

WP2 – R2.1

Overview of WP2 middleware as present in EDG 2.1 release and outlook to demo functionality





Overview

- ◆ Changes from R2.0 to R2.1
- ◆ Interfaces
 - Usage of info services
 - Usage of Storage services
- ◆ Known Issues and possible resolutions
- ◆ Outlook
- ◆ Discussion

Changes from R2.0 to R2.1

◆ **BUGFIXES**

- No major new functionalities
- Almost all open bugs fixed with priority < 4
- Complied with some enhancement requests

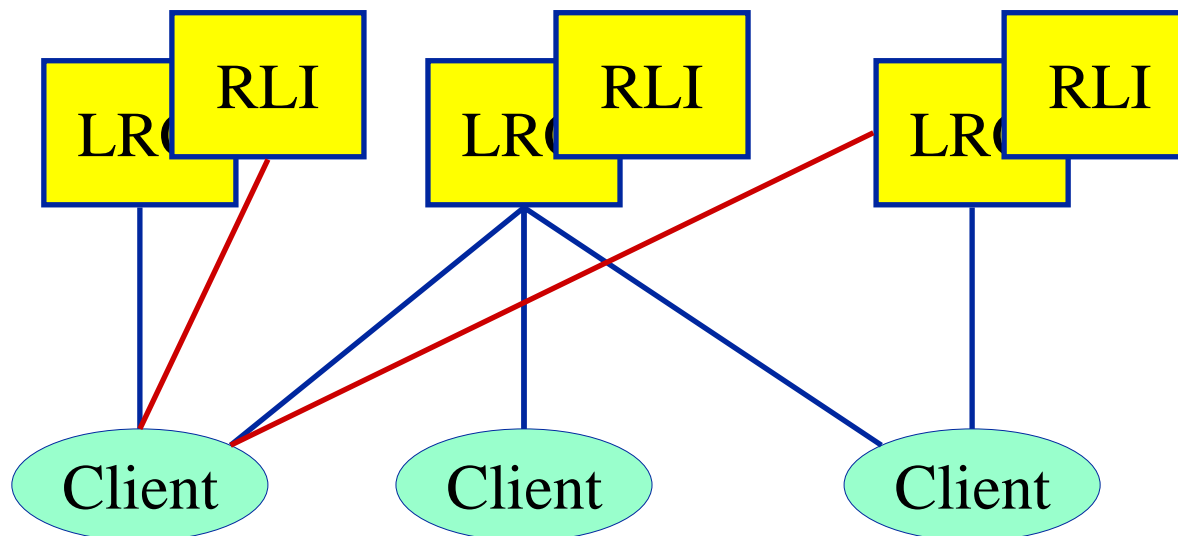
◆ Security

- VOMS certificate parsing
 - --vo flag is not necessary if user has issued voms-proxy-init
 - All services accept VOMS certs as well as old-style certs
- Unified tomcat configuration and installation through java-security
 - Easy maintenance, modular (even running on the same node, VO-specific services can be managed with no/little disruption to others)
 - Same infrastructure for WP2,WP3,WP5

Changes from R2.0 to R2.1

◆ Replication

- Full RLS, i.e. Introduction of Replica Location Indices
- No visible change to the user, all should be transparent
- Much better scalability: a local catalog is always available (LRC)
- Single central LRC model is still possible



Changes from R2.0 to R2.1

◆ Metadata

- RLS & RMC: enhanced user-defined metadata
 - Support for int, float in addition to string type
 - Support for complex query using SQL where clause i.e get all files with "I > 2 AND KIND LIKE '%TOP' OR Z <= 0.54"
- Spitfire: security support and much easier adaptability

◆ Conformance checking

- Logical file names, GUIDs and SURLS are checked for their 'correctness' i.e. adherence to their definitions

OK	WRONG
lfn: test guid: 3254-3a342-3768c-8899db srm: //lxshare0408.cern.ch/wpsix/data/myfile.dat	mytest guid: test srm: //path/data.txt

- SURL normalization: //, ../ resolution

Interfaces – Storage

- ◆ Full SE support. Issues:
 - Delete support – WP5 claims its there but we could not test yet
 - Fast file existence check – " "
 - Asynchronous operations – not for R2.1
 - SRM interface support – not for R2.1
- ◆ GridFTP server support ("Classic SE mode")
 - Fake GridFTP server to make it look like an SE
 - Works well with CASTOR-GridFTP which provides MSS access
- ◆ SRM support
 - SRM v1.1 support – integration tests with Fermilab, CASTOR
 - Not well tested

Interfaces – Infosystem

- ◆ Full RGMA integration
 - Publication to proper RGMA interfaces (producers)
 - Lookup through native RGMA interfaces
 - Cron entries provided to keep information fresh
- ◆ Full MDS integration for LCG
 - Info providers also exist for GLUE schema
 - MDS support through LDAP interface
- ◆ Replica Manager also configurable using a configuration file (Stub).

NOTE: The config file attributes have changed. Please see the release notes on the WP2 website
<http://cern.ch/edg-wp2/release2.1/index.html>

Known Issues : Performance

- ◆ Performance of edg-rm Command Line Interface
 - Due mainly to startup of full JVM per call
 - Java API is fast enough
 - C++ API is most efficient
- ◆ Example: listReplicas (insecure)
 - CLI: ~3.5 sec
 - Java API: ~0.2 sec
 - C++: ~30ms
- ◆ Resolution: provide C++ - based CLI for most critical operations (register, lookup)

Known Issues : enabling security



- ◆ All services have been tested in secure mode using certificates.
- ◆ All Java connectivity works, all tests pass in secure mode
 - Performance penalty: extra 200ms per call (or more, depending on latency. Nothing we can do about that, this is the nature of the SSL handshake mechanism, i.e. the price we pay for security.)
- ◆ Services do not provide authorization configuration options yet. This is something we want to provide for the last demo.
- ◆ C++ secure client has had a lot of issues, which were resolved only very late
 - Not enough testing has been done
 - Not advisable to deploy for production yet
 - Primary interface of RB to WP2
- ➔ Cannot enable security for the services on prod.TB NOW, but probably in a few weeks' time (which is probably too late)
- ➔ Solution is there and exists and we have deployed far worse stuff in the past, but now we're more cautious (and smarter).

Known Issues: Complexity

- ◆ Too much of the internals of the replication mechanism is exposed to the user.
- ◆ Replica Manager Service deferred to next project
 - Provide ease-of-use
 - Failsafety
 - Consistency checking
 - Reliability
 - Needs proper security delegation mechanism (through OGSA?)
- ◆ Solution to some extent: RLS Front-End (aka proxy) service, probable prototype to be demo'd at last review
- ◆ **NOTE:** POOL does not solve the problem but obscures it a bit, at least currently. It is a step in the right direction, but architecturally in the wrong layer.

Known Issues: Bulk operations

- ◆ Missing GDMP functionality: bulk transfer and subscription-based replication
 - Was planned for 2.1 but was cut off to keep features low
 - Experiments drove this feature cut
 - Effort was not available to keep GDMP supported or to adapt it to the new system
- ◆ RSS – Replication Subscription Service is an unfinished prototype. Its future is unclear.

Next steps

◆ **BUGFIXES**

◆ Make full RLS work

- It is essential to gather experience using a 'real' distributed Grid Middleware
- RGMA integration and automatic subscription can only be fully tested on the real system. There may be a few hiccups. Manual subscription is always possible and will probably have to be done.

◆ Work with applications to have a real Spitfire instance in production

- Gathering experience on this layer is also essential to be able to understand the requirements properly

For the demo:

◆ Provide RLS-FE service. This will lower the complexity on the client side and deal with the outbound connectivity issue at least for this service.

◆ Provide authorization options for RMC, RLS (i.e. Who may do what?) based on VOMS proxy content.



Conclusion and Outlook

- ◆ We have provided working services to EDG and LCG on a very aggressive timescale.
 - Gained a lot of expertise in the domain of data management and security
 - We established a lot of bonds to other organizations and have built good collaborations
- ◆ We have learned a LOT on how to NOT do things in a project
- ◆ Towards the end we have started to do things right but
 - People got overloaded while adapting the proper procedures and keeping the tight timelines at the same time
- ◆ I think EDG has provided the most solid Grid M/W there is today!
 - We should advertise it at least as much as our friends from other projects do.