Enabling Grids for E-science in Europe

www.eu-egee.org

PTF 16.06.2004

Requirements from Operations

subjective summary of input from several SA1 persons

EGEE is a project funded by the European Union under contract IST-2003-508833

Contents

Enabling Grids for E-science in Europe

- Requirements gathered at the JRA1/SA1 meetings
- Recent Input

- What's next?

JRA1-SA1 Meetings

Enabling Grids for E-science in Europe

- Meetings between JRA1-SA1
  - 06/02/2004
  - 30/03/2004
  - 02/04/2004
  - 20/04/2004
  - 09/06/2004
- Special emphasis on deployment issues:
  - Reference platforms
  - Dependencies (versions of compilers etc.)
  - Packaging
  - Configuration
  - Release and delivery
  - Support
  - Networks
- Reference:
  - SA1 requirements: https://edms.cern.ch/document/456865
  - Minutes of meetings: https://edms.cern.ch/document/451069

JRA1-SA1 Meetings

Enabling Grids for E-science in Europe

Id	Requirement	Status
1	**Mw delivery to SA1**	
1.1	Delivery process: get a tarball*, certify, feedback to JRA1, and iterate; when SA1 converges to a certified version, they will say so and JRA1 must package exactly the version that is certified, with no changes or additions	Agreed
1.2	Keep the granularity of components and packages as fine as possible	Agreed
2	**Release management**	
2.1	Quick bug releases and security patches must have a very quick turn-around	Agreed
2.2	Bug fixes will be provided to all versions being run by SA1	Agreed
3	**Deployment scenarios**	
3.1	JRA1 will deliver deployment recommendations for services as part of a release, and define the minimum running requirements for the entire system	Agreed
3.2	SA1 will give feedback to JRA1 about the most common deployment scenarios	Agreed
4	**Middleware configuration**	
4.1	Mw installation and configuration should be kept separated	Agreed
4.1	Provide simple and tool independent configuration mechanism	Mechanism not decided yet. To be discussed further.
4.2	JRA1 will provide a standard set of configuration files and documentation with examples that SA1 can use to design tools. Format to be agreed between SA1-JRA1	Agreed
4.3	It is the responsibility of SA1 to provide configuration tools to the sites	Agreed
4.4	Classify the configuration parameters into: - The ones that must be changed - The ones that might be changed at some time - The ones that are rarely changed - And always give valid/sensible default values	To be discussed further.
4.5	*Give configuration instructions covering all the given deployment use cases (see req 3.2). Or a tool that can do that*	*Depends on 3.1*
5	**Enforcement of the procedures**	
5.1	If bugs are found at any stage (certification, validation, deployment), bugs must be reported to JRA1 using Savannah. Bug fixing is JRA1 responsibility.	Agreed. Close collaboration needed between JRA1 integration/testing and SA1 operations
6	**Platforms to support**	

6	**Platforms to support**	
6.1	Primary platform: Red Hat Enterprise 3.0, gcc 3.2.3 and icc8 compilers (both 32 and 64-bits)	Agreed
6.2	Secondary platform: Windows (XP/2003), vc++ 7.1 compiler (both 32 and 64-bits)	Agreed
6.3	Secondary platform will be used to ensure portability	Agreed
6.4	Services needed in the Worker Nodes should not be constrained to fixed platforms	Agreed
6.5	Reduce to the minimum the services to be run in the Worker Nodes and make them easily portable	Agreed
7	**Versions for compilers, libraries, third party software etc**	
7.1	Avoid using multiple versions of the same libraries, tools, external programs	Agreed
7.2	Change of agreed version done through request to the Change Control Board	Agreed
7.3	JRA1 will define the 3rd party software versions supported. The EGEE mw will be tested against these versions.	Agreed
7.4	JRA1 will not deliver 3rd party software as part of a release, but will provide a list of dependencies and documentation.	Agreed
7.5	*The mw should not be dependent on any particular version of some external software (e.g. perl, openssl)*	*To be decided in advance. Further changes should be requested via the CCB.*
8	**Programming languages**	
8.1	Recommendation to simplify and minimize the number of programming languages used (if possible, only one)	Recommendation
9	**Packaging and software distribution**	
9.1	Packages should be relocatable (the whole tree, not separated branches)	Agreed
9.2	JRA1 will provide a release packaged in the native format of all supported platforms and a tarball. Both source and binaries would be provided.	Agreed
10	**Others**	

Enabling Grids for E-science in Europe

10	Others	
10.1	Sites must be allowed to organize the network as they wish, internal or external connectivity, NAT, firewall, etc, all must be possible, no special constraints. WNs must not require Outgoing IP connectivity; Not inbound connectivity either.	Agreed
10.2	Don't use mail as alarm system	To be discussed
10.3	The mw should not require to be installed/configured/run as root (or as a privileged user)	To be discussed

Input

Enabling Grids for E-science in Europe

- Input during the last few days (not in chorological order)
- Kostas Koumantaros
  - Would like the platform requirements to be revisited
    - Conclusion of the JRA1-SA1 meeting 09/06/2004
      - The reference platform term is discarded; there will be no reference platform as such.
      - Deployment platforms: platforms supported by SA1. SA1 has to support a range of platforms as wide as possible. For this year, the platform support will be mainly oriented to Linux based (RH Enterprise 3.x), 32-64 bits platforms. But the project has to get to a situation where a wide range of platforms is supported in a simple way.
      - The support responsibility inside SA1 is shared between the ROCs, which will support the different platforms deployed in their respective regions, and be involved in debugging and understanding platform issues for the platforms deployed on the region (according to TA) before passing them to JRA1.
      - Linux RH Enterprise 3.x-based systems are the most widely accepted by the sites (see results of SA1 platform survey). No flavor of it (RHEL30, Scientific Linux, CEL3, Fermi Linux, etc) will be enforced/recommended.
      - JRA1 testing testbed (distributed in 3 sites: RAL, NIKEF and CERN) will run each site production platform (e.g. CEL3 in the case of CERN).
      - SA1 certification testbed will also try to reproduce as much as possible the different deployed platforms.
      - Windows still remains as secondary platform; JRA1 will compile, build and test on this platform as a way to ensure portability.
      - Both SA1 and JRA1 will not support the platforms, will support the middleware.

Input

Enabling Grids for E-science in Europe

- Andrea Ferraro
  - Suggests to split along the lines of ROC and CIC
    - ROC
      - release : incremental m/w release able to run over a non dedicated machine
      - deployment : platform independent, backward compatible, no security holes- 1st level support : clear documentation
    - CIC
      - core infrastructure: Grid service 'sla' measurement and standard log files
      - comprehensive error messages
      - Job-Vo accounting:
      - policies rule and Vo-based priority queues
      - 2nd level support
      - clear documentation and reporting line to JRA1

Input

Enabling Grids for E-science in Europe

- Rolf Rumler
  - Brought up an important issue concerning operations:
  - Components (services) must support the redirection of workflow in case of problems.
    - Example: Problem to drain a RB by other means than waiting
  - Grid components should have administrative interfaces
    - for monitoring services
    - taking out a site smoothly

    ……

Summary

Enabling Grids for E-science in Europe

- Summary
  - In terms of platforms, packaging, config. releases some work has been done and we are in the process of refining this (JRA1-SA1 meetings)
  - **Operational aspects need more attention (some random selections):**
    - Phasing new services in and old out in a transparent way
      - clear state model for services has to be defined that allows this
    - Moving state:
      - state should be transportable (services have to be reborn)
        - » A service stopped on node A should be restart able on node B by moving the state between the nodes
    - ONE set of APIs for administrative and monitoring interfaces
      - not different ones for each service
      - probably with authorization to allow remote operations by ROC/CIC/OMC
    - Unique logging format for all services
      - needed for audit trails which are required by all sites
      - needed for debugging
    - Accounting interface
      - currently software filters multiple log files and joins them (not ideal)
    - Software has to be able to reflect resource usage policies as complex as the sites have implemented them locally
      - we got request that sites would like to have time depended policies

Summary

Enabling Grids for E-science in Europe

- **More things to consider**
  - Scalability
    - analysis of scalability of services needed
      - and tested as far as possible
  - Single point of failures have to be avoided
    - for each service a short risk analysis is needed:
    - the effect on the overall system when:
      - the service is not accessible for different duration
      - the underlying hardware died
      - the service is overloaded
    - for services that are essential it has to be clear how the service can be deployed in a redundant way
      - effect on scalability, performance,
  - Exception handling
    - services have to be prepared to handle non standard situations
      - confronted with old protocols/ corrupted data
      - resource exhaustion (memory, CPU)
    - currently services sometimes just crash or get blocked

The last points can be rephrased:

- The designers and developers have to be aware that any component can fail at any time

Summary

Enabling Grids for E-science in Europe

- **Services and VOs**
  - Adding and removing VOs has to be a light weight operation for a site
    - Currently many resources/services are done per VO (pathes, services, etc.)
      - **hard to manage**
      - **error prone to add new VO (needs modifications at many places)**
    - VO based policies needed for resource utilization
    - Adding a VO and defining the policy should be done at one place for a site and not for each service/resource independently

Next Steps

Enabling Grids for E-science in Europe

- **Inside SA-1**
  - Need some brain storming on operational requirements
- Continue meetings between JRA1 and SA1