SA3 Report

Markus Schulz
For EGEE-II SA3
IT Department, CERN
2nd EU Review of EGEE-II

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

Information Society and Media

Outline

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

SA3 in Numbers

SA3 Partners

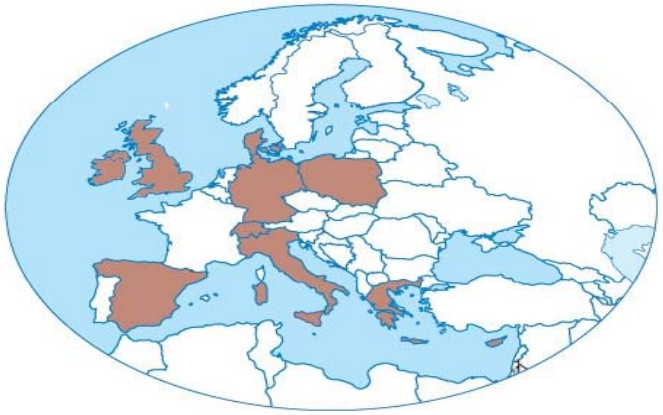

Manpower: 12 partners, 9 countries, 30 FTE

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

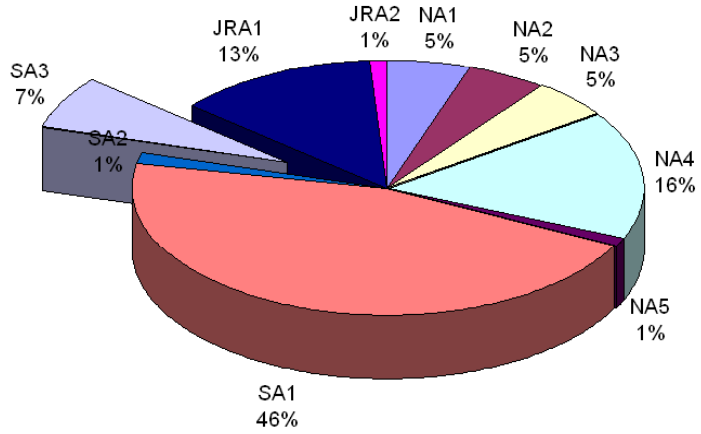

Activity Goals

- **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
- Capture requirements

Achievements: gLite-3.0 gLite-3.1

- **gLite-3.0: Integrated release of LCG-2.7 and gLite-1.5**
  - Different
    - build systems, configuration management, overlapping functionality
  - Different process……
    - LCG-2 tailored to production, gLite process tailored to development
- **Released on May 4th 2006**
  - 4 days later than planned 5 months before
- **gLite-3.1: Based on VDT-1.6, Scientific Linux 4, ETICS**
  - Using the new process components have been released incrementally
  - New major versions for major components
    - WMS, LB and CE
  - All clients and several services released for 64bit
  - Component based, modular configuration tool (YAIM 4)

Achievements: Process

- **Introduced new software lifecycle process**
  - Based on the gLite process and LCG-2 experience
    - 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
  - 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 has been added
  - Several process monitoring tools have been developed

Achievements: Testing

- **Test strategy, process,framework and external testbeds**
  - SAM framework for automated testing (SA1 product)
  - Central repository for tests
  - Formal follow-up on test development
  - Increased number of test cases ▪
    - Development of tests mostly by partners
  - Formal process for Patch certification
  - Extended test beds: 8 sites
    - about 100 nodes to cover additional deployment scenarios
  - Extensive use of *virtualized* test beds
    - Main mode of testing, significantly improved efficiency
  - Use of "Experimental Services"
    - Massive scalability tests can't be conducted on test infrastructures
  - Dedicated scalability testbeds for CEs

Achievements: Interoperability

- **For details see dedicated presentation**

- **Proof of concept demonstrated for: NAREGI**
- **Demonstrated interoperability with: UNICORE and ARC**
- **First steps towards interoperation with: ARC**
  - Pilot VO
  - Accounting, monitoring, support

- **Continuous production use with: OSG**
  - Added a interoperability testbed within the PPS

Achievements: Multiplatform Support

- **Based on ETICS for multi platform build support**
- **gLite clients for more platforms are now available with a short delay after new releases appear**
  - Still covering mainly different Linux distributions

Status

- **Integration and Release Management**

Handling Bugs and Patches (simplified)

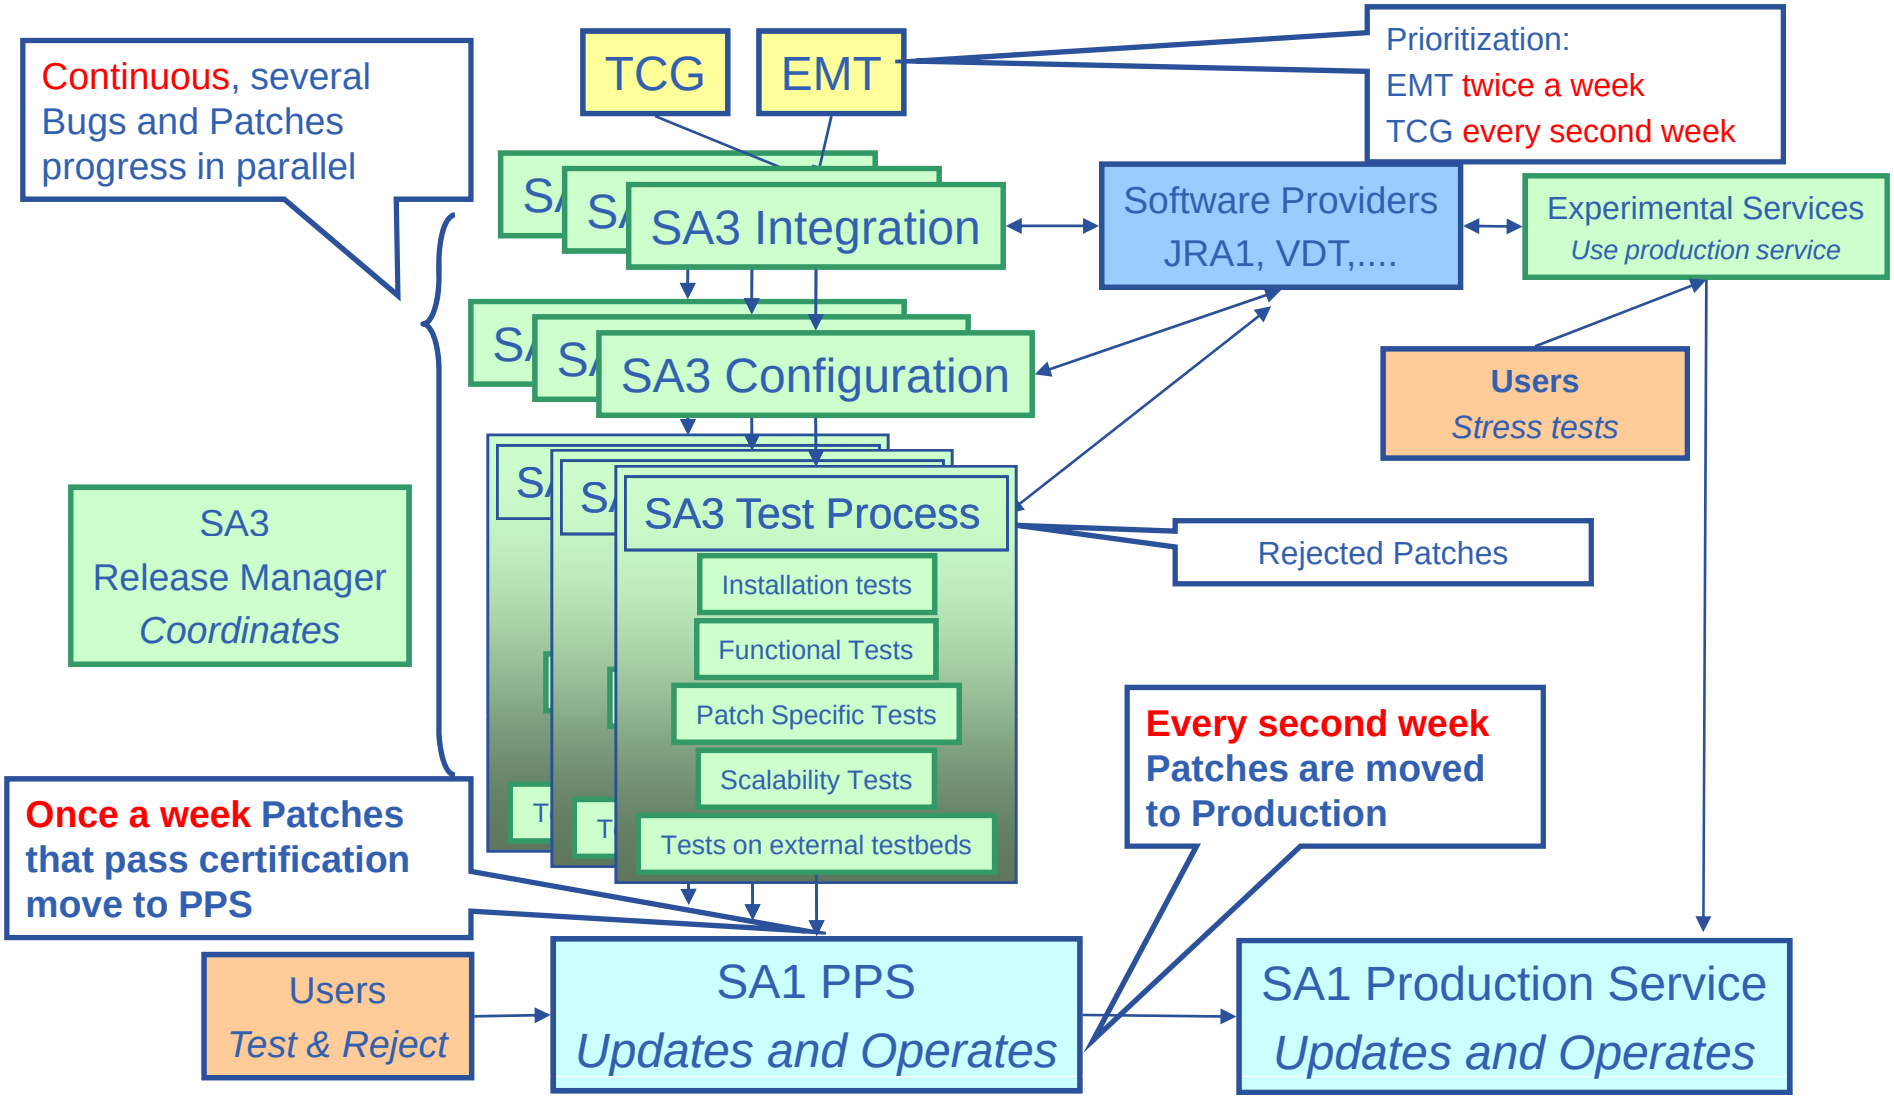

Usage

- **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 Bugs*
  - During EGEE-II 3099 bugs have been opened
    - 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**

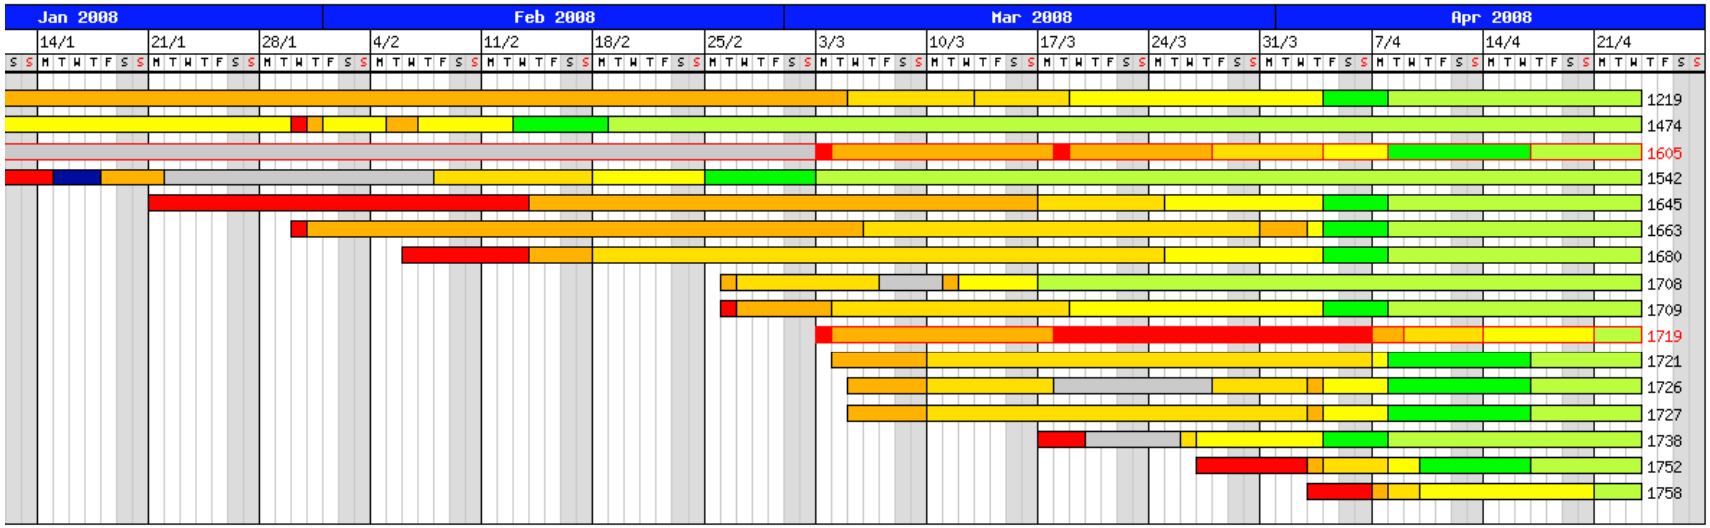

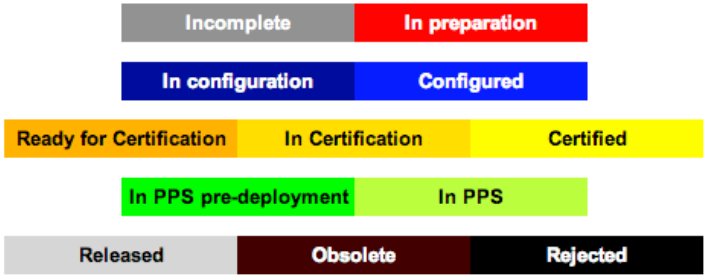

Process Monitors

- **Several web based tools to track status**

EGEE-II: SA3 Activity
Patch Watch

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

HTML 4.01 CSS

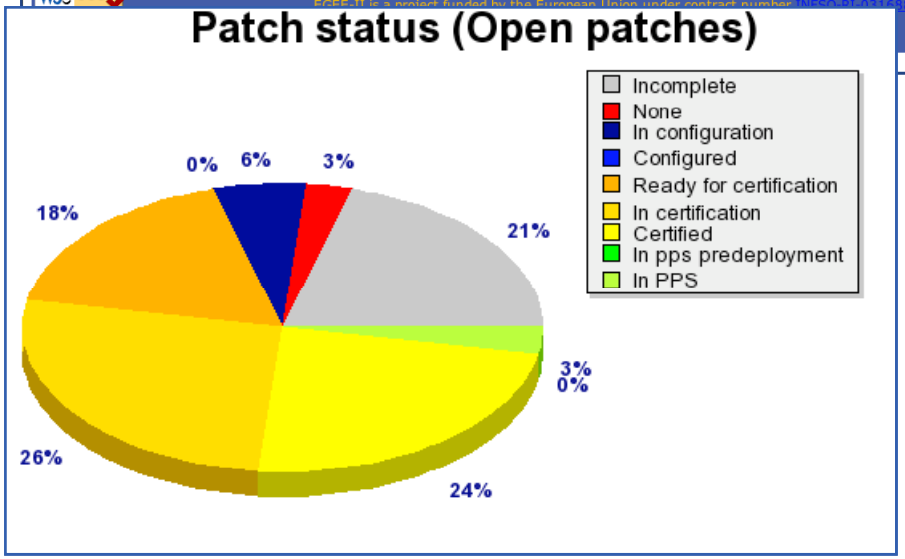

Patch processing times

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

Process

- **Processing patches moved progressively to partners**
  - Required improved tools for progress tracking
  - Partners tend to work on complex Patches that require some time
    - Approximately 10% of the patches have been handled outside CERN
    - Corresponds to about 20% of the certification effort

- **To improve efficiency we developed tools that can directly access to DB of the tracking tool (Savannah)**

Configuration Management

- **gLite-1 configuration: XML and Python**
- **LCG-2 configuration: Key-Value pairs + bash (YAIM)**
- **Site administrators preferred YAIM (result of survey)**
- **Moved all components to YAIM**
  - Initially monolithic architecture
  - Every configuration change required an update to all components

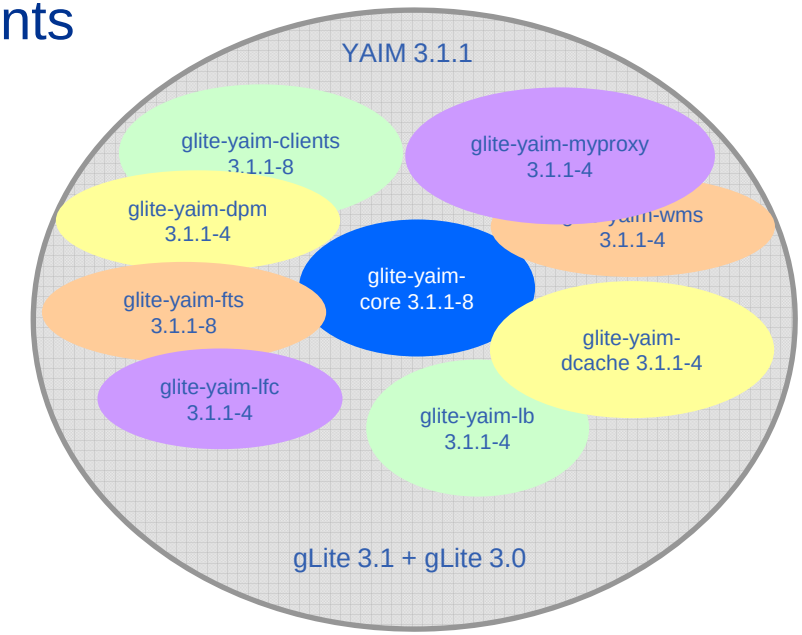

Configuration Management

- **YAIM 4**
  - Component based to ease independent frequent releases
  - 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
  - Complicates dependency management, release management

- **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 2007
  - Overall experience has been positive
    - Initial release lacked maturity
    - Functionality and performance has improved significantly over time
    - Multiplatform build support was very helpful

Testing

- **Test plans and process documented in MSA3.5**
- **Test strategy**
  - Multi level tests (from simple functional tests, to stress tests)
    - To abort as early as possible
  - As much steps in parallel as possible
    - Component based
      - *Install, configure, functional tests, first patch certification*
    - Requires many temporary testbeds
      - *We use virtualization (Xen based) to save time and resources*
    - We use the locally developed Vnode management system
  - First local then external testbeds
  - Testing relies now on a “Baseline Release” testbed
    - Required significant reorganization of the testbed operation

Testing Framework

- **We have chosen SAM as our framework for testing**
  - Maintained and used by SA1, sharing tests
  - Provides Web based, customizable views and history

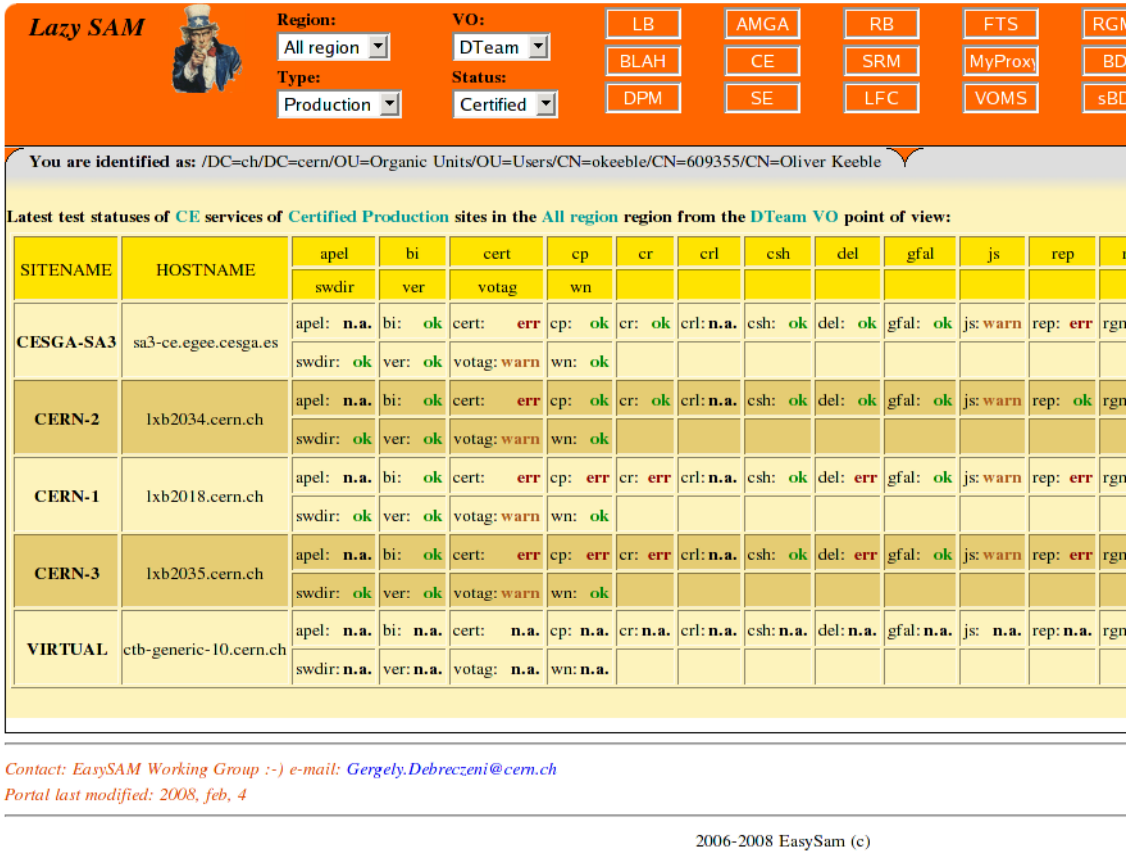

Lazy SAM

Region: All region | VO: DTeam | Type: Production | Status: Certified

LB | AMGA | RB | FTS | RGM
BLAH | CE | SRM | MyProxy | BD
DPM | SE | LFC | VOMS | sBD

You are identified as: /DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=okeeble/CN=609355/CN=Oliver Keeble

Latest test statuses of CE services of Certified Production sites in the All region region from the DTeam VO point of view:

SITENAME	HOSTNAME	apel	bi	cert	cp	cr	crl	csh	del	gfal	js	rep	
		swdir	ver	votag	wn								
CESGA-SA3	sa3-ce.egee.cesga.es	apel: n.a.	bi: ok	cert: err	cp: ok	cr: ok	crl: n.a.	csh: ok	del: ok	gfal: ok	js: warn	rep: err	rgn
		swdir: ok	ver: ok	votag: warn	wn: ok								
CERN-2	lxb2034.cern.ch	apel: n.a.	bi: ok	cert: err	cp: ok	cr: ok	crl: n.a.	csh: ok	del: ok	gfal: ok	js: warn	rep: ok	rgn
		swdir: ok	ver: ok	votag: warn	wn: ok								
CERN-1	lxb2018.cern.ch	apel: n.a.	bi: ok	cert: err	cp: err	cr: err	crl: n.a.	csh: ok	del: err	gfal: ok	js: warn	rep: err	rgn
		swdir: ok	ver: ok	votag: warn	wn: ok								
CERN-3	lxb2035.cern.ch	apel: n.a.	bi: ok	cert: err	cp: err	cr: err	crl: n.a.	csh: ok	del: err	gfal: ok	js: warn	rep: err	rgn
		swdir: ok	ver: ok	votag: warn	wn: ok								
VIRTUAL	ctb-generic-10.cern.ch	apel: n.a.	bi: n.a.	cert: n.a.	cp: n.a.	cr: n.a.	crl: n.a.	csh: n.a.	del: n.a.	gfal: n.a.	js: n.a.	rep: n.a.	rgn
		swdir: n.a.	ver: n.a.	votag: n.a.	wn: n.a.								

Contact: EasySAM Working Group :-) e-mail: Gergely.Debreczeni@cern.ch
Portal last modified: 2008, feb, 4

2006-2008 EasySam (c)

Test Status

- **Test development mainly by partners**
  - Progress is monitored and documented every 2 weeks
- **Many tests from external sources**
  - Volunteers, other projects
- **Security testing**
  - Done by Posznan
    - Code reviews (VOMS, R-GMA, DPM), penetration tests
    - Independent testbed
      - *Report to the grid vulnerability group*
- **Interoperability tests**
  - For OSG within the scope of the PPS

Test Status

- **Number of test cases available now >250**
  - During the second year we almost doubled the number of tests
- **Most progress has been achieved for the following components:**
  - Client tests
    - Many options
  - Data management tests
    - SRM
    - DPM
    - LFC
    - FTS
  - Stress tests for:
    - WMS/LB
    - CE
- **Suitable tests for regression tests have been identified**
  - Integration into the ETICS framework started

Test Beds

- Virtual testbeds for individual testers (> 10)
- Dynamical allocated test nodes (> 50 nodes)
- Central certification testbed (> 50 nodes)

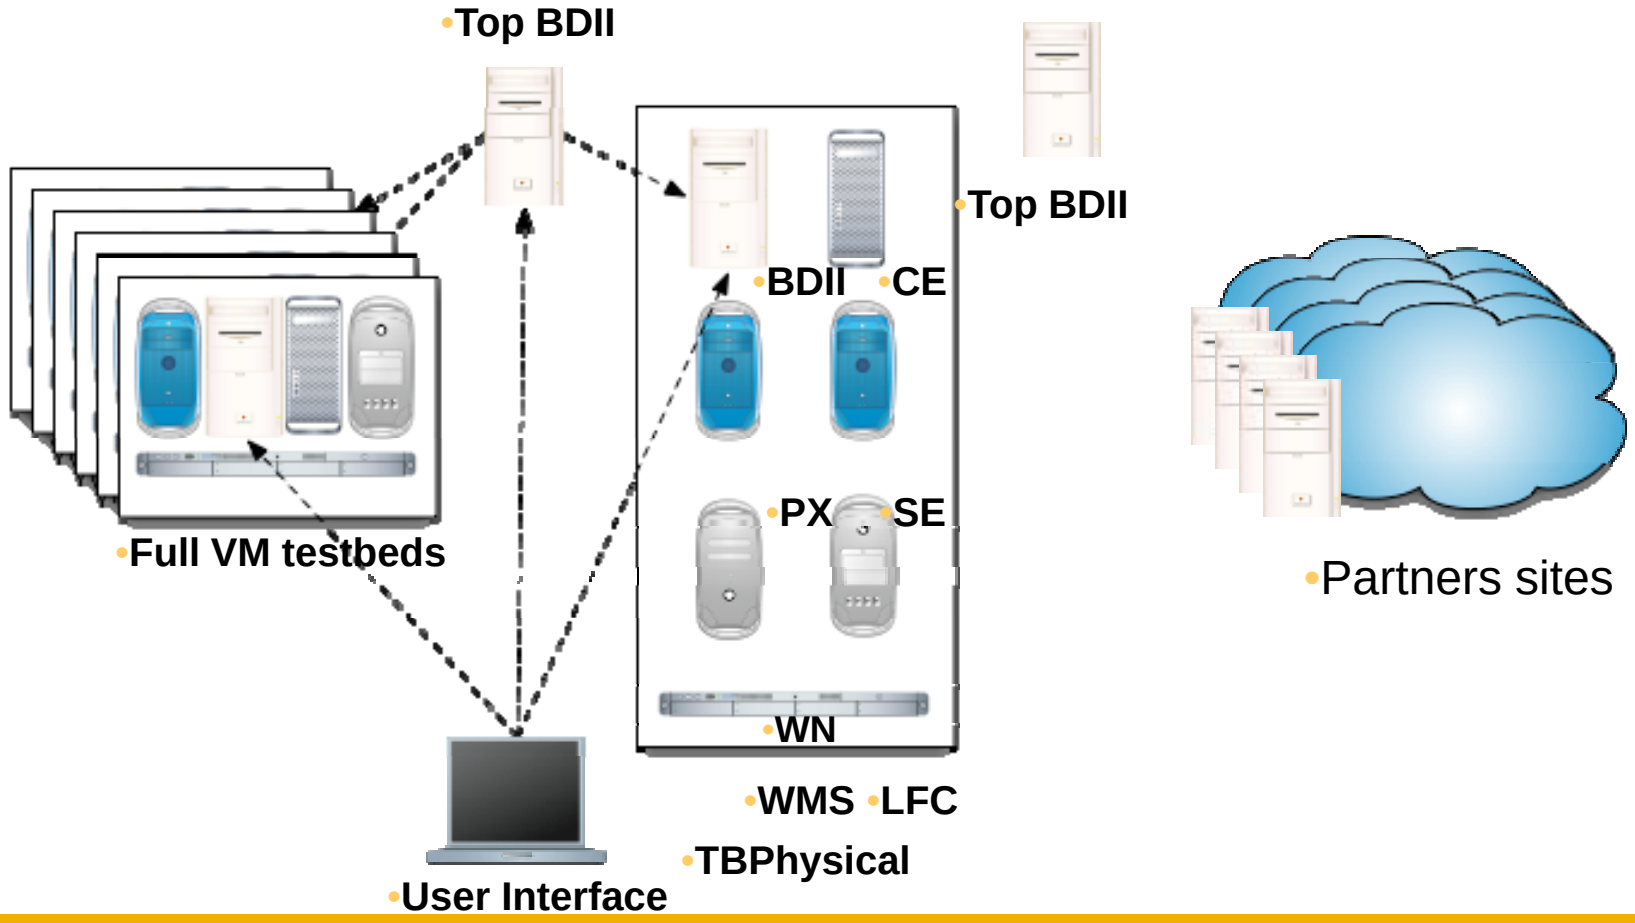

Test Beds

- **External testbeds linked to the certification testbed**
  - CESGA (SGE)
  - PIC (Condor)
  - GRNET (Torque)
  - UCY (Torque)
  - INFN (LSF)
  - LAL (DPM,LFC)
  - DESY (dcache)

Usage pattern has changed over time.

Partners carry out more independent Patch certification on their sites

- **Standalone testbeds**
  - Posznan (Security)
  - IMPERIAL (WMS)
  - TCD (Porting)

Interoperability

- **See dedicated presentation**

Standardization

- **Covered in the interoperability presentation**

Porting

- **Main partners are TCD and Posznan**

- **Problems with porting**
  - Software dependencies and interdependencies
    - Execution of the “Plan for glite restructuring” improved the situation
    - ETICS support for multiplatform build made the process more efficient
  - Up to now mainly “post release” porting
    - Difficult to follow change rate

- **TCD is moving to ETICS to close the gap**
  - Supports better concurrent multi platform build and tests
  - https://twiki.cern.ch/twiki/bin/view/EGEE/PortingWithEtics

- 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

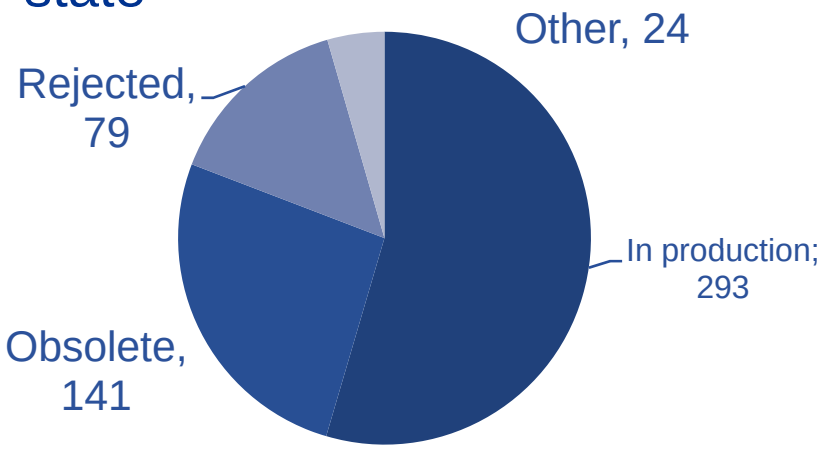

Issues

- **Testing**
  - Depends still too much on central team
  - For complex services tester require significant training
    - 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

•gLite

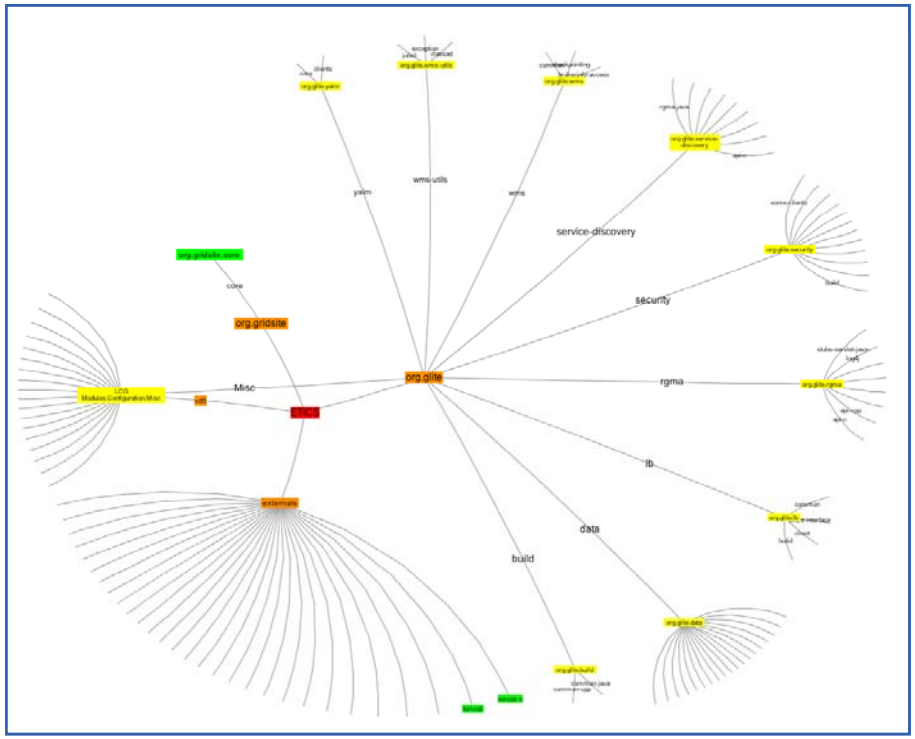

•Data management

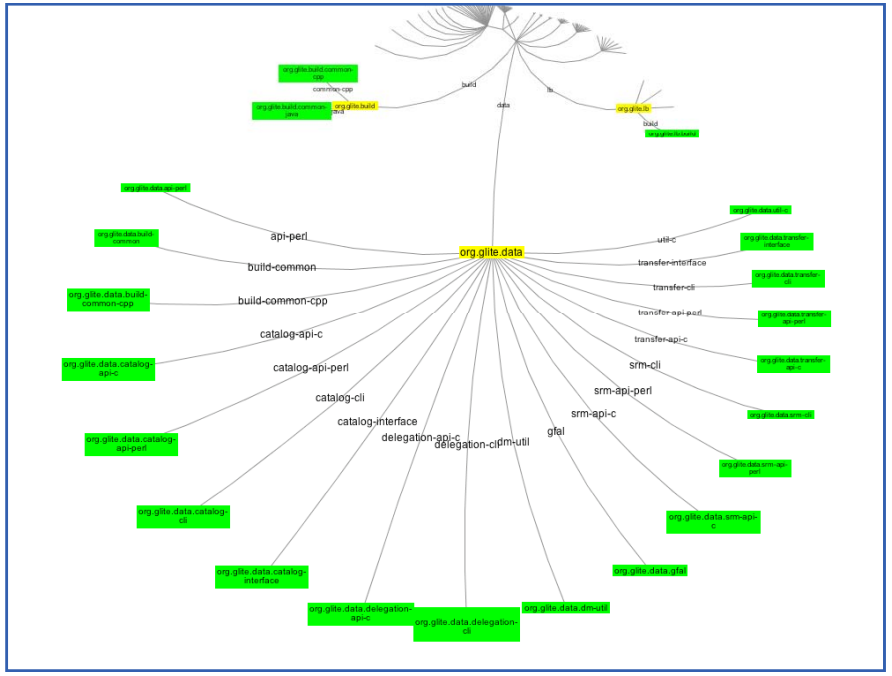

Plans

- **Automate more aspects of the process**
  - Testing
    - Regression tests, deployment tests
  - Patch handling
    - Ease the workload of the developers and integrators
      - *Tools for patch handling*
- **Distributed Patch processing**
  - Use experience of partners to increase throughput
- **Improve the process**
  - Patch iterations
  - Transition from certification to PPS to 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
- **Move to gLite-3.1 has been completed**
  - Uniform build system (ETICS)
- **Multiplatform support is now well 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