SA3 Report

Markus Schulz
For EGEE-II SA3
IT Department, CERN
Final EU Review of EGEE-II

www.eu-egee.org
www.glite.org

Information Society and Media

Outline

- **Activity Goals**
- **Status and Achievements**
  - Integration and Release Management
  - Testing
  - Multiplatform Support
- **Issues for SA3**
- **Future Plans**
- **Summary**

SA3 in Numbers

SA3 Partners

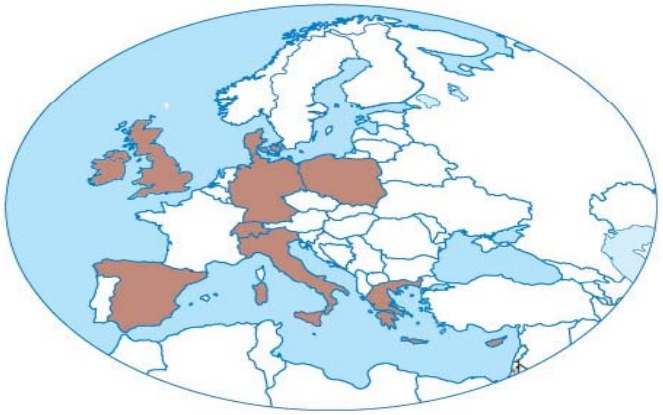

Manpower: 12 partners, 9 countries, 30 FTE

Short Name	Country	Total (PMs)
CERN	Switzerland	432
PSNC	Poland	36
TCD	Ireland	19
IMPERIAL	UK	24
INFN	Italy	60
UKBH	Denmark	12
UCY	Cyprus	34
GRNET	Greece	24
CSIC	Spain	12
PIC	Spain	24
CESGA	Spain	12
FZJ	Germany	36
Total		**725**

EGEE-II Budget

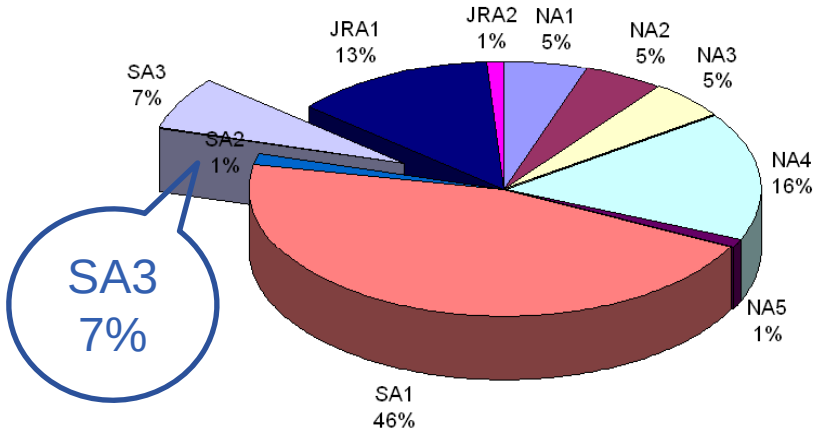

Activity Goals

- **Manage the process of building middleware distributions**
  - Integrating middleware components from a **variety** of sources
    - Based on TCG decisions
  - Define acceptance criteria for accepting components
    - Ensure:
      - ***reliability, robustness, scalability, security and usability***
  - Decouple middleware **distributions** from middleware **development**

Tasks

- Integration and Packaging
- Testing and Certification
  - ***Functional and Stress Testing***
  - ***Security, Vulnerability Testing***
  - ***Operate Certification and Testing Test Beds***
  - ***Project Testing Coordination***
- Debugging, Analysis, Support
- Interoperation
- Support for porting
- Participate in standardization efforts

Covered in other presentations

- **Interoperability:**
  - Proof of concept demonstrated for: NAREGI
  - Demonstrated interoperability with: UNICORE and ARC
  - First steps towards interoperation with: ARC
    - Accounting, monitoring, support
  - Continuous production use with: OSG
- **Standardization:**
  - GLUE-2
  - GIN-INFO
- **Software Metrics**

Integration and Release Management

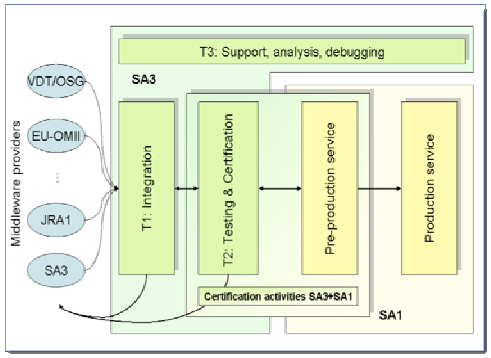

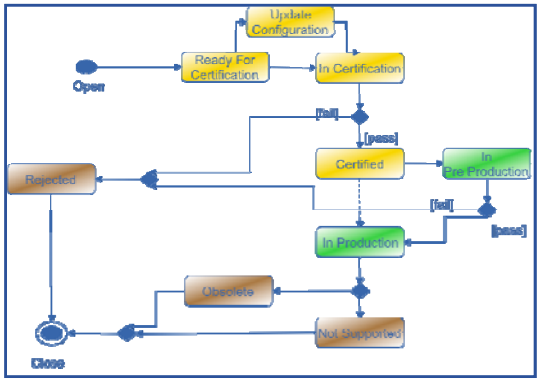

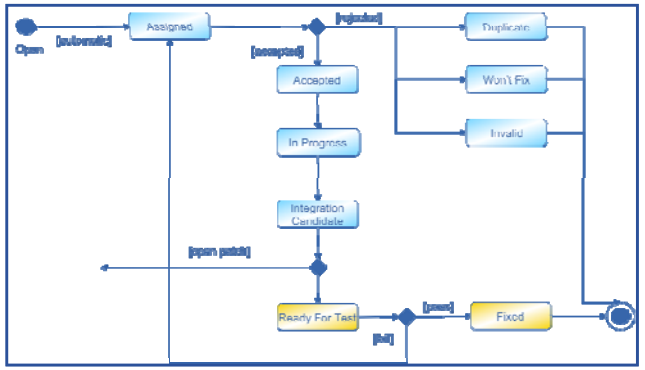

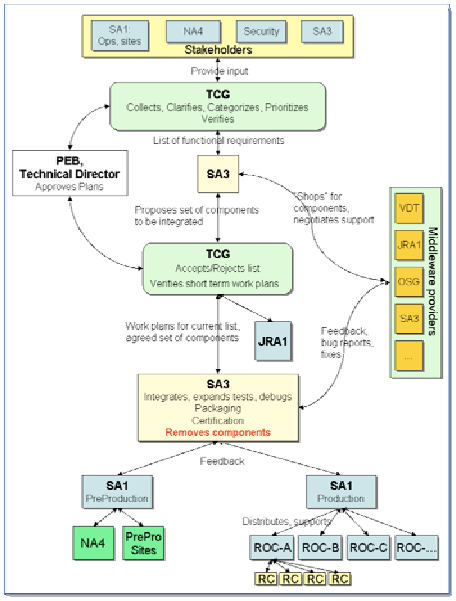

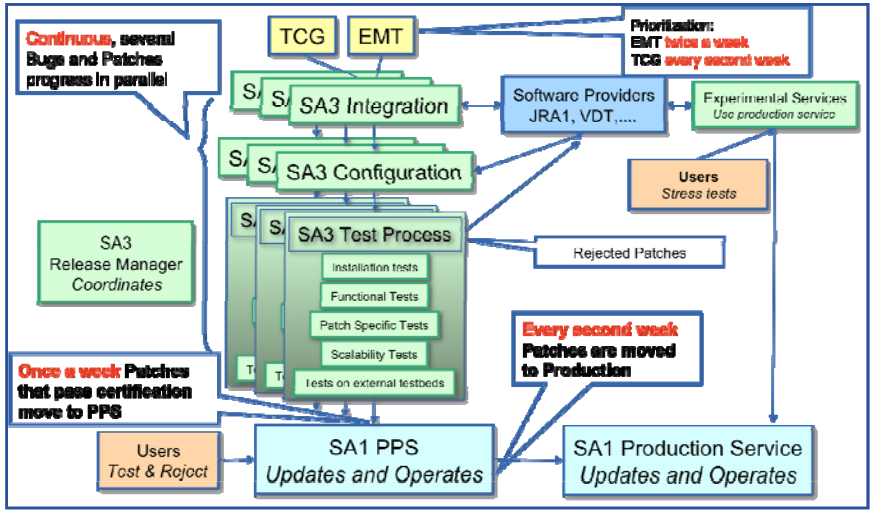

Link with SA1and JRA1

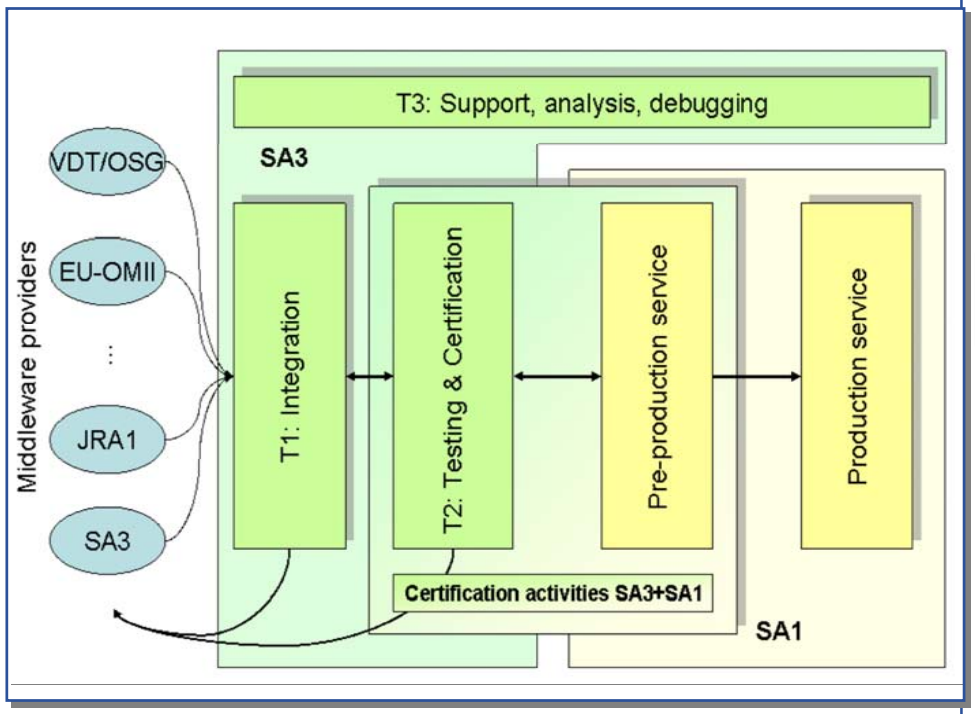

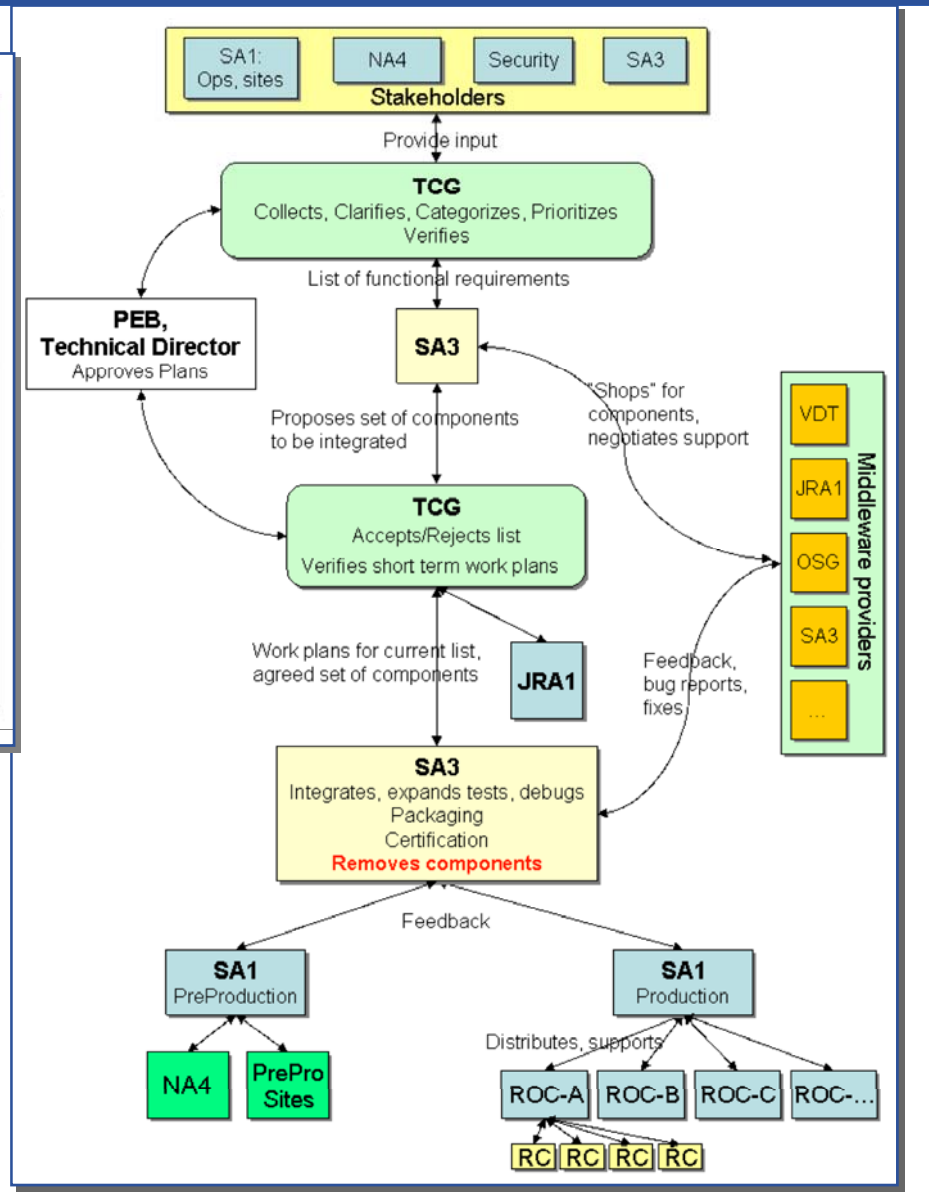

- Clear defined responsibilities

Process

- **Made full use of the software lifecycle process**
  - Documented in MSA3.2 and in use since July 2006
  - Components are updated independently
  - Updates are delivered on a ***weekly*** basis to the PPS
    - Move after **2 weeks** to production
  - Clear link between component ***versions***, ***Patches*** and ***Bugs***
    - Semi-automatic release note production
      - *Reducing the workload, improving the quality (one source)*
  - Clear prioritization by stakeholders
    - TCG for medium term (3-6 months) and EMT for short term goals
    - Clear definition of roles and responsibilities
- **Required only minor modifications in the second year**
  - One state was added
  - Several process monitoring tools were developed
  - More tasks were automated

Releases: gLite-3.0 gLite-3.1

- **gLite-3.0: Integrated release of LCG-2.7 and gLite-1.5**
  - Released on May 4th 2006
  - Phase out started (about 60 sites)
  - Has seen **49 updates**
    - A reflection of the dynamic evolution of the middleware

- **gLite-3.1: Based on VDT-1.6, Scientific Linux 4, ETICS**
  - Components have been released incrementally
  - New major versions for core components
    - WMS, LB, CE, FTS
  - All clients and several services released for 64bit
  - Component based, modular configuration tool (YAIM 4)
  - > 200 sites are running gLite-3.1

Usage

- **Process is in active use since July 2006**
  - Produced 26 sets of updates to the system in the first year
  - Second year:
    - Produced 23 sets of updates to gLite-3.0
    - Produced 17 sets of updates to gLite-3.1
  - Processed a total of **565** Patches
    - 361 for gLite-3.0, 204 for gLite-3.1
    - First year: **269** Patches
      - *Addressing 835 Change Requests*
  - During EGEE-II 3099 change requests have been opened
    - Increased usage and new use cases have uncovered more issues
    - 14% related to enhancements
    - 86% related to defects
    - Closed bugs: 1464 EGEE-II and 1002 EGEE-I

Process Monitors

- Several web based tools to track status
- Spot critical delays

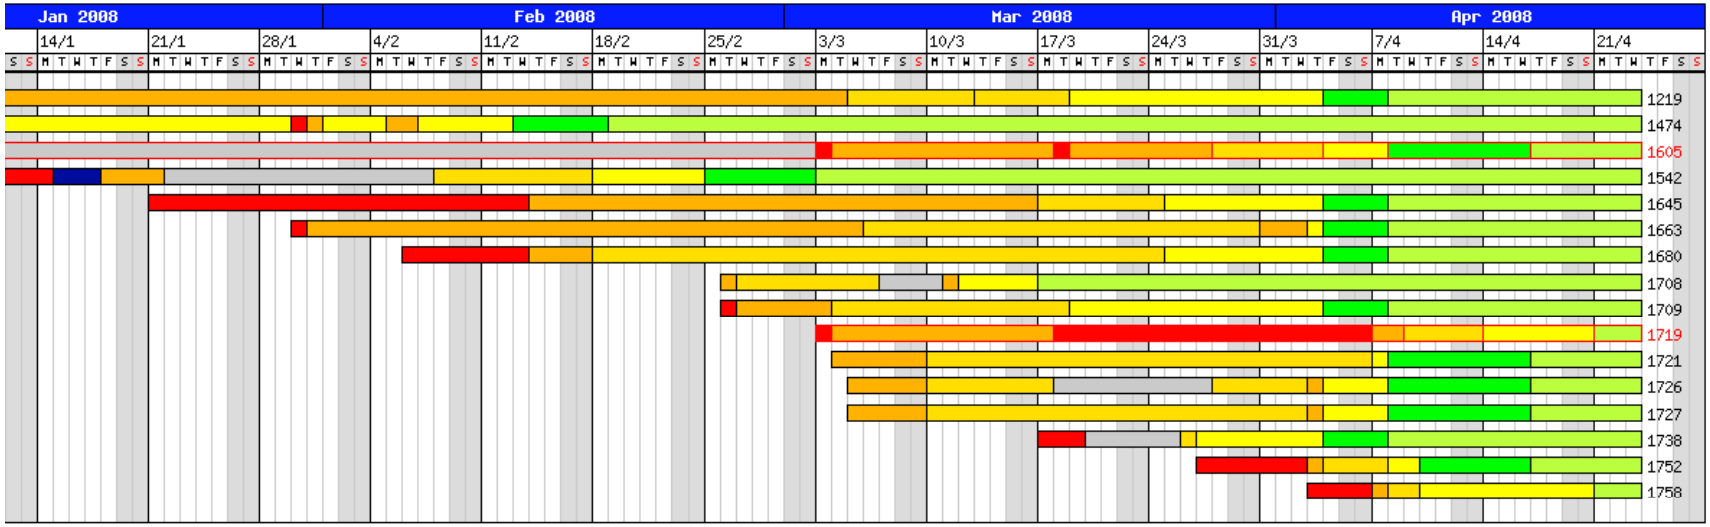

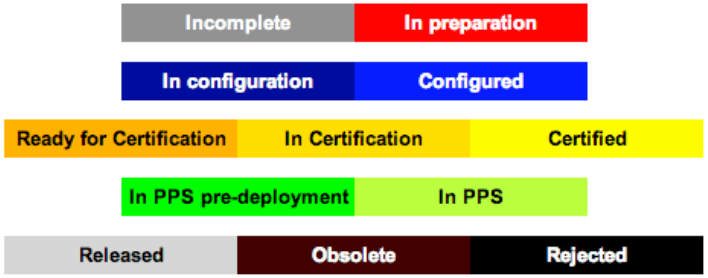

Process Monitors

- **Can create on demand complex reports**

EGEE-II: SA3 Activity Patch Watch

Enabling Grids for E-sciencE

Total:58

Queue Priority:High

#1760, 6days	#1594, 6days	#1850, 3days

Queue Priority:Normal

#1649, 146days	#1633, 122days	#1638, 122days	#1658, 122days	#1514, 122days	#1502, 122days	#1659, 122days
#1662, 120days	#1576, 76days	#1788, 43days	#1759, 43days	#1278, 43days	#1219, 43days	#1645, 43days
#1729, 43days	#1723, 43days	#1721, 43days	#1680, 43days	#1683, 43days	#1663, 43days	#1713, 42days
#1542, 41days	#1682, 41days	#1474, 41days	#1672, 41days	#1790, 31days	#1784, 27days	#1579, 24days
#1841, 23days	#1786, 16days	#1873, 15days	#1849, 13days	#1848, 13days	#1892, 13days	#1893, 13days
#1900, 10days	#1901, 10days	#1703, 8days	#1838, 8days	#1846, 8days	#1856, 8days	#1693, 7days
#1896, 7days	#1732, 3days	#1781, 3days	#1755, 3days	#1803, 2days	#1780, 2days	#1942, 2days
#1946, 2days	#1830, 1days	#1947, 1days	#1877, 1days	#1922, 1days	#1894, 1days	

Queue Priority:Low

Queue Priority:On Hold

W3C HTML 4.01 W3C CSS

EGEE-II is a project funded by the European Union under contract number INFSO-RI-031688

Patch status (Open patches)

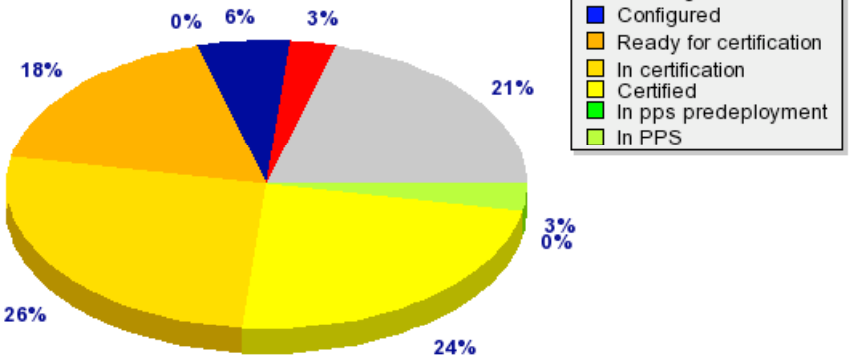

Patch processing times

Graph type: Distribution of total time for patch cycle

Patch priority: All priorities

gLite release series: All

Software area: All

Submit

Patch state distribution

Graph type: Distribution of open patch states (simplified)

Patch priority: All priorities

gLite release series: All

Software area: All

Submit

Patch Processing

- **Patch processing has seen strong partner participation**
  - Required advanced tools for progress tracking
  - Partners prefer to work on complex Patches
    - Reduced communication overhead
    - More flexible time management
    - Approximately **10%** have been handled outside CERN
    - Corresponds to about **20%** of the certification effort
- **To improve efficiency we developed tools that can directly access the DB of the tracking tool (Savannah)**
  - This is the basis for several automation efforts

Configuration Management

- **YAIM: Simplicity**
  - Key-Value pairs + bash
- **Popular with site administrators**
  - Result of a survey
  - Easy to integrate with local tools
  - Easy to modify

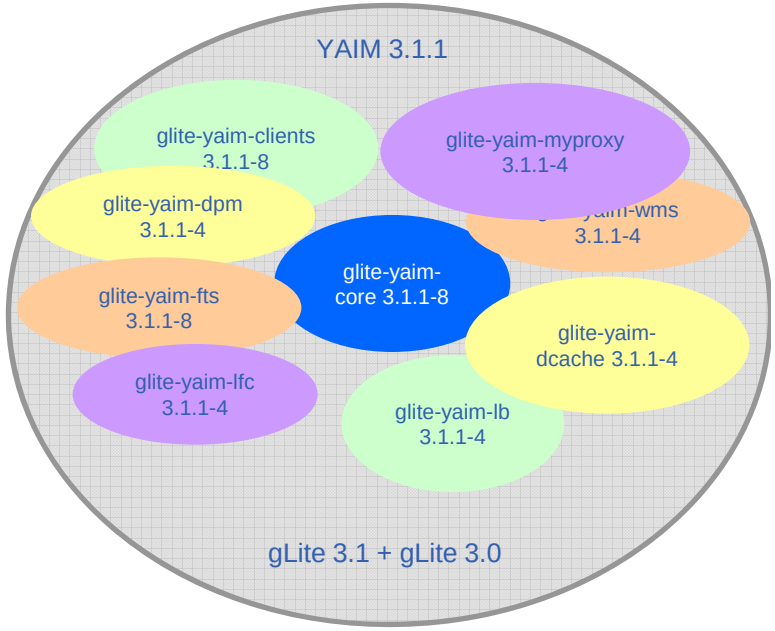

- **Moved all components to YAIM**
  - Initially monolithic architecture
  - Every configuration change required an update to all components

Configuration Management

- **YAIM 4**
  - Component based
    - Supports independent frequent releases of components
  - Allowed to distribute configuration effort
    - 25 contributors
    - Coordinated at CERN (quality control, testing)
  - Released October 2007
  - 33 modules released, 4 under development
- **Installation tool**
  - Started with APT for (semi) automatic RPM updates
    - Standard Debian tool, widely used
  - With SL4 we moved to YUM (comes with the release)
  - RPM lists for other tools
  - Tarballs for UIs and WNs

Build Systems

- **Started with 3 systems**
  - LCG, gLite, ETICS
  - Complicate dependency management, release management
- **Moved to 1**
- **ETICS**
  - Used for the gLite-3.1 branch
  - Migration process to ETICS started in early August 06
    - Finished for almost all components September 2007
    - Last component moved February 2008
  - Overall experience has been positive
    - Functionality and performance has improved significantly over time
    - Multiplatform build support was very helpful

Test strategy, framework

- **Test strategy:**
  - Test plans and process documented in MSA3.5
  - Multi level tests (from simple functional tests, to stress tests)
  - As much steps and components as possible are tested in parallel
- **SAM framework for automated testing**
  - Developed by SA1, sharing tests, customizable views and history

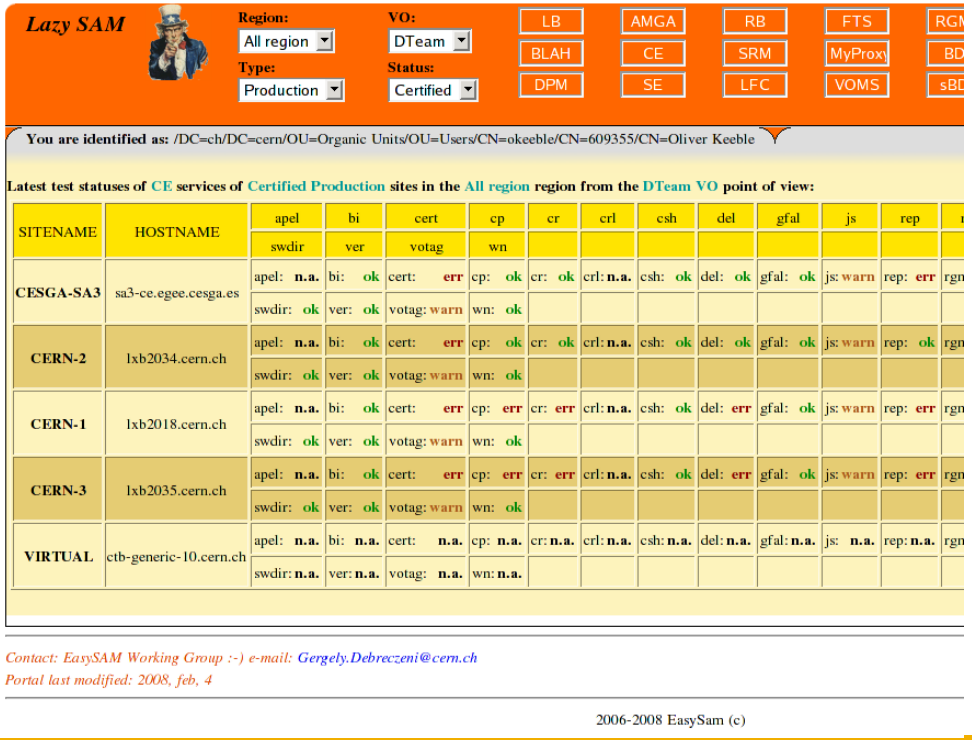

Testbeds

- **Central "Baseline Testbed" (> 50 nodes @CERN)**
- **Extended distributed test beds: 7 sites**
  - about 100 nodes to cover additional deployment scenarios
- **Virtualized test beds (>10 @CERN, each 1-5 nodes)**
  - Operation has been automated with the vNode tool
  - Main mode of testing, improved efficiency
- **Dedicated CE scalability test bed (> 25 nodes @CERN)**
- **Dynamical allocated test nodes (> 50 nodes @CERN)**
- **Use of "Experimental Services" (JRA1, SA1, SA3, NA4)**
  - Massive scalability tests can only be done in production
- **Standalone testbeds**
  - Posznan (Security), IMPERIAL (WMS), TCD (Porting)
- **Testbeds are expensive (hardware and humans)**

Test Beds

Usage pattern has changed over time.

Partners carry out more independent Patch certification on their sites

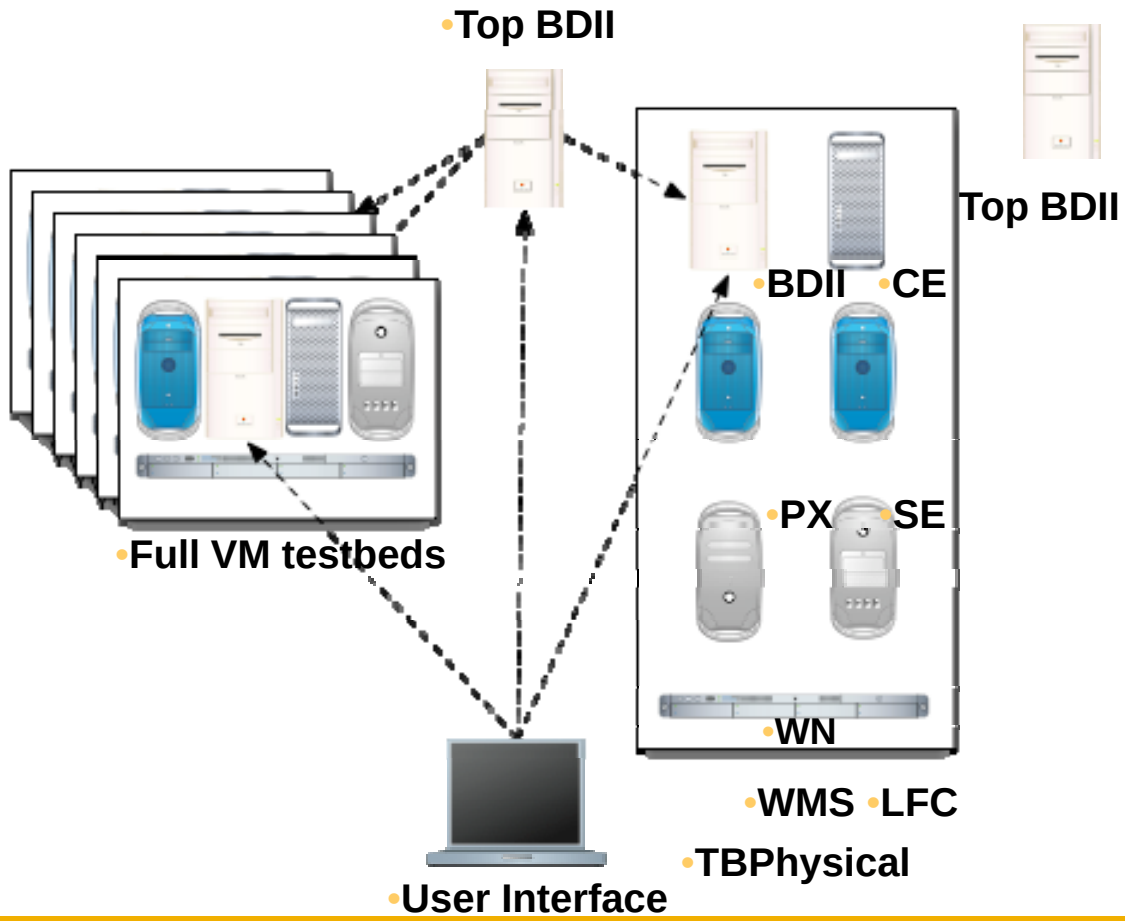

- Partners sites
- CESGA (SGE)
- PIC (Condor)
- GRNET (Torque)
- UCY (Torque)
- INFN (LSF)
- LAL (DPM,LFC)
- DESY (dcache)

Test Cases

- **Central repository for tests**
  - Contains more than 250 test cases
  - During the second year we almost doubled the number of tests
  - Most progress has been achieved for the following components:
    - Clients (many options, quite good coverage)
    - Data management tests: SRM, DPM, LFC, FTS
    - Stress tests: WMS/LB, CE
- **Test development is mainly done by partners**
  - Formal follow-up on test development
    - *Progress is monitored and documented every 2 weeks*
- **Many tests (about 30%) come from outside sources**
  - Volunteers, other projects,…

Test Cases

- **Security testing**
  - Done by Posznan
    - Code reviews (VOMS, R-GMA, DPM)
    - Penetration tests
    - Independent testbed
      - *Report to the Grid Security Vulnerability Group*
        - The GSVG classifies the vulnerabilities and does the followup
- **Interoperability tests**
  - For OSG within the scope of the PPS
- **Suitable tests for regression tests have been identified**
  - Integration into the ETICS framework started

Multi Platform Support

- **Main partners are Trinity College Dublin and Posznan**
- **Problems with porting**
  - Software dependencies and interdependencies
    - Execution of the “Plan for glite restructuring” improved the situation
  - Up to now mainly “post release” porting
    - Difficult to follow change rate
- **TCD moved to ETICS to close the gap**
  - Supports better concurrent multi platform build and tests
  - https://twiki.cern.ch/twiki/bin/view/EGEE/PortingWithEtics
- **Clients for several Linux versions are now available**

- Status table at TCD:
  - http://cagraidsvr06.cs.tcd.ie/autobuild

Builds using ETICS version: 1.3.6-1

Worker Node Build Status													
ARCH	OS TYPE	VERSION	DISTRO	torque	VDT	deps	GridIre	Basic	RGMA	VOMS	DM	gfal	WN-dev
ia32	CentOS	4	yum	3/3	0/1	30/30	2/2	12/12	41/41	13/13	17/17	21/23	107/109
	CentOS	5	yum	3/3	4/4	30/30	2/2	12/12	41/41	14/14	17/17	20/20	106/109
	Debian	4	debs	3/3	1/1	29/30	1/1	12/12	41/41	14/14	16/17	16/20	95/107
	Solaris	10	pkg/tarball	3/3	1/1	23/23	2/2	12/12	33/41	0/11	7/17	7/20	N/A
	SuSE	10	apt	3/3	4/4	30/30	1/1	12/12	41/41	13/13	17/17	18/20	N/A
x86_64	CentOS	4	yum	3/3	1/1	26/26	2/2	9/9	41/41	15/15	18/18	21/21	90/108
	CentOS	5	yum	3/3	4/4	24/30	N/A	N/A	N/A	N/A	N/A	N/A	N/A
	SuSE	10	apt	3/3	4/4	1/30	N/A	N/A	N/A	N/A	N/A	N/A	N/A
powerpc	Mac OS X	10.4	fink/tarball	3/3	1/1	1/30	1/1	12/12	33/41	0/11	0/17	11/29	57/109
	AIX	5	rpm/tarball	3/3	1/1	22/30	1/1	10/11	0/6	0/4	7/17	7/20	N/A
	Yellow Dog	6	yum	3/3	0/3	0/27	N/A	N/A	N/A	N/A	N/A	N/A	N/A

PSNC Build Results

Worker Node Build Status												
ARCH	OS TYPE	VERSION	DISTRO	torque	VDT	deps	Basic	RGMA	VOMS	DM	gfal	WN-dev
x86_64	Debian	4	debs	3/3	1/1	22/22	12/12	41/41	14/14	16/17	18/21	75/107

Obselete OS Build Results

Worker Node Build Status												
ARCH	OS TYPE	VERSION	DISTRO	torque	VDT	deps	GridIre	Basic	RGMA	VOMS	DM	WN-dev
ia32	SuSE	9	apt	3/3	3/3	28/28	2/2	12/12	41/41	14/14	17/17	107/107
x86_64	SLES	9	apt	3/3	1/1	24/24	1/1	9/9	37/41	10/15	12/18	74/109
powerpc	Mac OS X	10.3	fink/tarball	3/3	1/1	23/23	1/1	12/12	32/41	0/11	0/17	60/109

Legend	Colour	.	.	.	.
	Meaning	To be Started	Started	DONE	Not Applicable

Batch System Support

- **SA3 supports now:**
- **Torque/PBS -> reference platform**
  - LCG-CE, CREAM-CE
- **SGE**
  - LCG-CE, gLite-CE
- **Condor**
  - LCG-CE
- **LSF**
  - No direct support by a defined partner
  - LCG-CE, CREAM

Maintenance

- **SA3 ported LCG-CE to SL4**
  - Stop gap solution until CREAM-replaces the LCG-CE
- **SA3 improved the performance of the LCG-CE**
  - To cope with increased usage of the infrastructure
  - Speedup > 5 time

Issues: 2nd Year

- **Change management**
  - Move to SL4, VDT-1.6, globus-4
  - Move to ETICS
  - Many transitions in the infrastructure
  - While keeping changes flowing to production
- **Patch tracking reveals that SA3 can't handle the change rate**
  - Many Patches end in "Obsolete" state
  - We coped better than last year
    - Improved tools
    - Automation
    - Highly trained staff
  - Increased Patch latency

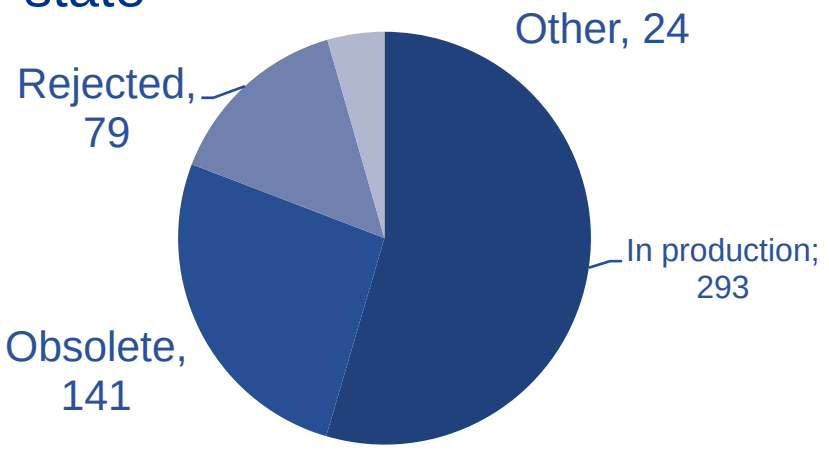

Issues

- **Testing**
  - Depends still too much on the central team
  - For complex services testers require significant training
    - Certifiers train Certifiers…. (NA3 is not involved)
    - Specialization → can result in patches being queued
  - We work towards more complete automation
    - Automation comes at a cost
    - Automation can't replace in depth understanding of the service

Issues

- **Multiplatform support**
  - Still suffers from complex dependencies

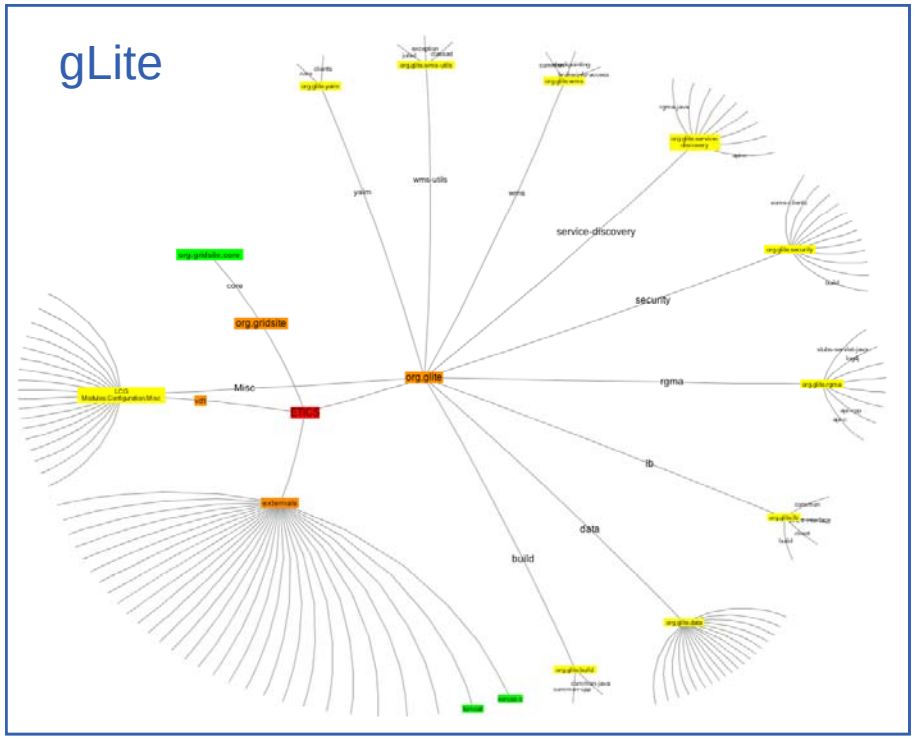

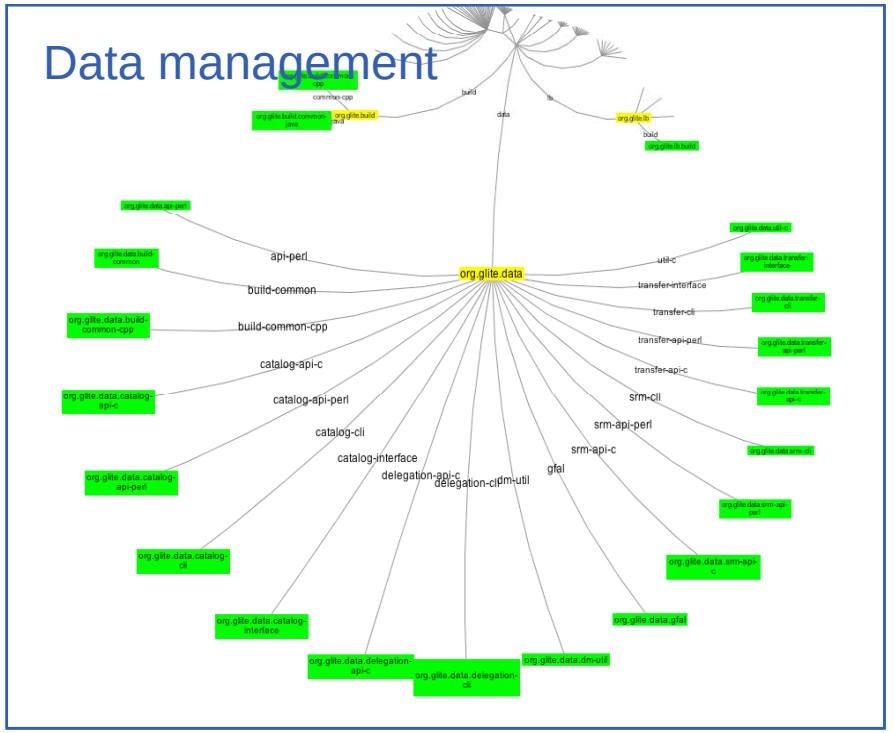

Plans

- **Automate more aspects of the process**
  - Testing
    - Regression tests, deployment tests (ETICS)
  - Patch handling
- **Distributed Patch processing**
  - Use experience of partners to increase throughput
- **Improve the process**
  - Patch iterations (adapt the process to reality)
  - Transition: development → certification
  - Transition: certification → Pre Production Service → Production
  - Goal: Reduced Patch latency
- **Alternative distribution of clients**
  - “push” multiple versions for user preview

Plans

- **Support at least 2 additional platforms for all releases**
  - To be defined by TCG (now TMB)
  - Can be restricted to some components (UIs, WN)

Summary

- **SA3 worked well as an activity**
- **We have a working Software Life Cycle process**
  - Component based updates work!
  - Very flexible, modular configuration tool, YAIM-4
- **Test process defined and implemented**
  - Many additional tests
  - Common framework with SA1 (SAM)
  - External testbeds to cover deployment scenarios
  - Virtualized testbeds improved efficiency (key technology)
- **Move to gLite-3.1 has been completed**
  - Uniform build system (ETICS)
- **Multiplatform support is now better understood**
  - Significant progress during the last year

Summary

- **Interoperability**
  - OSG is in production
  - ARC close to production
  - UNICORE demonstrated basic functionality
  - NAREGI demonstrated core functionality
    - Job level and data