GlossaryOfTerms

Documentation - Glossary Of Terms

Documentation

DocumentClassification	
DocumentType	

Table of Contents

1 Glossary of Terms

A-B-C-D-E-F-G-H-I-J-K-L-M-N-O-P-Q-R-S-T-U-V-W-X-Y-Z

1.1 A

A-B-C-D-E-F-G-H-I-J-K-L-M-N-O-P-Q-R-S-T-U-V-W-X-Y-Z

Access Anticipation Method
> Tries to foresee the I/O access characteristics of the application based on programmer's hints, anticipated knowledge or reaction to identified behavior.

Accounting (grid accounting)
> The OSG accounting system, called Gratia, tracks VO members' resource usage and presents that information in a consistent Grid-wide view, focusing in particular on CPU and Disk Storage utilization.

Administrative Domain
> One or more Resource Groups run under a single set of Policies; and, typically, run by a single team.

Agent
> A software component in OSG that operates on behalf of a User or Resource Owner or another Agent.

Alliance
> A collaboration of small application communities that develops systems and runs on the persistent grid infrastructure.

Application
> With respect to grid computing in general, "application" refers to a "layer" of grid components (above infrastructure and resources). An application is a name used to identify a set of software that will execute computational jobs, manage data (access, store, read...) and has many attributes. Any application when invoked (executed) includes information that allows tracing back to the individual who is responsible for the execution.

Application administrator
> A person designated by a VO who is charged with making sure that a particular application works on the participating grid resources.

Application community
> Providers of a particular end-to-end application and/or system that runs on the persistent grid infrastructure. Smaller application communities may collaborate as Alliances in developing systems and running them on the grid. These organizations will contribute by providing application requirements and interfaces to the grid services. An application community may consist of or span multiple VOs or VO groups.

Application layer
> Final of five layers to support grid applications. It includes all application-specific services, including Request Interpretation and part of Request Management.

Application Level Method
> Organizes the mappings of the application's main memory objects to respective disk objects to make disk accesses more efficient by exploiting data locality.

Application middleware
> Application Middleware is application-specific middleware which has some embedded capabilities and interfaces that are not general, e.g., information providers. This middleware depends on grid-wide middleware.

Application Programming Interface
> An Application Programming Interface (API) is a particular set of rules and specifications that a software program can follow to access and make use of the services and resources provided by another particular software program that implements that API. It serves as an interface between different software programs and facilitates their interaction, similar to the way the user interface facilitates interaction between humans and computers.

ARDA
> Architectural Roadmap towards Distributed Analysis (ARDA) is a project within LCG that seeks to coordinate the prototyping of distributed analysis systems for the LHC experiments. *See* http://lcg.web.cern.ch/LCG/activities/arda/arda.html.

Auditing (grid auditing)
> Grid auditing in OSG relates to resolving claims of challenged authentication and exposed risk on grid services which accept delegated credentials. The auditing system will use information from the accounting system and link it to information from other sources to allow full tracking and analysis of the actions and events related to a user's resource usage.

AUP
> Acceptable Use Policy. For example, *see* OSG User Acceptable Use Policy.

1.2 B

A-B-C-D-E-F-G-H-I-J-K-L-M-N-O-P-Q-R-S-T-U-V-W-X-Y-Z

BDII
> Berkeley Database Information Index (BDII) is an LCG implementation of Globus GIIS-like information index based on the Berkeley Database.

BeStMan
> Berkeley Storage Manager (BeStMan) is an Lawrence Berkeley National Laboratory (LBNL) implementation of Storage Resource Manager (SRM) based on Unix disks and the High Performance Storage System HPSS. *See also* https://sdm.lbl.gov/bestman/.

Brokering
> In the context of a Storage Resource Manager (SRM), brokering is contacting the source location to pin a file and then invoke a transfer service to get the file when the file is not locally available.

1.3 C

A-B-C-D-E-F-G-H-I-J-K-L-M-N-O-P-Q-R-S-T-U-V-W-X-Y-Z

CA
> *See* Certificate Authority.

Campus Grid
> A Grid operated within the context of a single Facility (such as a university of a national lab).

CE
> *See* Compute Element.

CEMon
> The CEMon service provides a common interface for publishing information about a Compute Element in your network. The framework can be configured with multiple sensors for collecting different kinds of data, such as the Generic Information Provider (GIP). The information can be published in multiple formats, such as the LDAP Data Interchange Format LDIF and Condor ClassAds. In OSG it is an important part if the Information Service (IS): the CEMons running on the resources collect information from the GIP and push it to the Resource Selection Service (ReSS) and to the CEMon Consumer at the Grid Operations Center GOC that in turn publishes the information to the both the OSG BDII and the WLCG Interoperability BDII. *See also* About CEMon.

CE-SE binding
> Job scheduling often requires both a compute element (CE) to run the job and a Storage Element (SE) to provide for an input or output storage extent. Currently , there are static relationships between individual CEs and SEs set by Site Admins. The CE-SE bind schema (part of GLUE schema) aims at providing the means for publishing such a relationship with eventual per-pair data. At the moment, the published information is limited to the local mount point on the CE pointing to the SE's storage space.

CE Storage
> Disk spaces of Storage Elements accessible from within a Compute Element.

Consumer
> A User or Agent who makes use of an available Resource or Agent or Service.

CPU Hour
> Used to measure processing delivered by Compute Elements. A Computer Processing Unit (CPU) hour is one CPU executing a job for one hour of wall clock time. To get a comparative measure, the hours must be normalized by the relative processing speeds of the CPUs.

Certificate
> A public-key certificate is a digitally signed statement from one entity (e.g., a certificate authority), saying that the public key (and some other information) of another entity (e.g., the grid user) has some specific value. The X.509 standard defines what information can go into a certificate, and describes how to write it down (the data format). Read more at CertificateWhatIs.

Certificate Revocation List (CRL)
> A **C**ertificate **R**evocation **L**ist is a list of Certificates that have been revoked and should not be relied upon. The list enumerates revoked certificates along with the reason(s) for revocation. The issuer and the issue date are also included. In addition, each list contains a proposed date for its next release. When a potential user attempts to access a server, the server allows or denies access based on the CRL entry for that particular user. For more information consult Wikipedia.

Certificate Authority
> An entity that issues certificates for use by other parties. The OSG recognizes certificates issued by a number of certificate authorities. *See* CertificateWhatIs. For more background, consult Wikipedia.

Cloud
> A set of Services, Providers, Resources and Policies providing a single point of access for all the computing needs of Consumers. The resources are not necessarily owned by the consumer, but may be leased or otherwise accessed.

Cluster
> A networked group of worker nodes (plus head node, if applicable) at a site. In the GLUE schema, a cluster is a container that groups together subclusters, or computer nodes. A cluster may be referenced by more then one computing element (CE).

Community (Cyber-)Infrastructure
> A set of services and software that has been established by a community to meet the needs of its members. The management of the distributed infrastructure is the responsibility of the community; and the resources are all, or nearly all, owned by the Virtual Organization and its members.

Compute Element (CE)
> Compute element is a term used in Grids to denote any kind of computing interface, e.g., a job entry or batch system. A compute element consists of one or more similar machines, managed by a single scheduler/job queue, which is set up to accept and run grid jobs. The machines do not need to be identical, but must have the same OS and the same processor architecture. In OSG, the CE runs the bulk of the OSG software stack. *See* Site Planning.

Compute Hours
> *See* CPU Hour

CRL
> *See* Certificate Revocation List.

CVMFS
> CernVM? File System (CVMFS) is a network file system developed by CERN, based on HTTP and optimized to deliver experiment software in a fast, scalable, and reliable way.

1.4 D

A-B-C-D-E-F-G-H-I-J-K-L-M-N-O-P-Q-R-S-T-U-V-W-X-Y-Z

DAG
> A Directed Acyclic Graph (DAG) is data structure used to represent dependencies. Each job is a node in a DAG and each node can have a number of parent or children nodes but no loops. A DAG is

defined by a .dag file, listing each of its nodes and their dependencies.

DAGman

Directed Acyclic Graph (DAG) Manager. A Grid middleware component to manage interdependent Grid jobs, provided by the Condor team. It allows you to specify dependencies between your Condor jobs (e.g., don't run B until A has completed successfully). It acts as a "meta-scheduler" managing the submission of your jobs to Condor, based on the DAG dependencies. DAGman is packaged as part of the Virtual Data Toolkit (VDT). Read more at http://www.cs.wisc.edu/condor/dagman/.

Data federation

The process of bringing data together in a single virtual location for on-demand application access.

Data Transfer Protocol

Protocols to to reliably transfer data in HPC and Grid environments.

dCache

A specific data file caching system that acts as an intelligent manager between the user and the data storage facilities. It optimizes the location of staged copies according to an access profile. It decouples the (potentially slow) network transfer rate from the (fast) storage media I/O rate in order to keep the mass storage system from bogging down. *See* http://www.dcache.org/.

Distinguished Name (DN)

The unique name of the entity whose public key the certificate identifies. The DN includes the subject's Common Name (first and last name), Organizational Unit (e.g., institution), Organization (e.g., VO), and 2-character Country code (and optionally locality and state or province).

Disk Cache

A disk cache can be viewed as an intermediate "relay station" between client applications and a Hierarchical Storage Manager (HSM). It decouples the potentially slow network transfer (to and from client machines) from the fast storage media I/O in order to keep the HSM from bogging down.

Durable file

A durable file has the behavior of a volatile file in that it has a lifetime associated with it, but also the behavior of a permanent file in that when the lifetime expires the file is not automatically eligible for removal.

Durable storage

Storage in which a "durable" physical replica of a file can only be created and removed by the administrator of the disk cache. A durable file is designated to stay in cache by the disk cache administrator (i.e,. not subject to dynamic removal) until he or she decides to change its status to "volatile".

Disk Resource Manager (DRM)

An SRM-compliant implementation on a disk-based unix file system. *See also:* SRM; TRM; HRM.

1.5 E

A-B-C-D-E-F-G-H-I-J-K-L-M-N-O-P-Q-R-S-T-U-V-W-X-Y-Z

EGEE
> Note: This project has ended and is replaced by EGI. The Enabling Grids for E-sciencE (EGEE) project was funded by the European Commission and aimed to provide a seamless Grid infrastructure for e-Science available to scientists 24 hours-a-day. *See* http://www.eu-egee.org/.

EGI
> The European Grid Infrastructure enables access to computing resources for European researchers from all fields of science, from High Energy Physics to Humanities. *See* http://www.egi.eu/.

1.6 F

A-B-C-D-E-F-G-H-I-J-K-L-M-N-O-P-Q-R-S-T-U-V-W-X-Y-Z

Fabric
> With respect to grid computing, "fabric" refers to a "layer" of grid components (underneath applications, tools, and middleware). Fabric encompasses local resource managers (e.g., operating systems, queuing systems, device drivers, libraries, etc.), and networked resources (e.g., compute and storage resources, data sources, etc.). Grid applications use grid tools and middleware components to interact with the fabric.

Fabric layer
> The first of five layers to support grid applications. It includes computing, storage and network resources; catalogs; code repositories; etc.

Facility
> A "facility" is a logical name denoting a collection of one or more Sites under the same administrative domain. Facilities are anticipated to provide other grid services - especially, but not exclusively, services in support of the major applications running for whom they provide computational resources.

Factory Type
> Factory Type is the name of the Job Manager in case of Globus GRAM-WS. A non-complete list of factory types is `Fork`, `ManagedFork`, `Condor`, `SGE`, `PBS`, `LSF`, `CCS`.

Federation Peer
> A level coupling of resources within OSG or of external grids to OSG. A federated union of resources or grids enables jobs to migrate between them.

Fully Qualified Attribute Name (FQAN)
> VOMS extended X.509 Attribute Certificate specification for defining extra attributes, based on RFC3281. FQANs contain Role and VO membership information for a Grid user.

1.7 G

A-B-C-D-E-F-G-H-I-J-K-L-M-N-O-P-Q-R-S-T-U-V-W-X-Y-Z

Gatekeeper
: A gatekeeper is a process used at a site to take incoming job requests and check the security to make sure each is allowed to use the associated computing resource(s). The gatekeeper process starts up the job-manager process after successful authentication.

Generic Information Provider (GIP)
: A configurable LDAP information provider that differentiates between static and dynamic information. OSG sites use GIP to advertise a variety of grid-related configuration data. GIP is interoperable with LCG.

GFAL
: *See* Grid File Access Layer.

GGF
: The Global Grid Forum (GGF) is a community-initiated forum of thousands of individuals from industry and research leading the global standardization effort for grid computing. GGF's primary objectives are to promote and support the development, deployment, and implementation of Grid technologies and applications via the creation and documentation of "best practices" - technical specifications, user experiences, and implementation guidelines. *See:* http://www.gridforum.org/.

GIP
: *See* Generic Information Provider.

gLExec
: gLExec, a close relative to suexec, is software that lets one user start processes as another user, by presenting a valid proxy certificate of the target user. gLExec is commonly used for what are referred to as "pilot" or "glidein" jobs (see GlideinWMS). *For more information, see:* About Glexec.

GlideinWMS
: A Workload Management System (WMS) that provides a simple way to access the Grid resources using the HTCondor system. Once setup, final users can submit regular HTCondor jobs to the local queue and the glidein factory will provide the computing resources behind the scenes. From the final user point of view, the HTCondor pool just magically grows and shrinks as needed. The user need not worry about grid entry points, managing queues, or provisioning worker nodes. See NavTechGlideinWMS for installation and reference documents.

GLUE schema
: *See* Grid Laboratory Uniform Environment schema.

GOC
: *See* Grid Operations Center.

gPLAZMA
>*See* grid-aware PLuggable AuthoriZation MAnagement.

GRAM
>*See* Grid Resource Allocation and Management.

Grid
>A named set of Services, Providers, Resources, and Policies, overlapping and/or including other Grids operating as a coherent infrastructure in support to the contracting Virtual Organizations.

Grid File Access Layer (GFAL)
>An LCG-provided interface for the normal file I/O operations (Open/Seek/Read/Write/Close) which includes a limited set of POSIX functions for file management. The GFAL interface is designed to hide the grid storage interactions (replica catalog, SRM and file access mechanism) from user applications.

Grid Laboratory Uniform Environment (GLUE) schema
>An abstract modeling for Grid resources and mapping to concrete schemas that can be used in Grid Information Services. It aims to define, publish and enable the use of common schemas for interoperability between the EU and US physics grid project efforts. *See also* GLUE Schema site.

Grid Monitor
>A part of Condor-G that replaces the monitoring (polling) duties previously done by jobmanagers. This is specific to the gt2 Grid Type. (*See also*
>http://www.cs.wisc.edu/condor/manual/v6.8/5_3Grid_Universe.html)

Grid Operations Center
>The help and support center for the entire OSG. *See also* Operations.

Grid Proxy
>A limited-life certificate signed by the user, or by another proxy, used for sending several jobs without having to reauthenticate with your password each time. Grid proxies provide a convenient alternative to constantly entering passwords, but are also less secure than the user's normal security credential.

Grid Resource Allocation and Management (GRAM or WS-GRAM)
>Service that provides a single interface for requesting and using remote system resources for the execution of 'jobs'. The most common use of GRAM is remote job submission and control. It is designed to provide a uniform, flexible interface to job scheduling systems.

Grid Security Infrastructure (GSI)
>A globus authorization and authentication schema that includes PKI (Public Key Infrastructure), single sign-on, and a time-limited proxy certificate. It is an extension of X.509 as specified by RFC-3820.

Grid User Management System (GUMS)

GUMS is a Grid Identity Mapping Service. It maps the credential for each incoming job at a site to an appropriate site credential, and communicates the mapping to the gatekeeper. GUMS is particularly well suited to a heterogeneous environment with multiple gatekeepers; it allows the implementation of a single site-wide usage policy, thereby providing better control and security for access to the site's grid resources. *See also* BNL's site.

grid-aware PLuggable AuthoriZation MAnagement (gPLAZMA)

An architecture that utilizes VOMS extended X.509 certificate specification for defining extra attributes (FQANs), based on RFC-3281, and provides role-based access control (RBAC) to dCache.

grid-mapfiles

A file-based Access Control List (ACL) mechanism that statically maps a certificate subject to a userID, authorizing them to access and use GSI-enabled services (i.e., OSG). It is a plain text file listing the subject name of the trusted certificate and the corresponding local user name. Each Distinguished Name (DN) can be listed in the file only once; but several DNs can be mapped to a single user name.

GridFtp

1. GridFTP is a Globus project that produces high-performance, secure, reliable data transfer technologies optimized for high-bandwidth wide-area networks. The project provides clients (`globus-url-copy`) and servers to transfer files on the Grid.
2. An extension to File Transfer Protocol (FTP) that includes strong authentication and encryption via Globus GSI; multiple, parallel data channels; third-party transfers; tunable network & I/O parameters; and server-side processing and command pipelining. GridFTP v2 is specified in GDF.47 and v1 is specified in GDF.20. The FTP protocol (RFC-959) is a common and extensible protocol for file transfer. GridFTP is using security extensions (RFC-2228), feature negotiation (RFC-2389) and alternate directory listing (RFC-3659). Systems can be "gridFTP-enabled". *See Also* GSIftp.

Gridmap Callout Interface

The Globus gatekeeper Gridmap callout interface allows for the replacement of the built-in Gridmap file mechanism with a component that can query the GUMS identity mapping service. The interface between PRIMA and the GUMS identity mapping service is based on the OGSA SAML Authorization Interface.

GSIftp

Method specifier used in URI to identify resources accessible using the GridFTP protocol (`gsiftp://myserver.domain: 2811/path/file`), sometime GSIftp is used as synonym of GridFTP to identify the protocol. It uses an authentication and authorization extension (RFC-2228) of FTP based on GSI, therefore the name. *See Also* GridFTP.

GUMS

See Grid User Management System.

1.8 H

A-B-C-D-E-F-G-H-I-J-K-L-M-N-O-P-Q-R-S-T-U-V-W-X-Y-Z

Head Node
>   This is an inexact term that generally refers to a node in an site cluster through which jobs are submitted. In a simple cluster, the head node is the one on which the gatekeeper and monitoring software are installed. In a more complicated cluster, there may be multiple gatekeepers on different nodes, and other issues, thus "head node" is not well defined. A head node would be networked to a group of worker nodes.

Hierarchical Resource Manager (HRM)
>   An SRM that manages access to both one or more disk caches and one or more tape archiving systems.

Hierarchical Storage Management (HSM)
>   HSM systems transparently migrate files from more expensive hard disk systems to cheaper optical disk and/or magnetic tape, usually robotically accessible. When a user requests the files, the HSM transparently moves them back to the hard disk system. *See also* HSM at Wikipedia.

High Performance Computing (HPC)
>   Computing environments that deliver large processing capacity via supercomputers and computer clusters. It often is used with a narrower meaning of supercomputers in the teraflops (Trillion Floating Operations Per Second) range.

High Throughput Computing (HTC)
>   Computing environments that deliver large amounts of processing capacity over long periods of time. The relevant measure is floating point operations per month rather than Floating Operations Per Second (FLOPS). *See* Condor High Throughput Computing. HTC is a form of High Performance Computing.

High Throughput Parallel Computing (HTPC)
>   HTPC is a computational paradigm for the class of applications where large ensembles (hundreds to thousands) of modestly parallel (4- to ~64- way) jobs are used to solve scientific problems ranging from chemistry, biophysics, weather and flood modeling, to general relativity.

HPC
>   *See* High Performance Computing.

HPSS
>   A commercial High Performance Storage System designed to scale beyond a petabyte and be highly reliable. *See* HPSS Collaboration.

HRM

    See Hierarchical Resource Manager.

HSM

    See Hierarchical Storage Manager.

HTC

    See High Throughput Computing.

HTPC

    See High Throughput Parallel Computing.

HTCondor

    HTCondor is a specialized workload management system for compute-intensive jobs. Like other full-featured batch systems, HTCondor provides a job queueing mechanism, scheduling policy, priority scheme, resource monitoring, and resource management. HTCondor is a grid middleware component developed at UW Madison, packaged and distributed by OSG. Read more at http://www.cs.wisc.edu/htcondor/.

HTCondor-G

    HTCondor-G is the Grid management part of HTCondor. HTCondor-G provides for submitting jobs to multiple, independently managed, local grids. OSG uses mostly the "gt2" (Globus) type of grid jobs but HTCondor-G can also support many other Grid and cloud protocols.

HTCondor-CE

    A gatekeeper implementation based on the HTCondor software.

1.9 I

A-B-C-D-E-F-G-H-I-J-K-L-M-N-O-P-Q-R-S-T-U-V-W-X-Y-Z

Information Provider

    Information Provider (IP) software interfaces to any data collection service, collects virtually any type of data it's asked to, and communicates the information for publishing to the grid.

Information server

    A read-only server that maintains detailed file and volume information.

Integration Test Bed (ITB)

    A distributed set of Resources that are provided, maintained, and operated by a selection of Virtual Organizations (VO) in order to provide a testing environment for new software releases. The ITB is used for integrating new services with existing services and core infrastructure for testing prior to deployment. An ITB release is defined as a set of functionalities that are provided by a site (gatekeeper services running on a computing element), a VO (VO specific services provided by that VO, such as a VOMS server), or another resource provider (such as a group providing a monitoring or

discovery service). All official OSG software releases are tested on the ITB before being approved for release. Individual VOs also use the ITB to test their software.

International Grid Trust Federation
> The International Grid Trust Federation (IGTF) is a body to establish common security policies and guidelines among its members.

ITB
> *See* Integration Test Bed.

1.10 J

A-B-C-D-E-F-G-H-I-J-K-L-M-N-O-P-Q-R-S-T-U-V-W-X-Y-Z

Job
> An executable set of code submitted to and run on the Resources provided by the Grid. A job is the fundamental unit of work in a Grid infrastructure.

Job Manager
> A Globus term that refers to a program used to interface to the batch system on the remote side. The spelling is case sensitive when used with Globus client tools. Supported job managers are `fork`, `managedfork`, `ccs`, `condor`, `lsf`, `pbs` and `sge`. The `fork` job manager is the default job manager. It executes the job directly on the gatekeeper.

1.11 K

A-B-C-D-E-F-G-H-I-J-K-L-M-N-O-P-Q-R-S-T-U-V-W-X-Y-Z

1.12 L

A-B-C-D-E-F-G-H-I-J-K-L-M-N-O-P-Q-R-S-T-U-V-W-X-Y-Z

LCG (Worldwide LHC Computing Grid)
> Computing Grid project at CERN. The goal of the LCG project is to meet the LHC experiments' unprecedented computing needs by deploying a worldwide computational grid service, integrating the capacity of scientific computing centers spread across Europe, America and Asia into a virtual computing organisation. *See* http://lcg.web.cern.ch/LCG/.

LDAP
> Lightweight Directory Access Protocol, a protocol used to locate resources in a network.

LDIF
> LDAP Data Interchange Format (LDIF) is a standard plain text data interchange format for representing LDAP (Lightweight Directory Access Protocol) directory content and update requests. LDIF conveys directory content as a set of records, one record for each object (or entry). It represents

update requests, such as Add, Modify, Delete, and Rename, as a set of records, one record for each update request.

LHC

The Large Hadron Collider at CERN. *See* http://public.web.cern.ch/public/en/lhc/lhc-en.html.

Lifetime (of a pin)

For storage, the period of time for which a file is guaranteed to remain available (remain "pinned") in a given storage area for a client. A pin lifetime is requested by a client and granted/denied by the Storage Resource Manager (SRM) at the time the file is moved in for that client. A separate, independent pin lifetime is set (or denied) for any other client who requests it once the file is already there. Client may release pin before lifetime expires. Action taken on file after last pin lifetime expires (or pin is released) is dependent on the SRM configuration.

Lifetime (of a space)

For storage, the period of time for which a (temporary, volatile or durable) storage area is guaranteed to remain available. Particular space may be requested by a client, in which case space lifetime may be negotiated. Or space and space lifetime may be assigned by default by the Storage Resource Manager (SRM). Any pin lifetime set for a file in the space may not exceed the remaining space lifetime.

Logical File Name (LFN)

A globally unique name for a file on the grid, that is location and machine independent; it may point to any Physical File Name (PFN) for the given file. LFNs are governed by the replica catalog and use the RLS.

Logistical Networking

Logistical Networking technology builds on a highly generic storage service that uses the Internet Backplane Protocol (IBP), which it deploys on storage servers called depots. *See* SPD SRM Parallel Depot.

1.13 M

A-B-C-D-E-F-G-H-I-J-K-L-M-N-O-P-Q-R-S-T-U-V-W-X-Y-Z

Mass Storage System

A high-capacity, large-scale data archive (usually tape) that is more intelligent than a normal storage system, and used to hold large amounts of infrequently accessed data. See also HRM, TRM.

Managed Fork

The ManagedFork jobmanager is a more scalable way to manage jobs which run on the compute element. Each fork job that is intended for execution on the compute element is queued in condor and run in the Condor "local universe" on the CE. This allows logging of what commands were run, and also a mechanism to throttle how many are running at once.

Match Maker (MM)

OSGMM is the Open Science Grid Match Maker, a service that sits on top of the OSG client software stack. It obtains site information from ReSS (Resource Selection Service) and uses a feedback system from/to Condor to publish the site information and keep job success history data. OSGMM can be used to schedule compute jobs across all resources available to a particular VO on the OSG, and to subsets of the resources with Condor's requirements expressions.

Matchmaking

The process of matching a job to a slot while maintaining site priorities.

Member organization of OSG

An institution, facility or VO that contributes to the development and resource pool of OSG.

Middleware

Middleware is software that connects two or more otherwise separate applications across the Internet or local area networks. More specifically, the term refers to an evolving layer of services that resides between the network and more traditional applications for managing security, access and information exchange. OSG recognizes two levels of middleware: Grid middleware (e.g., VDT, Grid3-gridmap, etc.) and VO- specific application middleware.

Message Passing

An approach to parallel computing where data and work is divided across the processors (striving for optimal load balancing) and communication between them is managed (striving for minimal and non-blocking communication) by explicitly calling specific library functions. The *de facto* standard for message passing is MPI.

Metronome

Metronome (formerly The NMI Build & Test System) is a distributed, multi-platform framework designed to provide automated software building and testing capabilities to a variety of grid computing projects. Software distributed by OSG is being moved into this build/test process. *See also* NMI Build and Test Lab.

Message Passing Interface (MPI)

A standard for message passing between processors. *See also* http://www.mpi-forum.org/docs/docs.html.

Multi-Grid Interoperability Now (GIN)

An initiative between 9 grid infrastructure projects to promote interoperability.

Monitoring (grid monitoring)

Grid monitoring involves collecting, analyzing and displaying information from the distributed production infrastructure in order to determine server status and application progress, and to log performance data of CPUs, networks and storage devices.

MSS

See Mass Storage System.

MyOSG

MyOSG is an one-stop location for various OSG information including, but not limited to, monitoring status, availability history, Virtual Organization information, contact information and accounting metrics. *See it in action:* http://myosg.grid.iu.edu/about.

1.14 N

A-B-C-D-E-F-G-H-I-J-K-L-M-N-O-P-Q-R-S-T-U-V-W-X-Y-Z

NGI
> National Grid Infrastructure. An individual national member of the European Grid Initiative

Network Element
> A network path or a set of network hops. This includes both end-to-end and hop-by-hop path information.

NMI
> *See* Metronome.

1.15 O

A-B-C-D-E-F-G-H-I-J-K-L-M-N-O-P-Q-R-S-T-U-V-W-X-Y-Z

OGSA
> *See* Open Grid Services Architecture.

OIM
> *See* Open Science Grid Information Management.

Open Grid Services Architecture (OGSA)
> A Services Architecture represents an evolution towards a Grid system architecture based on Web services concepts and technologies. *See also* Globus's Open Grid Services Architecture (OGSA) site.

Open Science Grid Information Management (OIM)
> A service that defines the topology used by various OSG systems and services; it is based on the OSG Blueprint Document available at http://osg-docdb.opensciencegrid.org/cgi-bin/ShowDocument?docid=18. For example, MyOSG, BDII, and Gratia all use topology defined in OIM.

Ownership
> A state of having absolute or well-defined partial rights and responsibilities for a Resource depending on the type of control. OSG considers two such types: actual Ownership and Ownership by virtue of a Contract/Lease. A Lessee is a limited Owner of the Resource for the duration of the Contract/Lease.

1.16 P

A-B-C-D-E-F-G-H-I-J-K-L-M-N-O-P-Q-R-S-T-U-V-W-X-Y-Z

Pacman

>A custom repository management system used by the Open Science Grid to install software on grid resources. Developed at and distributed by Boston University. *See also* Metronome.

Pacman Cache

>A Pacman Cache is a repository that provides software to be installed using the `pacman` program. A Pacman Cache is defined by a URL or a location on a local file system providing the Pacman File. Some Pacman Caches are *registered* and use simple names like *VDT*. Others are just URLs or file system locations.

Pacman file

>A .pacman file (e.g., xyz.pacman) contains instructions on how the software in the xyz package is fetched, installed, setup, uninstalled, what other packages it depends on, and so on.

Pacman package

>A software environment created by installation of a .pacman file.

PanDA

>A job scheduling and management system that provides an integrated service architecture with late binding of jobs, maximal automation through layered services, tight binding with ATLAS Distributed Data Management system, advanced error discovery and recovery procedures, and other features.

Partner

>Specifically, an individual or organization affiliated with a grid external to the Consortium with which the Open Science Grid interacts with through federation of resources. More broadly, the Open Science Grid maintains a partner relationship with many other organizations that are related grid infrastructures, other high performance computing infrastructures, international consortia, and certain project organizations that operate in the broad space of high throughput or high performance computing. *See also* Partner Organizations?.

Permanent storage

>A data storage system, or a data collection in a storage system, wherein a physical file can only be created and removed by the owner of the data collection.

Persistent storage

>See Permanent storage.

Physical File Name(PFN)

>The URL of a physical replica of a file, minus the protocol.

Pinning (a file)

>Pinning refers to the capability of an Storage Resource Manager (SRM) to keep a particular file in non-permanent storage space for a period of time set by the client, prior to making the file eligible for transfer or removal from the SRM. A pin is requested and released by a client. Pinning a file is a way of keeping the file in place, while not locking its content.

Policy

A statement of well-defined requirements, conditions or preferences put forth by a Provider and/or Consumer that is utilized to formulate decisions leading to actions and/or operations within the infrastructure. Policies of the Open Science Grid can be found at http://www.opensciencegrid.org/docdb_dashboard/index.php.

Pool (in dCache)

A virtual data partition in the dCache storage space. A pool is a cell responsible for storing retrieved files and for providing access to that data.

Pool (in Condor)

The set of resources controlled by a single Condor instance.

POOL

POOL is a Large Hadron Collider Computing Grid (LCG) software system and project, developed with US participation. It has been created to implement a common persistency framework for the LCG application area and to replace the Objectivity database system in the US LHC software. POOL is tasked to store experiment data and metadata in the multi Petabyte area in a distributed and grid enabled way. Read more at http://twiki.cern.ch/twiki/bin/view/Persistency.

Privilege

The Virtual Organization (VO) Services project, formerly known as the VO Privilege Project, implements finer-grained authorization for access to grid-enabled resources and services in order to improve user account assignment and management at grid sites, and reduce the associated administrative overhead. Depending on its implementation, privilege relies on, interfaces to and further develops at least some of the following independent pieces of VO-implemented and site-implemented authorization software: VOMS, VOMRS, Grid-map callout interface, PRIMA, GUMS, and SAZ. The project closed in June 2009.

PRIMA

PRIvilege Management and Authorization, a component of the privilege project for user authorization at a site, is used with GUMS and VOMS to implement dynamic, fine-grained, role-based identity mapping. PRIMA extracts the VOMS attributes containing the VO and role information from the user's proxy certificate, and queries GUMS for an appropriate local user account assignment.

Provider

Makes a Resource or Agent or Service available for access and use.

1.17 Q

A-B-C-D-E-F-G-H-I-J-K-L-M-N-O-P-Q-R-S-T-U-V-W-X-Y-Z

1.18 R

A-B-C-D-E-F-G-H-I-J-K-L-M-N-O-P-Q-R-S-T-U-V-W-X-Y-Z

RB (Resource Broker)

Grid middleware component that brokers the running of Grid jobs (making use of information services to obtain grid status information about available resources) and schedules jobs.

RBAC (Role-Based Access Control)

An infrastructure which provides a framework for role-based access to resources and services.

Release

An OSG Release refers to a set of software that is, or is expected to be, tested and available for download and installation. It contains a documented set of capabilities. Releases are sequentially numbered with a r.n.p format (e.g. 1.2.15) where r is the release, n is changed for major releases with significant change, and p is the point release number. Major releases occur infrequently (~yearly), while point releases are more frequent (~months). *For more information, see* DocReleaseProcess and ReleaseDocumentation.

Replica catalog

Provides mappings between logical names for files and one or more copies of the files on physical storage systems.

Resource

A resource is any physical or virtual entity of limited availability. In the OSG, all resources are represented by a unique DNS endpoint. The entity is available through the grid for use by researchers, typically a machine providing CPU cycles or storage capacity.

Resource Group

A named collection of resources for administrative purposes.

Resource Owner

An organization that has permanent specific control, rights and responsibilities for a Resource associated with ownership.

ReSS (Resource Selection Service)

ReSS is a component of the OSG Job Management Infrastructure that facilitates automated job/resource management components by publishing information about the OSG resource information in the Glue Schema format so grid users can easily access this published information and use it in their automated job management system.

RLI (Replica Location Interface)

A Grid middleware component used to distribute RLS information.

RLS (Replica Location Service)

A Grid middleware component provides information about location of data sets within the data grid.

RP (Resource Provider)

A facility offering resources (e.g., CPU, network, storage) to other parties (e.g., VOs) according to a specific Memorandum of Understanding (MOU).

RPM (Red Hat Package Manager)

A software package manager, developed by Red Hat, to manage, distribute, and update software as new versions are developed. RPMs are used by OSG in Release 3.0.

RSV (Resource and Service Validation) service
>	provides a scalable and easy to maintain resource/service monitoring infrastructure for an OSG site admin.

1.19 S

A-B-C-D-E-F-G-H-I-J-K-L-M-N-O-P-Q-R-S-T-U-V-W-X-Y-Z

SA
>	*See* Storage Area.

SAM
>	Service availability monitoring (SAM) displays resource and service availability metrics from the central RSV database.

SAML
>	Security Assertion Markup Language.

SAZ
>	*See* Site AuthoriZation service.

Schema
>	A schema refers to a description of objects and attributes needs to describe Grid resources, and the relationships between the objects.

SE
>	*See* Storage Element.

Security
>	Control of and reaction to intentional, unacceptable use of any part of the infrastructure.

Service
>	A method for accessing a Resource or Agent.

SIMD (Single Instruction, Multiple Data)
>	A class of parallel computers that describes computers with multiple processing elements that perform the same operation on multiple data simultaneously. Thus, such machines exploit data level parallelism.

SISD (Single Instruction, Single Data)
>	A term referring to a computer architecture in which a single processor, a uniprocessor, executes a single instruction stream, to operate on data stored in a single memory.

Site
>	A site is a logical name denoting a concrete, persistent, uniquely identifiable, and testable collection of Services, Providers and Resources for administrative purposes. A Facility is a collection of Sites under a single administrative domain. A site offers computing services, persistent storage services, or

both. A site offering computing services is identified by a one or more gatekeeper services (hostname and port) and a GSIftp service (hostname and port). Multiple sites at a facility may share certain services at that facility.

Site Authorization Service (SAZ)
Allows security authorities of the grid site to impose site-wide policy and to control access to the site; allows administrators to control user access to the site resources; and, provides means to retrieve the information about users and their access.

SPD (SRM Parallel Depot)
An integration of the Storage Resource Manager (SRM) with a technology for wide area data management, Logistical Networking (LN), which is now being used by several important OSG application communities. The strategy of combining SRM with Internet Backplane Protocol (IBP)-based depots is designed to provide an interoperable foundation for wide area storage infrastructure that can address the problems of in-transit data management through a familiar interface, but with maximum flexibility and performance.

Squid
Squid is a web caching service that speeds up downloads from http servers by locally caching files and serving the cached files rather than retrieving the files over the internet. Also referred to as a proxy, the typical use case would be for squid to be installed on a single server on a cluster with CE nodes using that server to proxy http and https requests.

Squid monitor
SNMP-based monitoring service for a Squid web cache.

SRM
See Storage Resource Manager.

SRM-dCache
An implementation of SRM as a "door" into dCache.

Storage Area
The storage area is a logical portion of storage extent assigned to a VO. Storage areas can overlap the same physical space, thus having contention over the free space among different VO s.

Storage Element
The group of services responsible for a storage resource on the Grid. Services may include data access (e.g. protocol for local data access, protocol for secure wide-area transfer), quota management or space management (providing the space available and allowing space reservation). Storage resources contributed to a Grid system can vary from simple disk servers managed via GridFTP to complex massive storage systems managed via SRM. The SE is a sufficiently flexible interface to grid storage units that allows interoperability of the grid without forcing local sites to change their existing software stack.

Storage Resource Manager (SRM)
Middleware components that manage shared storage resources on the Grid and provide: uniform

access to heterogeneous storage; ftp negotiation; dynamic TURL allocation; access to permanent and temporary types of storage; advanced space and file reservation; and, reliable transfer services. *See also:* DRM, TRM; HRM.

Subcluster

In the GLUE schema, a subcluster represents a "homogeneous" collection of nodes, where the homogeneity is defined by a collection whose required node attributes all have the same value. For example, a subcluster represents a set of nodes with the same CPU, memory, OS, network interfaces, etc. Subclusters provide a convenient way of representing useful collections of nodes.

Submit machine

The computer from which a job has been submitted.

1.20 T

A-B-C-D-E-F-G-H-I-J-K-L-M-N-O-P-Q-R-S-T-U-V-W-X-Y-Z

Temporary file

A file in temporary storage.

Temporary storage (space)

A shared space that is allocated to a user, but can be reclaimed by the file system (after some guaranteed amount of time has passed). If space is reclaimed, all the files in that space are removed by the file system. The implication is that files in these spaces are also temporary.

TeraGrid

An open scientific discovery infrastructure, launched by the National Science Foundation in August 2001, and now combining leadership class resources at eleven partner sites to create an integrated, persistent computational resource. TeraGrid integrates high-performance computers, data resources and tools, and high-end experimental facilities around the country. *See* https://www.teragrid.org/web/about/.

Tier 0

Initial tier in the grid hierarchy for a research project; it is the site at which raw data is taken. The experimental online system interfaces to the tier 0 resources. For the LHC high energy physics experiments, CERN is the tier 0 facility. Fermilab is the tier 0 facility for the Run II experiments at the Tevatron.

Tier 1

Next tier, after tier 0, in grid hierarchy. Tier 1 sites are connected to a Tier 0 site based on an MOU with the Tier 0 site. Typically a tier 1 site offers storage, analysis, and services, and represents a broad constituency (e.g., there may be a single tier 1 site per country or region which connects with multiple tier 2 sites in that country or region). In the US, tier 1 centers for the LHC high energy physics experiments ATLAS and CMS are BNL and Fermilab, respectively.

Tier 2

Tier 2 is the next level down in the grid hierarchy of sites, after tier 1. Tier 2 sites are typically regional computing facilities at University institutions providing a distributed Grid of facilities.

Tier 3

Tier 3 is typically a small to medium IT cluster/grid resource targeted at supporting a small group of scientists.

Tools

With respect to grid computing, "tools" refers to a "layer" of grid components (underneath applications, and above middleware and fabric). The tools layer encompasses resource brokers, monitoring tools, debuggers, etc.

Transfer File Name (TFN)

The IP address of a file, containing the transfer protocol name, the host and the Logical File Name (LFN).

TRM

A Tape Resource Manager (TRM) is a middleware layer that interfaces to systems that manage robotic tapes in a data grid. TRM is one type of a Storage Resource Manager (SRM).

TURL

A Transfer URL (TURL) is a URL used in the file transfer negotiation. A TURL includes the transfer protocol name and the name of the file server machine along with the path to the file.

1.21 U

A-B-C-D-E-F-G-H-I-J-K-L-M-N-O-P-Q-R-S-T-U-V-W-X-Y-Z

User

A person who makes a request of the Open Science Grid infrastructure. Typically, this is a person who submits jobs to run on OSG.

1.22 V

A-B-C-D-E-F-G-H-I-J-K-L-M-N-O-P-Q-R-S-T-U-V-W-X-Y-Z

Virtual Data Toolkit (VDT)

A collection of grid middleware upon which the OSG software stack is based. *See also* VDT site.

Virtual Organization (VO)

A dynamic collection of Users, Resources and Services for sharing of Resources. A VO is party to contracts between Resource Providers & VOs which govern resource usage & policies. A VO is a participating organization in a grid to which grid end-users must be registered and authenticated in order to gain access to the grid's resources. A VO must provide resources and/or establish resource-usage agreements with grid resource providers. Members of a VO may come from many different home institutions, may have in common only a general interest or goal (e.g., CMS physics analysis), and may communicate and coordinate their work solely through information technology

(hence the term virtual). An organization like a High Energy Physics experiment can be regarded as one VO. A subVO is a subset of the Users and Services within a VO which operates under the contracts of the parent.

Virtual Site
A set of Sites that agree to use the same policies in order to act as an administrative unit. Sites and Facilities negotiate a common administrative context to form a "virtual" site or facility.

Volatile File
A file that is temporary in nature, but has a (pin) lifetime guarantee.

Volatile Storage
Storage in which the "volatile" physical replica of a file is subject to removal from the Storage Resource Manager SRM or Distributed Resource Manager (DRM) according to preset policies.

Virtual Organization Management Registration Service (VOMRS)
VOMRS is a server that provides the means for registering members of a Virtual Organization, and coordination of this process among the various VO and grid resource administrators. VOMRS consists of a database to maintain user registration and institutional information, and a web user interface (web UI) for input of data into the database and manipulation of that data. The database is shared with the Virtual Organization Membership System. *See also* http://computing.fnal.gov/docs/products/vomrs/.

Virtual Organization Membership System (VOMS)
A system that manages real-time user authorization information for a VO. VOMS is designed to maintain, in a database shared with the VO Management Registration Service, only general information regarding the relationship of the user with his VO, e.g., groups he belongs to, certificate-related information, and capabilities he should present to resource providers for special processing needs. It maintains no personal identifying information besides the certificate. VOMS has 2 components: VOMS-Admin and VOMS-Server. Each officially registered VO in Open Science Grid is expected to have a functional VOMS system, either operated by itself or by arrangement with another VO.

VOMS-Admin
The components of Virtual Organization Membership System (VOMS) which is used for administration and book-keeping with database.

VOMS-Server
The component of Virtual Organization Membership System (VOMS) system which users contact to get a VOMS Proxy certificate. When a user issues a remote request to his VO's VOMS-Server (using voms-proxy-init command line client side tool), the VOMS-Server verifies user authenticity and grants an X.509 proxy certificate which contains extra attributes. Attributes returned by VOMS-Server are inserted into an RFC-3281 compliant Attribute Certificate, and thus compatible with other formats, e.g., Globus format. The AC has tokens which represent user's group membership(s) and role(s) in his VO.

VOMS Proxy
A Virtual Organization Membership System (VOMS) Proxy is an extended X.509 Grid Proxy certificate which uses extra attributes to define the group membership(s) and role(s) of a grid end-user. Grid services can be configured to read these attributes from a proxy certificate and perform

decisions based on their values.

VO

See Virtual Organization.

VO Manager

A person designated by a Virtual Organization (VO) to be in charge of defining policies for the VO. These policies may be related to, but not limited to, resource usage and VO membership. This person is also usually, but not always, in charge of the team operating the VOMS system of the VO.

1.23 W

A-B-C-D-E-F-G-H-I-J-K-L-M-N-O-P-Q-R-S-T-U-V-W-X-Y-Z

WN

See Worker Node.

Wall Hours

The number of hours taken by a computer to complete a task, including CPU hours + I/O time + any communication channel delay.

WLCG

The Worldwide Large Hadron Collider Computing Grid. *See* LCG.

Worker Node

The terminology used in the OSG to name the node that is executing jobs on the behalf of grid users. Each worker node is only associated with one Compute Element CE. A Compute Element can be associated with more than one Worker Node.

1.24 X

A-B-C-D-E-F-G-H-I-J-K-L-M-N-O-P-Q-R-S-T-U-V-W-X-Y-Z

XD

The National Science Foundation's eXtreme Digital program (XD) is making new infrastructure and next-generation digital services available to researchers and educators.

XRAC

The XSEDE Resource Allocation Committee (XRAC) awards allocations of compute resources in the XSEDE high performance computing environment, including allocations that go to the OSG. The committee meets quarterly to review requests based on merit and then award available resources on XD.

XSEDE

The Extreme Science and Engineering Discovery Environment (XSEDE) is a collection of integrated advanced digital resources and services, consisting of 16 supercomputers and high-end visualization and data analysis resources across the US. The five-year, $121-million XSEDE project in the XD program is supported by the National Science Foundation. It replaces and expands on the NSF TeraGrid? project.

1.25 Y

A-B-C-D-E-F-G-H-I-J-K-L-M-N-O-P-Q-R-S-T-U-V-W-X-Y-Z

1.26 Z

A-B-C-D-E-F-G-H-I-J-K-L-M-N-O-P-Q-R-S-T-U-V-W-X-Y-Z

1.27 Sources Used in the Glossary of Terms

- [GLUE13] Sergio Andreozzi, Stephen Burke, Flavia Donno, Laurence Field, Steve Fisher, Jens Jensen, Balazs Konya, Maarten Litmaath, Marco Mambelli, Jennifer M. Schopf, Matt Viljoen, Antony Wilson, Riccardo Zappi. 2007 "GLUE Schema Specification version 1.3".
- [BOEH] Chuck Boeheim, Andy Hanushevsky, David Leith, Randy Melen, Richard Mount, Teela Pulliam, Bill Weeks. 2006. "Scalla: Scalable Cluster Architecture for Low Latency Access Using xrootd and olbd Servers".
- [DEJ06] de Jong, Richard & Koot, Matthijs. 2006. RP2: Preparing the Worldwide LHC Computing Grid for MPI applications (PDF). University of Amsterdam. 2006.
- [HUB04] Hubovsky, Rainer. 2004. *Dealing with Massive Data: from Parallel I/O to Grid I/O*. Vienna: Institute for Applied Computer Science and Information Systems, University of Vienna. 10-11.
- [PER04] T. Perelmutov et al, "Storage Resource Manager", *CHEP'04*, Abstract 107, Interlaken, September 2004.
- [SSG02] Shoshani, Aire; Sim, Alexander; Gu, Junmin. 2002. "Storage Resource Managers: Middleware Components for Grid Storage". Nineteenth IEEE Symposium on Mass Storage Systems (MSS '02).
- [SSG03] Shoshani, Aire; Sim, Alexander; Gu, Junmin. 2003. "Storage Resource Managers: Essential Components for the Grid". In *Grid Resource Management: State of the Art and Future Trends*. Edited by Jarek Nabrzyski, Jennifer M. Schopf, Jan weglarz. Kluwer Academic Publishers.

See GlossaryToPDF for a procedure to produce a clean glossary in pdf format with working internal/external links.

JamesWeichel - 11 Feb 2011

This topic: Documentation > GlossaryOfTerms
Topic revision: r33 - 12 Aug 2014 - 15:27:16 - LotharBauerdick

TWiki | Report Bugs | Privacy Policy