



Enabling Grids for E-scienceE

SA3 Status and Plans

Markus Schulz

SA3

CERN-IT-GD

www.eu-egee.org



- **Integration, Testing and Releases**
 - Primary goals
 - Merging tools, processes and stacks of LCG-2.7 and gLite
 - Contribute to the improvement of the software's quality
 - Move to new foundation when required (▪ SL4, vdt-1.3, condor)
 - Include additional components driven by TCG's priority list
 - Support multiple platforms

- **Integration, Testing and Releases**
 - Derived Goals
 - Define a process that is agreed by all stakeholders
 - *Component based*
 - Common configuration system for all components
 - *Component based*
 - Detailed documentation of releases and updates
 - *Traceability (hard links between bugs and fixes)*
 - Speed up the test process
 - *Component based*
 - Move to ETICS as THE build system
 - *for all components and platforms*

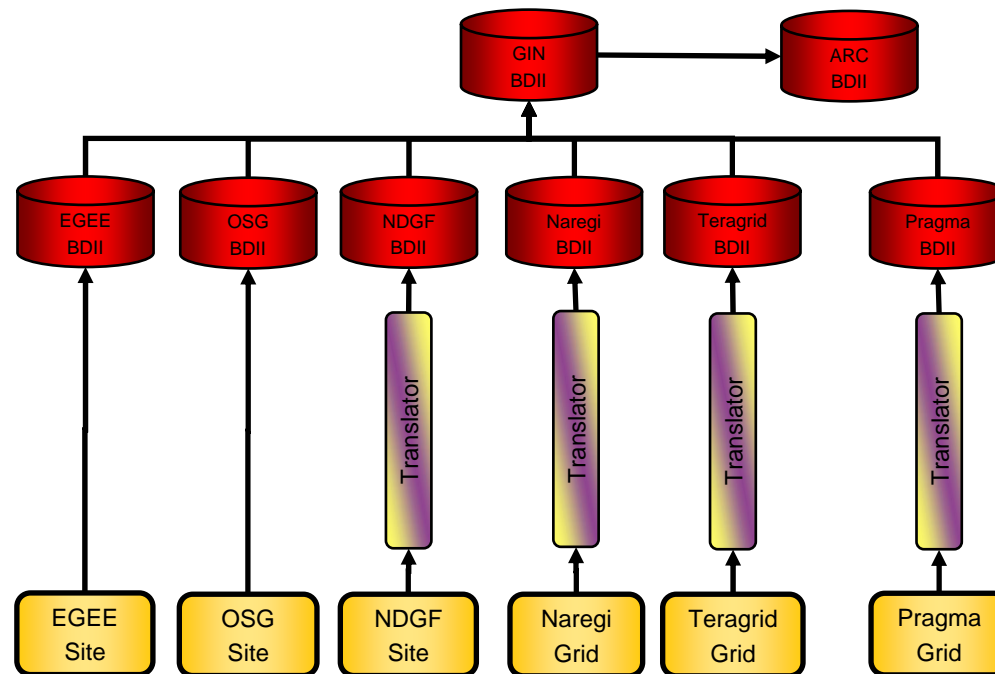
- **Integration, Testing and Releases**
 - **Derived Goals**
 - SL4 Port
 - Allow Condor, VDT and other core packages to be versioned by component
 - Improved coverage of deployment scenarios
 - *Batch systems, distributions, config tools*
 - *External testbeds*
 - Simultaneous releases for:
 - *64bit and 32bit (x86 64bit)*
 - *SL and Debian support*
 - Repository for automated tests
 - Vastly extended test coverage
 - *Including performance tests*
 - Regression tests
 - Common framework to present and archive test results

- **Interoperation**
 - Primary goals
 - Define plans to achieve interoperation with:
 - *ARC*
 - *Unicore*
 - Derived goals
 - Common strategy to achieve interoperation between grids
 - *Information systems*
 - *Schema evolution reflecting the need for interoperation*
- **Overall Goals**
 - Integrate the 13 SA3 partners into a working activity

- **First integration of LCG-2.7 and gLite on time for data challenges**
 - gLite-3.0 was not trouble free, but has been released
 - Uniform configuration management (for site managers)
- **PROCESS <----- The Achievement of the first year**
- **Definition of the software management process**
 - Iterative approach
 - Documented in milestone document
 - <https://edms.cern.ch/file/724371/1/EGEE-MSA3.2-724371-v1.5.doc>
 - Extension to “Experimental Services”
 - For massive scalability testing
 - For rapid turnaround with developers
 - Very lightweight
 - *Bugs and Patches are still tracked via Savannah*
 - *Changes on instance are tracked via wiki page*
 - *No configuration, certification, preproduction step*
 - Has been applied to WMS 3.0 and 3.1, gLite-CE

- **Implemented the Process**
 - Except the acceptance criteria, these have been defined, but have still to be implemented
- **The concept of splitting Bugs and Patches was really helpful**
- **The Savannah implementation worked very well**
 - Giving us traceability and accountability
 - Summaries for releases
- **EMT to prioritize bugs and patch processing**
 - Very compact overview
 - As a result we ended up with a few well aged patches
- **Processed since the introduction of the new process (August) > 70 patches**
 - And we know which bugs are fixed!!!
 - Currently 20 patches in progress

- **Interoperability work progressed beyond expectation**
 - OSG seamless for CMS since 8 months
 - After a slow start Unicore made real progress (Condor based)
 - Information system work in the scope of GIN
 - EGEE, OSG, Teragrid, NDGF, NGS, DEISA, Pragma, Naregi, APAC
 - Missing CrownGrid and Garuda, OurGrid (contact established)



- **Information System related work**
 - For historical reasons BDII and info providers are in JRA1 and SA3
 - In depth analysis of Site BDII performance problems in summer 2006
 - Provided solution
 - Additional caching and deployment recommendations
 - Significant contribution to evolution of GLUE
 - OGF Information System work

- **Providing “Experimental Production Service” WMS for CMS and Atlas**
 - Improved quality from the user perspective significantly
 - Uncovered operational shortcomings

- **Started same process for WMS-3.1 and gLite-CE**
 - Currently 3 CERN SA3 people are working 100% on WMS and CE
 - + 2*25% from EIS

- **Work with ETICS progressed relatively well (build related)**
 - System points into the right direction
 - Long term benefits can be expected
 - ETICS team responds well to feedback
 - System might need to be more adaptable to our way of working
 - Different projects will require different styles
 - There should not be “The ETICS Way”
 - Testing will require a second look
- **Testing**
 - Gap list of tests
 - Test plan with partners taking responsibility for tests
 - SAM as the backend for testing
 - Some test scripts have to be provided by partners
 - *but more has to come*
 - Multi step parallel testing with virtualized testbeds
 - *Has already started for configuration and installation tests*
 - First distributed testbeds

- **Tarball and RPM based SL3 WN that can be installed and used on SL4**
 - In pre-production since November 13
 - Experiments started over the weekend (Jan 12th) testing
 - Will be released as soon all 4 experiments have verified usability
- **Tarball based SL3 UI on SL4 nodes**
 - In tests since the beginning of January
 - Just started working on Monday
- **Better support for batch systems (Condor, SGE, LSF)**
 - See SA3 web page
 - Support provided by SA3 partners
- **Communication inside SA3**
 - 2 AllHands Meetings (last in November)
 - Weekly test coordination phone conference
 - Unfortunately we had to skip it quite often
- **The large number of bugs is related to the 16 fold increase in usage**
 - 2005: 1 Million Jobs 2006: 16 Million Jobs

- **More details in the breakout session**
- **SL4 port**
 - Hold back by build problems and the software
- **Only one platform supported**
 - TCD has additional ports, but not integrated in the release (one offs)
- **Testing**
 - Still lack of tests
 - Coverage and depth is not sufficient
 - We started performance testing for WMS and LFC
 - Regression testing very primitive
 - No growing set of tests
 - Poor test automatization
 - External tests not always reproducible
 - Testbed not reflecting component based approach

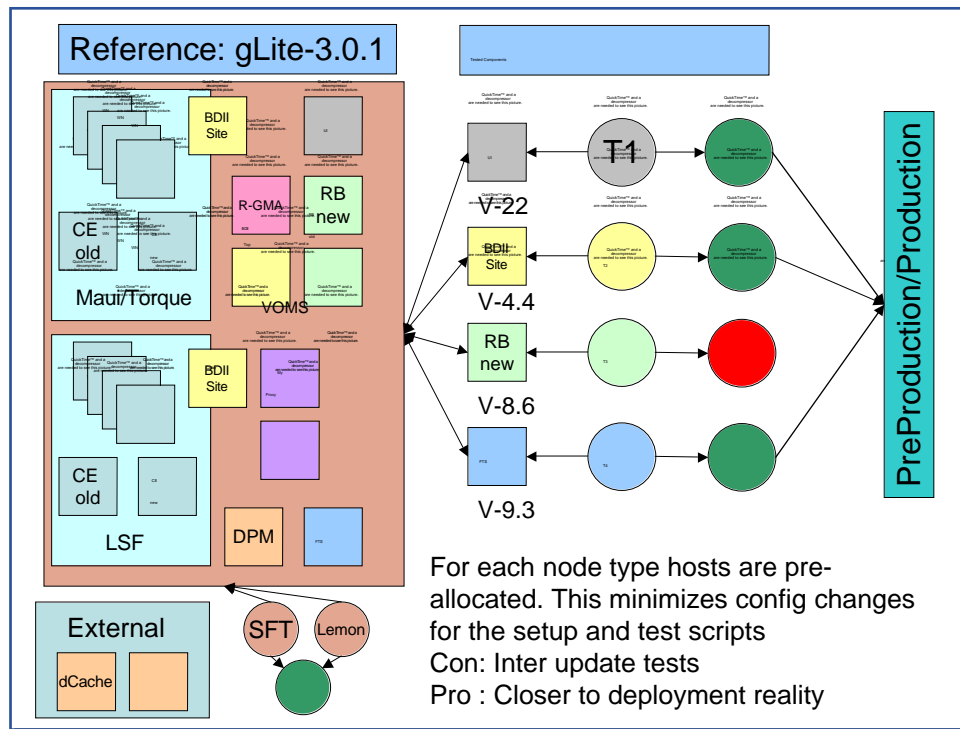
- **More details in the breakout session**
- **Team can only hold up with patches and platform moves (2-3 per week)**
 - No major new components integrated (AMGA, DGAS, glexec, GPBOX, ...)
 - However, plenty of change has been handled
- **Resources are not sufficient to introduce structural change while providing service**
 - Building
 - Testing
 - Configuration management
- **Configuration management (component based YAIM) is severely delayed**
 - Many patches required configuration changes
 - Linked are some config changes for Job Priority and DN based VO naming
 - For many patches this is the bottleneck

- **Configuration management needs more resources**
 - One of the bottlenecks of getting patches out
 - Due to high pressure many “trivial” errors are currently made
 - Affects the full chain
- **Restructuring and leaving the “Russian Doll” behind us**
 - We have to move to one layer, component based config.
 - Implementation of restructured YAIM is already advanced
 - First work on removing Python layer started
- **This will require reallocation of resources**
 - At least for 2 months

- **Moving all non legacy software to ETICS**
 - Finished within 3 months
- **The the real work starts** ▪
 - Restructuring the build (3rd quarter 2007)
 - Node type based
 - Component based
 - Cleaning client dependencies
 - Individual versioning of server dependencies
 - Continuous work until end of the project
 - Providing support for more platforms
 - X86 64 bits (July)
 - Debian (soon after build via ETICS works)
 - All official porting via ETICS (end of the 2007)
- **Process**
 - More formal approach to “Experimental Services”
 - Implement transition acceptance criteria as described in the process document
 - First for new components
 - AMGA as a test case by April

- **Testing**
 - Implementing the test plan
 - Described in the test plan milestone document
 - SAM as the backend for testing
 - *Started, continuous work*
 - *3rd quarter 2007*
 - Tests provided by partners as agreed
 - *2nd quarter 2007*
 - Multi step parallel testing with virtualized testbeds
 - *Has already started for configuration testing*
 - *In use for installation tests*
 - *2nd quarter 2007*
 - Proper regression tests
 - *Not before end of the project*
 - *Needs complete automation of testing to be efficient*

- **Restructure testbed management**
 - Reflecting component and node type based upgrades
 - Automation -----> improve reproducibility
 - Slowed down by tests needed for SL4 and VDT-1.3.x move
 - End 2007
- **ETICS for unit tests (end 2007??)**



- **Interoperation**

- ARC and UNICORE
 - Follow the planning document
 - Unicore we have some hopes to be faster
- Continue to work in the GIN activity
- Establish contacts to CrownGrid, Garuda and OurGrid
- Finish first release and evolve “Interoperation Cookbook” (March 2007)

- **Information Systems**

- In depth measurement of the scalability of the current system
 - 1st quarter 2007
- Propose scalability strategy (2nd quarter 2007)
- Continue Glue and OGF work

- **Introducing new components**
 - Driven by the TCG prioritized list
 - Not before transition to ETICS, SL4 and vdt-1.3 has been finished
 - Not more than 2 components at a time (resource limitation)

- **Continue with gLite-3.1 component stress testing and integration**
 - 1st quarter 2007

- **Repository for middleware provided by VOs**
 - Many interesting user space and service components are available
 - Too many to certify on our testbeds
 - Lack of expertise of testers
 - Providers should be encouraged to use ETICS (2nd quarter 2007)
 - Building
 - *Multiple platforms*
 - Basic testing
 - *coexistence testing*
 - *Additional unit tests if provided by developers*
 - SourceForge instance + APT repository via ETICS (March 2007)
 - We have to discuss what the minimal requirements are

- **Activity Coordination**
 - Two AllHands Meetings (not at CERN)
 - Partner reviews
 - Discussed at the November meeting
 - Draft by Zdenek Sekera
 - First early March
 - More regular phone conferences



Enabling Grids for E-scienceE

Review of SA3 Process Support for VO provided Software

Breakout Session

www.eu-egee.org



- **Discuss the problems mentioned in the “SA3 Status and Plans” presentation**
- **Discuss what should be done with the software process**
 - Explain why this is not the bottleneck
- **Understand better how we can make the middleware available that is used by the VOs**
 - Explain our limitations
 - Understand what the VOs require
- **Discuss a bit some general problems**
 - What means short term planning with full work plans
- **Mostly discussions**

- **Not complete**
- **No specific order**
- **Not meant to excuse our bad performance**
- **Should trigger discussions**

- **Tell us what we should do how and when**
 - **But be aware:**
 - **You might have to do it**

- **Applies to SL4, new components, bug fixes, and any change**
 - Complex builds
 - Complex dependencies
 - Complex configuration
 - Poorly tested software
 - Some small scale functional tests should have been done by the developers
- **The Process has not slowed us down (2-3 patches per week)**
 - Trivial patches have been moved by the EMT directly to production
 - Prioritization by EMT reflects projects overall needs
 - As a result some specific patches stay longer in the queue
 - For very active components the “Experimental Services” have been used (patches applied within hours)
 - **The traceability and accountability is indispensable**

- **gLite build concept:**
 - Assuming reflective dependencies:
 - You build WMS
 - You rebuild security too, because this might be affected
 - *The build environment has changed*
 - This is very save
 - *Even small inconsistencies are hidden*
 - This is very dangerous
 - *You loose the ability to separate out building blocks*
 - *You loose track of the real dependencies*

- **SL4 (and other) builds are delayed because:**
 - Many interdependencies between middleware components
 - Structure of the software:
 - Project, subsystem, component
 - *Back references make builds very complex*
 - *No clear split between clients with minimal requirements and services*
 - Component updates require updates of not related nodes
 - Have a look at our release page
 - Structure of the gLite build (which has to be exported to ETICS)
 - The build intelligence in the gLite system tried to achieve consistency
 - This was achieved by the way builds are done
 - *Costly subsystem builds*
 - This can work despite missing explicit information on component level
 - *Which made component based releases very hard*
 - gLite build concept:
 - *Assuming reflective dependencies:*
 - You build WMS
 - You rebuild security too, because this might be affected (shared dependencies)

- **The effects are severe**
- **An example**
 - Failing rebuilds of the WMS helper package holds up interoperation
 - This part of the new WMS writes the condor submit file
 - *Condor is ready to submit to Unicore and Arc*
 - This RPM has stopped progress for almost 6 month!!!!
- **What about ETICS???**
- **ETICS is in that respect more flexible**
 - But has to be used correctly (requires reworking of the builds)
 - Importing naively the gLite build information creates same and new problems
- **Then why didn't you move long time ago??**

- **We tried:**
 - First trials with ETICS in **March** 2006
 - Expected to be ready in **April**
 - No integrated system
 - **June** basic functionality was there
 - Missing edit functionality (ETICS team had to edit meta data)
 - **August** edit available
 - First tutorial EGEE-06 in **September**
 - Since **September**
 - People got trained
 - User feedback
 - Production bugs
 - Started with monolithic gLite-3.1 build
 - Moving to node type driven build
 - *Not easy because of the software structure*
 - People active in moving gLite to ETICS : ~3

- **Why does SA3 not drive the simplification?**
- **We started “House Cleaning” of WNs and UIs**
 - Announced at the EMT
 - Wiki-Page <https://twiki.cern.ch/twiki/bin/view/EGEE/WnCleanup>
 - Setup November 15th
 - Current status:
 - *Some feedback from developers*
 - *But no feedback on: 139 RPMs*
 - *We lack resources to verify them on a trial and error basis*
 - To be fair, we didn't push very hard

- **JRA1 - SA3 can work together to restructure the way the stack can be decomposed**
 - Component based release affects:
 - Build
 - Dependencies
 - Interfaces between clients and services
 - *Example: Client may depend on server code*
 - Client, server and common code have to be cleanly separated
 - Clients should have absolutely minimal dependencies
 - *No “exotic” material permitted*
 - Avoid obscure 3rd party package dependencies
 - *Deadly for porting*
 - Static building might help deployment
 - *Doesn't address the build problem*
 - *Doesn't address portability problem*
- **How to move towards portable middleware?**
 - Maybe related to the above?
- **We need to plan this work....**
- **What is then dropped from the work plan**

- What is expected NA4?
- **Repository for middleware provided by VOs**
 - Many interesting user space and service components are available
 - Too many to certify on our testbeds
 - Lack of expertise of testers
 - Providers should be encouraged to use ETICS (2nd quarter 2007)
 - Building
 - *Multiple platforms*
 - Basic testing
 - *coexistence testing*
 - *Additional unit tests if provided by developers*
 - SourceForge instance + APT repository via ETICS (March 2007)
 - We have to discuss what the minimal requirements are

- **Introducing structural change is almost impossible**
 - Personnel already at full throttle with day to day work
 - This includes project overhead (communication)
 - Change has to be introduced while providing services
 - Bug fixes
 - Significant performance enhancements
 - Security fixes
 - Support (tickets from sites, roc)
- **More bugs have shown up**
 - 16 fold increase in usage of the infrastructure during 2006
- **SA3 resources split over 13 partners**
 - Difficult coordination and communication
 - Some partners don't have ▪ "critical mass"

- **Distributed Partners**
 - Some are very independent
 - What means certification of a component
 - *Without scripts and results archived*
 - Work is done, but impossible to asses
 - *Results are communicated highly summarized*
 - External Testbeds
 - Effort is going into it, but communication overhead too high for day to day updates

- **PreProduction Service needs a redefinition**
 - Experiments are not using it (see SL4 WNs)
 - Focus on deployment testing????
 - Merging with some SA3 testbeds?

- **TCG has turned into a “strange” meeting**
 - Work plans of JRA1 and SA3 are full for at least another 6 months
 - What short term planning can be done now?
 - Nothing is ever dropped...
 - EMT does the week to week juggling of priorities
 - Many significant changes are queued up
 - Data management, job priority, etc.
- **New requirements should reflect the experience gained with the new components**
 - Which can arrive after PPS ...