

Balázs Kónya (editor), Lund University, ARC
Aleksandr Konstantinov, Oslo University, ARC
Laurence Field, CERN, gLite
Moreno Marzolla, INFN, gLite
Morris Riedel, Julich, UNICORE
Bernd Schuller, Julich, UNICORE

UMD, EGI:
<http://web.eu-egi.eu/>

June xxx, 2009

Survey of potential Universal Middleware Distribution (UMD) components

Status of the document

The content of the document should be considered as a continuous work in progress. The current version presents a completed inventory with respect to the ARC, gLite and the Unicore candidate components. Unfortunately the inventory is incomplete with respect to components of “3rd party” middleware teams. Furthermore, the standard landscape overview sections would also need more work.

Version 0.84

Last modified: 15:00, 26 June

Scope and purpose of the document

Technical experts of the three middleware consortia (i.e. ARC, gLite, UNICORE) have been jointly preparing a survey of the currently available middleware components. The survey intends to provide a functional decomposition, taxonomy and status evaluation of the components of the middleware stacks that have a potential relevance for UMD. Due to limited resources and time constraints the technical team initially can only focus on the components of the ARC/gLite/Unicore middleware stacks, nevertheless wherever it was possible third party solutions were considered in the survey.

The main purpose of this middleware component survey is to be used as the technical background document during the UMD component identification and selection process. The document itself contains no recommendations of particular components: instead its purpose is to present a technically sound state-of-the-art modular decomposition and inventory of existing and under development middleware solutions. In particular, the document provides information in terms of numerous categories along functionalities in the areas of computing, data, information, accounting and security.

Contents

1	Introduction	5
1.1	Component description template	6
2	Computing functionality	7
2.1	Standards landscape	7
2.2	Overview table: computing functionality	8
2.3	Inventory of components	8
2.3.1	ARC: Grid Manager	8
2.3.2	ARC: A-REX	9
2.3.3	ARC: libarcclient (job management module)	11
2.3.4	ARC: ng* job management CLI	12
2.3.5	ARC: arc* job management CLI	13
2.3.6	ARC: Janitor	14
2.3.7	gLite: WMS	14
2.3.8	gLite: WMS Clients	15
2.3.9	gLite: LB	16
2.3.10	gLite: CREAM Service	17
2.3.11	gLite: CREAM Command Line Tools	18
2.3.12	gLite: LCG-CE	19
2.3.13	UNICORE: TSI	20
2.3.14	UNICORE: XNJS	21
2.3.15	UNICORE: UAS (Job Submission and Management)	22
2.3.16	UNICORE: OGSA-BES Adoption	23
2.3.17	UNICORE: Workflow Engine	24
2.3.18	UNICORE: Service Orchestrator	25
2.3.19	UNICORE: Rich Client	25
2.3.20	UNICORE: Command-line Client	26
2.3.21	UNICORE: HILA	27
3	Data functionality	28
3.1	Standards landscape	28
3.2	Overview table: data functionality	28
3.3	Inventory of components	29
3.3.1	ARC: Classic SE	29
3.3.2	ARC: libarcdata2	30
3.3.3	ARC: ng* data management CLI	31
3.3.4	ARC: arc* data management CLI	32
3.3.5	ARC: Chelonia	33
3.3.6	ARC: Chelonia client	34
3.3.7	gLite: DPM	35
3.3.8	gLite: lcg_utils	36
3.3.9	gLite: htcp	36
3.3.10	gLite: SRM client	37
3.3.11	gLite: DPM client	38
3.3.12	gLite: LFC client	39
3.3.13	gLite: FTS client	39
3.3.14	gLite: DPM/DICOM interface	40
3.3.15	gLite: LCG File Catalogue (LFC)	41
3.3.16	gLite: File Transfer Service (FTS)	42
3.3.17	gLite: Hydra	43
3.3.18	gLite: AMGA	44
3.3.19	gLite: AMGA Client API	45
3.3.20	UNICORE: UAS (Storage Management and File Transfer)	45
3.3.21	other: dCache	46
3.3.22	Other: StoRM	47
4	Information functionality	48

4.1	Standards landscape.....	48
4.2	Overview table: information area.....	49
4.3	Inventory of components.....	49
4.3.1	ARC: Classic Infoserver.....	49
4.3.2	ARC: Classic Infoindex.....	50
4.3.3	ARC: ALIS.....	51
4.3.4	ARC: ISIS.....	52
4.3.5	ARC: libarcclient (infosys modules).....	53
4.3.6	ARC: Grid Monitor.....	54
4.3.7	gLite: BDII.....	55
4.3.8	gLite: lcg-info and lcg-infosites.....	56
4.3.9	gLite: Service Discovery API.....	57
4.3.10	UNICORE: CIS.....	58
4.3.11	UNICORE: Service Registry.....	58
5	Security functionality.....	59
5.1	Standards landscape.....	59
5.1.1	Authentication (AuthN).....	60
5.1.2	Authorization (AuthZ).....	60
5.2	Overview table: Security area.....	61
5.3	Inventory of components.....	62
5.3.1	ARC: HED security API.....	62
5.3.2	ARC: Shibbridge.....	63
5.3.3	ARC: arcproxy.....	64
5.3.4	ARC: Charon Authorization Service.....	65
5.3.5	ARC: Proxy Store.....	65
5.3.6	ARC: Fruittly credential service.....	66
5.3.7	gLite: VOMS.....	67
5.3.8	gLite: VOMS-Admin.....	68
5.3.9	gLite: voms-proxy-*.....	69
5.3.10	gLite: ARGUS.....	70
5.3.11	gLite: Short Lived Credential Service.....	71
5.3.12	gLite: SCAS.....	72
5.3.13	gLite: LCAS.....	73
5.3.14	gLite: LCMAPS.....	74
5.3.15	gLite: gLite-Exec (glExec).....	74
5.3.16	gLite: gLite-Delegation Client.....	76
5.3.17	gLite: Shibboleth Short Lived Credential Service Client.....	76
5.3.18	gLite: Gridsite.....	77
5.3.19	gLite: CGSI_gSoap.....	78
5.3.20	gLite: delegation-Java.....	79
5.3.21	gLite: Trustmanager.....	80
5.3.22	UNICORE: Gateway.....	81
5.3.23	UNICORE: XUADB.....	82
5.3.24	UNICORE: XACML Entity.....	83
5.3.25	UNICORE: UVOS.....	84
5.3.26	other: MyProxy.....	85
6	Accounting functionality.....	86
6.1	Standards landscape.....	86
6.2	Overview table: Accounting area.....	86
6.3	Inventory of components.....	86
6.3.1	ARC: JURA.....	86
6.3.2	gLite: Accounting Service.....	87
6.3.3	gLite: APEL.....	88
6.3.4	other: SGAS.....	89
7	Other functionality.....	90
7.1	Overview table: Other components.....	90

7.2 Inventory of components 90
7.2.1 ARC: HED 90
8 Effort estimates 92
9 References 96

1 Introduction

The Universal Middleware Distribution (UMD) is a collection of compiled software components with integrated configuration from a single repository that add a layer of Grid functionality on top of operating systems (similar in concept to a Linux distribution). As such it is composed of multiple production quality software components originating from different providers each with their own separate software development infrastructure. UMD is the proposed approach to provide and manage a high quality middleware stack for the future European grid infrastructure, the EGI¹. The UMD software management process is defined in separate document "EGI: Managing the software process"²

This survey presents a functional decomposition and taxonomy of the existing middleware solutions ARC, gLite and UNICORE. The middleware solutions are decomposed into components which are categorized by their provided functionality. The document classifies the middleware components by the following functionality areas:

- Computing functionality
- Data functionality
- Information functionality
- Accounting functionality
- Security functionality
- and other.

Within each functionality area a brief overview of the current standard landscape was prepared, where it was possible, an initial set of specifications to be followed by UMD components was recommended. This standard landscape overview also tries to identify gaps within a certain functionality area where further standardization work is critically needed. In this particular context, the UMD Tech team also takes the work of the OGF Production Grid Infrastructure (PGI) group into account, but focuses more in-depths on the demands given by EGI.

The middleware components are surveyed and presented within a particular functionality area and were evaluated against a recommended set of standards. A rough estimate of efforts per components needed to converge towards the recommended minimum set of standards is provided.

A middleware component is defined as either being a service, a development library or a client implementing one of the identified grid functionalities. In the context of this document services represent self-contained functionality which can be remotely accessed by Grid users or other services. Libraries are publicly supported development interfaces which can be used in the implementation of services or user level applications and can be thought of as part of a Standard Development Kit (SDK). Clients are binary tools aimed at user to be able to carry out some Grid functionality.

The UMD candidate components from the middleware stacks were evaluated against a set of criteria (see component template below).

¹ EGI Blueprint: <http://www.eu-egi.eu/blueprint.pdf>

² UMD Process document:

http://knowledge.eu-egi.eu/knowledge/index.php/UMD#Process_working_group

1.1 Component description template

Name

Description

A short paragraph describing the component

Basic information

Component type: *{service, library, client}*

Developer: *the team providing development/maintenance/support for the module*

Documentation:

Code repository:

Development language:

Build environment:

Download:

Licence:

Status

Maturity: *{planned, development, pre-production, production}*

Usage: *Infrastructures and/or user communities depending on the component*

Interface maturity: *{development, pre-production, production}*

Support status: *{abandoned, best effort, short term, long term}*

Support team: *person/project/organization providing support for the component*

Development plans:

Distribution/Availability

Middleware stack: *The name and the release of the middleware stack distributing the component*

Standalone usability: *Can the component alone be part of UMD?*

Component dependencies: *components listed in this survey on which this component depends*

External software dependencies: *the major external software dependencies*

Component consumers: *list of other components relying on the component*

Supported platforms: *e.g. 32/64 linux, mac-os x, windows, RHEL4, ubuntu, ...*

Planned to be supported platforms: *e.g. 32/64 linux, mac-os x, windows, RHEL4, ubuntu, ...*

Service component information *(only for a service component)*

Public Interfaces:

Aligned security: *security component or security information transferred with the requests*

Capabilities: *capabilities offered by the service according to the Capability_t type of Glue2*

Library component information *(only for a library component)*

Consumers: *list of other components (services/clients) relying on the library*

Language binding: *programming languages in which the API is available*

Interoperability/Standards

Standard-compliance: *{none, non-UMD, early-UMD, full-UMD}*

Interoperable components:

Additional standards: *implemented and relevant non-UMD recommended standards*

Cost estimates³

Maintenance/support effort: *in FTEs*

Standard convergence effort: *in FTEs*

New development effort: *in FTEs*

Total effort: *in FTEs*

³ Cost estimates are given in a summarized form in the dedicated Section 8

Other

Any other kind of information

2 Computing functionality

2.1 Standards landscape

There are currently two major standard specifications which apply to job management services, namely the Basic Execution Service (BES) and the Job Submission Description Language (JSDL) specifications. Besides the core BES and JSDL specifications there exist several extensions and profiles created by different OGF activities.

The JSDL [1] is an XML-based notation for describing the requirements of computational jobs for submission to Grid environments. The aim of JSDL is to provide a notation for describing the structure and requirements of individual jobs. The core JSDL is deliberately narrow scoped.

The BES specification [2] describes a basic Web Service interface for creation, monitoring and control of computational jobs. In the BES terminology, jobs are called activities, and are described using the JSDL notation. While JSDL is used to describe the static structure of an activity, BES specifies a set of operations which can be executed on activities: creation, termination, obtaining the current status of an activity or a set of activities and so on. In general, a BES service acts as a frontend to one or more resources, where a resource is a generic term to denote anything from a supercomputer, to a pool of workstations managed through a batch system such as LSF, PBS or Torque, or individual computers.

Although both the BES and JSDL specifications have been finalized, they are not suitable for production use because of several limitations and shortcomings. Some of the most relevant limitations are the following:

- Lack of a common security profile. Security considerations are outside the scope of the BES specification; the result is that different BES implementations might adopt incompatible security settings, which result in the services being not interoperable.
- Lack of common information model. Both BES and JSDL employ a limited-scope information model. The information model exposed through the BES interface is not satisfactory in production environment. Not all significant features of the contained resources are exposed, and there are no efficient ways to query the information model.
- Insufficient treatment of data staging and application environments.

For these reasons, the BES and JSDL specifications need to be suitably extended and reengineered in order to be usable in production infrastructures. Furthermore, the core specifications should be harmonized with GLUE v2.0 and security models. Currently the OGF Production Grid Infrastructures (PGI) Working Group is working towards defining a production profile based on refinements, extensions or redesign of BES, JSDL and GLUE. The outcome of the PGI WG will likely constitute one of the building blocks for the implementation of interoperable job submission and management components in the context of UMD.

It is recommended that UMD through the middleware consortia increases its involvement in the standardization work of PGI working group in order to define a production-ready common specification for job management and job description within one year. Once the PGI specification is released it is suggested that all computing functionality services included in UMD should comply with the PGI profile for job management. The standard convergence work related to PGI conformance within the computing area is estimated to require substantial development efforts from the middleware teams.

2.2 Overview table: computing functionality

	Component name	Type	Maturity	Standard compliance		
2.3.1	ARC: Grid Manager	Service	Production	none		
2.3.2	ARC: A-REX	Service	Pre-production	Early-umd		
2.3.3	ARC: libarcclient (job management module)	Library	Pre-production	Early-umd		
2.3.4	ARC: ng* job management CLI	Client	production	None		
2.3.5	ARC: arc* job management CLI	Client	Pre-production	Early-umd		
2.3.6	ARC: Janitor	Service	development	none		
2.3.7	gLite: WMS	service	production	None		
2.3.8	gLite: WMS Clients	client	production	None		
2.3.9	gLite: LB	service	production	None		
2.3.10	gLite: CREAM Service	service	production	Non-umd		
2.3.11	gLite: CREAM Command Line Tools	client	production	None		
2.3.12	gLite: LCG-CE	service	production	None		
2.3.13	UNICORE: TSI	service	production	None		
2.3.14	UNICORE: XNJS	library	production	Non-umd		
2.3.15	UNICORE: UAS (Job Submission and Management)	service	production	Non-umd		
2.3.16	UNICORE: OGSA-BES Adoption	service	Pre-production	Non-umd		
2.3.17	UNICORE: Workflow Engine	service	Pre-production	none		
2.3.18	UNICORE: Service Orchestrator	service	Pre-production	None		
2.3.19	UNICORE: Rich Client	client	production	None		
2.3.20	UNICORE: Command-line Client	client	production	Non-umd		
2.3.21	UNICORE: HILA	library	Pre-production	???		

Comment [BK1]: ???

2.3 Inventory of components

2.3.1 ARC: Grid Manager

Name ARC Grid Manager

Description

The ARC Grid Manager is a powerful realization of a pre Web Service computing element. It implements job execution capability over a large variety of computational resources. The computing element is built around the Grid Manager. The main features of the Grid Manager include the transparent, automatic and integrated data staging and caching capability, support for large number of batch systems, session directory management, support for RuntimeEnvironments, logging capability, flexible access control, etc...The pre-web service

computing element exposes the Grid Manager's functionality through a custom GridFTP- and XRSL-based interface.

Basic information

Component type: service
Developer: Oslo University, NorduGrid Collaboration
Documentation:
<http://www.nordugrid.org/documents/GM.pdf>
Code repository:
<http://svn.nordugrid.org/trac/nordugrid/browser/arc0>
Development language: c++, shell scripts, perl
Build environment: GNU autotools
Download: download.nordugrid.org
Licence: Apache License Version 2.0

Status

Maturity: production
Usage: as the core server-side ARC component the GM is deployed at every ARC site installation
Interface maturity: production
Support status: long term
Support team: NorduGrid Collaboration
Development plans: only bugfixes

Distribution/Availability

Middleware stack: ARC production releases (latest version 0.6.5)
Standalone usability: yes
Component dependencies: none
External software dependencies: globus pre-WS libraries, gsoap
Component consumers: ARC: ng*, ARC: arc*, ARC: libarcclient
Supported platforms: major linux flavours
Planned to be supported platforms: none

Service component information *(only for a service component)*

Public Interfaces: custom gridftp and XRSL based interface
Aligned security:
Capabilities:

Interoperability/Standards

Standard-compliance: none
Interoperable components: none
Additional standards: gridftp protocol is used for job submission/management

Other

As part of the development of a next generation ARC computing element the gridftp-based custom interface layer and the security framework of the Grid Manager are being replaced by standard compliant, WS-based technology solution. The new component, called A-REX, will gradually wrap all the existing functionality of the Grid Manager thus replacing the pre-WS technology computing element. All the new functionality development is already taking place as part of the A-REX. Support for Grid Manager will be offered as long as the service is deployed in production.

2.3.2 ARC: A-REX

Name ARC Resource-coupled EXecution service A-REX

Description

A-REX is the next generation computing element of ARC offering WS-interfaces and advanced security solutions. The powerful computing element implements job execution capability over a large variety of computational resources. A-REX is built around the Grid Manager of the pre-WS ARC computing element. A-REX interprets standard job description (JSDL with extensions), offers OGF-compliant WS interface, features transparent, automatic and integrated data staging and caching capability, support for large number of batch systems, session directory management, logging capability, RuntimeEnvironments. A-REX offers WS-based local information interface serving GLUE2 information. A-REX is capable working together with community approved security frameworks.

Basic information

Component type: service
Developer: Oslo University, NIIF
Documentation:
Code repository: <http://svn.nordugrid.org/trac/nordugrid/browser/arc1>
Development language: c++, shell scripts, perl
Build environment: GNU aototools
Download: download.nordugrid.org
Licence: Apache License Version 2.0

Status

Maturity: pre-production
Usage: as the core server-side ARC component A-REX will be deployed at every ARC site installation
Interface maturity: production, but open for changes as the PGI profile emerges
Support status: short term support by KnowARC project, long term support by NorduGrid collaboration
Support team: Oslo University, NIIF
Development plans: add support for PGI profile

Distribution/Availability

Middleware stack: part of the technology preview releases (0.9.x) and will be included in the next production ARC release (0.8.0) scheduled for May 2009
Standalone usability: no
Component dependencies: ARC: HED
External software dependencies:
Component consumers: ARC: libarcclient, ARC: arc* CLI
Supported platforms: major linuxes, solaris, mac-osx
Planned to be supported platforms: windows

Service component information *(only for a service component)*

Public Interfaces: extended BES, extended JSDL, Glue v2.0 through WSRF-RP
Aligned security:
Capabilities:

Interoperability/Standards

Standard-compliance: early-UMD
Interoperable components: Unicore client
Additional standards: -

Other

A-REX is gradually replacing the pre-WS ARC computing element by wrapping and preserving the powerful functionality offered by the Grid Manager component into novel WS-based solution enriched by an enhanced security framework. All new development related to server-side ARC computing element takes places in A-REX.

2.3.3 ARC: libarcclient (job management module)

Name libarcclient (job management module)

Description

The next generation general purpose grid library (libarcclient) offers transparent access to grid resources through an intuitive and easy-to-use Multilanguage API. The libarcclient comes loaded with plenty of functionality including Grid job management, resource selection and brokering. It has a pluggable architecture that allows it to implement cross-grid job management. The current implementation comes with plugins for ARC Grid Manager (pre-WS solution), A-REX (extended BES solution), Unicore and CREAM.

Libarcclient is targeted for grid developers and advanced application integrators and provides a general toolkit for high-level cross-middleware, multiplatform client side development.

The library is the core of the highly successful and powerful, nevertheless lightweight ARC clients, grid portals and client-side job managers.

Basic information

Component type: library

Developer: Uppsala University, Copenhagen University, NIIF of NorduGrid Collaboration

Documentation:

http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/doc/tech_doc/client/client_technical.pdf

Code repository:

<http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/src/hed/acc> ,

<http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/src/hed/libs/client>

Development language: C++ (language bindings are offered for Python, Java see below)

Build environment: GNU autotools

Download: download.nordugrid.org

Licence: Apache License Version 2.0

Status

Maturity: pre-production

Usage: as the underlying library of most of the ARC client-side tools it'll be used by every ARC user community

Interface maturity: pre-production

Support status: short term support by KnowARC project, long term support by NorduGrid collaboration

Support team: Uppsala University

Development plans: enhancing maturity of implementation, add support for the emerging PGI-profile

Distribution/Availability

Middleware stack: technology preview release of ARC (version 0.9.3)

Standalone usability: no

Component dependencies: ARC HED

External software dependencies: openssl, glibmm, libxml2

Component consumers: ARC: arc* CLI, ARC:LAP

Supported platforms: all flavours of Linux, MacOS, Solaris

Planned to be supported platforms: Windows

Library component information *(only for a library component)*

Consumers: ARC client tools

Language binding: C++, Python, Java (not tested)

Interoperability/Standards

Standard-compliance: early-UMD

Interoperable components: the library is able to handle CREAM and UNICORE computing elements

Additional standards: already supports many flavors of JSDL and BES with extensions and uses GLUE v2.0 as internal data representation

Other

The libarcclient development library is scheduled to be part of the next production ARC release scheduled for 3rd quarter of this year as a replacement of the aging previous generation ARCLIB.

2.3.4 ARC: ng* job management CLI

Name ARC ng* job management CLI

Description

The ARC ng* job management tools (ngsub, ngget, ngcat, ngstat, ngclean, ngkill, ngresub, ngrenew, ngresume, ngsync) are command line utilities for managing Grid jobs on computing elements equipped with the pre-WS Grid Manager ARC component. Their functionality includes resource selection according to criteria specified in job description, job submission, monitoring, re-submission and result retrieval. The ng* CLI provides a light-weight but powerful grid client including an integrated broker.

Basic information

Component type: client

Developer: Uppsala University of NorduGrid collaboration

Documentation: <http://www.nordugrid.org/documents/ui.pdf>

Code repository:

<http://svn.nordugrid.org/trac/nordugrid/browser/arc0/trunk/cli>

Development language: C++

Build environment: GNU autotools

Download: <http://download.nordugrid.org>

Licence: Apache License Version 2.0

Status

Maturity: production

Usage: one of the most often used tool of the broad NorduGrid community

Interface maturity: production

Support status: long term support

Support Team: Uppsala University from NorudGrid Collaboration

Development plans: none (see other comments)

Distribution/Availability

Middleware stack: ARC production releases (latest version 0.6.5)

Standalone usability: Yes

Component dependencies: ARC: Grid Manager

External software dependencies: pre-ws globus packages (globus-ftp-client, globus-rsl, globus-rsl-client, LFC-client-libs), gsoap, openldap

Component consumers: none

Supported platforms: all the major Linuxes

Planned to be supported platforms: none

Interoperability/Standards

Standard-compliance: none,

Interoperable components: none

Additional standards: makes use of gridftp-based custom interface to submit/manages jobs

Other

The ng* CLI is going to be complemented by the arc* CLI. The transition is scheduled for the 3rd quarter of this year.

2.3.5 ARC: arc* job management CLI

Name arc* job management client tools

Description

The arc* job management client tools implement Grid job management functionality – selection of resource aka brokering, submission of job, monitoring job status, retrieval of results – in a cross-grid environment. They are based on the libarcclient library. The new client offers the same functionality as the old ng* client with significantly improved performance. It also offers new features such as job migration, reliable job resubmission and a plugin structure for custom brokering algorithms.

Basic information

Component type: client

Developer: Uppsala University, Copenhagen University, NIIF from NorduGrid Collaboration

Documentation:

http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/doc/tech_doc/client/client_technical.pdf

Code repository:

<http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/src/clients/arclib>

Development language: C++

Build environment: GNU autotools

Download: <http://download.nordugrid.org/>

Licence: Apache License Version 2.0

Status

Maturity: pre-production

Usage: as the main command line user interface of ARC it'll be used by every ARC user

Interface maturity: pre-production

Support status: short term support by KnowARC project, long term support by NorduGrid collaboration

Support team: Uppsala University

Development plans: enhancing maturity of implementation

Distribution/Availability

Middleware stack: ARC Version 0.9.x

Middleware stack: technology preview releases of ARC (0.9.x)

Standalone usability: no

Component dependencies: libarcclient, ARC:HED

External software dependencies: openssl, glibmm, libxml2

Component consumers: none, it is an end-user tool

Supported platforms: all the major Linuxes, MacOS, Solaris

Planned to be supported platforms: windows

Interoperability/Standards

Standard-compliance: early-UMD

Interoperable components: can manage jobs on CREAM and Unicore computing elements

Additional Standards: supports JSDL, BES and uses GLUE 2.0 as internal data representation

Other

The arc* commands are phasing out the ng* commands. The transition is scheduled for the 3rd

quarter of this year.

2.3.6 ARC: Janitor

Name Janitor

Description

Software installation & management component with a WS interface

Basic information

Component type: service

Developer: Nordugrid Collaboration

Documentation: <http://svn.nordugrid.org/trac/nordugrid/browser/arc1/branches/janitor/doc/janitor>

Code repository: <http://svn.nordugrid.org/trac/nordugrid/browser/arc1/branches/janitor>

Development language: perl

Build environment: gnu autotools

Download: not yet distributed

Licence: Apache License Version 2.0

Status

Maturity: development

Usage: heavily needed e.g. by the bioinformatics community

Interface maturity: development

Support status: best effort

Support team: NorduGrid Collaboration

Development plans: finish the original design

Distribution/Availability

Middleware stack: will be distributed as part of the technology preview releases of ARC (0.9.x)

Standalone usability: no

Component dependencies: ARC:A-REX

External software dependencies: couple of perl modules

Component consumers: none

Supported platforms: all linuxes

Planned to be supported platforms: mac-osx, windows

Service component information (*only for a service component*)

Public Interfaces: custom WS interface for managing application software environment

Aligned security: under development

Capabilities:

Interoperability/Standards

Standard-compliance: none

Interoperable components: none

Additional standards: none

2.3.7 gLite: WMS

Name gLite Workload Management System (WMS)

Description

The Workload Management System (WMS) is a service responsible for the distribution and management of computational jobs across Grid resources, in such a way that applications are conveniently, efficiently and effectively executed. A job is expressed in a flexible Job Description

Language, based on Condor ClassAds. Different types of jobs are supported: simple batch jobs, DAGs, collections, parametric. On top of them other types of jobs, such as MPI and interactive, can be implemented.

Basic information

Component type: *Service*
Developer: EGEE3 JRA1/Workload Management System team
Documentation: <https://edms.cern.ch/document/572489/1>
Code repository: gLite code repository <http://jra1mw.cvs.cern.ch:8180/>
Development language: C++
Build environment: ETICS
Download: <http://glite.web.cern.ch/glite/packages/R3.1/>
Licence: Apache license

Status

Maturity: *production*
Usage: EGEE sites
Interface maturity: *production*
Support status: short term (until the end of the EGEE3 project)
Support team: EGEE3/JRA1 team
Development plans: maintenance

Distribution/Availability

Middleware stack: gLite 3.1 distribution
Standalone usability: no
Component dependencies: LB, BDII, depending on job request (MyProxy, LFC)
External software dependencies: Full list available on request
Supported platforms: SLC4/x86
Planned to be supported platforms: SLC5/x86_64

Service component information

Public Interfaces: WMPProxy (WSDL-based)
Aligned security: *user proxy*
Capabilities:

Library component information

Consumers: gLite WMS User Interface (UI)
Language binding: Java, C++, Python

Interoperability/Standards

Standard-compliance: *none*
Interoperable components: *none*
Additional Standards: gridftp-based transfer of input files

2.3.8 gLite: WMS Clients

Name gLite Workload Management System (WMS) Clients

Description

The WMS clients allow a user or a higher-level application to access the job management services made available by the WMS

Basic information

Component type: *Client*
Developer: EGEE3 JRA1/Workload Management System team

Documentation: <https://edms.cern.ch/document/572489/1>
Code repository: gLite code repository <http://jra1mw.cvs.cern.ch:8180/>
Development language: C++/Python
Build environment: ETICS
Download: <http://glite.web.cern.ch/glite/packages/R3.1/>
Licence: Apache license

Status

Maturity: *production*
Usage: EGEE sites
Interface maturity: *production*
Support status: short term (until the end of the EGEE3 project)
Support team: EGEE3 JRA1 team
Development plans: maintenance

Distribution/Availability

Middleware stack: gLite 3.1 distribution
Standalone usability: no
Component dependencies: gLite Workload Management System
External software dependencies: Full list available on request.
Component consumers: end user tool
Supported platforms: SLC4/x86
Planned to be supported platforms: SLC5/x86_64

Interoperability/Standards

Standard-compliance: none
Interoperable components: none
Additional Standards: none

2.3.9 gLite: LB

Name gLite Logging and Bookkeeping (LB) Service

Description

The primary purpose of the Logging and Bookkeeping service (L&B) is tracking Grid jobs as they are processed by various Grid middleware components. The information that is gathered is used to inform Grid users on the job state, can be used for debugging and detecting problems, and can also serve as an input for various Grid monitoring and statistics tools. The data are gathered in the form of L&B events that are logged actively by middleware components. The L&B infrastructure concentrates the events on a per-job basis and processes the information to give a high-level view on the job state that is presented to the user.

Basic information

Component type: service
Developer: EGEE3 JRA1/LB team
Documentation: <http://egee.cesnet.cz/en/JRA1/LB/>
Code repository: gLite code repository <http://jra1mw.cvs.cern.ch:8180/>
Development language: C/C++
Build environment: ETICS
Download: <http://glite.web.cern.ch/glite/packages/R3.1/>
Licence: Apache license

Status

Maturity: production
Usage: EGEE sites

Interface maturity: production
Support status: short term (until the end of the EGEE3 project)
Support team: EGEE3/JRA1 team
Development plans: maintenance

Distribution/Availability

Middleware stack: gLite 3.1
Standalone usability: no
Component dependencies: gSOAP-plugin
External software dependencies: MySQL, gSOAP
Component consumers: gLite WMS
Supported platforms: SLC4/x86
Planned to be supported platforms: SLC5/x86_64

Service component information

Public Interfaces: Legacy, includes C/C++ bindings for client applications
Aligned security: *none*
Capabilities: *none*

Interoperability/Standards

Standard-compliance: none
Interoperable components: *none*
Additional Standards: none

2.3.10 gLite: CREAM Service

Name CREAM Computing Resource Execution and Management (CREAM) Service

Description

The CREAM (Computing Resource Execution And Management) Service is a web service for job management operations at the Computing Element (CE) level. Its interface is well-defined using the Web Service Description Language (WSDL). The main functionality of CREAM is job submission: CREAM jobs are specified using the same Job Description Language (JDL) used to specify the characteristics and requirements of jobs submitted to the gLite WMS. Other operations supported by CREAM include job cancellation, job status (with configurable level of verbosity and filtering based on submission time and/or job status), job list, job suspension and resume, job purge for terminated jobs. CREAM also allows automatic disabling of new job submissions, when certain conditions are met, e.g. when the number of active jobs is higher than a specified threshold or if the CE administrator explicitly requires this. The interface with the underlying Local Resource Management System (LRMS) is implemented via Batch Local Ascii Helper (BLAH).

Basic information

Component type: Service
Developer: EGEE3 JRA1/CREAM team
Documentation:
CREAM User's Guide <https://edms.cern.ch/document/595770>;
CREAM JDL Attributes Specification <https://edms.cern.ch/document/592336>;
CREAM Java CLI Documentation
<http://grid.pd.infn.it/cream/field.php?n=CREAM.CREAMJavaCLIDocumentation>
Code repository: gLite code repository <http://jra1mw.cvs.cern.ch:8180/>
Development language: Java
Build environment: ETICS
Download: <http://glite.web.cern.ch/glite/packages/R3.1/>
Licence: Apache License

Status

Maturity: Production

Usage: EGEE sites

Interface maturity: Production

Support status: short term (until the end of the EGEE3 project)

Support team: EGEE3/JRA1 team

Development plans: Implement failover mechanisms to enhance resiliency; implement additional standards-compliant job management interface.

Distribution/Availability

Middleware stack: gLite 3.1

Standalone usability: Yes

Component dependencies: BLAH, LCAS/LCMAPS, gLite Delegation Service

External software dependencies: Apache Axis 1.4; Apache Tomcat application server

Component consumers: gLite: CREAM Command Line Tools, ICE (Interface to CREAM Environment)

Supported platforms: SLC4/x86

Planned to be supported platforms: SLC5/x86_64

Service component information

Public Interfaces: Legacy (WSDL-based), includes Java and C++ bindings

Aligned security: TLS/SSL with Globus-style X509 proxies with VOMS Extensions

Capabilities: *none*

Interoperability/Standards

Standard-compliance: non-UMD (OGSA-BES/JSDL only in development branch, not used in production)

Interoperable components: BES interface in development

Additional Standards: JSDL, GLUE

Other

A prototype standard BES/JSDL interface has been implemented in a development branch of the gLite CREAM service. This branch is currently not maintained and unsupported, as the standard BES/JSDL specifications proved not to be appropriate for production use. We plan to enhance the gLite CREAM service by providing an additional standards-compliant interface based on the outcome of the OGF Production Grid Infrastructures (PGI) Working Group.

2.3.11 gLite: CREAM Command Line Tools

Name gLite: CREAM CLI Command Line Tools

Description

The gLite CREAM CE has its own set of CLI tools, which can be used for direct job submission (as opposed to submission through the gLite Workload Management System, for which other command-line tools must be used). The CREAM CLI is made of a collection of commands implemented in C++ using the gSOAP engine for implementing the client-side Web Service interface. Using the CREAM CLI users can access all the operations of the CREAM interface, and submit and manage their jobs. Note that interactions using the CREAM CLI are subject to the same authentication/authorization mechanisms which are employed for job management through the gLite WMS. Authentication and authorization are enforced by CREAM itself, so they are in place regardless of the specific clients which are being used.

Basic information

Component type: Client

Developer: EGEE3 JRA1/CREAM team
Documentation:
CREAM Client API C++ Tutorial and Documentation
<http://grid.pd.infn.it/cream/field.php?n=Main.CREAMClientAPITutorialAndDocumentation>
Code repository: gLite code repository <http://jra1mw.cvs.cern.ch:8180/>
Development language: C++
Build environment: ETICS
Download: <http://glite.web.cern.ch/glite/packages/R3.1/>
Licence: Apache License

Status

Maturity: Production
Usage: EGEE sites
Interface maturity: Production
Support status: short term (until the end of the EGEE3 project)
Support team: EGEE3/JRA1 team
Development plans: maintenance

Distribution/Availability

Middleware stack: gLite 3.1
Standalone usability: No
Component dependencies: gLite CREAM, gSoap-plugin
External software dependencies: gSoap
Component consumers: end user tool
Supported platforms: SLC4/x86
Planned to be supported platforms: SLC5/x86_64

Interoperability/Standards

Standard-compliance: none
Interoperable components: none
Additional Standards: none

2.3.12 gLite: LCG-CE

Name gLite LCG-CE

Description

LCG-CE is a native computing resource access service with Globus Gatekeeper. LCG has modified some of its component to improve its performance like job manager.

Basic information

Component type: Service
Developer: EGEE3 JRA1
Documentation: <http://glite.web.cern.ch/glite/packages/R3.1/deployment/lcg-CE/lcg-CE.asp>
Code repository: gLite code repository <http://jra1mw.cvs.cern.ch:8180/>
Development language: c, perl
Build environment: ETICS
Download: <http://glite.web.cern.ch/glite/packages/R3.1/>
Licence: Apache License

Status

Maturity: Production
Usage: EGEE sites
Interface maturity: Production
Support status: short term (until the end of the EGEE3 project)

Support team: EGEE3/JRA1 team
Development plans: The LCG-CE will be replaced by the CREAM CE

Distribution/Availability

Middleware stack: gLite 3.1
Standalone usability: Yes
Component dependencies: globus
External software dependencies: *globus*
Component consumers: gLite WMS (in particular, the gLite Job Controller/Log Monitor modules)
Supported platforms: SLC4/x86
Planned to be supported platforms: SL4, SL5 ...

Service component information

Public Interfaces: Legacy
Aligned security: TLS/SSL with *Globus-style X509 Proxy Certificates with VOMS extensions*
Capabilities: *none*

Interoperability/Standards

Standard-compliance: none
Interoperable components: *none*
Additional Standards: none

2.3.13 UNICORE: TSI

Name UNICORE Target System Interface (TSI)

Description

The UNICORE TSI is the lowest layer of the UNICORE middleware stack and responsible for interactions with one local resource management system (RMS) that performs the scheduling on a specific computational resource. For each of the supported RMSs (e.g. Torque, Slurm, PBSPro, LoadLeveler, etc.) a dedicated TSI exists. Most of them are implemented in Perl. UNICORE jobs are forwarded by the enhanced Network Job Supervisor (XNJS) to the respective TSI for the execution on a computational resource.

Basic information

Component type: *service*
Developer: UNICORE SourceForge Community
Documentation: http://www.unicore.eu/unicore/architecture/system-layer.php#anchor_tsi (2009-03)
Code repository: <http://unicore.svn.sourceforge.net/svnroot/unicore/tsi/>
Development language: Perl
Build environment: UNIX
Download:
<http://downloads.sourceforge.net/unicore/unicore-servers-6.2.0-p1.tgz>
Licence: BSD License

Status

Maturity: *production*
Usage: *DEISA, D-Grid, SKIF-Grid, commercial partners*
Interface maturity: *production*
Support status: *long term*
Support team: supported by UNICORE SourceForge Community
Development plans: Refactoring for one DRMAA-based Java TSI, improved batch system support

Distribution/Availability

Middleware stack: UNICORE 6.2
Standalone usability: Yes
Component dependencies: none
Software dependencies from external providers: *Perl*
Component consumers: *UNICORE XNJS*
Supported platforms: all systems that can run Perl
Planned to be supported platforms: *None*

Service component information

Public Interfaces: Proprietary interface
Aligned security: *none*
Capabilities: *executionmanagement.jobexecution*

Interoperability/Standards

Standard-compliance: *None*
Interoperable components: *None*
Additional Standards: *DRMAA (development)*

2.3.14 UNICORE: XNJS

Name The UNICORE enhanced Network Job Supervisor (XNJS)

Description

The UNICORE enhanced Network Job Supervisor (XNJS) is the execution backend of the UNICORE service layers (UNICORE Atomic Services or OGSA-BES). It tracks the job progress and interacts with a specific TSI on a respective computational resource. The XNJS is capable of transforming an abstract JSDL-based job submission to a concrete site-specific execution using the UNICORE Incarnation Database (IDB).

Basic information

Component type: *library*
Developer: UNICORE SourceForge Community
Documentation: http://www.unicore.eu/unicore/architecture/service-layer.php#anchor_xnjs
Code repository: <http://unicore.svn.sourceforge.net/svnroot/unicore/xnjs/>
Development language: Java
Build environment: Apache Maven
Download:
<http://downloads.sourceforge.net/unicore/unicore-servers-6.2.0-p1.tgz>
Licence: BSD License

Status

Maturity: *production*
Usage: *DEISA, D-Grid, SKIF-Grid, commercial partners*
Interface maturity: *production*
Support status: *long-term*
Support team: *UNICORE SourceForge Community*
Development plans:

Comment [BK2]: ???

Distribution/Availability

Middleware stack: UNICORE 6.2
Standalone usability: *no*
Component dependencies: *UNICORE TSI.*
External software dependencies: *none*

Component consumers: The UNICORE Atomic Services (UAS) and UNICORE OGSA-BES rely on this particular execution backend.

Supported platforms: All Java-supported platforms

Planned to be supported platforms: None

Library component information

Consumers: *UNICORE Atomic Services, UNICORE OGSA-BES, UNICORE Workflow service*

Language binding: *Java*

Interoperability/Standards

Standard-compliance: *non-umd (JSDL)*

Interoperable components: *None*

Additional Standards: *File transfers support HTTP, FTP, GridFTP*

2.3.15 UNICORE: UAS (Job Submission and Management)

Name The UNICORE Atomic Services (UAS) elements TargetSystemFactory service, TargetSystem service, and JobManagement service

Description

The WS-RF-based UNICORE Atomic Services (UAS) are the proprietary Web service interfaces of the UNICORE middleware. Among other WS related to file transfers and storage management, three services are related to job submission and management. Firstly, the TargetSystemFactory service acts as a factory for user-specific TargetSystem service instances that represent computational resources. One TargetSystem service in turn allows for the submission of JSDL-compliant jobs that can be handled via JobManagement service instances. The UAS rely on the XNJS for optimal execution. For managing the job's working directory and handling file transfers, the StorageManagement and FileTransfer service are used.

Basic information

Component type: *service*

Developer: *UNICORE SourceForge Community*

Documentation:

http://www.unicore.eu/unicore/architecture/service-layer.php#anchor_uas (2009-03)

<http://www.unicore.eu/documentation/manuals/unicore6/unicorex/index.html>

Code repository: <http://unicore.svn.sourceforge.net/svnroot/unicore/unicorex/trunk/uas-core/>

Development language: *Java*

Build environment: *Apache Maven*

Download:

<http://downloads.sourceforge.net/unicore/unicore-servers-6.2-0-p1.tgz>

Licence: *BSD License*

Status

Maturity: *production*

Usage: *DEISA, D-Grid, SKIF-Grid, commercial partners*

Interface maturity: *production*

Support status: *long-term*

Support team: *UNICORE SourceForge Community*

Development plans: [

Comment [BK3]: ???

Distribution/Availability

Middleware stack: *UNICORE 6.2*

Standalone usability: *Yes* Component dependencies: *UNICORE/X service container, UNICORE*

storage and filetransfer services

 External software dependencies: *none*

Supported platforms: *All Java-based platforms*

Planned to be supported platforms: None

Service component information

Public Interfaces: TargetSystemFactory service, TargetSystem service, JobManagement service
Aligned security: TLS/SSL (optional with proxies) and SAML Assertions in the SOAP Header
Capabilities: executionmanagement.jobmanager

Interoperability/Standards

Standard-compliance: non-umd (JSDL)
Interoperable components: *None*
Additional Standards: WS(RF)

2.3.16 UNICORE: OGSA-BES Adoption

Name UNICORE OGSA – Basic Execution Services (BES) Adoption

Description

The OGSA-BES services of UNICORE represent the standard alternative to the job submission and management elements of the UNICORE UAS.

Basic information

Component type: *service*
Developer: FZJ, UNICORE SourceForge Community
Documentation: http://www.unicore.eu/unicore/architecture/service-layer.php#anchor_ogsa
Code repository: <http://unicore.svn.sourceforge.net/viewvc/unicore/unicorex/trunk/ogsabes-core/>
Development language: Java
Build environment: Apache Maven
Download:
<http://downloads.sourceforge.net/unicore/unicore-servers-6.2.0-p1.tgz>
Licence: BSD License

Status

Maturity: pre-production
Usage: WISDOM, Virtual Physiological Human, EUFORIA
Interface maturity: production
Support status: long-term
Support team: supported by FZJ and the UNICORE SourceForge Community
Development plans: Push for real production usage

Distribution/Availability

Middleware stack: UNICORE 6.2
Standalone usability: No
Component dependencies: UNICORE/X container
External software dependencies: none
Component consumers:
Supported platforms: All Java-based platforms
Planned to be supported platforms: None

Comment [BK4]: ???

Service component information

Public Interfaces: BESFactory service, BESManagement service, BESActivity service
Aligned security: TLS/SSL (optional with proxies) and SAML assertions in SOAP header
Capabilities: executionmanagement.jobmanager

Interoperability/Standards

Standard-compliance: non-umd (OGSA-BES JSDL)

Interoperable components: *interoperability with component CREAM-BES of gLite and A-REX of ARC, GridSAM of OMII-UK, and other OGSA-BES components*
Additional Standards: *HPC-Basic Profile, WS(RF)*

2.3.17 UNICORE: Workflow Engine

Name UNICORE Workflow Engine

Description

The workflow support in UNICORE 6 is implemented as a two-layered architecture consisting of a workflow engine and the service orchestrator layers. Besides enhanced scalability this also has the benefit of a pluggable workflow engine. Different workflow description dialects, which comply to domain-specific requirements, can be plugged in. Due to the service oriented approach, even plugging in a complete different workflow engine is possible. The currently available workflow engine is based on the open-source Shark XPDL engine and was developed within the Chemomomentum project.

Basic information

Component type: *service*
Developer: FZJ and UNICORE SourceForge Community
Documentation: http://www.unicore.eu/unicore/architecture/service-layer.php#anchor_workflow
Code repository: <http://unicore.svn.sourceforge.net/svnroot/unicore/workflow>
Development language: Java
Build environment: Apache Maven
Download:
<http://downloads.sourceforge.net/unicore/unicore-workflow-6.1.3.tgz>
Licence: BSD License

Status

Maturity: *pre-production*
Usage: *DEISA, D-Grid, SKIF-Grid, commercial partners* Interface maturity: production
Support status: *long-term*
Support team: *support by FZJ and UNICORE SourceForge Community*
Development plans: feature enhancements, production hardening

Distribution/Availability

Middleware stack: UNICORE 6
Standalone usability: yes
Component dependencies: UNICORE Service Orchestrator
External software dependencies: none
Component consumers:
Supported platforms: All Java-based platforms
Planned to be supported platforms: None

Comment [BK5]: ???

Service component information

Public Interfaces: WorkflowFactoryservice, WorkflowManagement service
Aligned security: *TLS/SSL (optional proxies) and SAML assertions in SOAP Headers*
Capabilities: executionmanagement.executionandplanning

Interoperability/Standards

Standard-compliance: *None*
Interoperable components: *None*
Additional Standards: WS(RF)

2.3.18 UNICORE: Service Orchestrator

Name UNICORE Service Orchestrator

Description

The service orchestrator layer is responsible to executing the individual tasks in a workflow, handling job execution and monitoring on the Grid. Different brokering strategies are implemented to find the best suited resources for each workflow step. Other brokering strategies can easily be plugged-in. To increase performance and scalability of the service orchestrator layer, multiple instances can be installed.

Basic information

Component type: *service*

Developer: FZJ & UNICORE SourceForge Community

Documentation: http://www.unicore.eu/unicore/architecture/service-layer.php#anchor_servorch

Code repository: [\[redacted\]](#)

Comment [BK6]: ???

Development language: Java

Build environment: Apache Maven

Download:

http://downloads.sourceforge.net/unicore/unicore-workflow-6.1.3.tgz?use_mirror=switch

Licence: BSD License

Status

Maturity: *pre-production*

Usage: *DEISA, D-Grid, SKIF-Grid, commercial partners* Interface maturity: *production*

Support status: *long-term*

Support team: *supported by FZJ & UNICORE SourceForge Community*

Development plans: More brokering algorithms, integration of different brokers, standards support (e.g. OGSA-BES)

Distribution/Availability

Middleware stack: UNICORE Workflow System Version 6

Standalone usability: *Yes*

Component dependencies: UNICORE Atomic services

External Software dependencies: *none*

Component consumers: *UNICORE Workflow service*

Supported platforms: *All Java-based platforms*

Planned to be supported platforms: *None*

Service component information

Public Interfaces: *ServiceOrchestrator service*

Aligned security: *TLS/SSL (optional proxies) and SAML assertions in SOAP Headers*

Capabilities: *executionmanagement.executionandplanning*

Interoperability/Standards

Standard-compliance: *None*

Interoperable components: *None*

Additional Standards: *WSRF*

2.3.19 UNICORE: Rich Client

Name UNICORE Rich Client

Description

The Eclipse-based UNICORE Rich Client (URC) offers the full set of functionalities to the users in

a graphical representation. It targets a wide range of users with varying Grid and IT experience. It provides a useful graphical view of the Grid, which can be filtered in order to find specific resources, services or files. Detailed resource requirements for jobs (e.g. required number of CPUs, amount of RAM) can be specified. Most notably, it allows for powerful mechanisms to construct directed acyclic graphs (DAGs) used for Grid workflows.

Basic information

Component type: *client*

Developer: *UNICORE SourceForge Community*

Documentation: http://www.unicore.eu/unicore/architecture/client-layer.php#anchor_urc

Code repository: <http://unicore.svn.sourceforge.net/viewvc/unicore/eclipseclient/>

Development language: Java

Build environment: Apache Ant, Apache Maven

Download: http://sourceforge.net/project/showfiles.php?group_id=102081&package_id=263955

Licence: BSD License

Status

Maturity: *production*

Usage: *DEISA, D-Grid, SKIF-Grid, commercial partners* Interface maturity: *production*

Support status: *long-term*

Support team: supported by FZJ and UNICORE SourceForge Community

Development plans: Improve and extend functionality

Distribution/Availability

Middleware stack: *UNICORE 6*

Standalone usability: Yes

Component dependencies: UNICORE Atomic services

Supported platforms: *All Java-based platforms (except Mac OS/X)*

Planned to be supported platforms: *Mac OS/X*

Interoperability/Standards

Standard-compliance: *None*

Interoperable components: *None*

Additional Standards: *Web Service Resource Framework (WS-RF), Eclipse rich client platform*

2.3.20 UNICORE: Command-line Client

Name UNICORE Command-line Client (UCC)

Description

The UNICORE command line client (UCC) is a very versatile command-line tool that allows users to access all features of the UNICORE service layer in a shell or scripting environment. It allows to run jobs, monitor their status and retrieve generated output, both in single job mode or in a powerful and flexible batch mode for multiple jobs. Additionally, workflows can be submitted and controlled with the UCC. Via extensions it also allows for OGSA-BES-based job submissions.

Basic information

Component type: *client*

Developer: *UNICORE SourceForge Community*

Documentation:

http://www.unicore.eu/unicore/architecture/client-layer.php#anchor_ucc (2009-03)

Code repository:

<http://unicore.svn.sourceforge.net/viewvc/unicore/ucc/>

Development language: Java

Build environment: Java

Download: http://sourceforge.net/project/showfiles.php?group_id=102081&package_id=263954
Licence: BSD

Status

Maturity: *production*
Usage: *DEISA, D-Grid, SKIF-Grid, commercial partners*
Interface maturity: *production*
Support status: *long-term*
Support team: *FZJ and UNICORE SourceForge community*
Development plans: *keep up with server-side developments, new features*

Distribution/Availability

Middleware stack: *UNICORE 6*
Standalone usability: *Yes*
Component dependencies: *UNICORE Atomic services*
External software dependencies: *none*
Component consumers:
Supported platforms: *All Java-based platforms*
Planned to be supported platforms: *None*

Comment [BK7]: ???

Interoperability/Standards

Standard-compliance: *non-umd (OGSA-BES)*
Interoperable components: *ARC A-REX and CREAM-BES*
Additional Standards: *Web Services Resource Framework (WS-RF)*

2.3.21 UNICORE: HILA

Name UNICORE High level API

Description

HiLA is a High Level API for Grid Applications, that allows simple development of clients with just a few lines of code for otherwise complex functionality. It can also be used to integrate UNICORE 6 in user applications to e.g. directly submit simulation jobs to Grid resources without using UNICOREs command line or graphical clients. HiLA provides a single interface with multiple implementations for UNICORE 5, UNICORE 6 and OGSA-BES. HiLA is used to integrate UNICORE 6 access into DESHL (including SAGA support), as part of the DEISA portal, and to connect GAT with UNICORE 6. The nature of the HiLA API leads to a concise coding toolkit for building collective tier Grid services and client interfaces. Following the UNICORE principle of seamlessness, the design of the API models the Grid through a object-oriented facade, presenting abstract representations of the underlying resources. Importantly, this includes encapsulating security configuration behind well defined interfaces, further enhancing the uncluttered API.

Basic information

Component type: *library*
Developer: *UNICORE SourceForge Community*
Documentation:
http://www.unicore.eu/unicore/architecture/client-layer.php#anchor_hila (2009-03)
Code repository: <http://unicore.svn.sourceforge.net/svnroot/unicore/hila>
Development language: *Java*
Build environment: *Apache Maven*
Download: http://sourceforge.net/project/showfiles.php?group_id=102081&package_id=225019
Licence: *BSD License*

Status

Maturity: *pre-production*
 Usage: *DE/ISA*
 Interface maturity: *pre-production*
 Support status: *long-term*
 Support team: FZJ and the UNICORE SourceForge Community
 Development plans: improve for production use, new features

Distribution/Availability

Middleware stack: *UNICORE 6*
 Standalone usability: *Yes*
 Component dependencies: *UNICORE Atomic Services*
 External software dependencies:
 Component consumers:
 Supported platforms: *All Java-based Platforms*
 Planned to be supported platforms: *None*

Comment [BK8]: ???

Library component information (*only relevant for a library*)

Consumers: *JavaGAT, DESHL*
 Language binding: *Java*

Interoperability/Standards

Standard-compliance:
 Interoperable components:
 Additional Standards: *WS-RF*

Comment [BK9]: ???

3 Data functionality

3.1 Standards landscape

To be provided later...

Overview of what existing, followed data related relevant specifications.. standards in this area have to be evolve...or we have to drive them more...even SRM was firstly out of OGF and has been just pushed to OGF later when it was already adopted by many...

Discussion whether there are already existing set of standards, which can be used to achieve std-based interoperability of UMD data components.

Recommendation of what an UMD data component should comply with

3.2 Overview table: data functionality

	Name	Type	Maturity	Standard compliance		
3.3.1	ARC: Classic SE	service	production	None-umd		
3.3.2	ARC: libarcdata2	library	Pre-production	Non-umd		
3.3.3	ARC: ng* data management CLI	client	production	Non-umd		
3.3.4	ARC: arc* data management CLI	client	Pre-production	Non-umd		
3.3.5	ARC: Chelonia	service	Pre-production	none		
3.3.6	ARC: Chelonia client	client	Pre-production	none		
3.3.7	gLite: DPM	service	production	Non-umd		
3.3.8	gLite: lcg_utils	client	production	none		
3.3.9	gLite: htcp	client	production	Non-umd		

3.3.10	gLite: SRM client	client	production	Non-umd		
3.3.11	gLite: DPM client	client	production	None		
3.3.12	gLite: LFC client	client	production	None		
3.3.13	gLite: FTS client	client	production	None-umd		
3.3.14	gLite: DPM/DICOM interface	client	production	None		
3.3.15	gLite: LCG File Catalogue (LFC)	service	production	None		
3.3.16	gLite: File Transfer Service (FTS)	service	production	Non-umd		
3.3.17	gLite: Hydra	service	production	None		
3.3.18	gLite: AMGA	service	production	None		
3.3.19	gLite: AMGA Client API	library	production	None		
3.3.20	UNICORE: UAS (Storage Management and File Transfer)	service	production	Non-umd		
3.3.21	Other: dCache	service	production	Non-umd		
3.3.22	Other: StoRM	service	production	Non-umd		

3.3 Inventory of components

3.3.1 ARC: Classic SE

Name ARC Classic SE

Description

A Gridftp server developed using the Globus libraries. The Classic SE comes with powerful backend plugins presenting various information using FTP interface. Those include: jobplugin for Grid job submission and management, gaclplugin for storing ordinary files with assigned access policies in GACL language, fileplugin for ordinary FTP interface.

Basic information

Component type: service
Developer: Oslo University
Documentation: The NorduGrid Grid Manager And GridFTP Server,
<http://www.nordugrid.org/documents/GM.pdf>
Code repository: <http://svn.nordugrid.org/trac/nordugrid/browser/arc0/>
Development language: c++
Build environment: GNU autotools
Download: download.nordugrid.org
Licence: Apache License Version 2.0

Status

Maturity: production
Usage: a key ARC component as the gridftp job plugin provides the job submission interface for ARC computing element. The ftp server is also used as a simple storage element.
Interface maturity: production
Support status: long term
Support team: NorduGrid Collaboration
Development plans: only bugfixes

Distribution/Availability

Middleware stack: ARC production releases (latest version 0.6.5)
Standalone usability: yes
Component dependencies: none
External software dependencies: globus pre-WS libraries, gsoap
Component consumers: ARC Grid Manager, ARC data clients
Supported platforms: major linux flavours
Planned to be supported platforms: none

Service component information (*only for a service component*)

Public Interfaces: Gridftp
Aligned security:
Capabilities:

Interoperability/Standards

Standard-compliance: non-UMD (GridFTP)
Interoperable components: any gridftp-capable data client
Additional standards: -

Other

With the new ARC:A-REX component the ARC Gridftp server soon will not be needed as a job submission interface module, nevertheless the ARC Classic SE is to be supported as long as there is a need for a simple GridFTP server.

3.3.2 ARC: libarcdata2

Name ARC data management library (libarcdata2)

Description

The component implements a general purpose data library with support for multiple data transfer and management protocols. Due to its pluggable architecture it is relatively easy to add support for new protocols. Currently supported protocols are: HTTP(S,G), (Grid)FTP, LDAP, LFC, RLS, SRM and Chelonia.

The libarcdata is targeted for grid developers and advanced application integrators and provides a general toolkit for high-level cross-middleware, multiplatform client side development. The library provides the core of the data management layer of the highly successful and powerful, nevertheless lightweight ARC clients, grid portals and client-side job managers.

Basic information

Component type: library
Developer: Uppsala University and University of Oslo from NorduGrid Collaboration
Documentation:
http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/doc/internal_notes/dmc/dmc.pdf
Code repository:
<http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/src/hed/dmc>
<http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/src/hed/libs/data>
Development language: C++
Build environment: GNU autotools
Download: download.nordugrid.org
Licence: Apache License version 2.0

Status

Maturity: pre-production
Usage: as the underlying library of all ARC data management the library will be widely deployed
Interface maturity: pre-production
Support status: short term support by KnowARC project, long term support by NorduGrid

collaboration
Support team: Uppsala University
Development plans: enhance maturity, decouple from gsoap

Distribution/Availability

Middleware stack: technology preview releases of ARC (0.9.x)
Standalone usability: no
Component dependencies: ARC:HED
External software dependencies: globus data libraries (globus-ftp-client, globus-rsl-client), LFC-client-libs, gsoap, openldap
Component consumers: ARC: arc* CLI
Supported platforms: all the major flavours of linuxes, mac-osx
Planned to be supported platforms: windows (most of the effort needed is to port the globus dependencies)

Library component information *(only for a library component)*

Consumers: ARC: arc* CLI
Language binding: Python, java

Interoperability/Standards

Standard-compliance: non-UMD (gridftp, HTTP(S), SRM storage services are supported)
Interoperable components: any storage system with SRM interface
Additional standards: -

Other

The library is scheduled for inclusion into production ARC release for the 3rd quarter of this year.

3.3.3 ARC: ng* data management CLI

Name ng* data management CLI

Description

The ng* data management tools (ngls, ngcp, ngrm, ngtransfer, ngstage) are based on older version of libarcdata library and feature several data protocols. They are powerful command line tools support listing, removing and copying data.

Basic information

Component type: client
Developer: Uppsala University and Oslo University of NorduGrid Collaboration
Documentation: <http://svn.nordugrid.org/trac/nordugrid/browser/arc0/trunk/doc/ui/ui.pdf>
Code repository: <http://svn.nordugrid.org/trac/nordugrid/browser/arc0/trunk/grid-manager>
Development language: C++
Build environment: GNU autotools
Download: download.nordugrid.org
Licence: Apache License version 2.0

Status

Maturity: production
Usage: widely used as the client-side data handling utility
Interface maturity: production
Support status: long term
Support team: Oslo University
Development plans: none

Distribution/Availability

Middleware stack: production ARC releases (latest version 0.6.5)
Standalone usability: yes
Component dependencies: none
External software dependencies: globus-ftp-client, globus-rls-client, LFC-client-libs, gsoap
Component consumers: none, it is a standalone client
Supported platforms: all the major linux flavours
Planned to be supported platforms: none

Interoperability/Standards

Standard-compliance: non-UMD (supports SRM, gridftp)
Interoperable components: all the SRM, gridftp-based storage solutions
Additional standards:

Other

This component is being gradually phased out by the next generation arc* data management clients.

3.3.4 ARC: arc* data management CLI

Name arc* data management client tools

Description

The arc* data management tools (arccp, arcls, arcrm) are based on the libarcdata2 library and feature support for several data protocols. Additional protocols can be added by writing corresponding plugins. These command line tools support listing, removing and copying of data.

Basic information

Component type: client
Developer: Uppsala University and Oslo University from NorduGrid Collaboration
Documentation:
http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/doc/tech_doc/client/client_technical.pdf
Code repository: <http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/src/clients/data>
Development language: C++
Build environment: GNU autotools
Download: download.nordugrid.org
Licence: Apache License version 2.0

Status

Maturity: pre-production
Usage: it will be widely used as the main client-side data handling utility replacing the ng* data CLI
Interface maturity: pre-production
Support status: short term support by KnowARC project, long term support by NorduGrid collaboration
Support team: Uppsala University
Development plans: support for third party transfers

Distribution/Availability

Middleware stack: technology preview releases of ARC (0.9.x)
Standalone usability: yes
Component dependencies: ARC: libarcdata2
External software dependencies: no direct dependency but through libarcdata2 it has optional dependencies on globus data libraries (globus-ftp-client, globus-rsl-client), LFC-client-libs, gsoap, and openldap
Component consumers: none, it is a standalone client

Supported platforms: Linux, MacOS, Solaris
Planned to be supported platforms: windows

Interoperability/Standards

Standard-compliance: non-UMD (supports SRM, gridftp, http(s))
Interoperable components: any SRM-based storage system
Additional standards: -

Other

This component is gradually phasing out the old ng* data management clients. The tools are scheduled for inclusion into production ARC release for the 3rd quarter of the year.

3.3.5 ARC: Chelonia

Name ARC Chelonia storage system

Description

Chelonia is a user-friendly, grid enabled and replicated storage system with a hierarchical global namespace. Users of Chelonia can easily upload and download files and collections to the storage cloud managed by Chelonia and share their files (data, photos, movies etc) with any other user of the system. Contrary to dCache, which has a set of centralized core services, Chelonia is distributed by design which allows for a flexible setup and the elimination of single points of failure. It is possible to integrate the namespace of third-party storage systems (e.g. dCache) into the Chelonia namespace in a similar way as mounting remote filesystems into a local filesystem. Chelonia is non-intrusive by design.

Basic information

Component type: service
Developer: NIIF Institute, Uppsala University and Oslo University from NorduGrid Collaboration
Documentation:
http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/doc/tech_doc/storage_system/arc-storage-documentation.pdf
http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/doc/tech_doc/storage_system/arc-storage-manual.pdf
Code repository:
<http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/src/services/storage>
Development language: python
Build environment: GNU autotools
Download: download.nordugrid.org
Licence: Apache License version 2.0

Status

Maturity: pre-production
Usage: first will be deployed in the Hungarian ClusterGrid
Interface maturity: production
Support status: short term via KnowARC project and long term via NorduGrid Collaboration
Support team: NIIF Institute from NorduGrid Collaboration
Development plans: performance and stability improvements, add support for SRM interface

Distribution/Availability

Middleware stack: technology preview releases of ARC (0.9.x)
Standalone usability: no
Component dependencies: ARC: libarcdata2; ARC: HED
External software dependencies: SWIG (1.3.28 or later), Python dev libraries (2.4 or later)
Component consumers: ARC: arc* data cli

Supported platforms: Linux, Mac OS X, Solaris
Planned to be supported platforms: Windows

Service component information

Public Interfaces: proprietary WS-based interface
Aligned security: TLS/SSL, proxies are supported, VOMS attributes are supported
Capabilities: data.access.flatfile, data.management.replica, data.management.storage,
data.transfer, data.naming.resolver

Interoperability/Standards

Standard-compliance: none
Interoperable components: none
Additional Standards: bytelO, HTTP(s)

Other

Chelonia is scheduled to be included into the production ARC releases for 1st quarter of 2010.

3.3.6 ARC: Chelonia client

Description

Client solutions for the Chelonia storage system, including a command line tool and a module for FUSE enabling transparent access to Chelonia-enabled Grid storage resources.

Basic information

Component type: client
Developer: NIIF Institute, Uppsala University and Oslo University from NorduGrid Collaboration
Documentation:
http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/doc/tech_doc/storage_system/arc-storage-documentation.pdf
http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/doc/tech_doc/storage_system/arc-storage-manual.pdf
Code repository
<http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/src/services/storage/cli>
Development language: Python
Build environment: GNU autotools
Download: download.nordugrid.org
Licence: Apache License version 2.0

Status

Maturity: pre-production
Usage: The client tools will be tested by the Hungarian ClusterGrid users
Interface maturity: production
Support status: short term via the KnowARC project, long term through NorduGrid Collaboration
Support team: NIIF, Oslo from NorduGrid Collaboration
Development plans: increase user-friendliness and ease-of-use, GUI tools

Distribution/Availability

Middleware stack: technology preview releases of ARC (0.9.x)
Standalone usability: no
Component dependencies: ARC: libarcdata2, ARC:HED
External software dependencies: SWIG (1.3.28 or later), Python dev libraries (2.4 or later)
Component consumers: none (end users)
Supported platforms: Linux, Mac OS X, Solaris
Planned to be supported platforms: Windows

Interoperability/Standards

Standard-compliance: none
Interoperable components: none
Additional standards: HTTP(s)

Other

Chelonia clients are scheduled for ARC production release inclusion for 1st quarter of 2010.

3.3.7 gLite: DPM

Name gLite Disk Pool Manager (DPM)

Description

The light-weight Disk Pool Manager (DPM) offers a simple solution for a disk-only Storage Element. It is easy to install and configure and requires very low maintenance effort. The DPM consists of a set of servers with their own client interface. All the interfaces work only in secure mode using X509 certificates or Kerberos 5 tokens and can use VOMS based ACLs. The DPM is deployed in about 120 sites within EGEE.

Basic information

Component type: Service
Developer: EGEE3 JRA1
Documentation: <https://twiki.cern.ch/twiki/bin/view/LCG/DpmGeneralDescription>
Code repository: gLite code repository <http://jra1mw.cvs.cern.ch:8180/>
Development language: C
Build environment: ETICS
Download: <http://glite.web.cern.ch/glite/packages/R3.1/>
Licence: Apache License

Status

Maturity: Production
Usage: Disk only storage solution used in production
Interface maturity: Production
Support status: short term (until the end of the EGEE3 project)
Support team: EGEE3/JRA1 team
Development plans: maintainance

Distribution/Availability

Middleware stack: gLite 3.1
Standalone usability: yes
Component dependencies: none
External software dependencies: full list available on request
Component consumers:
Supported platforms: SLC4/x86
Planned to be supported platforms: SLC5/x86_64

Service component information

Public Interfaces: SRM
Aligned security: X509 certificates or Kerberos 5 tokens and can use VOMS based ACLs
Capabilities: *none*

Interoperability/Standards

Standard-compliance: non-umd (SRM, GridFTP)
Interoperable components: *SRM interface, GridFTP*
Additional Standards: rfiio, HTTP(S), xrootd

3.3.8 gLite: lcg_utils

Name gLite lcg_utils

Description

In the lcg_utils a mixture of GFAL and Globus functionality was used to create Grid equivalents of UNIX copy commands and registration that spans two storage systems and catalogues of choice (i.e. LFC, EDG). Other functions have been implemented, such as getting the list of file replicas from many sources, performing lookups on LFNs or GUIDs, changing the status of the file or finally removing it from the Grid.

Basic information

Component type: Client

Developer: EGEE3 JRA1

Documentation: <https://edms.cern.ch/file/454439/LCG-2-UserGuide.html#SeDMClientTools>

Code repository: gLite code repository <http://jra1mw.cvs.cern.ch:8180/>

Development language: C

Build environment: ETICS

Download: <http://glite.web.cern.ch/glite/packages/R3.1/>

Licence: Apache License

Status

Maturity: Production

Usage: End user tool on the WN and UI

Interface maturity: Production

Support status: short term (until the end of the EGEE3 project)

Support team: EGEE3/JRA1 team

Development plans: maintainance

Distribution/Availability

Middleware stack: gLite 3.1

Standalone usability: *SRM and LFC*

Component dependencies: none

External software dependencies: full list on request

Component consumers: none

Supported platforms: SLC4/x86

Planned to be supported platforms: SLC5/x86_64

Interoperability/Standards

Standard-compliance: none

Interoperable components: none

Additional Standards: *none*

3.3.9 gLite: htcp

Name gLite htcp

Description

htcp is a client to fetch files or directory listings from remote servers using HTTP or HTTPS, or to put or delete files or directories onto remote servers using HTTPS. htcp is similar to the scp command, but uses HTTP/HTTPS rather than ssh as its transfer protocol.

Basic information

Component type: Client
Developer: EGEE3 JRA1
Documentation: <http://www.gridsite.org/doc/1.1.x/htcp.1.html>
Code repository: gLite code repository <http://jra1mw.cvs.cern.ch:8180/>
Development language: C
Build environment: ETICS
Download: <http://glite.web.cern.ch/glite/packages/R3.1/>
Licence: Apache License

Status

Maturity: Production
Usage: UI and WN
Interface maturity: Production
Support status: short term (until the end of the EGEE3 project)
Support team: EGEE3/JRA1 team
Development plans: maintainance

Distribution/Availability

Middleware stack: gLite 3.1
Standalone usability: *yes*
Component dependencies: none
External software dependencies: *none*
Component consumers: used for data management operations
Supported platforms: SLC4/x86
Planned to be supported platforms: SLC5/x86_64

Interoperability/Standards

Standard-compliance: non-umd
Interoperable components: none
Additional Standards: http

3.3.10 gLite: SRM client

Name gLite SRM client

Description

Command line Client for the SRM interface

Basic information

Component type: Client
Developer: EGEE3 JRA1
Documentation: <https://edms.cern.ch/file/454439/LCG-2-UserGuide.html>
Code repository: gLite code repository <http://jra1mw.cvs.cern.ch:8180/>
Development language: Java
Build environment: ETICS
Download: <http://glite.web.cern.ch/glite/packages/R3.1/>
Licence: Apache License

Status

Maturity: Production
Usage: UI and WN
Interface maturity: Production
Support status: short term (until the end of the EGEE3 project)
Support team: EGEE3/JRA1 team
Development plans: maintenance

Distribution/Availability

Middleware stack: gLite 3.1
Standalone usability: *yes*
Component dependencies: none
External software dependencies: full list available on request
Component consumers: end users
Supported platforms: SLC4/x86
Planned to be supported platforms: SLC5/x86_64

Interoperability/Standards

Standard-compliance: non-umd (SRM)
Interoperable components: none
Additional Standards: -

3.3.11 gLite: DPM client

Name gLite DPM client

Description

Command line Client for the DPM interface

Basic information

Component type: Client
Developer: EGEE3 JRA1
Documentation: <http://twiki.cern.ch/twiki/bin/view/LCG/DpmGeneralDescription>
Code repository: gLite code repository <http://jra1mw.cvs.cern.ch:8180/>
Development language: C
Build environment: ETICS
Download: <http://glite.web.cern.ch/glite/packages/R3.1/>
Licence: Apache License

Status

Maturity: Production
Usage: Sites running the DPM
Interface maturity: Production
Support status: short term (until the end of the EGEE3 project)
Support team: EGEE3/JRA1 team
Development plans: maintenance

Distribution/Availability

Middleware stack: gLite 3.1
Standalone usability: *DRM*
Component dependencies: none
External software dependencies: *full list available on request*
Component consumers: end users
Supported platforms: SLC4/i386
Planned to be supported platforms: SLC5/x86_64

Interoperability/Standards

Standard-compliance: none
Interoperable components: none
Additional Standards: none

3.3.12 gLite: LFC client

Name gLite LFC client

Description

Command line Client for LFC

Basic information

Component type: Client
Developer: EGEE3 JRA1
Documentation: <https://twiki.cern.ch/twiki/bin/view/LCG/LfcAdminGuide>
Code repository: gLite code repository <http://jra1mw.cvs.cern.ch:8180/>
Development language: C
Build environment: ETICS
Download: <http://glite.web.cern.ch/glite/packages/R3.1/>
Licence: Apache License

Status

Maturity: Production
Usage: UI and WN
Interface maturity: Production
Support status: short term (until the end of the EGEE3 project)
Support team: EGEE3/JRA1 team
Development plans: maintenance

Distribution/Availability

Middleware stack: gLite 3.1
Standalone usability: yes
Component dependencies: none
External software dependencies: full list available on request
Component consumers: end user
Supported platforms: SLC4/x86
Planned to be supported platforms: SLC5/x86_64

Interoperability/Standards

Standard-compliance: none
Interoperable components: none
Additional Standards: none

3.3.13 gLite: FTS client

Name gLite FTS client

Description

Command line Client for FTS

Basic information

Component type: Client
Developer: EGEE3 JRA1
Documentation: <https://edms.cern.ch/file/454439//LCG-2-UserGuide.html>
Code repository: gLite code repository <http://jra1mw.cvs.cern.ch:8180/>
Development language: Java
Build environment: ETICS
Download: <http://glite.web.cern.ch/glite/packages/R3.1/>
Licence: Apache License

Status

Maturity: Production
Usage: UI
Interface maturity: Production
Support status: short term (until the end of the EGEE3 project),
Support team: EGEE3/JRA1 team
Development plans: maintainance

Distribution/Availability

Middleware stack: gLite 3.1
Standalone usability: yes
Component dependencies: FTS
External software dependencies: *full list available on request*
Component consumers: end user
Supported platforms: SLC4/x86
Planned to be supported platforms: SLC5/x86_64

Interoperability/Standards

Standard-compliance: none-umd (SRM, GridFTP)
Interoperable components:none
Additional Standards: -

3.3.14 gLite: DPM/DICOM interface

Name gLite DPM/DICOM interface

Description

The Digital Imaging and Communications in Medicine (DICOM) is a standard medical image storage system used by Biomedical researchers, an important user community in EGEE. gLite includes an interface between DICOM and the DPM storage element. This DPM/DICOM interface is a command line client that takes anonymized medical images from a DICOM system, enters the image metadata into a metadata storage system such as AMGA, uses EDS and Hydra to encrypt the image, stores the resulting file in DPM and stores storage metadata in LFC. These encrypted files may then be analyzed or studied in a Grid environment. The planned support effort for DPM/DICOM is driven by the requests from the user community, currently the Biomedical researchers.

Basic information

Component type: Client
Developer: EGEE3 JRA1
Documentation: <https://twiki.cern.ch/twiki/bin/view/EGEE/DMEDS>
Code repository: gLite code repository <http://jra1mw.cvs.cern.ch:8180/>
Development language: C
Build environment: ETICS
Download: <http://glite.web.cern.ch/glite/packages/R3.1/>
Licence: Apache License

Status

Maturity: Production
Usage: biomed communtiy
Interface maturity: mature
Support status: short term (until the end of the EGEE3 project)
Support team: EGEE3/JRA1 team
Development plans: maintainance

Distribution/Availability

Middleware stack: gLite 3.1
Standalone usability: *no*
Component dependencies: DPM
External software dependencies: *full list available on request*
Component consumers: none
Supported platforms: SLC4/x86
Planned to be supported platforms: SLC5/x86_64

Interoperability/Standards

Standard-compliance: none
Interoperable components: none
Additional Standards: none

3.3.15 gLite: LCG File Catalogue (LFC)

Name LCG File Catalogue (LFC)**Description**

The LFC offers a hierarchical view of files to users and provides Logical File Name (LFN) to Storage URL (SURL) mappings (via a GUID) with authorization on its logical namespace. Read-only distributed catalog can be achieved by using Oracle streams. The LFC is deployed at most EGEE European and Asian sites, in addition to the CERN site. It supports several deployment scenarios; central, local and local readonly replica of global.

Basic information

Component type: Service
Developer: EGEE3 JRA1
Documentation: <https://twiki.cern.ch/twiki/bin/view/LCG/LfcAdminGuide>
Code repository: gLite code repository <http://jra1mw.cvs.cern.ch:8180/>
Development language: C
Build environment: ETICS
Download: <http://glite.web.cern.ch/glite/packages/R3.1/>
Licence: Apache License

Status

Maturity: Production
Usage: most EGEE European and Asian sites, in addition to the CERN site
Interface maturity: Production
Support status: short term (until the end of the EGEE3 project)
Support team: EGEE3/JRA1 team
Development plans: maintenance

Distribution/Availability

Middleware stack: gLite 3.1
Standalone usability: yes
Component dependencies: none
External software dependencies: full list available on request
Component consumers: data management tools
Supported platforms: SLC4/x86
Planned to be supported platforms: SLC5/x86_64

Service component information

Public Interfaces: LFC interface

Aligned security: *none*
Capabilities: none

Interoperability/Standards

Standard-compliance: none
Interoperable components: *none*
Additional Standards: DLI

3.3.16 gLite: File Transfer Service (FTS)

Name gLite File Transfer Service (FTS)

Description

The gLite File Transfer Service (FTS) is a low-level data movement service for transferring files between Storage Elements. In addition, it provides features for administration and monitoring of these transfers. The FTS exposes an interface to submit asynchronous bulk requests and performs the transfers using either third-party GridFTP or SRM Copy. The FTS servers are typically deployed at (large) sites where there are large amounts of data to be transferred. From the security point of view, the interactions with these external services (mainly SRMs and Storage Elements) always use the client proxy credentials either retrieved from MyProxy or delegated by the GridSite delegation components.

Basic information

Component type: Service
Developer: EGEE3 JRA1
Documentation: http://glite.web.cern.ch/glite/packages/R3.1/deployment/glite-FTS_oracle/glite-FTS_oracle.asp
Code repository: gLite code repository <http://jra1mw.cvs.cern.ch:8180/>
Development language: Java
Build environment: ETICS
Download: <http://glite.web.cern.ch/glite/packages/R3.1/>
Licence: Apache License

Status

Maturity: Production
Usage: Tier 1 sites
Interface maturity: Production
Support status: short term (until the end of the EGEE3 project),
Support team: EGEE3/JRA1 team
Development plans: maintenance

Distribution/Availability

Middleware stack: gLite 3.1
Standalone usability: *yes*
Component dependencies: none
External software dependencies: full list available on request
Component consumers: data management tools
Supported platforms: SLC4/i386
Planned to be supported platforms: SLC5/x86_64

Service component information

Public Interfaces: FTS interface
Aligned security: *none*
Capabilities: none

Interoperability/Standards

Standard-compliance: none-umd (SRM, GridFTP)

Interoperable components: none

Additional Standards:-

3.3.17 gLite: Hydra**Name** gLite Hydra**Description**

The symmetric encryption keys for encrypted data (files) storage are stored in a specific set of servers called "Hydra". Hydra provides controlled access to these keys (through certificate DN and VOMS attributes based ACLs) and secured communication to the requester. The Hydra service is a Java Web Service, which can be deployed in a J2EE container, such as Tomcat. It requires a database back-end, and communicates via the HTTPS protocol with its clients.

In addition, Hydra exploits the Shamir secret-sharing scheme to improve security and reliability of this service. Shamir's scheme consists of splitting keys into N fragments stored in different places. Only $M < N$ fragments are needed to reconstruct a complete key. However, owning less than M key fragments, does not give any information on the complete key. Thus, the system is both resistant to attacks (at least M key stores need to be compromised for an attacker to be able to reconstruct a key) and reliable (the disconnection of a limited number of servers does not prevent the key reconstruction).

Basic information

Component type: Service

Developer: EGEE3 JRA1

Documentation: <https://twiki.cern.ch/twiki/bin/view/EGEE/DMEDS>

Code repository: gLite code repository <http://jra1mw.cvs.cern.ch:8180/>

Development language: Java

Build environment: ETICS

Download: <http://glite.web.cern.ch/glite/packages/R3.1/>

Licence: Apache License

Status

Maturity: Production

Usage: biomed community

Interface maturity: Production

Support status: short term (until the end of the EGEE3 project)

Support team: EGEE3/JRA1 team

Development plans: maintenance

Distribution/Availability

Middleware stack: gLite 3.1

Standalone usability: yes

Component dependencies: none

External software dependencies: Tomcat, MySQL

Component consumers: *data management tools*

Supported platforms: SLC4/x86

Planned to be supported platforms: SLC5/x86_64

Service component information

Public Interfaces: none

Aligned security: HTTPS, certificate DN and VOMS attributes based ACLs

Capabilities: *none*

Interoperability/Standards

Standard-compliance: none
Interoperable components: none
Additional Standards: none

3.3.18 gLite: AMGA

Name AMGA**Description**

AMGA is a metadata service which provides database access for Grid applications. AMGA is implemented as a thin layer between a Grid application and the underlying database and acts as an interface between them, providing a Grid style authentication mechanism. AMGA also provides functionality to make Grid applications interoperable with different database backends. Finally, AMGA's support for replication of relational data can be used to build scalable and reliable Grid applications.

Basic information

Component type: Service
Developer: EGEE3 JRA1
Documentation: https://documents.eumedgrid.org/record/160/files/istanbul-AMGA_srvinst-060706.ppt
Code repository: gLite code repository <http://jra1mw.cvs.cern.ch:8180/>
Development language: Java
Build environment: ETICS
Download: <http://glite.web.cern.ch/glite/packages/R3.1/>
Licence: Apache License

Status

Maturity: Production
Usage: various VOs
Interface maturity: Production
Support status: short term (until the end of the EGEE3 project)
Support team: EGEE3/JRA1 team
Development plans: maintenance

Distribution/Availability

Middleware stack: gLite 3.1
Standalone usability: yes
Component dependencies: none
External software dependencies: *full list available on request*
Component consumers: VO tools
Supported platforms: SLC4/x86
Planned to be supported platforms: SLC5/x86_64

Service component information

Public Interfaces: AMGA interface
Aligned security: none
Capabilities: none

Interoperability/Standards

Standard-compliance: none
Interoperable components: *none*
Additional Standards: none

3.3.19 gLite: AMGA Client API

Name AMGA Client API

Description

Client API for contacting the AMGA service.

Basic information

Component type: API

Developer: EGEE3 JRA1

Documentation: https://documents.eumedgrid.org/record/160/files/istanbul-AMGA_srvinst-060706.ppt

Code repository: gLite code repository <http://jra1mw.cvs.cern.ch:8180/>

Development language: Java

Build environment: ETICS

Download: <http://glite.web.cern.ch/glite/packages/R3.1/>

Licence: Apache License

Status

Maturity: Production

Usage: various VOs

Interface maturity: Mature

Support status: short term (until the end of the EGEE3 project)

Support team: EGEE3/JRA1 team

Development plans: maintenance

Distribution/Availability

Middleware stack: gLite 3.1

Standalone usability: yes

Component dependencies: none

External software dependencies: full list available on request

Component consumers: VO tools

Supported platforms: SLC4/x86

Planned to be supported platforms: SLC5/x86_64

Library component information

Consumers: VO tools

Language binding: *Java*

Interoperability/Standards

Standard-compliance: none

Interoperable components: none

Additional Standards: none

3.3.20 UNICORE: UAS (Storage Management and File Transfer)

Name The UNICORE Atomic Services (UAS) elements StorageManagement service and File transfer services

Description

The WS-RF-based UNICORE Atomic Services (UAS) are the proprietary Web service interfaces of the UNICORE middleware. Among other WS related to computation a various of services are related to storage management and file transfers. Firstly, the StorageManagement service acts as a representation of various types of storages and the various flavours of the file transfer

services allow different types of transfers (e.g. plain HTTP, ByteIO). The UAS rely on the XNJS for optimal execution.

Basic information

Component type: *service*

Developer: UNICORE SourceForge Community

Documentation: http://www.unicore.eu/unicore/architecture/service-layer.php#anchor_xnjs

Code repository: <http://unicore.svn.sourceforge.net/viewvc/unicore/unicorex/trunk/uas-core/>

Development language: Java

Build environment: Apache Maven

Download:

<http://downloads.sourceforge.net/unicore/unicore-servers-6.2.0-p1.tgz>

Licence: BSD License

Status

Maturity: *production*

Usage: *DEISA, D-Grid, SKIF-Grid, commercial partners*

Interface maturity: *production*

Support status: *long-term*

Support team: UNICORE SourceForge Community

Development plans: Add and enhance support for various storage backends (e.g. iRODS, Apache Hadoop)

Distribution/Availability

Middleware stack: UNICORE 6.2

Standalone usability: Yes

Component dependencies: *UNICORE/X container*

External software dependencies: none

Component consumers: UNICORE Atomic services, UNICORE OGSA-BES

Supported platforms: *All Java-based platforms*

Planned to be supported platforms: *none*

Service component information

Public Interfaces: Proprietary StorageManagement Service and File Transfer Services

Aligned security: *TLS/SSL (proxies optional), SAML assertions in the SOAP header*

Capabilities: data.management.storage; data.management.transfer

Interoperability/Standards

Standard-compliance: non-umd (HTTP *FileTransfer Service*)

Interoperable components: None

Additional Standards: OGSA ByteIO, *WSRF*

3.3.21 other: dCache

Name dCache

Description

The goal of dCache is to provide a system for storing and retrieving huge amounts of data, distributed among a large number of heterogeneous server nodes, under a single virtual filesystem tree with a variety of standard access methods. Depending on the Persistency Model, dCache provides methods for exchanging data with backend (tertiary) Storage Systems as well as space management, pool attraction, dataset replication, hot spot determination and recovery from disk or node failures. Connected to a tertiary storage system, the cache simulates unlimited direct access storage space. Data exchanges to and from the underlying HSM are performed automatically and invisibly to the user. Filesystem namespace operations may be performed

through a standard NFS(2) interface.

Basic information

Component type: {service, library, client}

Developer:

Documentation:

Code repository:

Development language:

Build environment:

Download:

Licence:

Status

Maturity: {planned, development, pre-production, production}

Usage:

Interface maturity: {development, pre-production, production}

Support status: {abandoned, best effort, short term, long term}

Support team:

Development plans:

Distribution/Availability

Middleware stack:

Standalone usability:

Component dependencies:

External software dependencies:

Component consumers:

Supported platforms:

Planned to be supported platforms:

Service component information *(only for a service component)*

Public Interfaces:

Aligned security:

Capabilities:

Library component information *(only for a library component)*

Consumers:

Language binding:

Interoperability/Standards

Standard-compliance: {none, non-UMD, early-UMD, full-UMD}

Interoperable components:

Additional standards:

Other

Any other kind of information

3.3.22 Other: StoRM

Name StoRM

Description

StoRM is a storage resource manager service for disk based storage systems implementing the SRM interface version 2.2. StoRM provides a flexible, configurable, scalable and high performance SRM solution. It supports standard Grid access protocols as well as direct access (native POSIX I/O call) on data, fostering the integration of non Grid aware application providing local access on shared storage. StoRM takes advantage from high performance parallel and

cluster file systems, but it works on top of any standard POSIX file system with ACL (Access Control List) support. Indeed, StoRM uses the ACLs provided by the underlying file systems to provide a strong and efficient security model.

Basic information

Component type: Service
Developer: INFN-CNAF
Documentation: <http://storm.forge.cnaf.infn.it>
Code repository: <https://svn.forge.cnaf.infn.it/svnroot/storm>
Development language: C/C++ and Java
Build environment: GNU autotools/ant
Download: <http://storm.forge.cnaf.infn.it/repository/RPMS.storm>
Licence: Apache 2

Status

Maturity: Production
Usage: EGEE installations
Interface maturity: Production
Support status: Short term
Support team: INFN-CNAF

Development plans:

Distribution/Availability

Middleware stack: INFNGRID--gLite 3.1
Standalone usability: Yes
Component dependencies:
External software dependencies:
Component consumers:
Supported platforms: SLC4/x86
Planned to be supported platforms: SLC5/x86_64

Service component information

Public Interfaces:
Aligned security: X509 certificates with VOMS attributes
Capabilities:

Interoperability/Standards

Standard-compliance: non-umd
Interoperable components:
Additional Standards: SRM, GridFTP

4 Information functionality

4.1 Standards landscape

The information system components of the three middlewares are still utilizing different information models: ARC relies on the NorduGrid schema, gLite makes use of the GLUE1.3 schema while Unicore relies on its own schema as well.

Comment [BK10]: Unicore to confirm it

Fortunately, the OGF GLUE working group after several years of intensive work has released the GLUE v2.0 specification which provides a common information model embraced by all the major middlewares and grid infrastructures. However, GLUE v2.0, as an information model, does not regulate how the information is obtained, queried or exchanged. There is no common such mechanisms and no ongoing relevant standardization work is addressing this requirement.

It is proposed that every information functionality UMD component should comply with the GLUE v2.0 specification. Concerning information publishing/query it is recommended that the currently used mechanisms are to be harmonized by agreeing on some best practices. To satisfy the GLUE v2.0 compliance requirement all the UMD candidate components require considerable development efforts invested into standard convergence.

4.2 Overview table: information area

	Name	Type	Maturity	Standard compliance		
4.3.1	ARC: Classic Infoserver	service	Production	None		
4.3.2	ARC: Classic Infoindex	Service	Production	None		
4.3.3	ARC: ALIS	service	Pre-production	Early-umd		
4.3.4	ARC: ISIS	service	Development	None		
4.3.5	ARC: libarcclient (infosys modules)	library	Pre-production	Early-umd		
4.3.6	ARC: Grid Monitor	client	production	None		
4.3.7	gLite: BDII	service	production	none		
4.3.8	gLite: lcg-info and lcg-infosites	client	production	none		
4.3.9	gLite: Service Discovery API	library	production	None		
4.3.10	UNICORE: CIS	service	Pre-production	Early-umd		
4.3.11	UNICORE: Service Registry	service	production	None		

4.3 Inventory of components

4.3.1 ARC: Classic Infoserver

Name ARC Classic Infoserver

Description

The Infoserver generates and stores detailed resource and activity information about grid resources. The ARC Infoserver is responsible for resource (computing or storage) description and characterization. The local information is generated on the resource, optionally cached and served through a query interface. The Infoserver implementation is based upon LDAP server with BDII module. The information generation is done via powerful infoprovider scripts. The local information is expressed according to the NorduGrid LDAP schema.

Basic information

Component type: service

Developer: Lund, Uppsala and Oslo Universities from NorduGrid Collaboration

Documentation: The NorduGrid/ARC Information System,

http://www.nordugrid.org/documents/arc_infosys.pdf

Code repository: <http://svn.nordugrid.org/trac/nordugrid/browser/arc0>

Development language: mostly perl scripts

Build environment: GNU autotools

Download: download.nordugrid.org

Licence: Apache License Version 2.0

Status

Maturity: production

Usage: it is a central component used by every ARC client

Interface maturity: production

Support status: long term

Support team: NorduGrid Collaboration

Development plans: none, with a low priority add support for the GLUE v2 LDAP schema

Distribution/Availability

Middleware stack: ARC production releases (latest version 0.6.5)

Standalone usability: no

Component dependencies: ARC: Grid Manager, glite:BDII

External software dependencies: OpenLDAP, perl

Component consumers: ARC clients (ng*, arc*, libarcclient)

Supported platforms: major linux flavours

Planned to be supported platforms: none

Service component information *(only for a service component)*

Public Interfaces: LDAP with NorduGrid information schema

Aligned security: none, anonymous access

Capabilities:

Interoperability/Standards

Standard-compliance: none

Interoperable components: every LDAP client

Additional standards: LDAP

Other

Recently replaced the old globus GRIS-ldap modul dependency by native LDAP and BDII

4.3.2 ARC: Classic Infoindex

Name ARC Classic Infoindex

Description

The ARC Classic Infoindex component is responsible for indexing grid services. The Infoindices maintain a dynamic list of resources (LDAP URLs of the local InfoServers) and other Index Services. The Infoindices are further coupled together implementing a specific topology. Client use the Infoindices to discover resource information endpoints.

The implementation of the Infoindex component is based upon LDAP server and the registrations are carried out via ldapadd operations.

Basic information

Component type: service

Developer: Uppsala University from NorduGrid Collaboration

Documentation: The NorduGrid/ARC Information System,

http://www.nordugrid.org/documents/arc_infosys.pdf

Code repository:

http://svn.nordugrid.org/trac/nordugrid/browser/arc0/branches/v_0_8/infosystem-index

Development language: c++

Build environment: GNU autotools

Download: download.nordugrid.org

Licence: Apache License Version 2.0

Status

Maturity: production

Usage: it is a central component used by every ARC client. Every country or organization is expected to operate couple of Infoindex services

Interface maturity: production

Support status: long term

Support team: NorduGrid Collaboration

Development plans: none, only bugfixes

Distribution/Availability

Middleware stack: ARC production releases (latest version 0.6.5)

Standalone usability: no

Component dependencies: none

External software dependencies: OpenLDAP

Component consumers: ARC clients (ng*, arc*, libarcclient)

Supported platforms: major linux flavours

Planned to be supported platforms: none

Service component information *(only for a service component)*

Public Interfaces: LDAP with special NorduGrid schema

Aligned security: none, anonymous access is provided

Capabilities:

Interoperability/Standards

Standard-compliance: none

Interoperable components: none

Additional standards: LDAP

Other

The Infoindex service implementation has recently migrated from the old globus GIS ldap backend to a Nordugrid backend written on top of native ldap.

4.3.3 ARC: ALIS

Name ARC Local Information System (ALIS)

Description

The ALIS is a component implementing the "local information system" functionality for ARC services. It is based upon a HED library which provides WSRF interface for information query (service and client side) and XML container for storing information. It also features registration functionality for registering services to information collectors/indexing services. It's internal design is made as much independent of external interfaces. That makes it easy to extend supported interfaces without significant changes in services and client tools.

Basic information

Component type: library

Developer: Oslo University, Lund University, NIIF from NorduGrid Collaboration

Documentation:

https://www.knowarc.eu/wiki/images/9/90/Knowarc_D1.2-1_07.pdf (Section 2)

<http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/doc/ARC1-API.pdf>

(Arc::InfoRegisterContainer and related classes for API reference)

Code repository: <http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/src/hed/libs/infosys>

Development language: C++

Build environment: GNU autotools

Download: download.nordugrid.org

Licence: Apache License Version 2.0

Status

Maturity: pre-production

Usage: used by every ARC services which are implemented in the HED framework

Available Documentation: https://www.knowarc.eu/wiki/images/9/90/Knowarc_D1.2-1_07.pdf (Section 2) and <http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/doc/ARC1-API.pdf>

Interface maturity: development

Support status: short term support by KnowARC project, long term support by NorduGrid collaboration

Support team: NIIF and Oslo University from NorduGrid

Development plans: enhancing maturity and interoperability of implementation, adding caching functionality

Distribution/Availability

Middleware stack: technology preview releases of ARC (version 0.9.x)

Standalone usability: no

Component dependencies: ARC:HED

External software dependencies: libxml

Component consumers: ARC:A-REX, ARC:Chelonia (ARC services use this library to provide information about themselves)

Supported platforms: Linux, MacOS, Solaris, Windows (early support)

Planned to be supported platforms: Windows

Service component information

Public Interfaces: WSRF (Selected functionality), proprietary registration interface

Aligned security: Evaluation of ARC policies embedded into informational document

Capabilities:

Library component information

Consumers: ARC services and client tools

Language binding: C++, Python, Java (not tested)

Interoperability/Standards

Standard-compliance: early-umd (preliminary GLUE v2.0 support)

Interoperable components: None

Additional Standards: WSRF selected functionality

4.3.4 ARC: ISIS

Name ARC Information System Indexing Service (ISIS)

Description

The ISIS is a service which stores information pushed to it by another services and serves as resource registration and discovery point. Multiple ISIS services form peer-to-peer network to provide fault-tolerant information system.

Basic information

Component type: service

Developer: NIIF from NorduGrid Collaboration

Documentation:

<http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/doc/infosys/TechnicalHandbook/ISIS%20-%20TechnicalHandbook.pdf>

Code repository: <http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/src/services/isis>

Development language: C++

Build environment: GNU autotools
Download: <http://download.nordugrid.org/software/nordugrid-arc1/releases/latest/>
Licence: Apache License Version 2.0

Status

Maturity: development
Usage: once completed it'll be a key component of ARC-based grids as the ISIS will form the information backbone connecting resources grid-enabled by ARC components
Interface maturity: pre-production
Support status: short term support by KnowARC project, long term support by NorduGrid collaboration
Support team: NIIF
Development plans: enhancing maturity of implementation.

Distribution/Availability

Middleware stack: technology preview releases of ARC (version 0.9.x)
Standalone usability: no
Component dependencies: ARC:HED
External software dependencies: libxml, Glibm
Component consumers: ARC:libarcclient, ARC:A-REX, ARC: Chelonia, etc..
Supported platforms: Linux, MacOS, Solaris, Windows (early)
Planned to be supported platforms: Windows

Service component information

Public Interfaces: WSRF (Selected functionality), proprietary registration interface
Aligned security: Authorization of registering and neighbor services, evaluation of ARC policies embedded into informational document
Capabilities:

Interoperability/Standards

Standard-compliance: none
Interoperable components: None
Additional Standards: WSRF selected functionality

Other

ISIS will gradually replace the ARC: Classic Infoindex component.

4.3.5 ARC: libarcclient (infosys modules)

Name The information system modules of libarcclient

Description

The next generation general purpose grid library (libarcclient) offers transparent access to grid resources through an intuitive and easy-to-use Multilanguage API. The libarcclient comes loaded with plenty of functionality including resource discovery, information system interrogation and resource selection and brokering as well. It has a pluggable architecture that allows it to implement cross-grid information system query. The current implementation comes with plugins for ARC Classic Infoserver and Infoindex (ldap-based pre-WS solution), A-REX (WSRF-based solution), Unicore (WSRF-based solution) and CREAM (mixture of ldap-based BDII and a proprietary WS-based solution).

Libarcclient is targeted for grid developers and advanced application integrators and provides a general toolkit for high-level cross-middleware, multiplatform client side development.

The library is the core of the highly successful and powerful, nevertheless lightweight ARC clients, grid portals and client-side job managers.

Basic information

Component type: library

Developer: Uppsala University, Copenhagen University, NIIF of NorduGrid Collaboration

Documentation:

http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/doc/tech_doc/client/client_technical.pdf

Code repository:

<http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/src/hed/acc> ,

<http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/src/hed/libs/client>

Development language: C++ (language bindings are offered for Python, Java see below)

Build environment: GNU autotools

Download: download.nordugrid.org

Licence: Apache License Version 2.0

Status

Maturity: pre-production

Usage: as the underlying library of most of the ARC client-side tools it'll be used by every ARC user community

Interface maturity: pre-production

Support status: short term support by KnowARC project, long term support by NorduGrid collaboration

Support team: Uppsala University from NorduGrid

Development plans: enhancing maturity of implementation, add support for the emerging PGI-profile, improved support for Unicore resource discovery

Distribution/Availability

Middleware stack: technology preview release of ARC (version 0.9.3)

Standalone usability: no

Component dependencies: ARC HED

External software dependencies: openssl, glibmm, libxml2

Component consumers: ARC: arc* CLI

Supported platforms: all flavours of Linux, MacOS, Solaris

Planned to be supported platforms: Windows

Library component information *(only for a library component)*

Consumers: ARC client tools

Language binding: C++, Python, Java (not tested)

Interoperability/Standards

Standard-compliance: early-UMD (preliminary GLUE v2.0 support)

Interoperable components: CREAM, BDII, Unicore information services

Additional standards: LDAP, WSRF

Other

The libarclient development library is scheduled to be part of the next production ARC release scheduled for 3rd quarter of 2009 as a replacement of the aging previous generation ARCLIB.

4.3.6 ARC: Grid Monitor

Name ARC Grid Monitor**Description**

The Grid Monitor is a Web client tool for the ARC Information System, allowing browsing all the published information about the system. It makes use of the hierarchical information organization and the PHP LDAP module to provide a real-time monitoring and primary debugging for ARC-based grids.

Basic information

Component type: Client
Developer: Lund University from NorduGrid Collaboration
Documentation: "The Grid Monitor: Usage Manual". [NORDUGRID-MANUAL-5] (March 2009)
Code repository: <http://svn.nordugrid.org/arc0/trunk/monitor>
Development language: PHP
Build environment: N/A
Download: download.nordugrid.org
Licence: Apache License Version 2.0

Status

Maturity: Production
Usage: deployed for monitoring as part of the ARC-based grids
Interface maturity: Production
Support status: Best effort
Support team: Lund University from NorduGrid
Development plans: Has to be adapted for GLUE-based information systems and WS interfaces

Distribution/Availability

Middleware stack: ARC production releases (latest v0.6.5)
Standalone usability: YES
Component dependencies: none
External software dependencies: ldap, php
Component consumers: none, it is an end-user tool
Supported platforms: All
Planned to be supported platforms: none

Interoperability/Standards

Standard-compliance: none
Interoperable components: none
Additional Standards: LDAP

Other

Is localizable in any human language (currently, 10 languages are available)

4.3.7 gLite: BDII

Name gLite Berkeley Database Information Index (BDII)

Description

The BDII is a core component of the EGEE information system. It is used to index information sources and provides an up-to-date cache of the information from these sources. It is used at the resource, site and top level in the infrastructure.

Basic information

Component type: Service
Developer: EGEE3 JRA1
Documentation: <https://twiki.cern.ch/twiki/bin/view/EGEE/BDII>
Code repository: gLite code repository <http://jra1mw.cvs.cern.ch:8180/>
Development language: python
Build environment: ETICS
Download: <http://glite.web.cern.ch/glite/packages/R3.1/>
Licence: Apache License

Status

Maturity: Production
Usage: Infrastructure
Interface maturity: Production
Support status: short term (until the end of the EGEE3 project),
Support team: EGEE3/JRA1 team
Development plans: maintenance

Distribution/Availability

Middleware stack: gLite 3.1
Standalone usability: yes
Component dependencies: none
External software dependencies: OpenLDAP
Component consumers: Grid Services end users tools
Supported platforms: SLC4/i386 and x86_64
Planned to be supported platforms: SLC5/x86_64

Service component information

Public Interfaces: LDAP
Aligned security: *none*
Capabilities: none

Interoperability/Standards

Standard-compliance: none
Interoperable components: none
Additional Standards: LDAP, GLUE v1.x

4.3.8 gLite: lcg-info and lcg-infosites

Name gLite: lcg-info and lcg-infosites

Description

lcg-info and lcg-infosites are two command line clients which are used to query the information system.

Basic information

Component type: Client
Developer: EGEE3 JRA1
Documentation: <https://edms.cern.ch/file/454439//LCG-2-UserGuide.html>
Code repository: gLite code repository <http://jra1mw.cvs.cern.ch:8180/>
Development language: perl
Build environment: ETICS
Download: <http://glite.web.cern.ch/glite/packages/R3.1/>
Licence: Apache License

Status

Maturity: Production
Usage: UI
Interface maturity: Production
Support status: short term (until the end of the EGEE3 project),
Support team: EGEE3/JRA1 team
Development plans: maintenance

Distribution/Availability

Middleware stack: gLite 3.1
Standalone usability: no
Component dependencies: BDII
External software dependencies: none
Component consumers: none
Supported platforms: SLC4/i386
Planned to be supported platforms: SLC5/x86_64

Interoperability/Standards

Standard-compliance: none
Interoperable components: *none*
Additional Standards: LDAP

4.3.9 gLite: Service Discovery API

Name gLite Service Discovery API

Description

Service Discovery is a set of libraries that provide an abstraction for finding Grid services and discovering their basic attributes, independently of the underlying Information System infrastructure.

Basic information

Component type: library
Developer: EGEE3 JRA1
Documentation: none
Code repository: gLite code repository <http://jra1mw.cvs.cern.ch:8180/>
Development language: C++
Build environment: ETICS
Download: <http://glite.web.cern.ch/glite/packages/R3.1/>
Licence: Apache License

Status

Maturity: Development
Usage: Development
Interface maturity: Production
Support status: short term (until the end of the EGEE3 project),
Support team: EGEE3/JRA1 team
Development plans: Developing

Distribution/Availability

Middleware stack: gLite 3.1
Standalone usability: *yes*
Component dependencies: SAGA
External software dependencies: full list available on request
Component consumers: Grid Service and end user tools
Supported platforms: SLC4/i386
Planned to be supported platforms: SLC5/x86_64

Library component information

Consumers: Grid Service and end user tools
Language binding: *C, C++, Java, python*

Interoperability/Standards

Standard-compliance: SAGA

Interoperable components: *SAGA plugin engine*
Additional Standards: *Glue*

4.3.10 UNICORE: CIS

Name UNICORE Common Information Service (CIS)

Description

The Common Information Service (CIS) is the information service of UNICORE 6. It gathers both static and dynamic information from all connected UNICORE sites, which are then displayed either in raw XML or human-readable text form in some GUI. For instance, as longitude and latitude information is also stored, a Google maps view allows a geographical representation of the Grid infrastructure. CIS is based on the GLUE 2.0 standard from OGF. Like with the UNICORE registry, a central CIS is necessary in an operational UNICORE infrastructure.

Basic information

Component type: *service*
Developer: FZJ and UNICORE SourceForge Community
Documentation: http://www.unicore.eu/unicore/architecture/service-layer.php#anchor_cis
Code repository: <http://unicore.svn.sourceforge.net/viewvc/unicore/CommonInformationService/>
Development language: Java
Build environment: Apache Maven
Download: http://sourceforge.net/project/showfiles.php?group_id=102081&package_id=269054
Licence: BSD License

Status

Maturity: *pre-production*
Usage: *D-Mon as part of D-Grid*
Interface maturity: *production*
Support status: *long-term*
Support team: *UNICORE SourceForge Community*
Development plans: *Full GLUE2 support*

Distribution/Availability

Middleware stack: UNICORE 6
Standalone usability: *Yes*
Component dependencies: *CIS Infoprovider deployed on UNICORE sites*
External software dependencies: none
Component consumers: None
Supported platforms: *All Java-based platforms*
Planned to be supported platforms: *None*

Service component information

Public Interfaces: *CIS proprietary interface*
Aligned security: *TLS/SSL*
Capabilities: *information.provenance*

Interoperability/Standards

Standard-compliance: *early-umd (GLUE2)*
Interoperable components: *None*
Additional Standards: *None*

4.3.11 UNICORE: Service Registry

Name UNICORE Service Registry

Description

Like many service-oriented environments, a global service registry is available, where the different services can register once they are started. A single service registry is necessary to build-up and operate a distributed UNICORE infrastructure. This service registry is contacted by the clients in order to "connect to the Grid". Therefore it is the central service of a UNICORE 6 Grid. Like the XNJS, the service registry runs in UNICORE's WS-RF hosting environment.

Basic information

Component type: *Web service*
Developer: UNICORE SourceForge Community
Documentation: http://www.unicore.eu/unicore/architecture/service-layer.php#anchor_registry
Code repository: <http://unicore.svn.sourceforge.net/viewvc/unicore/unicorex/trunk/uas-core/>
Development language: Java
Build environment: Apache Maven
Download:
http://downloads.sourceforge.net/unicore/unicore-quickstart-6.1.3-p1.tgz?use_mirror=ovh
Licence: BSD License

Status

Maturity: *production*
Usage: *DEISA, D-Grid, SKIF-Grid, and commercial partners*
Available Documentation:
http://www.unicore.eu/unicore/architecture/service-layer.php#anchor_registry (2009-03)
Interface maturity: *production*
Support status: *long-term and supported by UNICORE SourceForge Community*
Development plans: *None*

Distribution/Availability

Middleware stack: UNICORE 6
Standalone usability: *Yes*
Component dependencies: *UNICORE/X container*
External software dependencies: none
Component consumers: UNICORE clients, UNICORE atomic services
Supported platforms: *All Java-based platforms*
Planned to be supported platforms: *None*

Service component information

Public Interfaces: WS-RF ServiceGroups
Aligned security: *TLS/SSL*
Capabilities: *information.provenance*

Interoperability/Standards

Standard-compliance: *None*
Interoperable components: *none*
Additional Standards: WS-RF ServiceGroups

5 Security functionality

5.1 Standards landscape

The standards landscape in terms of middleware security can be basically differentiated between authentication and authorization.

5.1.1 Authentication (AuthN)

Authentication of end-users refers to the check whether certificates of end-users have been signed by a trusted CA, are valid (date), and not revoked by Certificate Revocation List (CRL) mechanisms. In terms of authentication the middleware should support commonly used schemes that use certificate-based identity information from the transport layer security (TLS)/Secure Socket Layer(SSL) protocols. In this context, TLS connections can be initiated using full X.509 end-entity certificates [SEC1] or X.509 proxy certificates [SEC2].

We emphasize that UMD components have to either support one or the other! Because of modern handler designs it is possible to even combine different solutions without breaking the other correspondent other model in interoperability use cases.

5.1.2 Authorization (AuthZ)

Authorization can be differentiated in basically two approaches. The purely identity-based authorization uses information of X.509 certificates to grant end-users access to the systems. However, in terms of UMD, we require a more flexible solution and raise the demand to support attribute-based authorization.

In this context, attributes indicate the position of an end-user within a virtual organization (VO) and group membership as well as role possession. These attributes are released from so-called Attribute Authorities (AA), which in terms of UMD are typically represented by either the Virtual Organization Membership Service (VOMS) or the UNICORE Virtual Organization Service (UVOS). To ensure the correctness of attributes a signing mechanism must be used and over time we have seen two approaches of signing and then conveying these attributes to middleware systems.

The first approach uses Attribute Certificates (AC) [SEC3] that are typically conveyed within the extension of X.509 certificates. This can be done with full end-entity X.509 certificates or X.509 proxies on the transport level (i.e. authentication). One AA that supports this approach is the non Web service-based 'classic VOMS' system.

The second approach uses Security Assertion Markup Language (SAML) [SEC4] assertions that are typically conveyed within the SOAP header during Web service calls (e.g. OGSA-BES systems). This can be also done with using either full end-entity X.509 certificates or X.509 proxies on the transport level (i.e. authentication). One AA that supports this approach is the Web service-based 'OMII-Europe/EGEE-III VOMS system' and UVOS. Although this approach seems to be limited in terms of using it only with Web service technologies there is a trend towards the use of this approach since it is much more flexible than the first approach. For instance, the second approach allows for the transport of n different SAML assertions maybe obtained from n different AAs. But most notably, the second approach encapsulates the transport level (i.e. TLS/SSL) from the transport of attributes (SOAP header instead of extension within a certificate used for TLS/SSL).

We emphasize that UMD components have to either support one or the other transport mechanism! Because of modern handler designs it is possible to even combine these different approaches without breaking the other correspondent other model in interoperability use cases. Also, many systems, like the 'classic VOMS' and 'OMII-Europe/EGEE-III VOMS' are based on the same VO database and thus release the same attributes in a different manner.

However, since at the end, only the attributes matter in terms of authorization decisions it is important that within UMD an agreement is achieved upon the security attributes and their semantics. The common UMD security semantics maybe based on Extensible Access Control Markup Language (XACML) standards [SEC5] policies or using rather proprietary mechanisms. Work should be focused on an agreement of the attribute semantics and their structure. This can

be done in close collaboration with the OGF Production Grid Infrastructure Working Group that works on this topic.

An UMD security compliant component should be able to 1) validate users' credentials either in full X509 certificate or X509 proxy certificate format; 2) fetch attributes (roles) of the users either from well-defined VOMS extensions or from SAML assertions within the SOAP header of the messages. 3) be able to understand the semantics of the common UMD set of security attributes. As a result of the above three requirements UMD compliant components will be able to base their security decision (accept/reject the request) on the gathered information. In order to fulfill the requirements considerable amount of work is necessary in terms of both standardization and code development.

5.2 Overview table: Security area

	Name	Type	Maturity	Standards compliance		
5.3.1	ARC: HED security API	library	Pre-production	Non-umd		
5.3.2	ARC: Shib Bridge	service	development	Non-umd		
5.3.3	ARC: arcproxy	client	Pre-production	Non-umd		
5.3.4	ARC: Charon Authorization Service	service	development	none		
5.3.5	ARC: Proxy Store	service	development	none		
5.3.6	ARC: Fruitfly credential service	service	development	none		
5.3.7	gLite: VOMS	service	production	Non-umd		
5.3.8	gLite: VOMS-Admin	service	production	None		
5.3.9	gLite: voms-proxy-*	client	production	None		
5.3.10	gLite: ARGUS	service	development	Non-umd		
5.3.11	gLite: Short Lived Credential Service	service	Pre-production	None		
5.3.12	gLite: SCAS	service	Pre-production	None		
5.3.13	gLite: LCAS	library	production	None		
5.3.14	gLite: LCMAPS	library	production	None		
5.3.15	gLite: gLite-Exec (glExec)	client	production	None		
5.3.16	gLite: gLite-Delegation Client	client	production	None		
5.3.17	gLite: Shibboleth Short Lived Credential Service Client	client	production	None		
5.3.18	gLite: Gridsite	library	production	None		
5.3.19	gLite: CGSI_gSoap	library	production	None		
5.3.20	gLite: delegation-Java	library	production	None		
5.3.21	gLite: Trustmanager	Library	production	None		
5.3.22	UNICORE: Gateway	service	production	Non-umd		
5.3.23	UNICORE: XUADB	service	production	None		
5.3.24	UNICORE: XACML Entity	other	production	None		
5.3.25	UNICORE: UVOS	service	production	Non-umd		
5.3.26	other: MyProxy	service	production	None		

5.3 Inventory of components

5.3.1 ARC: HED security API

Name

HED security API

Description

This component offers libraries for developing services, clients and plugins within the ARC HED framework. It includes: the interface for security plugins, the interface for access control policy processing (libarcsecurity), the class for credential delegation (libarcdelegation), the class for credential processing (libarccredential), the class for xml security processing (libxmlsec) and the class for WS-Security processing (libws-security)

Basic information

Component type: library

Developer: University of Oslo from NorduGrid

Documentation:

<http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/doc/sec/SecurityFrameworkofARC1.pdf>

Code repository:

Code repository:

<http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/src/hed/libs/credential>

<http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/src/hed/libs/delegation>

<http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/src/hed/libs/security>

<http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/src/hed/libs/ws-security>

<http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/src/hed/libs/xmlsec>

Development language: C++

Build environment: GNU autotools

Download: download.nordugrid.org

Licence: Apache License Version 2.0

Status

Maturity: pre-production

Usage: A key library of the new WS-based ARC framework.

Interface maturity: pre-production

Support status: short term through the KnowARC, long term through NorduGrid Collaboration

Support team: Oslo University

Development plans:

Distribution/Availability

Middleware stack: Technology preview releases of ARC (version 0.9.x)

Standalone usability: no

Component dependencies: ARC:HED

External software dependencies: libxmlsec1, openssl, libxml2

Component consumers: ARC: A-REX, ARC: Chelonia

Supported platforms: Linux, MacOS, Solaris, Windows (preliminary)

Planned to be supported platforms: complete windows

Library component information *(only for a library component)*

Consumers: ARC: A-REX, ARC: Chelonia

Language binding: c++, Python

Interoperability/Standards

Standard-compliance: non-UMD

Interoperable components: none

Additional standards:

5.3.2 ARC: Shibbridge

Name Shibbridge

Description

This component provides the SAML2 Web SSO profile implementation in ARC. It can be used for authenticating client using username/password rather than using client's X.509 certificate, i.e., using username/password plus SSL/TLS (without client certificate) authentication instead of mutual SSL/TLS authentication.

The component bridges ARC services to External Identity Providers (IdP) that are compatible to standard SAML2 Web SSO profile (such as Shibboleth IdP, or any other standard-compatible IdP implementations).

Basic information

Component type: mostly service but together with a client

Developer: University of Oslo from NorduGrid

Documentation:

<http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/doc/sec/SecurityFrameworkofARC1.pdf>

Code repository:

<http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/src/services/saml/SPService.cpp>

<http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/src/hed/libs/client/ClientSAML2SSO.cpp>

Development language: C++

Build environment: GNU autotools

Download: download.nordugrid.org

Licence: Apache License Version 2.0

Status

Maturity: development

Usage: the component is yet under development

Interface maturity: development

Support status: short term by the KnowARC project

Support team: Oslo University

Development plans:

Distribution/Availability

Middleware stack: technology preview releases of ARC (version 0.9.x)

Standalone usability: no

Component dependencies:

External software dependencies:

Component consumers:

Supported platforms:

Planned to be supported platforms:

Service component information *(only for a service component)*

Public Interfaces:

Aligned security:

Capabilities:

Interoperability/Standards

Standard-compliance: non-UMD

Interoperable components:

Additional standards: SAML2 Web SSO profile

Other

Any other kind of information

5.3.3 ARC: arcproxy

Name arcproxy

Description

The client utility which can be used to create: RFC-compatible proxy certificate; RFC-compatible proxy certificate with delegation policy embedded for fine-grained restriction; to contact voms server and create RFC-compatible proxy certificate with VOMS AC embedded (the same as voms-proxy-init); to contact myproxy server, delegate a credential to myproxy server (the same as myproxy-init) and get a delegated credential from myproxy server (the same as myproxy-get-delegation).

Basic information

Component type: client

Developer: Oslo University from NorduGrid

Documentation:

<http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/doc/sec/SecurityFrameworkofARC1.pdf>

Code repository:

<http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/src/clients/credentials/arcproxy.cpp>

Development language: C++

Build environment: GNU autotools

Download: download.nordugrid.org

Licence: Apache License Version 2.0

Status

Maturity: pre-production

Usage: once it'll be part of the production release it will be used as the default proxy generation and management utility of ARC

Interface maturity: pre-production

Support status: short term by KnowARC project and long term by the NorduGrid collaboration

Support team: Oslo University

Development plans:

Distribution/Availability

Middleware stack: technology preview releases of ARC (version 0.9.x)

Standalone usability: yes

Component dependencies:

External software dependencies:

Component consumers:

Supported platforms:

Planned to be supported platforms:

Interoperability/Standards

Standard-compliance: non-UMD

Interoperable components: gLite: VOMS, other: myproxy

Additional standards:

Other

The ARC: arcproxy utility is a good candidate for a unified, cross-middleware proxy-generating client.

5.3.4 ARC: Charon Authorization Service

Name Charon Authorization Service

Description

The Charon ARC authorization web service accepts request for access control decision and returns the access control decision. Currently the ARC policy schema (simplified XACML policy) is supported while the XACML policy schema is planned to be supported.

Basic information

Component type: service

Developer: Oslo University from Nordugrid Collaboration

Documentation:

<http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/doc/sec/SecurityFrameworkofARC1.pdf>

Code repository:

<http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/src/services/charon>

Development language: C++

Build environment: GNU autotools

Download: <http://download.nordugrid.org/software/nordugrid-arc1/releases/latest/>

Licence: Apache License Version 2.0

Status

Maturity: development

Usage:

Interface maturity: development

Support status: short term by the KnowARC project

Support team: Oslo University

Development plans: add XACML support

Distribution/Availability

Middleware stack: technology preview releases of ARC (version 0.9.x)

Standalone usability:

Component dependencies:

External software dependencies:

Component consumers:

Supported platforms:

Planned to be supported platforms:

Service component information *(only for a service component)*

Public Interfaces:

Aligned security:

Capabilities:

Interoperability/Standards

Standard-compliance: none

Interoperable components: none

Additional standards:

Other

Any other kind of information

5.3.5 ARC: Proxy Store

Name Proxy Store

Description

Proxy Store is a web service to which client can delegates its credential and then from which client can fetch the delegated credential. Functionality-wise it is similar to myproxy server while using SOAP communication rather than GSI. The Proxy Store does not use username/password to protect the delegated credential in the service. It is more similar to the gridsite delegation service.

Basic information

Component type: service

Developer: Oslo University from NorduGrid Collaboration

Documentation:

<http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/doc/sec/SecurityFrameworkofARC1.pdf>

Code repository:

<http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/src/services/delegation>

Development language: C++

Build environment: GNU autotools

Download: download.nordugrid.org

Licence: Apache License Version 2.0

Status

Maturity: development

Usage:

Interface maturity: {development

Support status: short term by the KnowARC project

Support team: Oslo University

Development plans:

Distribution/Availability

Middleware stack: technology preview releases of ARC (version 0.9.x)

Standalone usability:

Component dependencies:

External software dependencies:

Component consumers:

Supported platforms:

Planned to be supported platforms:

Service component information *(only for a service component)*

Public Interfaces:

Aligned security:

Capabilities:

Interoperability/Standards

Standard-compliance: none

Interoperable components:

Additional standards:

Other

Any other kind of information

5.3.6 ARC: Fruitfly credential service

Name ARC: Fruitfly credential service

Description

Short Lived Credential Service

Fruitfly is a web service which generates short-lived X.509 credential based on username/password. The SAML attribute assertion returned from IdP is embedded as certificate extension, so that afterwards when client uses this short-lived credential to access other generic Web Services, the SAML attributes can be pushed to service side and used for access control. The service comes with a client utility.

Basic information

Component type: service with simple client

Developer: Oslo University from NorduGrid

Documentation:

<http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/doc/sec/SecurityFrameworkofARC1.pdf>

Code repository:

<http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/src/services/slcs>

<http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/src/clients/credentials/slcs>

Development language: C++

Build environment: GNU autotools

Download: download.nordugrid.org

Licence: Apache License Version 2.0

Status

Maturity: development

Usage:

Interface maturity: development

Support status: short term by Knowarc project

Support team: Oslo University

Development plans:

Distribution/Availability

Middleware stack: technology preview releases of ARC (version 0.9.x)

Standalone usability:

Component dependencies:

External software dependencies:

Component consumers:

Supported platforms:

Planned to be supported platforms:

Service component information *(only for a service component)*

Public Interfaces:

Aligned security:

Capabilities:

Interoperability/Standards

Standard-compliance: none

Interoperable components:

Additional standards:

Other

Any other kind of information

5.3.7 gLite: VOMS

Name gLite Virtual Organization Membership Service (VOMS)

Description

VOMS is an Attribute Authority service that issues attributes (in the form of signed assertions)

expressing membership information of a subject within a VO. VOMS allows VO users to be partitioned in groups and to assign them roles and other free-form attributes that are then used to drive authorization decisions.

Basic information

Component type: Service

Developer: EGEE3 JRA1

Documentation:

VOMS Communication Protocol: <https://twiki.cnaf.infn.it/cgi-bin/twiki/view/VOMS/VOMSProtocol>,
<https://twiki.cnaf.infn.it/cgi-bin/twiki/view/VOMS/WebHome?>

Code repository: gLite code repository <http://jra1mw.cvs.cern.ch:8180/>

Development language: C++

Build environment: ETICS

Download: <http://glite.web.cern.ch/glite/packages/R3.1/>

Licence: Apache License

Status

Maturity: Production

Usage: all EGEE installations, plus other major Grid infrastructures worldwide

Interface maturity: Production

Support status: short term (until the end of the EGEE3 project)

Support team: EGEE3/JRA1 team

Development plans: maintenance

Distribution/Availability

Middleware stack: gLite 3.1

Standalone usability: Yes

Component dependencies: none

External software dependencies: full list available on request

Component consumers: voms-proxy-* command line tools, gLite VOMS-admin, gLite WMS, gLite WMS UI, gLite CREAM among others

Supported platforms: SLC4/x86, other Linux distributions

Planned to be supported platforms: SLC5/x86_64

Service component information

Public Interfaces: Legacy, SAML

Aligned security: *none*

Capabilities: *none*

Library component information

Consumers: *Grid Service and end user tools*

Language binding: C/C++/Java

Interoperability/Standards

Standard-compliance: non-umd (SAML)

Interoperable components: *none*

Additional Standards: *x509*

5.3.8 gLite: VOMS-Admin

Name VOMS-Admin

Description

VOMS-Admin is the administrative application that manages and populates the VOMS database with VO-specific membership information. It is implemented as a Web Service that can be

accessed either through a web interface or using a command line client.

Basic information

Component type: Service
Developer: EGEE3 JRA1
Documentation: <https://twiki.cern.ch/twiki/view/VOMS/WebHome?>
Code repository: gLite code repository <http://jra1mw.cvs.cern.ch:8180/>
Development language: Java
Build environment: ETICS
Download: <http://glite.web.cern.ch/glite/packages/R3.1/>
Licence: Apache License

Status

Maturity: Production
Usage: VO managers
Interface maturity: Production
Support status: short term (until the end of the EGEE3 project)
Support team: EGEE3/JRA1 team
Development plans: maintenance

Distribution/Availability

Middleware stack: gLite 3.1
Standalone usability: No
Component dependencies: none
External software dependencies: full list available on request
Component consumers: VOMS
Supported platforms: SLC4/x86
Planned to be supported platforms: SLC5/x86_64

Service component information

Public Interfaces: Web Interface
Aligned security: *none*
Capabilities: *none*

Interoperability/Standards

Standard-compliance: none
Interoperable components: none
Additional Standards: *none*

5.3.9 gLite: voms-proxy-*

Name gLite voms-proxy-*

Description

Commands to generate and manipulate X509 proxy certificates containing Attribute Certificates ACs by contacting a VOMS server.

Basic information

Component type: Client
Developer: EGEE3 JRA1
Documentation: <https://edms.cern.ch/file/454439/LCG-2-UserGuide.html>
Code repository: <http://glite.cvs.cern.ch:8180/cgi-bin/glite.cgi/>
Development language: C++
Build environment: ETICS
Download: <http://glite.web.cern.ch/glite/packages/R3.1/>

Licence: Apache License

Status

Maturity: Production
Usage: UI (all end users)
Interface maturity: Production
Support status: short term (until the end of the EGEE3 project)
Support team: EGEE3/JRA1 team
Development plans: maintenance

Distribution/Availability

Middleware stack: gLite 3.1
Standalone usability: No
Component dependencies: gLite VOMS
External software dependencies: full list available on request
Component consumers: end users
Supported platforms: SLC4/i386
Planned to be supported platforms: SLC5/x86_64

Interoperability/Standards

Standard-compliance: none
Interoperable components: none
Additional Standards: *none*

5.3.10 gLite: ARGUS

Name gLite ARGUS Authorization Service

Description

The goal of the Authorization Service is to provide a consistent means of distributing, evaluating, and enforcing authorization policies across grid services. Policies, which are stored and transported as XACML documents, may be authored by various parties (e.g. VO administrators, site administrators) and combined to form an effective policy. It offers a consistent approach to authorization so that all components employ the same code, offer the same functionality and evaluate policies in the same manner. The Authorization Service is composed of four components: Policy Enforcement Point (PEP), Policy Decision Point (PDP), Policy Administration Point (PAP), and Execution Environment Service (EES).

Basic information

Component type: Service
Developer: EGEE3 JRA1
Documentation: N/A
Code repository: N/A
Development language: N/A
Build environment: N/A
Download: N/A
Licence: Apache License

Status

Maturity: Development
Usage: Site level
Interface maturity: Development
Support status: short term (until the end of the EGEE3 project)
Support team: EGEE3/JRA1 team

Development plans: Developing

Distribution/Availability

Middleware stack: gLite 3.1

Standalone usability: yes

Component dependencies: N/A

External software dependencies: N/A

Component consumers: N/A

Supported platforms: none at the moment (the component is in development)

Planned to be supported platforms: SLC5/x86_64

Service component information

Public Interfaces: yes

Aligned security: N/A

Capabilities: none

Interoperability/Standards

Standard-compliance: non-umd (XACML)

Interoperable components: none

Additional Standards: none

5.3.11 gLite: Short Lived Credential Service

Name Short Lived Credential Service (SLCS)

Description

The Short Lived Credential Service (SLCS) is a service that issues short lived X.509 credentials based upon successful authentication at a Shibboleth Identity Provider. This service presents from the user's point of view an X.509 certificate factory that issues him/her a certificate with which he/she can access grid resources. The certificate is accepted by the international Grid Policy Management Authorities (IGTF), and therefore the SLCS certificates can be used on most existing Grid infrastructures. The SLCS server relies on Shibboleth to do the user authorization. It uses the Shibboleth attributes of the user to generate the required X.509 subject DN and extensions. An external online certificate authority server is used as a backend to sign the certificate signing requests sent by the client. The server also enforces the policies defined in the SLCS CP/CPS.

Basic information

Component type: Service

Developer: EGEE3 JRA1

Documentation: N/A

Code repository: N/A

Development language: N/A

Build environment: N/A

Download: N/A

Licence: Apache License

Status

Maturity: Pre-Production

Usage: Site level

Interface maturity: Pre-Production

Support status: short term (until the end of the EGEE3 project),

Support team: EGEE3/JRA1 team

Development plans: Developing

Distribution/Availability

Middleware stack: gLite 3.1
Standalone usability: *yes*
Component dependencies: N/A
External software dependencies: N/A
Component consumers: N/A
Supported platforms: SLC4/x86
Planned to be supported platforms: SLC5/x86_64

Service component information

Public Interfaces: *yes*
Aligned security: *none*
Capabilities: *none*

Interoperability/Standards

Standard-compliance: *none*
Interoperable components: *none*
Additional Standards: *none*

5.3.12 gLite: SCAS

Name gLite Site Central Authorization Service (SCAS)

Description

A centralized Local Centre Authorization Service (LCAS) implemented as a server front-end. This will allow LCAS to be invoked by any node other service in a site, providing centralized site-wide authorization policies.

Basic information

Component type: Service
Developer: EGEE3 JRA1
Documentation: <https://www.nikhef.nl/pub/projects/grid/gridwiki/index.php/SCAS>
Code repository: <http://glite.cvs.cern.ch:8180/cgi-bin/glite.cgi/>
Development language: C++
Build environment: ETICS
Download: <http://etics-repository.cern.ch:8080/repository/download/registered/org.glite>
Licence: Apache License

Status

Maturity: Pre-Production
Usage: Site Level
Interface maturity: Pre-Production
Support status: short term (until the end of the EGEE3 project)
Support team: EGEE3/JRA1 team
Development plans: Developing

Distribution/Availability

Middleware stack: gLite 3.1
Standalone usability: *yes*
Component dependencies: *none*
External software dependencies: Full list available on request
Component consumers:
Supported platforms: SLC4/x86
Planned to be supported platforms: SLC5/x86_64

Service component information

Public Interfaces: yes
Aligned security: *none*
Capabilities: *none*

Interoperability/Standards

Standard-compliance: none
Interoperable components: none
Additional Standards: *none*

5.3.13 gLite: LCAS

Name gLite Local Centre for Authorization Service (LCAS)

Description

The Local Centre Authorization Service (LCAS) is a framework that makes a binary authorization decision based on the credentials provided by the Grid user and is dynamically loaded by services for each incoming request. It provides authorization decisions in a fast way before getting to the mapping stage (LCMAPS). The LCAS framework does not need to communicate with any other service.

Basic information

Component type: Library
Developer: EGEE3 JRA1
Documentation: <http://www.dutchgrid.nl/DataGrid/wp4/lcas/>
Code repository: <http://glite.cvs.cern.ch:8180/cgi-bin/glite.cgi/>
Development language: C++
Build environment: ETICS
Download: <http://glite.web.cern.ch/glite/packages/R3.1/>
Licence: Apache License

Status

Maturity: Production
Usage: EGEE installations
Interface maturity: Production
Support status: short term (until the end of the EGEE3 project)
Support team: EGEE3/JRA1 team
Development plans: maintenance

Distribution/Availability

Middleware stack: gLite 3.1
Standalone usability: yes
Component dependencies: none
External software dependencies: full list available on request.
Component consumers: gLexec
Supported platforms: SLC4/x86
Planned to be supported platforms: SLC5/x86_64

Library component information

Consumers: *Grid Services*
Language binding: C++, *Java*

Interoperability/Standards

Standard-compliance: none
Interoperable components: none

Additional Standards: *none*

5.3.14 gLite: LCMAPS

Name gLite Local Credential MAPPING Service (LCMAPS)

Description

The Local Credential MAPPING Service (LCMAPS) is a framework implemented as a library which is loaded dynamically by a service needing to map the Grid credentials presented by a client into UNIX UID and GIDs. The LCMAPS framework also executes some sub-tasks in order to optimize the execution of these evaluations by providing internally the certificate chain of the users, the JDL, a list of VOMS data and a placeholder for storing data such as the gathered UNIX UID and GIDs. This information can be used by the plugins directly from LCMAPS in order to reduce calls to the VOMS API and/or the GridSite core library.

Basic information

Component type: Library
Developer: EGEE3 JRA1
Documentation: <http://www.dutchgrid.nl/DataGrid/wp4/lcmaps/edg-lcmaps-0.0.16/apidoc/html/index.html>
Code repository: <http://glite.cvs.cern.ch:8180/cgi-bin/glite.cgi/>
Development language: C++
Build environment: ETICS
Download: <http://glite.web.cern.ch/glite/packages/R3.1/>
Licence: Apache License

Status

Maturity: Production
Usage: EGEE installations
Interface maturity: Production
Support status: short term (until the end of the EGEE3 project)
Support team: EGEE3/JRA1 team
Development plans: maintenance

Distribution/Availability

Middleware stack: gLite 3.1
Standalone usability: *yes*
Component dependencies: none
External software dependencies: full list available on request
Component consumers: gLexec
Supported platforms: SLC4/x86
Planned to be supported platforms: SLC5/x86_64

Library component information

Consumers: *Grid Services*
Language binding: *C++, Java*

Interoperability/Standards

Standard-compliance: none
Interoperable components: *none*
Additional Standards: *none*

5.3.15 gLite: gLite-Exec (glExec)

Name gLite-Exec (glExec)

Description

glExec (gLite Exec) is a component that switches the identity of Grid user processes based on the Fully Qualified Attributes presented by the user credential (usually a VOMS proxy). It was initially designed to operate on a Computing Element to run the specific LRMS command with appropriate rights. Lately its scope has been extended to operate also on a Worker Node, to enable jobs to run other processes under different identities, with the purpose of safely supporting job management frameworks based on pilot jobs. The development of glExec has been forked from gsexec, which in turn was forked from the suexec program from the Apache project. This development decision was made since the predecessor programs have a proven history to be safe.

The glExec component operates with strong security checking in mind, in that a valid user proxy certificate chain must be presented for a mapping to take place. glExec itself cannot determine the mapping from Grid user to uid and gid itself and relies on the LCAS/LCMAPS system to do this.

In order to do this, LCAS/LCMAPS needs to have a valid user proxy containing the full certificate chain and the private key generated in the last delegation step to the machine running glExec. glExec checks the ownership and file permissions of this certificate before it calls out to LCAS/LCMAPS. If glExec determines that the conditions are incorrect, it will halt, stopping the whole identity switching process. Since changing uid and gid usually requires root privileges, glExec comes as a binary executable program with the suid bit set.

Basic information

Component type: Client

Developer: EGEE3 JRA1

Documentation: http://www.nikhef.nl/pub/projects/grid/gridwiki/index.php/Site_Access_Control

Code repository: <http://glite.cvs.cern.ch:8180/cgi-bin/glite.cgi/>

Development language: C++

Build environment: ETICS

Download: <http://glite.web.cern.ch/glite/packages/R3.1/>

Licence: Apache License

Status

Maturity: Production

Usage: Site Level

Interface maturity: Production

Support status: short term (until the end of the EGEE3 project)

Support team: EGEE3/JRA1 team

Development plans: maintenance

Distribution/Availability

Middleware stack: gLite 3.1

Standalone usability: Yes

Component dependencies: LCAS, LCMAPS

External software dependencies: *full list available on request*

Component consumers: BLAH, gLite CREAM and others

Supported platforms: SLC4/x86

Planned to be supported platforms: SLC5/x86_64

Interoperability/Standards

Standard-compliance: none

Interoperable components: none

Additional Standards: *none*

5.3.16 gLite: gLite-Delegation Client

Name gLite-Delegation Client

Description

This component also provides command-line utilities to manage (delegate, check and destroy) credentials on a service implementing the GridSite delegation protocol. The delegation client component depends on Service Discovery, gridsite-core and CGSI_gSOAP.

Basic information

Component type: Client
Developer: EGEE3 JRA1
Documentation: none
Code repository: <http://glite.cvs.cern.ch:8180/cgi-bin/glite.cgi/>
Development language: C++
Build environment: ETICS
Download: <http://glite.web.cern.ch/glite/packages/R3.1/>
Licence: Apache License

Status

Maturity: Production
Usage: EGEE installations
Interface maturity: Production
Support status: short term (until the end of the EGEE3 project)
Support team: EGEE3/JRA1 team
Development plans: maintenance

Distribution/Availability

Middleware stack: gLite 3.1
Standalone usability: *yes*
Component dependencies: none
External software dependencies: full list available on request
Component consumers: gLite CREAM, gLite WMPProxy
Supported platforms: SLC4/x86
Planned to be supported platforms: SLC5/x86_64

Interoperability/Standards

Standard-compliance: none
Interoperable components: none
Additional Standards: *none*

5.3.17 gLite: Shibboleth Short Lived Credential Service Client

Name Shibboleth Short Lived Credential Service Client

Description

Two shell commands are available as part of the SLCS client package. The main command is `slcs-init`, which authenticates the user with the given IdP, generates locally a private key and certificate signing request, then receives the resulting short lived X.509 certificate from the SLCS server. The information command `slcs-info` gives details about the configured Shibboleth identity providers.

Basic information

Component type: Client
Developer: EGEE3 JRA1

Documentation: http://www.switch.ch/grid/slcs/about/about_long.html
Code repository: N/A
Development language: N/A
Build environment: N/A
Download: N/A
Licence: Apache License

Status

Maturity: Production
Usage: Evaluation
Interface maturity: Production
Support status: short term (until the end of the EGEE3 project)
Support team: EGEE3/JRA1 team
Development plans: matenance

Distribution/Availability

Middleware stack: gLite 3.1
Standalone usability: *yes*
Component dependencies: none
External software dependencies: *N/A*
Component consumers: *N/A*
Supported platforms: SLC4/x86
Planned to be supported platforms: SLC5/x86_64

Interoperability/Standards

Standard-compliance: none
Interoperable components: none
Additional Standards: *none*

5.3.18 gLite: Gridsite

Name gLite Gridsite

Description

GridSite delegation is a proxy certificate delegation protocol, based on standard Web Service interfaces. The development of this tool was fostered by the need of replacing the GSI-specific HTTPG protocol by the standard HTTPS one, in order that off-the-shelf components (such as client libraries and service containers) could be used more easily for Grid development. The simple Web Services-based interface is adjusted to match the functionality of the WSRF-based Globus delegation protocol, thus making the transition between these systems easier for developers. This package also provides the SQL schema definitions for MySQL and Oracle implementations, which serve as an interface between the delegation service and other services using the delegated credentials. It also provides a programming language independent documentation of the file and database based back-ends for service implementers who wish to use the delegated credentials.

Basic information

Component type: library
Developer: EGEE3 JRA1
Documentation: <http://www.gridsite.org/>
Code repository: <http://glite.cvs.cern.ch:8180/cgi-bin/glite.cgi/>
Development language: C
Build environment: ETICS
Download: <http://glite.web.cern.ch/glite/packages/R3.1/>

Licence: Apache License

Status

Maturity: Production
Usage: Infrastructure
Interface maturity: Production
Support status: short term (until the end of the EGEE3 project)
Support team: EGEE3/JRA1 team
Development plans: maintenance

Distribution/Availability

Middleware stack: gLite 3.1
Standalone usability: *yes*
Component dependencies: none
External software dependencies: *full list available on request*
Component dependencies: none
Supported platforms: SLC4/x86
Planned to be supported platforms: SLC5/x86_64

Library component information

Consumers: many
Language binding: C, C++

Interoperability/Standards

Standard-compliance: none
Interoperable components: *none*
Additional Standards: *none*

5.3.19 gLite: CGSI_gSoap

Name gLite CGSI_gSoap

Description

CGSI-gSOAP is a library which allows gSOAP clients and services to use the Globus Security Infrastructure (GSI) authentication. By making minimal modifications to the C or C++ source code, a developer can enable the usage of HTTPS and HTTPG protocols via the GSI libraries. The server-side version of the CGSI-gSOAP library also aids the usage of VOMS credentials with the help of the VOMS library.

Basic information

Component type: Library
Developer: EGEE3 JRA1
Documentation: <http://glite.web.cern.ch/glite/documentation/default.asp>
Code repository: <http://glite.cvs.cern.ch:8180/cgi-bin/glite.cgi/>
Development language: C
Build environment: ETICS
Download: <http://glite.web.cern.ch/glite/packages/R3.1/>
Licence: Apache License

Status

Maturity: Production
Usage: Grid Services
Interface maturity: Production
Support status: short term (until the end of the EGEE3 project),

Support team: EGEE3/JRA1 team
Development plans: maintenacne

Distribution/Availability

Middleware stack: gLite 3.1
Standalone usability: *yes*
Component dependencies: none
External software dependencies: *full list available on request*
Component consumers: Grid Services
Supported platforms: SLC4/x86
Planned to be supported platforms: SLC5/x86_64

Library component information

Consumers: *Grid Services*
Language binding: *C*

Interoperability/Standards

Standard-compliance: none
Interoperable components: *none*
Additional Standards: *none*

5.3.20 gLite: delegation-Java

Name gLite delegation-Java

Description

The delegation-java component is a server-side Java library that implements a GridSite delegation service. It provides implementation for both the file and database based back-ends. The delegation-java library can be used by service developers who wish to add proxy delegation functionality to their Java-based web service. The delegation-service-java component is a packaging of the delegation-java library into a Java Web Service, which can be deployed in a J2EE container such as Tomcat. It is possible to use a database back-end. The delegation-service-java communicates via the HTTPS protocol with its clients. The delegation-java component depends on Util-java and Trustmanager. The delegation-service-java component depends on Tomcat.

Basic information

Component type: Library
Developer: EGEE3 JRA1
Documentation: <http://glite.web.cern.ch/glite/documentation/default.asp>
Code repository: <http://glite.cvs.cern.ch:8180/cgi-bin/glite.cgi/>
Development language: Java
Build environment: ETICS
Download: <http://glite.web.cern.ch/glite/packages/R3.1/>
Licence: Apache License

Status

Maturity: Production
Usage: Grid Services
Interface maturity: Production
Support status: short term (until the end of the EGEE3 project),
Support team: EGEE3/JRA1 team
Development plans: maintenacne

Distribution/Availability

Middleware stack: gLite 3.1
Standalone usability: *yes*
Component dependencies: none
External software dependencies: *full list available on request*
Component consumers: Grid Services
Supported platforms: SLC4/x86
Planned to be supported platforms: SLC5/x86_64

Library component information

Consumers: *Grid Services*
Language binding: *Java*

Interoperability/Standards

Standard-compliance: none
Interoperable components: *none*
Additional Standards: *none*

5.3.21 gLite: Trustmanager

Name gLite Trustmanager

Description

The Trustmanager is a Java library to check the validity of an X509 certificate chain and is interfaced with the Java extensible SSL engine. It contains the machinery to configure the Java SSL engine, to read the certificates in the format used in Grids, to support Certificate Revocation Lists (CRL) and to automatically reload at runtime the CRLs or the user credentials when they change. The package also contains plugins for integration with Tomcat web server and Axis web service clients. Trustmanager uses the bouncycastle Java cryptographic routines and depends on the presence of a servlet container.

Basic information

Component type: Library
Developer: EGEE3 JRA1
Documentation: <http://glite.web.cern.ch/glite/documentation/default.asp>
Code repository: <http://glite.cvs.cern.ch:8180/cgi-bin/glite.cgi/>
Development language: Java
Build environment: ETICS
Download: <http://glite.web.cern.ch/glite/packages/R3.1/>
Licence: Apache License

Status

Maturity: Production
Usage: Grid Services
Interface maturity: Production
Support status: short term (until the end of the EGEE3 project),
Support team: EGEE3/JRA1 team
Development plans: maintenacne

Distribution/Availability

Middleware stack: gLite 3.1
Standalone usability: *yes*
Component dependencies: none
External software dependencies: *full list available on request*
Component consumers: Grid Services
Supported platforms: SLC4/x86

Planned to be supported platforms: SLC5/x86_64

Library component information

Consumers: *Grid Services*

Language binding: *Java*

Interoperability/Standards

Standard-compliance: *none*

Interoperable components: *none*

Additional Standards: *none*

5.3.22 UNICORE: Gateway

Name UNICORE Gateway

Description

The Gateway component acts as the entry point to a UNICORE site and performs the authentication of all incoming requests. From the UNICORE perspective it is the "door to the outside world" in a site firewall and may serve several resources/target systems behind it. In more detail the UNICORE Gateway checks whether a certificate is signed by a trusted CA and valid (date check) as well as not revoked using Certificate Revocation List (CRL) principles. The UNICORE Gateway performs purely identity-based authentication using the certificate information from the transport level connection (i.e. TLS/SSL), optionally accepting Globus proxy certificates. It provides basic HTTP proxy capabilities, i.e. it will forward HTTP GET, PUT and POST requests to the target sites behind it.

Basic information

Component type: *service*

Developer: UNICORE SourceForge Community

Documentation: http://www.unicore.eu/unicore/architecture/service-layer.php#anchor_gateway

Code repository: <http://unicore.svn.sourceforge.net/svnroot/unicore/gateway/>

Development language: *Java*

Build environment: *Apache Maven*

Download:

<http://downloads.sourceforge.net/unicore/unicore-servers-6.2.0-p1.tgz>

Licence: *BSD License*

Status

Maturity: *production*

Usage: *DEISA, D-Grid, SKIF-Grid, commercial partners*

Interface maturity: *production*

Support status: *long-term*

Support team: *supported by UNICORE SourceForge Community*

Development plans: *Full HTTP support*

Distribution/Availability

Middleware stack: *UNICORE Core Server Bundle Version 6.2*

Standalone usability: *yes*

Component dependencies: *None*

External software dependencies: *None*

Component consumers: *None*

Supported platforms: *All Java-based platforms*

Planned to be supported platforms: *None*

Service component information

Public Interfaces: purely SSL (i.e. https)
Aligned security: TLS/SSL (optional proxy support)
Capabilities: security.authentication

Interoperability/Standards

Standard-compliance: non-umd (TLS/SSL, proxy certificates)
Interoperable components: None
Additional Standards: HTTP

5.3.23 UNICORE: XUADB

Name UNICORE enhanced UNICORE User Database (XUADB)

Description

For mainly identity-based authorization of users the XNJS uses the XUADB user database to perform the mapping from X.509 certificates to the actual users' logins and roles. The XUADB component is a Web service in itself, so it can be used from multiple UNICORE installations, e.g. within the same computing centre. Full X.509 certificates are used as base line, while the access control is based on XACML policies. Proxy certificates are optionally supported in UNICORE 6 e.g. to use GridFTP.

Basic information

Component type: *service*
Developer: UNICORE SourceForge Community
Documentation: http://www.unicore.eu/unicore/architecture/service-layer.php#anchor_xuadb
Code repository: <http://unicore.svn.sourceforge.net/viewvc/unicore/xuadb/>
Development language: Java
Build environment: Apache Maven
Download:
<http://downloads.sourceforge.net/unicore/unicore-servers-6.2.0-p1.tgz>
Licence: BSD License

Status

Maturity: *production*
Usage: *DEISA, D-Grid, SKIF-Grid, commercial partners*
Interface maturity: *production*
Support status: *long-term*
Support team: supported by UNICORE SourceForge Community
Development plans: *None*

Distribution/Availability

Middleware stack: UNICORE 6.2
Standalone usability: *yes*
Component dependencies: *None*
External software dependencies: *None*
Component consumers: UNICORE Atomic Services, UNICORE OGSA-BES
Supported platforms: *All Java-based platforms*
Planned to be supported platforms: *None*

Service component information

Public Interfaces: proprietary XUADB WS interface
Aligned security: *TLS/SSL and optional proxies, identity-based authorization PIP*
Capabilities: security.authorization; security.identitymapping

Interoperability/Standards

Standard-compliance: *None*
Interoperable components: *None*
Additional Standards: *WS*

Comment [BK11]: ???

5.3.24 UNICORE: XACML Entity

Name UNICORE Extensible Access Control Markup Language (XACML) Entity

Description

The XACML Entity in UNICORE performs authorization decisions based on XACML policy files. This is used as an PDP in UNICORE that allows for identity-based authorization in conjunction with the XUADB PIP and attribute-based authorization decisions if, and only if, attributes have been used in conjunction with UNICORE (e.g. using UVOS).

Although this component is not directly a full UNICORE component it is essential in understanding the architecture of UNICORE with respect to security.

Basic information

Component type: other (*non Web service functionality*)

Comment [BK12]: ???

Developer: UNICORE SourceForge Community

Documentation: http://www.unicore.eu/unicore/architecture/service-layer.php#anchor_xuadb

Code repository: <http://unicore.svn.sourceforge.net/viewvc/unicore/xnjs/>

Development language: Java

Build environment: Apache Maven

Download:

http://downloads.sourceforge.net/unicore/unicore-quickstart-6.1.3-p1.tgz?use_mirror=fastbull

Licence: BSD License

Status

Maturity: production

Usage: DEISA, D-Grid, SKIF-Grid, commercial partners

Interface maturity: production

Support status: long-term

Support team: UNICORE SourceForge Community

Development plans: None

Distribution/Availability

Middleware stack: UNICORE Core Server Bundle Version 6.1.3

Standalone usability: No

Component dependencies: None

Component consumers: UNICORE Atomic Services, UNICORE OGSA-BES

Supported platforms: *All Java-based platforms*

Planned to be supported platforms: *None*

Service component information

Public Interfaces: None

Aligned security: certificate-based identity information and security attributes are used to perform authorization decisions

Capabilities: security.authorization

Library component information

Consumers: None

Language binding: None

Interoperability/Standards

Standard-compliance: None
Interoperable components: None
Additional Standards: XACML

5.3.25 UNICORE: UVOS

Name UNICORE Virtual Organisation Service

Description

As an alternative or as an addition to the XUADB UVOS, this Virtual Organisation (VO) Service using the SAML standard can be used to authorize users. It acts as an attribute-authority (AA) and releases signed attributes for end-users stating the position of an end-user within a VO as well as group membership or role possession. Released attributes are later used in conjunction with the UNICORE XACML entity to perform attribute-based authorization decisions.

Basic information

Component type: service
Developer: UNICORE SourceForge Community
Documentation: http://www.unicore.eu/unicore/architecture/service-layer.php#anchor_uvos
Code repository: <http://unicore.svn.sourceforge.net/svnroot/unicore/uvos> (currently being migrated)
Development language: Java
Build environment: Apache Maven
Download:
<http://downloads.sourceforge.net/unicore/VOManager-1.3.2-linux.gtk.x86.tar.gz>
Licence: BSD License

Status

Maturity: production
Usage: Chemomentum project
Interface maturity: production
Support status: long term
Support team: ICM Warsaw and UNICORE SourceForge Community
Development plans:

Distribution/Availability

Middleware stack: UNICORE VO Service 1.3.2
Standalone usability: yes
Component dependencies: *None*
External software dependencies: None
Component consumers: UNICORE atomic services, UNICORE OGSA-BES
Supported platforms: All Java-based platforms
Planned to be supported platforms: None

Service component information

Public Interfaces: SAML
Aligned security: Username/Password, TLS/SSL, attributes of end-users conveyed via SAML assertions
Capabilities: security.attributeauthority

Interoperability/Standards

Standard-compliance: non-umd (SAML)
Interoperable components: *None*
Additional Standards: None

5.3.26 other: MyProxy

Name MyProxy

Description

MyProxy is used for managing X.509 Public Key Infrastructure (PKI) security credentials (certificates and private keys). Within EGEE this service is used to overcome the problem of short-lived proxies and long jobs. It allows the user to create and store a long-term proxy in a dedicated server (a MyProxy server). The WMS, CREAM CE, VOBOX and FTS will then be able to use this long-term proxy to periodically renew the proxy for a submitted job before it expires and until the job ends (or the long-term proxy expires).

Basic information

Component type: Service
Developer: EGEE3 JRA1
Documentation: <http://grid.ncsa.uiuc.edu/myproxy/>
Code repository: gLite code repository <http://jra1mw.cvs.cern.ch:8180/>
Development language: C++
Build environment: ETICS
Download: <http://glite.web.cern.ch/glite/packages/R3.1/>
Licence: Apache License

Status

Maturity: Production
Usage: Infrastructure
Interface maturity: Production
Support status: short term (until the end of the EGEE3 project),
Support team: EGEE3/JRA1 team
Development plans: maintenance

Distribution/Availability

Middleware stack: VDT
Standalone usability: yes
Component dependencies: none
External software dependencies: none
Component consumers: WMS, CE
Supported platforms: many
Planned to be supported platforms: many more

Service component information *(only for a service component)*

Public Interfaces: MyProxy interface
Aligned security: none
Capabilities: none

Interoperability/Standards

Standard-compliance: none
Interoperable components: none
Additional standards: none

Other

Any other kind of information

6 Accounting functionality

6.1 Standards landscape

There are currently two OGF specifications (still in their draft stage) dealing with representation and manipulation of accounting information. These specifications are the Usage Record (UR) [5] and Resource Usage Service (RUS) [6]. The Usage Record specification describes an XML notation for describing resource usage information, while the Resource Usage Service specification describes the Web Service interface for a standards-compliant accounting system, which enables grid resource usage auditing and accounting as well as grid economic model. The RUS service is needed to provide interface for UR insertion to accounting database and to provide query interface for information about resource usage to a variety of Grid entities: service manager who wishes to examine utilization across their resources; a service that wishes to charge for the use of the consumed resources; and a virtual organization that wishes to monitor resource activity within their Grid. Usage information is made available in UR format.

The main limitations of the RUS/UR specifications, which make them unsuitable for production use, are their immature status and the fact that the RUS interface does not allow aggregate resource usage information to be accessed. This means that, for example, it is not possible to get the average, per-VO utilization of a given computational resource over a given time span. Only individual accounting records can be provided and the client must then aggregate the data. This is highly inefficient, as the RUS service must provide large amount of information to clients, with serious performance problems.

Unfortunately, neither RUS nor the UR in its current form is ready for production usage. Currently the RUS specification is being improved by the OGF RUS-WG. Either a new version will be available in a reasonable timescale (~1 year) so that it can be evaluated and eventually implemented in production Grids or the existing RUS specification needs to be heavily profiled to remove its limitations and be used in production.

The currently available accounting-related standards do not provide useful recommendation for public accounting interfaces and associated data exchanged, nor is such a common standard expected to emerge in a reasonable timescale. However, given that accounting is usually very middleware-specific (and often also VO-specific), it is acceptable that each accounting system will maintain a different interface as long as these interfaces are open and properly published as well as well documented. Thus, within UMD it is suggested that accounting components should come with properly documented open interfaces.

6.2 Overview table: Accounting area

	Name	Type	Maturity	Standard compliance		
6.3.1	ARC: JURA	client	development	Non-umd		
6.3.2	gLite: Accounting Service	service	production	None		
6.3.3	gLite: APEL	client	production	None		
6.3.4	other: SGAS	service	production	Non-umd		

6.3 Inventory of components

6.3.1 ARC: JURA

Name Job Usage Reporter of ARC (JURA)

Description

The Job Usage Reporter of ARC (JURA) is an agent for reporting resource usage of jobs to accounting services. Its input consists of job log files created by the execution manager from metering data provided by the LRMS. These job logs are parsed, transformed into standard Usage Record format, and then submitted to accounting services for persistent recording. Currently ARC Job Usage Reporter interacts with the SGAS Logging and Usage Tracking Service.

Basic information

Component type: client

Developer: Lund University & NIF of NorduGrid

Documentation:

<http://svn.nordugrid.org/repos/nordugrid/arc1/trunk/doc/jura/jura-tech-draft.pdf>

Code repository:

<http://svn.nordugrid.org/repos/nordugrid/arc1/trunk/src/clients/jura>

Development language: c++

Build environment: autotools

Download: download.nordugrid.org

Licence: Apache Licence version 2.0

Status

Maturity: development

Usage: Hungarian Clustergrid, Nordic Grid deployments

Interface maturity: production

Support status: short term provided by the KnowARC project

Support team: NIF

Development plans: add support for other accounting services such as the gLite Accounting Service once their interfaces are well-defined and documented.

Distribution/Availability

Middleware stack: will be part of the next technology preview release of ARC (0.9.4)

Standalone usability: none

Component dependencies: ARC A-REX, ARC HED

External software dependencies: libxml, Glibm, openssl

Component consumers: ARC A-REX, SGAS

Supported platforms: all flavors of linuxes

Planned to be supported platforms: windows, mac-osx

Interoperability/Standards

Standard-compliance: non-UMD

Interoperable components: none

Additional standards: makes use of OGF UR to report to SGAS

Other

The JURA component is an important module undergoing rapid development and expected to be deployed in production as a replacement of the currently used SGAS JARM module already around the 3rd quarter of 2009.

6.3.2 gLite: Accounting Service

Name gLite Accounting Service

Description

The accounting information is sent to a central accounting database where it is processed to generate statistical summaries that are available through the EGEE/WLCG Accounting Portal.

These statistics are then available for use, in different views, by: all scientific communities using the named Grid Infrastructures; the owners and administrators of the constituent resources, and the management of the Grid Infrastructures. For example, they include the total CPU consumed by each VO.

Accounting portal: http://www3.egee.cesga.es/gridsite/accounting/CESGA/egee_view.php

Basic information

Component type: service
Developer: CESGA
Documentation: *N/A*
Code repository: *N/A*
Development language: *N/A*
Build environment: *N/A*
Download: *N/A*
Licence: *N/A*

Status

Maturity: Production
Usage: Infrastructure
Interface maturity: mature
Support status: Fully supported
Development plans: maintenance

Distribution/Availability

Middleware stack: gLite/EGEE
Standalone usability: yes
Component dependencies: none
External software dependencies: *N/A*
Component consumers: end users and VO managers
Supported platforms: *N/A*
Planned to be supported platforms: *N/A*

Service component information

Public Interfaces: *N/A*
Aligned security: *none*
Capabilities:

Interoperability/Standards

Standard-compliance: none
Interoperable components: none
Additional Standards: none

6.3.3 gLite: APEL

Name APEL (Accounting Processor for Event Logs)

Description

APEL is a component of EGEE CPU accounting distributed as part of the gLite middleware. APEL parses batch system and Grid gatekeeper logs at each Grid site to generate CPU usage records and publishes them into a centralised repository at a GOC (Grid Operations Centre).

Basic information

Component type: Accounting Client
Developer: RAL

Documentation: <https://edms.cern.ch/document/726137>
Code repository: <http://glite.cvs.cern.ch:8180/cgi-bin/glite.cgi/>
Development language: Java, perl
Build environment: ETICS
Download: <http://glite.web.cern.ch/glite/packages/R3.1/>
Licence: Apache

Status

Maturity: Production
Usage: CE
Available Documentation: None
Interface maturity: Production
Support status: Fully Supported
Development plans: maintenance

Distribution/Availability

Middleware stack: gLite
Standalone usability: yes
Component dependencies: none
External software dependencies: full list available on request
Component consumers: Accounting Service
Supported platforms: SL4
Planned to be supported platforms: SL5

Service component information

Public Interfaces: none
Aligned security: *none*
Capabilities: none

Library component information

Consumers: none
Language binding: none

Interoperability/Standards

Standard-compliance: none
Interoperable components: none
Additional Standards: none

6.3.4 other: SGAS

Name SweGrid Accounting System (SGAS)

Description

The SweGrid Accounting System (SGAS) is a Java implementation of a usage tracking and credit allocation enforcement system for the Grid. Services included are the Logging and Usage Tracking Service (LUTS) for recording usage data in an XML database after execution, and the Bank service for usage credit management. For these services, client applications called Job Account Reservation Manager (JARM) components are available for the computing element of an old version of the production ARC middleware and for Globus GRAM.

Basic information

Component type: collection of services, corresponding clients and management utilities
Developer: Umea University, NDGF
Documentation: <http://www.cs.umu.se/research/grid/sgas/>
Code repository: <svn://svn.cs.umu.se/proj/gird/sgas/trunk>

Development language: java
 Build environment: recently changed to Apache buildr from Maven1
 Download: <http://www.cs.umu.se/research/grid/sgas/releases/sgas-2.2.0.tar.gz>
 Licence: Apache License Version 2.0

Status

Maturity: production
 Usage: NDGF sites
 Interface maturity: production
 Support status: long term
 Support team: Umea University
 Development plans: replacement of the database backend, interface cleanup and optional standardization, add support for aggregations

Distribution/Availability

Middleware stack: SGAS distribution version 2.0
 Standalone usability: yes
 Component dependencies: none
 External software dependencies: java, globus java-ws-core 4.0.8
 Component consumers: ARC Usage Reporter
 Supported platforms: linuxes
 Planned to be supported platforms: none

Service component information *(only for a service component)*

Public Interfaces: WSRF Resource Properties-based custom interface to add/query usage records
 Aligned security: TLS/SSL proxies are supported, VOMS attributes are not understood, file-based and XACML-based access control policies within the servers
 Capabilities: security.accounting

Interoperability/Standards

Standard-compliance: non-UMD
 Interoperable components:
 Additional standards: incompletely supports an early version of OGF Usage Record

Other

SGAS is currently undergoing a major redesign and reimplementatation.

7 Other functionality

7.1 Overview table: Other components

	Name	Type	Maturity	Standard conformance		
7.2.1	ARC: HED	library	Pre-production	Non-umd		

7.2 Inventory of components

7.2.1 ARC: HED

Name ARC Hosting Environment (Daemon) Framework (HED)

Description

The ARC HED is a collection of libraries and plugins for building lightweight client/service applications. It is XML based and has minimalistic SOAP support. One of the most important feature of HED is it's extremely flexible architecture due to pluggable functionality. Current implementation of HED comes with modules for developing SOAP and Web Services. It serves as base infrastructure for all ARC components being developed.

Basic information

Component type: library (accompanied by plugins and a hosting executable)
Developer: NIIF, Oslo University from NorduGrid
Documentation: http://www.nordugrid.org/documents/ARCHED_article.pdf
Code repository: <http://svn.nordugrid.org/trac/nordugrid/browser/arc1/trunk/src/hed>
Development language: C++
Build environment: GNU autotools
Download: download.nordugrid.org
Licence: Apache License Version 2.0

Status

Maturity: pre-production
Usage: HED is a central component being part of most ARC deployments either client or server-side
Available Documentation: http://www.nordugrid.org/documents/ARCHED_article.pdf
Interface maturity: production (current features stable, new features added in compatible way)
Support status: short term support by KnowARC project, long term support by NorduGrid collaboration
Support team: Oslo University
Development plans: continuously adding new plugins. implementing new features

Distribution/Availability

Middleware stack: technology preview releases of ARC (v 0.9.x) and to be part of the production ARC release starting from the coming v 0.8.0 scheduled for May 2009
Standalone usability: Yes
Component dependencies: none
External software dependencies: libxml, Glibm, openssl, plus several optional dependencies
Component consumers: all the newer ARC components
Supported platforms: Linux, MacOS, Solaris, Windows (preliminary)
Planned to be supported platforms: improve support for windows

Library component information

Consumers: recently developed ARC services and client tools
Language binding: C++, Python, Java (not tested)

Interoperability/Standards

Standard-compliance: non-umd (plugins for handling TCP, TLS, GSI, SOAP, WSRF selected functionality, WS-Security selected functionality, etc.
Interoperable components: None
Additional Standards: see above

Other

HED is scheduled to be included in the next production release of ARC together with the first HED-based service, the A-REX.

8 Effort estimates

		FTEs per year			
		Maintenance, Support	Standard Convergence	New development	Total
ARC components		9.100	5.800	8.000	22.900
Compute	Grid Manager	0.500	0.000	0.000	0.500
Compute	A-REX	1.000	1.000	0.500	2.500
Compute	libarcclient (job management)	0.500	0.500	0.250	1.250
Compute	ng* job CLI	0.250	0.000	0.000	0.250
Compute	arc* job CLI	0.500	0.000	0.250	0.750
Compute	Janitor	0.250	0.000	0.750	1.000
Data	Classic SE	0.250	0.000	0.000	0.250
Data	libarcdata2	1.000	0.500	0.500	2.000
Data	ng* data CLI	0.250	0.000	0.000	0.250
Data	arc*data CLI	0.500	0.000	0.500	1.000
Data	Chelonia	0.250	0.500	0.500	1.250
Data	Chelonia client	0.250	0.000	0.250	0.500
Information	Classic Infoserver	0.250	0.000	0.000	0.250
Information	Classic Infoindex	0.250	0.000	0.000	0.250
Information	ALIS	0.250	0.500	1.000	1.750
Information	ISIS	0.250	0.000	0.750	1.000
Information	libarcclient (infosys)	0.500	0.500	0.500	1.500
Information	Grid Monitor	0.250	0.250	0.250	0.750
Security	HED security API	0.250	0.500	0.250	1.000
Security	Shibbridge	0.100	0.100	0.250	0.450
Security	arcproxy	0.100	0.000	0.000	0.100
Security	Charon	0.100	0.100	0.250	0.450
Security	Proxy Store	0.100	0.000	0.250	0.350
Security	Fruitfly Credential Service	0.100	0.200	0.000	0.300
Accounting	JURA	0.100	0.150	0.500	0.750
Other	HED	1.000	1.000	0.500	2.500
Associated components					
Data	dCache	4.000	1.000	1.000	6.000
Accounting	SGAS	0.500	0.750	0.750	2.000

		FTEs per year			
		Maintenance, Support	Standard Convergence	New development	Total
UNICORE components		3.535	3.215	4.910	11.660
Compute	TSI	0.125	0.125	0.250	0.500
Compute	XNJS	0.250	0.250	0.500	1.000
Compute	UAS-Compute	0.250	0.500	0.250	1.000
Compute	OGSA-BES	0.250	0.250	0.000	0.500
Compute	Workflow-Engine	0.250	n/a	0.250	0.500
Compute	Service Orchestrator	0.125	0.250	0.250	0.625
Compute	URC	0.500	0.500	0.500	1.500
Compute	UCC	0.125	0.125	0.250	0.500
Compute	HILA	0.125	0.125	0.250	0.500
Data	UAS-Data	0.250	0.500	1.000	1.750
Information	CIS	0.250	0.250	0.500	1.000
Information	Service Registry	0.125	n/a	n/a	0.125
Security	Gateway	0.250	n/a	0.250	0.500
Security	XUADB	0.125	n/a	0.125	0.250
Security	XACML Entity	0.080	n/a	0.080	0.160
Security	UVOS	0.250	0.250	0.250	0.750
Other	WSRFlite	0.125	0.000	0.125	0.250
Other	UNICORE/X	0.080	0.090	0.080	0.250

		FTEs per year			
		Maintenance, Support	Standard Convergence	New development	Total
gLite components		0.000	0.000	0.000	0.000
Compute	gLite: WMS				0.000
Compute	gLite: WMS Clients				0.000
Compute	gLite: LB				0.000
Compute	gLite: CREAM Service				0.000
Compute	gLite: CREAM CLI				0.000
Compute	gLite: LCG-CE				0.000
Data	gLite: DPM				0.000
Data	gLite: lcg_utils				0.000
Data	gLite: htcp				0.000
Data	gLite: SRM client				0.000
Data	gLite: DPM client				0.000
Data	gLite: LFC client				0.000
Data	gLite: FTS client				0.000
Data	gLite: DPM/DICOM interface				0.000
Data	gLite: LFC				0.000
Data	gLite: FTS				0.000
Data	gLite: Hydra				0.000
Data	gLite: AMGA				0.000
Data	gLite: AMGA Client API				0.000
Information	gLite: BDII				0.000
Information	gLite: lcg-info and lcg- infosites				0.000
Information	gLite: Service Discovery API				0.000
Security	gLite: VOMS				0.000
Security	gLite: VOMS-Admin				0.000
Security	gLite: voms-proxy-*				0.000
Security	gLite: ARGUS				0.000
Security	gLite: Short Lived Credential Service				0.000
Security	gLite: SCAS				0.000
Security	gLite: LCAS				0.000
Security	gLite: LCMAPS				0.000
Security	gLite: gLite-Exec (gExec)				0.000
Security	gLite: gLite-Delegation Client				0.000
Security	gLite: Shibboleth Short Lived Credential Service Client				0.000
Security	gLite: Gridsite				0.000
Security	gLite: CGSI_gSoap				0.000

Security	gLite: delegation-Java				0.000
Security	gLite: Trustmanager				0.000
Accounting	gLite: Accounting Service				0.000
Accounting	gLite: APEL				0.000
Other					0.000
Associated components					
Data	StoRM				0.000
Security	myProxy				0.000

ARC subtotals with 20% overhead	
Security	3.180
Data	6.300
Compute	7.500
Information	6.600
Accounting	0.900
Other	3.000
<u>TOTAL</u>	<u>27.480</u>

Unicore subtotals	
Security	1.660
Data	1.750
Compute	6.625
Information	1.125
Other	0.500
<u>TOTAL</u>	<u>11.660</u>

gLite subtotals with overhead	
Security	0.000
Data	0.000
Compute	0.000
Information	0.000
Accounting	0.000
Other	0.000
<u>TOTAL</u>	<u>0.000</u>

9 References

- [1] Ali Anjomshoaa, Fred Brisard, Michel Drescher, Donal Fellows, An Ly, Stephen McGough, Darren Pulsipher, and Adreas Savva. Job Submission Description Language (JSDL) Specification, Version 1.0, November 1 2005. OGF Specification GFD-R.056, <http://www.gridforum.org/documents/GFD.56.pdf>.
- [2] I. Foster, A. Grimshaw, P. Lane, W. Lee, M. Morgan, S. Newhouse, S. Pickles, D. Pulsipher, C. Smith, and M. Theimer. OGSA Basic Execution Service Version 1.0, August 2007. OGF Specification GFD.108, <http://www.ogf.org/documents/GFD.108.pdf>.
- [3] Blair Dillaway, Marty Humphrey, Chris Smith, Marvin Theimer, and Glenn Wasson. HPC Basic Profile, Version 1.0, August 28 2007. OGF Specification GFD-R-P.114, <http://www.ogf.org/documents/GFD.114.pdf>.
- [4] Marty Humphrey, Chris Smith, Marvin Theimer, Glenn Wasson, JSDL HPC Profile Application Extension, Version 1.0.
- [5] Usage Record Working Group: <https://forge.gridforum.org/projects/ur-wg/>
- [6] Resource Usage Service Working Group: <https://forge.gridforum.org/projects/rus-wg>
- [SEC1] R. Housley, W. Ford, W. Polk, D.Solo: Internet X.509 Public Key Infrastructure (PKI) Certificate and CRL Profile, January 1999. <http://www.ietf.org/rfc/rfc2459.txt>
- [SEC2] S. Tuecke, V. Welch, D. Engert, L. Pearlman, M. Thompson: Internet X.509 Public Key Infrastructure (PKI) Proxy Certificate Profile, Internet Engineering Task Force, June 2004. <http://www.ietf.org/rfc/rfc3820.txt>
- [SEC3] S. Farrel, R. Housley: An Internet Attribute Certificate Profile for Authorization, Internet Engineering Task Force, April 2002. <http://www.ietf.org/rfc/rfc3281>
- [SEC4] OASIS Security Services (SAML) Technical Committee, <http://www.oasis-open.org/committees/security/>
- [SEC5] OASIS XACML Technical Committee, http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=xacml