics and qualities of service. While they each could be implemented using a general-purpose RDBMS for an underlying data store, this would most likely result in trading off some desirable, specialized qualities of service for unneeded functionality. For example, a log service that must support fast writes would most likely find the performance and footprint costs of a full-function RDBMS to be prohibitive. Similar observations could be made about notification services.

3.9.3.4 *Publish/Subscribe Pattern*

In this pattern publishers disseminate information to consumers without any prior knowledge about them. Subscribers specify which published messages are of interest to them. For example, assuming the WSRF Basic Profile [OGSA WSRF BP], the WS-Notification specifications (see also §3.3.7) provide a core set of interfaces that could be leveraged by the OGSA information services. Features include a publish/subscribe model, a mechanism to organize items of interest to a subscriber known as “Topics”, the ability to filter messages based on the properties of the publisher or the message content, and management interfaces for publications and subscriptions.

3.9.3.5 *Subscribe/Publish Pattern*

In this pattern messages are created in response to a subscription. This allows the subscriber to choose what gets published in the first instance. For example, information dissemination (INFOD)¹⁵ allows subscribers not only to influence the selection of publishers but also to specify what is considered to be an event and what message gets published in response to an event. Publishers in INFOD are able to select which consumers are eligible as well as appropriate consumers for each message. INFOD complements WS-Notification by allowing more control on what information gets published in the first instance, and adds additional mutual filtering between publishers and consumers.

3.9.4 **Functional Capabilities**

We define discovery, message delivery, logging, and monitoring capabilities.

3.9.4.1 *Discovery*

One universally needed capability is service and resource discovery. A directory (or registry) is an obvious solution, but not the only one. A directory is distinguished from other possible solutions in that it has persistent storage for the “latest” information and is optimized for searches. Low latency response to a high volume of queries is required. The directory may be replicated for scalability.

Alternatively, a compilation or guide of information can be stored in an *index* (such as Google). Unlike a registry, which tends to be centrally controlled, anyone can create an index.

Another alternative is peer-to-peer discovery, where a web service is a node in a network of peers and dynamically queries its neighbors in search of a suitable match. The query propagates through the network from one node to another until a match is found, a particular hop count is reached, or some other termination criterion is satisfied.

Yet another alternative is dynamic announcement and discovery using multicast protocols in ad-hoc networks where no registry or index is available. Services and resources announce their arrival and departure from the network. To locate services and resources on the network messages are sent to a multicast group and group members that match the request return a response directly to the requestor. The ability to discover services and resources dynamically, assuming the existence of the network only, is very important as it enables the bootstrapping of their management.

3.9.4.2 *Message Delivery*

Producers and consumers interact by exchanging messages, and this can be handled by a common messaging infrastructure. This infrastructure is only concerned with how to distribute copies of messages to interested parties, not how these messages are constructed in the first place. Producers either send messages directly to relevant consumers or make use of an intermediary (message broker) that decouples producers from consumers. In the latter case the producers publish their messages to the broker, which takes the responsibility for forwarding the message to interested parties. A producer (or the intermediary) may provide notification capabilities and additional function such as a finder service (allowing its producers and consumers to find each other). Message brokers may store and forward messages in stable storage—not necessarily for persistency, but for reliable message delivery purposes.

¹⁵ See <https://forge.gridforum.org/projects/infod-wg/>

3.9.4.3 Logging

An OGSA logging service keeps log records in a persistent store for a period of time. It acts as an intermediary between log artifact producers and consumers. Producers write log artifacts sequentially, and consumers may read (but not update) the log records. This message exchange is optimized for performance. To ensure the general acceptance of the basic log semantics and to enable the exploitation of existing implementations, OGSA logging services should support key features found in existing logging implementations. There may be multiple producers and consumers for a given logging intermediary, and both may set filters for records.

3.9.4.4 Monitoring

Information that carries a field for ordering purposes (e.g., a time stamp and sequence number) can be used for monitoring. An OGSA monitoring service could be equally used for applications or resources. Some situations (e.g., real-time applications) might impose strict requirement on the monitoring service (e.g., high update rates and high performance). In such a case a special-purpose service might be needed.

3.9.4.5 General Information and Monitoring Service

A general OGSA service that provides a combination of the above capabilities can provide more flexibility to the end user. For Grid resources in general (including services and applications), the amount of available information about resources could be large, dispersed across the network, and updated frequently. Searches in this space may have unacceptable latencies. In order to manage such information in a controllable way, it is important to separate information source discovery from information delivery. Searches should only be used to locate information sources or sinks. A special-purpose directory is used to hold metadata about the resources. In this case the directory must cope with high rates of updates that are expected in the dynamic OGSA environment. Individual producer/consumer pairs can limit the amount of data flowing between them to that satisfying the consumer query. This model differs from a message broker that combines the mechanisms for finding sources and sinks of information and its delivery into a single searchable channel. The merits of this approach are described in the GMA document [GMA].

A user of such a general service should be able to put any information, irrespective of its intended use (e.g., discovery, monitoring), into it without needing to understand the complexity of the system. He first must specify what information is to be made available in OGSA, where the type and structure of that information should be well-defined. The user might also wish to specify certain properties and policies. This could include how the information is held, retention period, guarantee of delivery, persistency, and access control. A consumer could filter information of interest using a subscription topic, for example, and it could support a query to further refine the events delivered to it through predicates defined in the query expression.

More advanced users may require a deeper understanding of the internal workings of the service. Following a request for information in a general (GMA-based) service, expected behavior is that a “mediator” capability is used to perform a registry/schema lookup and locate suitable sources of information. For long-term queries the mediator ensures that, as sources are dropped or new relevant sources come online, the subscribed consumers are updated. Mediation is also concerned with planning the distributed queries to the relevant producers, as well as merging the results.

3.9.5 Properties

3.9.5.1 Security

Authentication and authorization rules for consumers and producers allow them to exchange information in a secure fashion. Discovering metadata information about producers and consumers (e.g., their existence or the type of information they produce or consume) could also be subjected to security rules. Some services (e.g., metering, authentication, authorization) require that messages be made secure (e.g., encrypted) for delivery.

3.9.5.2 Quality of Service

Various levels of QoS can be provided by OGSA information services, such as reliable, guaranteed delivery. WS-Reliability, defined by the OASIS WS-RM TC, provides these reliable message delivery properties [WS-Reliability].

3.9.5.3 Availability, Performance, and Scalability

Information systems play a critical role in OGSA. Since almost every other capability in OGSA makes use of them, they need to be available at all times and to be especially tolerant of partial failure. Many clients of the information systems expect to receive information at high rates and cannot afford to wait long periods of time. High-performance systems are needed in this case. Because a large number of resources, services and applications may wish to produce and consume information, the system must also be scalable across wide-area networks.

3.9.6 Interactions with Other OGSA Services

Standard event data models facilitate the transfer of information from producers to consumers. The use of recognized standards event schema permits the discovery and processing of events from a large variety of potentially different sources across a VO. Ideally, to avoid mediation costs, the resource event model and the infrastructure event model should be closely aligned. Also, to accommodate domain-specific events, event schemas should be extensible. For example, assuming the WSRF Basic Profile, the event schema developed in the OASIS WSDM-TC, which also accommodates the existing CIM event model, may be used as the basis for the OGSA event data model.

Notification mechanisms, or an extension of them, could be used to deliver events from producers to consumers once they have discovered each other, and between other components of the system.

Security services are needed for authentication/authorization between different components.

Data Federation is required in the query-planning phase of the mediation between producers and consumers of information in GMA.

Replication mechanisms are needed so that repositories of metadata (directories or registries) can be distributed and replicated to avoid single points of failure and improve scalability. Some temporary inconsistency between replicated copies might be acceptable in some situations; in this case the information system needs to be robust in the event of inconsistent or out-of-date metadata.

4 Security Considerations

Security considerations are discussed in the Security services section (§3.7).

Editor Information

Ian Foster
Distributed Systems Laboratory
Mathematics and Computer Science Division
Argonne National Laboratory
Argonne, IL 60439
Email: foster@mcs.anl.gov

Hiro Kishimoto
Information Technology Core Laboratories
Fujitsu Laboratories
4-1-1, Kamikodanaka, Nakahara, Kawasaki City, Japan
Email: hiro.kishimoto@jp.fujitsu.com

Andreas Savva
IT Architecture Lab.
Fujitsu Laboratories
4-1-1, Kamikodanaka, Nakahara, Kawasaki City, Japan
Email: andreas.savva@jp.fujitsu.com

Contributors

We gratefully acknowledge the contributions made to this document by Jeffrey Frey, Takuya Mori, Jeffrey Nick, Chris Smith, David Snelling, Latha Srinivasan, and Jay Unger.

Acknowledgements

We are grateful to numerous colleagues for discussions on the topics covered in this document, in particular (in alphabetical order, with apologies to anybody we have missed) Mario Antonioletti, Takuya Araki, Michael Behrens, Jamie Bernardin, Shel Finkelstein, Steve Fisher, Dennis Gannon, Kate Keahey, Carl Kesselman, Takashi Kojo, Peter Kunszt, Simon Laws, Tan Lu, James Magowan, Susan Malaika, Tom Maguire, David Martin, Nataraj Nagaratnam, Steven Newhouse, Dave Pearson, Benny Rochwerger, John Rofrano, Ellen Stokes, Paul Taylor, Steve Tuecke, Sachiko Wada, James Warnes, Philipp Wieder, and Ming Xu.

We would also like to thank the people who took the time to read and comment on earlier drafts. Their comments were valuable in helping us improve the readability and accuracy of this document.

Glossary

The glossary is provided as a companion document, the OGSA Glossary of Terms [OGSA Glossary].

Intellectual Property Statement

The OGF takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the OGF Secretariat.

The OGF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this recommendation. Please address the information to the OGF Executive Director.

Disclaimer

This document and the information contained herein is provided on an “As Is” basis and the OGF disclaims all warranties, express or implied, including but not limited to any warranty that the use of the information herein will not infringe any rights or any implied warranties of merchantability or fitness for a particular purpose.

Full Copyright Notice

Copyright © Open Grid Forum (2002-2006). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the OGF or other organizations, except as needed for the purpose of developing Grid Recommendations in which case the procedures for copyrights defined in the OGF Document process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the OGF or its successors or assignees.

References

- [ACS] Fukui, K. (ed.) Application Contents Service Specification 1.0, Global Grid Forum, Lemont, Illinois, U.S.A., GWD-R, November 2005.
http://www.ggf.org/Public_Comment_Docs/Documents/Feb-2006/draft-ggf-acs-spec-1.0.pdf
- [ByteIO] Morgan, M. (ed.) ByteIO Specification 1.0. Global Grid Forum, Lemont, Illinois, U.S.A., GWD-R, October 2005.
http://www.ggf.org/Public_Comment_Docs/Documents/Feb-2006/draft-byteio-rec-doc-v1-1.pdf

- [CAS] Pearlman, L., Welch, V., Foster, I., Kesselman, C., Tuecke, S. A Community Authorization Service for Group Collaboration. *Proceedings of the IEEE 3rd International Workshop on Policies for Distributed Systems and Networks*, 2001.
- [Condor] Tannenbaum, T., Wright, D., Miller, K., Livny, M., “Condor—A Distributed Job Scheduler” in Sterling, T. (ed) *Beowulf Cluster Computing with Linux*, The MIT Press, 2002.
- [CDDL] D. Bell, T. Kojo, P. Goldsack, S. Loughran, D. Milojicic, S. Schaefer, J. Tatemura, and P. Toft, Configuration Description, Deployment, and Lifecycle Management (CDDL) Foundation Document, Global Grid Forum, Lemont, Illinois, U.S.A., GFD-I.050, August 2005. <http://www.ggf.org/documents/final.htm>
- [Dynamic Access Control] Lepro, R., Cardea: Dynamic Access Control in Distributed Systems. *NAS Technical Report NAS-03-020*, 2003.
- [DAIS] Atkinson, M., Dialani, D., Guy, L., Narang, I., Paton, N., Pearson, D., Storey, T., Watson, P., “Grid Database Access and Integration: Requirements and Functionalities”, Global Grid Forum, Lemont, Illinois, U.S.A., GFD-I.13, March 2003. <http://www.ggf.org/documents/final.htm>
- [VOMS-1.1] EU DataGrid, VOMS Architecture version 1.1. 2003. http://grid-auth.infn.it/docs/VOMS-v1_1.pdf
- [Fine-Grain Auth] Keahey, K., Welch, V., Lang, S., Liu, B., Meder, S. Fine-Grain Authorization Policies in the GRID: Design and Implementation. *1st International Workshop on Middleware for Grid Computing*, 2003.
- [Fine-Grain Auth RM] Keahey, K., Welch, V. Fine-Grain Authorization for Resource Management in the Grid Environment. *Proceedings of Grid2002 Workshop*, 2002.
- [Globus MDS] K. Czajkowski, S. Fitzgerald, I. Foster, C. Kesselman. Grid Information Services for Distributed Resource Sharing. *Proceedings of the Tenth IEEE International Symposium on High-Performance Distributed Computing (HPDC-10)*, IEEE Press, August 2001.
- [Grid Auth Framework] Lorch, M., Cowles, B. (eds.) Conceptual Grid Authorization Framework and Classification. GGF Working Group on Authorization Frameworks and Mechanisms, 2004.
- [Grid AAA Req] Mullen, S. Crawford, M., Lorch, M, Skow, D. Grid Authentication Authorization and Accounting Requirements. GGF Site Authentication, Authorization, and Accounting working group. January, 2004.
- [Grid Anatomy] 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.
- [Grid Physiology] Foster, I., Kesselman, C., Nick, J. and Tuecke, S. The Physiology of the Grid: An Open Grid Services Architecture for Distributed Systems Integration. Globus Project, 2002. <http://www.globus.org/alliance/publications/papers/ogsa.pdf>

- [GS Security] Welch, V., et. al. Security for Grid Services, *Twelfth International Symposium on High Performance Distributed Computing (HPDC-12)*, IEEE Press, to appear June 2003.
- [GMA] B. Tierney, R. Aydt, D. Gunter, W. Smith, M. Swany, V. Taylor and R. Wolski, *A Grid Monitoring Architecture*, Global Grid Forum, Lemont, Illinois, U.S.A., GFD.I.7, January 2002. <http://www.ggf.org/documents/final.htm>
- [JobQueue] Katramatos, D., Humphrey, M., Grimshaw, G., and Chapin, S. J. JobQueue: A Computational Grid-Wide Queueing System. Proceedings of the Second International Workshop on Grid Computing. Lecture Notes in CS, 2001.
- [JSDL] Anjomshoaa, A., Brisard, F., Drescher, M., Fellows, D., Ly, A., McGough, S., Pulsipher, D., and Savva, A. (ed.) Job Submission Description Language (JSDL) Specification, Version 1.0, Global Grid Forum, Lemont, Illinois, U.S.A., GFD-R-P.056, November 2005. <http://www.ggf.org/documents/final.htm>
- [Kerberos V5] Kohl, J., and Neuman, C. The Kerberos Network Authentication Service (V5), RFC 1510, IETF, 1993.
- [Legion] S.J. Chapin, D. Katramatos, J.F. Karpovich, and A.S. Grimshaw, "Resource Management in Legion," *Journal of Future Generation Computing Systems*, vol. 15, pp. 583-594, 1999.
- [Liberty] Introduction to the Liberty Alliance Identity Architecture. Revision 1.0, Liberty Alliance Project. <http://www.projectliberty.org/>, March 2003.
- [MOWS] Sedukhin, I. (ed.) Web Services Distributed Management: Management of Web Services (WSDM-MOWS) 1.0, OASIS Standard, March 2005, <http://docs.oasis-open.org/wsdm/2004/12/wsdm-mows-1.0.pdf>
- [MUWS part 1] Vambenepe, W. (ed.) Web Services Distributed Management: Management using Web Services (MUWS 1.0) Part 1, OASIS Standard, March 2005, <http://docs.oasis-open.org/wsdm/2004/12/wsdm-muws-part1-1.0.pdf>
- [MUWS part 2] Vambenepe, W. (ed.) Web Services Distributed Management: Management using Web Services (MUWS 1.0) Part 2, OASIS Standard, March 2005, <http://docs.oasis-open.org/wsdm/2004/12/wsdm-muws-part2-1.0.pdf>
- [Nimrod] Abramson, D., Giddy, J. and Kotler, L., High Performance Parametric Modeling with Nimrod/G: Killer Application for the Global Grid? Proceedings of the International Parallel and Distributed Processing Symposium (IPDPS), Cancun, Mexico, 2000, 520-528.
- [OGSA BES] Grimshaw, A., Newhouse, N., Pulsipher, D., Morgan, M. OGSA Basic Execution Service Version 1.0, Global Grid Forum OGSA-BES-WG, Draft 16, February 2006. <https://forge.gridforum.org/projects/ogsa-bes-wg/document/ogsa-bes-draft-v16/en/1>
- [OGSA Glossary] Treadwell, J. (ed.) Open Grid Services Architecture Glossary of Terms Version 1.5. Global Grid Forum, Lemont, Illinois, U.S.A., GFD-I.081, July 2006. <http://www.ggf.org/documents/final.htm>

- [OGSA Profile Definition] Maguire, T., Snelling, D. OGSA Profile Definition. Global Grid Forum, Lemont, Illinois, U.S.A., GFD-I.059, January 2006.
<http://www.ggf.org/documents/final.htm>
- [OGSA RM] Maciel, F. B. (ed.) Resource Management in OGSA. Global Grid Forum, Lemont, Illinois, U.S.A., GFD-I.045, 2005, <http://www.ggf.org/documents/final.htm>
- [OGSA Roadmap] Kishimoto, H., and Treadwell, J. (eds.) Defining the Grid: A Roadmap for OGSA™ Standards 1.0. Global Grid Forum, Lemont, Illinois, USA, GFD-I.053, September 2005. <http://www.ggf.org/documents/final.htm>
- [OGSA BSP—Core] Mori, T., and Siebenlist, F., OGSA Basic Security Profile 1.0— Core. Global Grid Forum, Lemont, Illinois, U.S.A., GWD-R, February 2006.
http://www.ggf.org/Public_Comment_Docs/Documents/Apr-2006/draft-ggf-ogsa-basic-security-profile-core-08.pdf
- [OGSA BSP—Secure Channel] Mori, T., and Siebenlist, F., OGSA Basic Security Profile 1.0— Secure Channel. Global Grid Forum, Lemont, Illinois, U.S.A., GWD-R, February 2006.
http://www.ggf.org/Public_Comment_Docs/Documents/Apr-2006/draft-ggf-ogsa-basic-security-profile-secure-channel-19.pdf
- [OGSA Use Cases] Foster, I., Gannon, D., Kishimoto, H and J. Von Reich, Jeffrin. (eds.) Open Grid Services Architecture Use Cases. Global Grid Forum OGSA-WG, GFD-I.029, October 2004. <http://www.ggf.org/documents/final.htm>
- [OGSA Use Cases Tier 2] J. Von Reich, Jeffrin. (ed.) Open Grid Services Architecture: Second Tier Use Cases. Global Grid Forum OGSA-WG, Draft, March 2004.
<https://forge.gridforum.org/projects/ogsa-wg/document/draft-ggf-ogsa-usecase-tier2-20/en/1>
- [OGSA WSRF BP] Foster, I., Maguire, T., and Snelling, D. OGSA WSRF Basic Profile 1.0. Global Grid Forum, Lemont, Illinois, U.S.A., GWD-R, September 2005.
http://www.ggf.org/Public_Comment_Docs/Documents/Oct-2005/draft-ggf-ogsa-wsrf-basic-profile-v43.pdf
- [RFC3280] Housley, R., Polk, W., Ford, W., and Solo, D., Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile. RFC 3280, IETF, April 2002.
- [RFC2903] de Latt, C., Gross, G., Gommans, L., Vollbrecht, J., Spence, D. Generic AAA Architecture, RFC 2903, IETF, August 2000.
- [R-GMA] Cooke A, Gray AJG, Ma LS, Nutt W, Magowan J, Oevers M, Taylor P, Byrom R, Field L, Hicks S, Leake J, Soni M, Wilson A, Cordenonsi R, Cornwall L, Djaoui A, Fisher S, Podhorszki N, Coghlan B, Kenny S, O'Callaghan D, *R-GMA: An information integration system for grid monitoring*. Lecture Notes in Computer Science 2888: 462-481 2003
- [PRIMA] Lorch, M., Adams, D., Kafura, D., Koneni, M., Rathi, A., and Shah, S. The PRIMA System for Privilege Management, Authorization and Enforcement in Grid Environments, *4th Int. Workshop on Grid Computing - Grid 2003*, 17 November 2003 in Phoenix, AR, US

- [Role-Based VO] Canon, S., Chan, S., Olson, D., Tull, C., Welch, V. Using CAS to Manage Role-Based VO Sub-Groups. *Proceedings of Computing in High Energy Physics (CHEP '03)*, 2003.
- [SAZ] Sekhri, V. and Mandrichenko, I. Site Authorization Service (SAZ) at Fermilab. *Proceedings of Computing in High Energy Physics (CHEP '03)*, 2003.
- [VO Security] Siebenlist, F., Nagaratnam, N., Welch, V., Neuman, B.C. Security for Virtual Organizations: Federating Trust and Policy Domains. In *The Grid: Blueprint for a New Computing Infrastructure (2nd Edition)*, 2004.
- [WS-Addressing] Gudgin, M., Hadley, M., Rogers, T. (eds.). Web Services Addressing 1.0 – Core (WS-Addressing), W3C Recommendation, 9 May 2006. <http://www.w3.org/TR/2006/REC-ws-addr-core-20060509>
- [WS-Agreement] Andrieux, A., Czajkowski, K., Dan, A., Keahey, K., Ludwig, H., Nakata, T., Pruyne, J., Rofrano, J., Tuecke, S. and Xu, M. Web Services Agreement Specification (WS-Agreement). Global Grid Forum, Lemont, Illinois, U.S.A., GWD-R, September 2005. http://www.ggf.org/Public_Comment_Docs/Documents/Oct-2005/WS-AgreementSpecificationDraft050920.pdf
- [WS-Architecture] Booth, D., Haas, H., McCabe, F., Newcomer, E., Champion, M., Ferris, C. and Orchard, D. Web Services Architecture. W3C, Working Group Note, February 2004. <http://www.w3.org/TR/2003/WD-ws-arch-20040211/>
- [WS-I BP 1.1] K. Ballinger, D. Ehnebuske, C. Ferris, M. Gudgin, C.K. Liu, M. Nottingham, and P. Yendluri (ed.) Basic Profile Version 1.1, Web Services Interoperability Organization Final Material, 24 August 2004. <http://www.ws-i.org/Profiles/BasicProfile-1.1.html>
- [WS-I BSP 1.0] A. Barbir, M. Gudgin, M. McIntosh, and K.S. Morrison (ed.) Basic Security Profile Version 1.0, Web Services Interoperability Organization, Working Group Draft, 29 August 2005. <http://www.ws-i.org/Profiles/BasicSecurityProfile-1.0-2005-08-29.html>
- [WS-Naming] Grimshaw, A. (ed). WS-Naming Specification. Draft, March 2006. <http://forge.gridforum.org/projects/ogsa-naming-wg/>
- [WS-Reliability] Iwasa, K. (ed.) WS-Reliability 1.1, OASIS standard, November 2004. <http://www.oasis-open.org/committees/download.php/9330/WS-Reliability-CD1.086.zip>
- [WS-Resource] Graham, S., Karmakar, A., Mischkinsky, J., Robinson, I., Sedukhin, I. Web Services Resource, Version 1.2, OASIS Standard, April 2006. http://docs.oasis-open.org/wsr/wsr/wsr-ws_resource-1.2-spec-os.pdf