

SA3's Response to “Recommendation 40”

Markus Schulz
CERN IT

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

- **The reviewers gave one recommendation to SA3**
 - Implicit recommendations deal with:
 - Multiplatform Support
 - Interoperability
 - These will be addressed as part of the status
- **“40. SA3 and JRA1. Change the classification of bugs within the bugs list and separate true bugs (i.e. defects in the code) from feature change requests.”**
- **Reaction:**
 - SA3 introduced mandatory severity field in the bug tracker
 - Total number of bugs opened since introduction (June 21th):
 - 662 Defects 90%
 - 74 Enhancements 10%



Enabling Grids for E-scienceE

SA3 Status and Plans

Markus Schulz

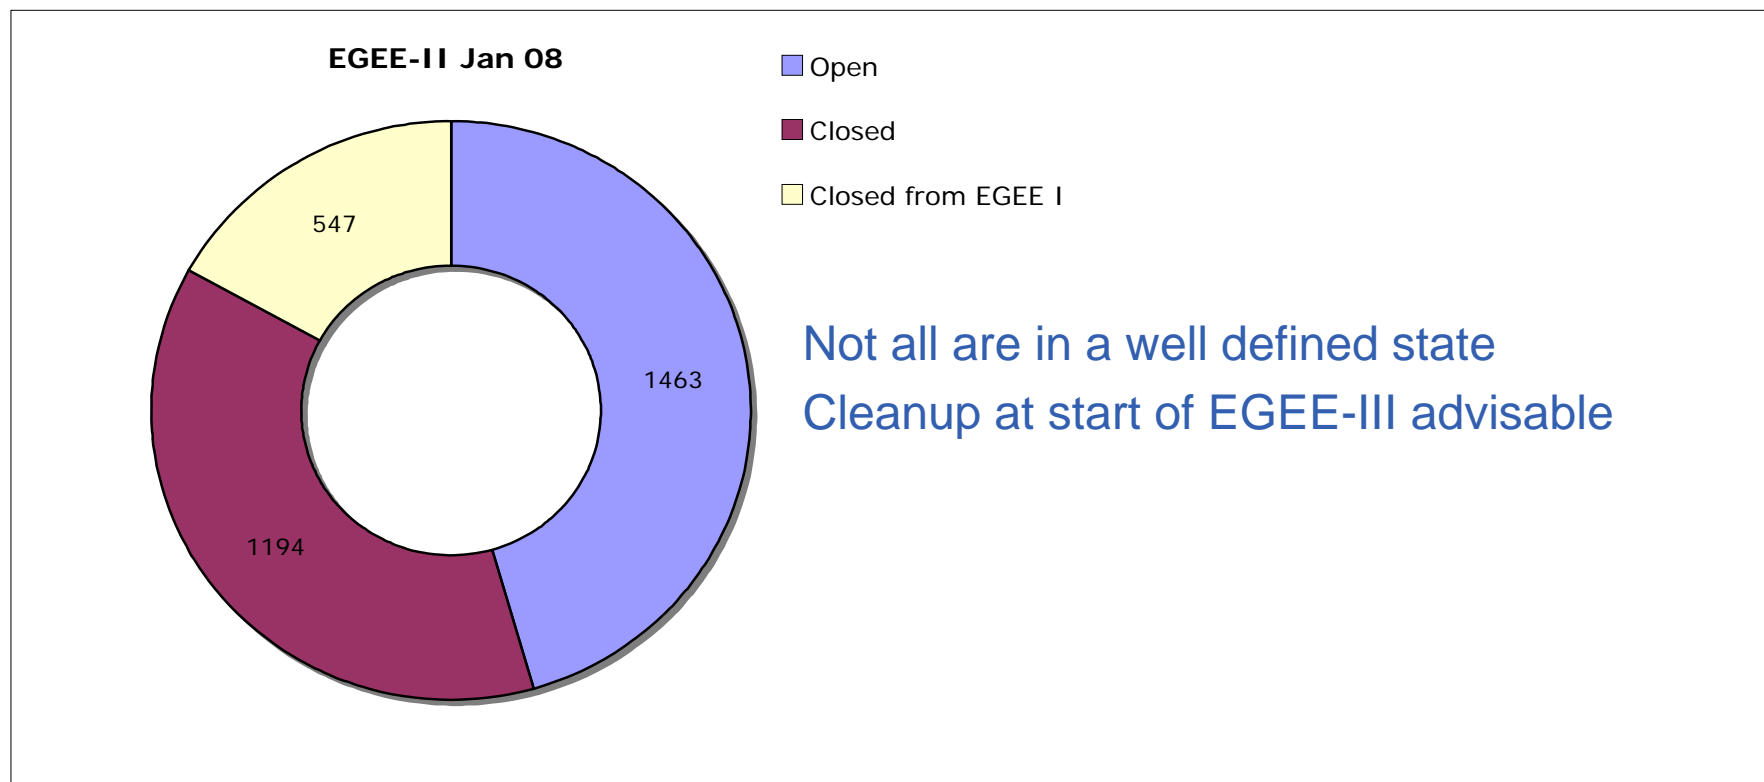
CERN IT

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

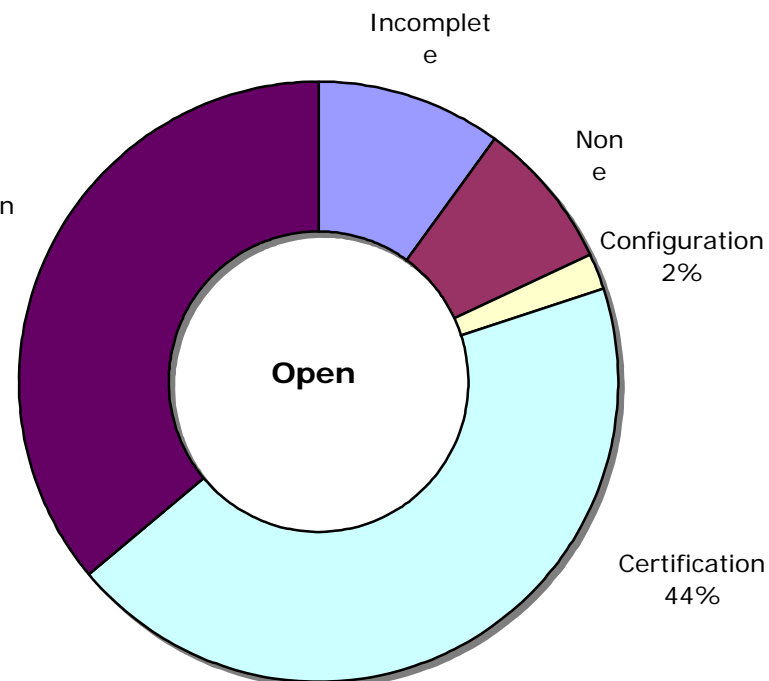
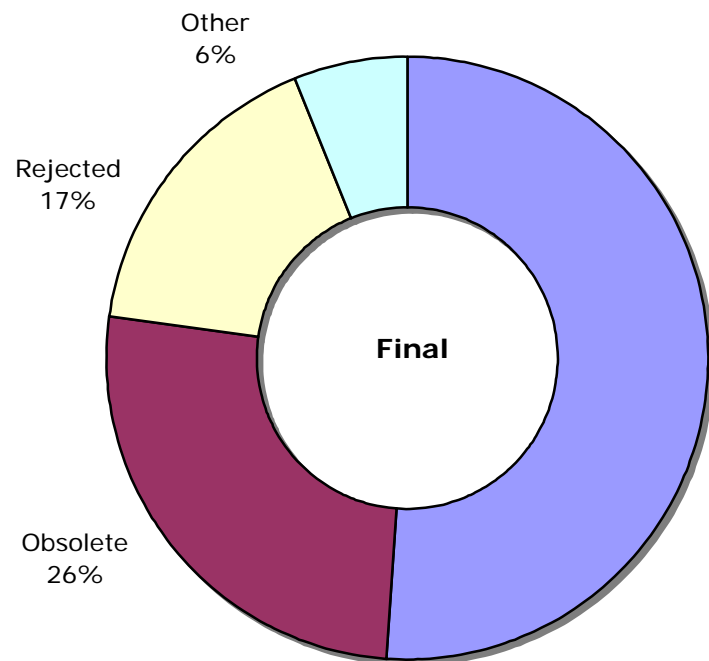


- **Focus on a few topics**
 - Bugs and patch statistics
 - High level view
 - *Detailed analysis in the “Process” part*
 - Status of move to SL4/VDT 1.6
 - Status of multi platform support
 - Status of batch system support
 - Status of Interoperability
 - Plans for the last 3 months of the project
 - Some issues
- **Nothing about YAIM (yaim is doing fine)**

- Bugs come rolling in: 182 since November
- We close them faster: 221 closed since November
- Bugs in EGEE-II: 2657
- Open bugs from EGEE-I: 771 (547 closed in EGEE II)



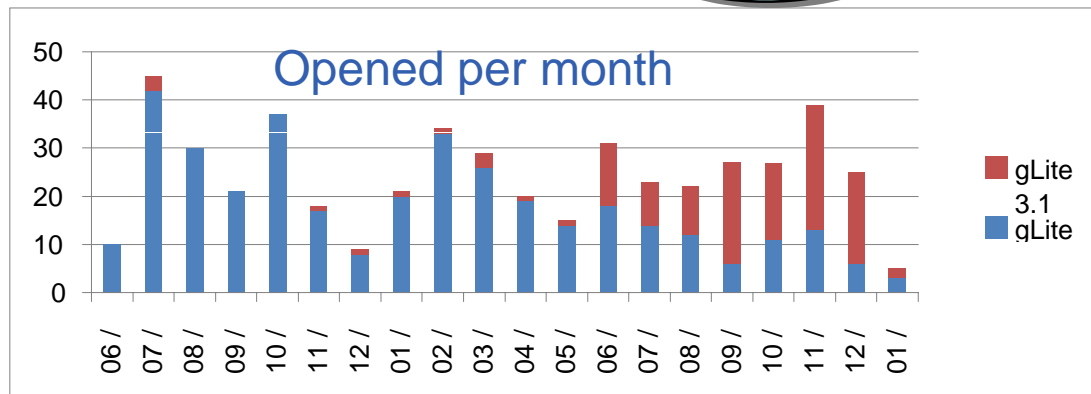
- Total: 482 Closed: 432 Open: 50



Outside CERN:

Certified: 28

In progress: 8



Experimental
Services
@
INFN

service	SL4/SLC4/I386	SL4/SLC4/x86_64
glite-WN	Released	Certification
glite-UI	Released	Build
glite-AMGA_postgres	PPS	Integration
glite-BDII	Released	Configuration
lcg-CE	Released	Build
glite-CREAM	Build	Build
glite-FTA_oracle	Integration	Build
glite-FTS_oracle	Integration	Integration
glite-FTM	Released	Integration
glite-LB	Certification	Build
glite-LFC_mysql	Released	Certification
glite-LFC_oracle	Released	Certification
glite-MON	Certification	Build
glite-PX	PPS	Integration
glite-SE_classic	Certification	Build
glite-SE_dcache_*	PPS	Integration
glite-SE_dpm_disk	Released	Certification
glite-SE_dpm_mysql	Released	Certification
glite-SE_dpm_oracle	Integration	Integration
glite-TORQUE_utils	Released	Build
glite-TORQUE_client	Released	Build
glite-TORQUE_server	Released	Build
glite-VOMS_oracle	Released	Integration
glite-VOMS_mysql	Released	Integration
glite-VOBOX	PPS	Integration
glite-WMS	Certification	Build

- We are about a year late
- Critical Components
 - CREAM-CE
 - Expected within a month
 - WMS/LB
 - Currently in pre-certification at Imperial
 - Might be ready for certification within a week

- **Focus on most critical node types**
 - WN and UI
- **Next platform: Debian 4**
 - Regular builds in ETICS
- **SL5 port should start soon to be ready in time**
- **Some progress related to building during the last 4 months**
 - ETICS matured and starts to pay off
 - Restructuring of gLite improved the situation
- **Testing of other platforms didn't start**
- **TCD produced a 10 step porting guide:**
 - Covering building of:
 - ETICS, globus, dependencies, Base Modules, VOMS, RGMA, DM/GFAL and the WN meta-package

<https://twiki.cern.ch/twiki/bin/view/EGEE/EGEEPortingGuide>

Builds using ETICS version: 1.2.1-1

Worker Node Build Status														
ARCH	OS TYPE	VERSION	DISTRO	VDT	deps	GridIre	Basic	VOMS	RGMA	DM	gfal	WN	Full	torque
ia32	CentOS	4	yum	3/3	30/30	2/2	9/9	14/14	45/45	17/17	20/20	108/108	344/374	3/3
	SuSE	9	apt	3/3	28/28	2/2	14/14	15/15	45/45	17/17	20/20	106/106	312/365	3/3
	CentOS	5	yum	1/1	24/24	2/2	9/9	10/15	24/45	0/17	0/19	63/99	226/368	3/3
x86_64	Solaris	10	pkg/tarball	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
	CentOS	4	yum	1/1	5/16	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
powerpc	SuSE	9	apt	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
	Mac OS X	10.4	fink/tarball	2/2	17/23	N/A	8/9	1/15	36/46	0/17	11/29	42/61	182/361	N/A
	Mac OS X	10.3	fink/tarball	2/2	18/23	N/A	8/9	1/15	36/46	0/17	11/29	44/62	186/354	N/A
	AIX	5	rpm/tarball	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A

Legend	Colour	Meaning
		To be Started
		Started
		DONE
		Vendor Built
		Not Possible

NOW!!!

Builds using ETICS version: 1.3.2-1

INFSO-RI-508833

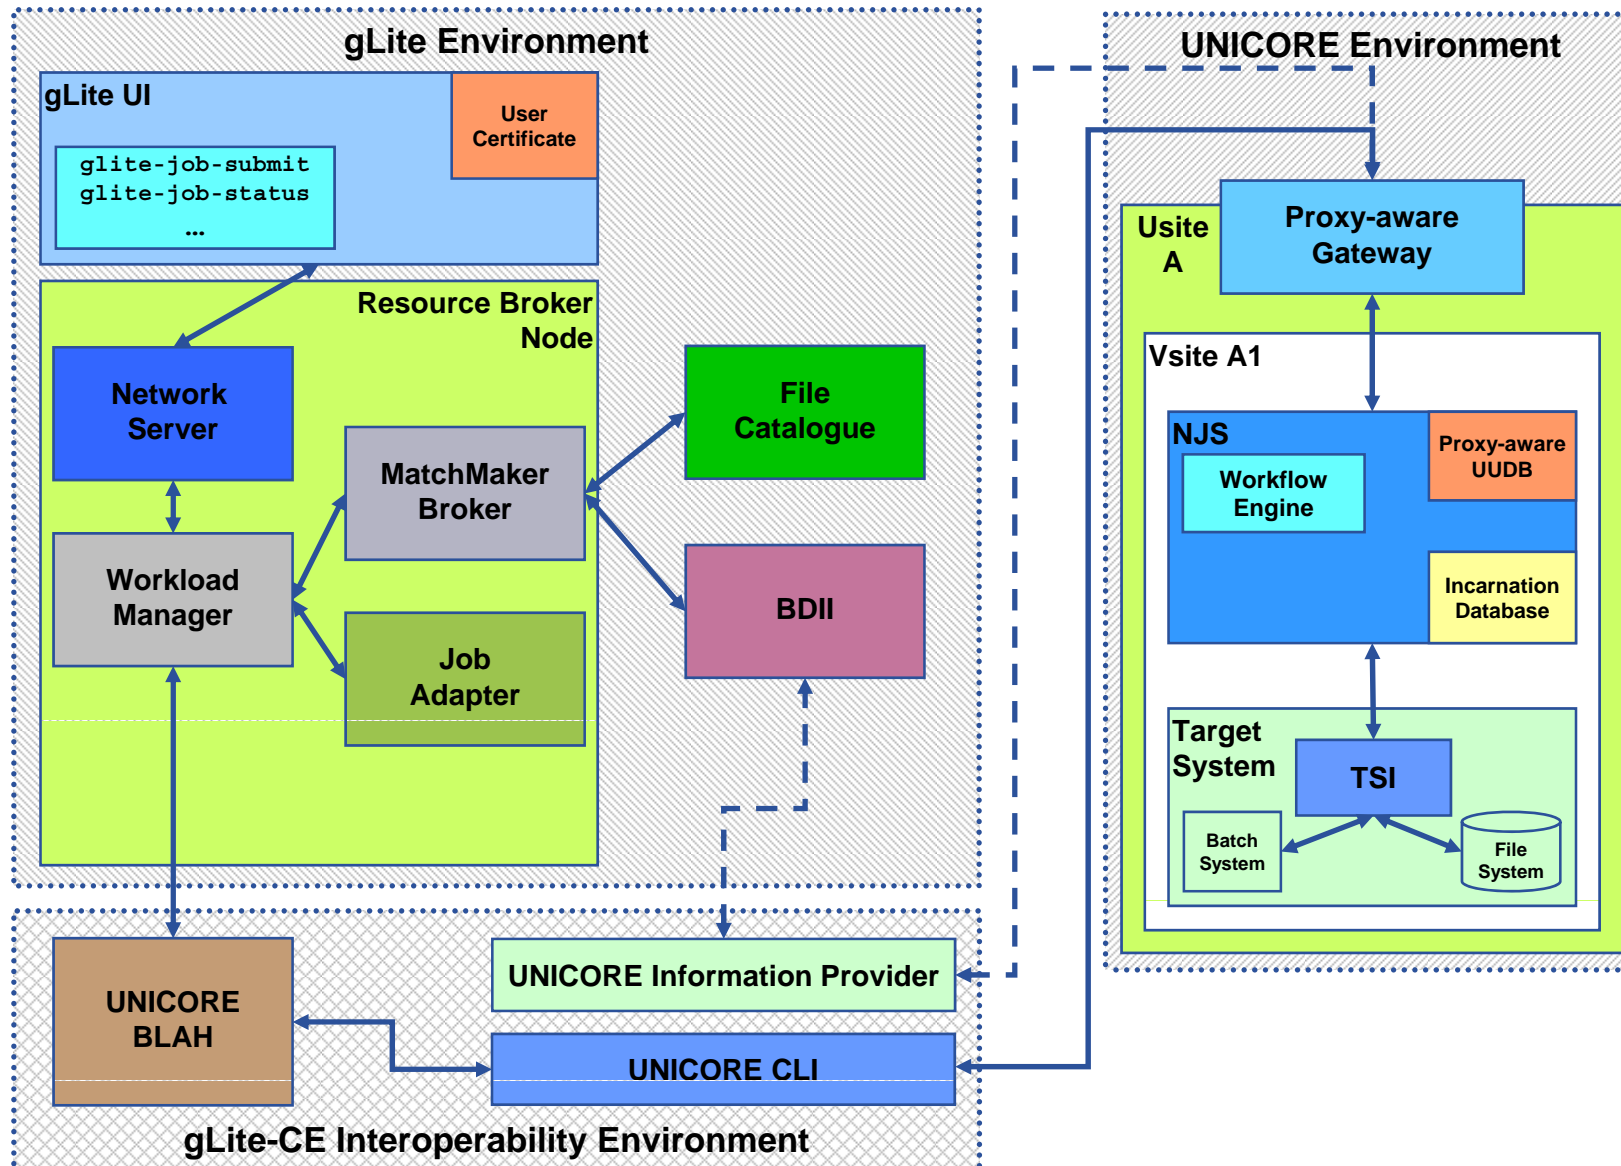
Worker Node Build Status														
ARCH	OS TYPE	VERSION	DISTRO	VDT	deps	GridIre	Basic	RGMA	VOMS	DM	gfal	WN	Full	torque
ia32	CentOS	4	yum	3/3	30/30	2/2	12/12	41/41	14/14	17/17	20/20	108/108	343/379	3/3
	SuSE	9	apt	3/3	28/28	2/2	12/12	41/41	14/14	17/17	20/20	106/106	313/374	3/3
	CentOS	5	yum	1/1	24/24	2/2	9/9	41/41	14/14	16/17	18/21	63/99	234/378	3/3
x86_64	Solaris	10	pkg/tarball	1/1	23/23	2/2	0/11	14/42	0/11	7/17	7/20	N/A	N/A	3/3
	CentOS	4	yum	1/1	26/26	2/2	9/9	41/41	15/15	18/18	19/21	66/105	0/1	3/3
powerpc	SLES	9	apt	1/1	24/24	1/1	9/9	24/41	10/15	11/18	N/A	N/A	239/375	3/3
	Mac OS X	10.4	fink/tarball	1/1	23/23	1/1	0/11	13/27	1/15	0/17	11/29	42/61	179/372	3/3
	Mac OS X	10.3	fink/tarball	1/1	23/23	1/1	0/11	13/27	1/15	0/17	11/29	44/62	183/374	3/3
	AIX	5	rpm/tarball	1/1	21/21	1/1	0/11	0/6	0/4	7/17	7/20	N/A	N/A	3/3

- **New voluntary partner started working on MS-Windows port**
- **Grid2Win by the Consorzio COMETA**
 - UI/WN and CE
 - Support for MS batch systems
 - UI Gui
- **Port based on Cygwin**
- **Much available for LCG-2.7**
- **Port to gLite-3.1 and ETICS started**

- **Batch Systems in production use:**
 - Condor, Torque, PBS, SGE, BQS and LSF
- **SA3 effort on Torque (reference), SGE and Condor**
 - LSF currently orphaned
 - Large sites provide local adaptation
 - These adaptations are needed, but default LSF is a must!
- **SGE guide and LCG-CE have been released (close to PPS)**
 - Port to gLite-CE stopped, move to CREAM-CE in progress
- **Condor:**
 - Still some issues with LCG-CE support
 - Authorization, user mapping....
 - CREAM-CE progresses thanks support from JRA1 developers

- **Up to date status has been documented:**
 - DSA3.2 Report on Progress of Interoperability
 - <https://edms.cern.ch/file/870050/2/EGEE-II-DSA3.2-870050-v0.6.doc>
- **Here only the most compact summary**
- **gLite -> OSG**
 - In production, PPS integrated an interoperability testbed
- **gLite -> UNICORE**
 - Prototype supporting simple job submission has been demonstrated, clear improvements during the last 3 months

- **gLite -> ARC**
 - Activity is behind the “Revised Work Plan”
 - Working on gLite-WMS -> ARC job submission interface
 - Cooperation between WMS developers and Christian
 - Bug that summarizes the situation:
<https://savannah.cern.ch/bugs/?13494>
 - Maybe we see a first prototype before the end
- **gLite -> Naregi**
 - Voluntary effort from Naregi
 - No production system available to verify Naregi's prototype
 - Production systems are planned for end 2008
- **GLUE**
 - Now an OGF working group co-chaired by an SA3 member



- **Release Management and Certification**
 - Provide the support required for LCG CCRC-08
 - Prepare improved version of the release process
 - Increase the number of patches processed by partners
 - Improve SA3 <-> SA1 transition in the process
 - Rationalize distributed test beds and PPS
 - Release: 3.1 WMS/LB and Cream-CE

- **Multiplatform Support**
 - End of March target:
 - 4 non-SL4 supported Linux platforms (including Debian x86/x86_64) + Mac OS X. building in ETICS
 - Debian WN in certification

- **Batch Systems**
 - Condor for LCG-CE and CREAM in certification
 - SGE for CREAM in certification
- **Interoperability**
 - Include more functionality in the UNICORE prototype
 - Get a proof of principle ARC prototype
- **Issues:**
- **Experimental services (WMS/LB)**
 - Expected fast transition to certification and production
 - Material was not ready, tests at Imperial are still ongoing
 - New LB RPMS, new time out parameters,....
 - This is debugging, not certification!
- **Interoperability with ARC**
 - Too little, too late

SA3 Transition to EGEE-III “Anticipated Changes”

Markus Schulz
CERN IT

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

Manpower: EGEE II 12 partners, 30 FTE

Short Name	Total (PMs)	NEW Total (PMs)
CERN	432	396
PSNC	36	24
TCD	19	36
IMPERIAL	24	0
INFN	60	96
UKBH	12	0
UCY	34	12
GRNET	24	30
CSIC	12	8
PIC	24	12
CESGA	12	12
CESNET	0	24
FOM	0	24
UH.HIP	0	12
JINR	0	10
PNPI RAS	0	10
SINP MSU	0	10
STFC	0	36
ASGC	0	40
FZJ	36	0
Total	725	792

EGEE III 17 partner 33 FTE

8 new partners

Significant resources co-located with JRA1 effort (including CERN DM)

This will require more coordination

Process has to ensure that:

Partners can contribute in Patch processing

JRA1 co-allocated partners issue complete and pre-tested Patches of good quality

These partners have to certify patches from other areas

- **Distributed Coordination**
 - CERN
 - release process, certification, interoperability, activity
 - NIKHEF
 - Batch system porting and support
 - TCD
 - Multi platform support
- **Invited new partners to the “SA3 All Hands” in Dublin**
 - Discussed an initial draft of the execution plan
- **Issues:**
 - Some staff are leaving before the end of the project
 - At the start of EGEE-III we will be very short of resources

Certification and Release Process Analysis and Improvements

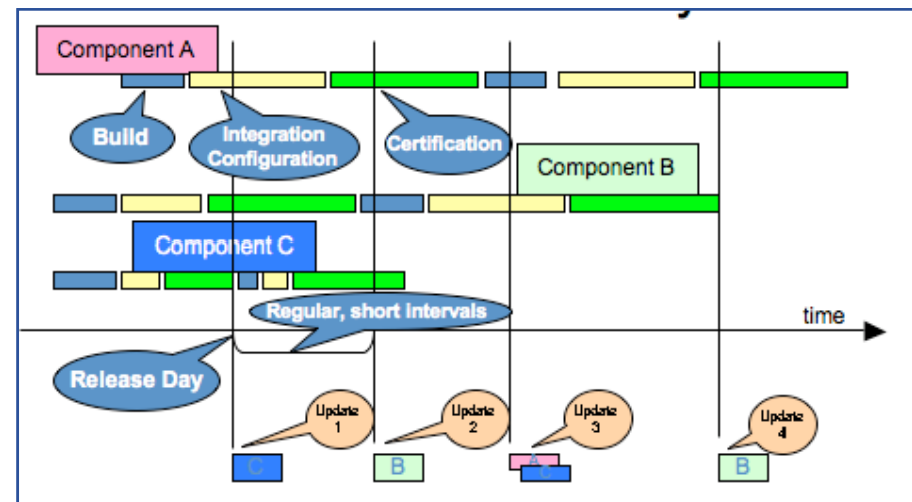
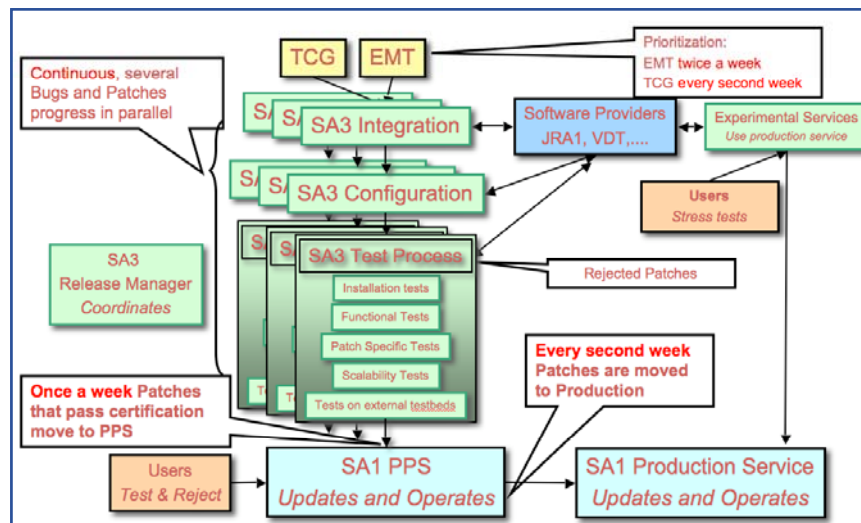
Markus Schulz
CERN IT

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

- **Not an introduction to the current certification/release process**
- **Summary of some issues observed**
- **Suggestions of improvements**

- **Why do we have to change?**
- **New conditions to operate in**
 - More platforms in different “flavors” (32bit/64bit)
 - 1 change == n patches?
- **Shift of resources closer to developers**
 - More strict differentiation between testing, Patch preparation, certification and release preparation
- **New services don't fit the current model**
 - Many iterations are needed
- **The process is perceived as too slow**
 - Latency and throughput
- **We have to work in a more distributed environment**

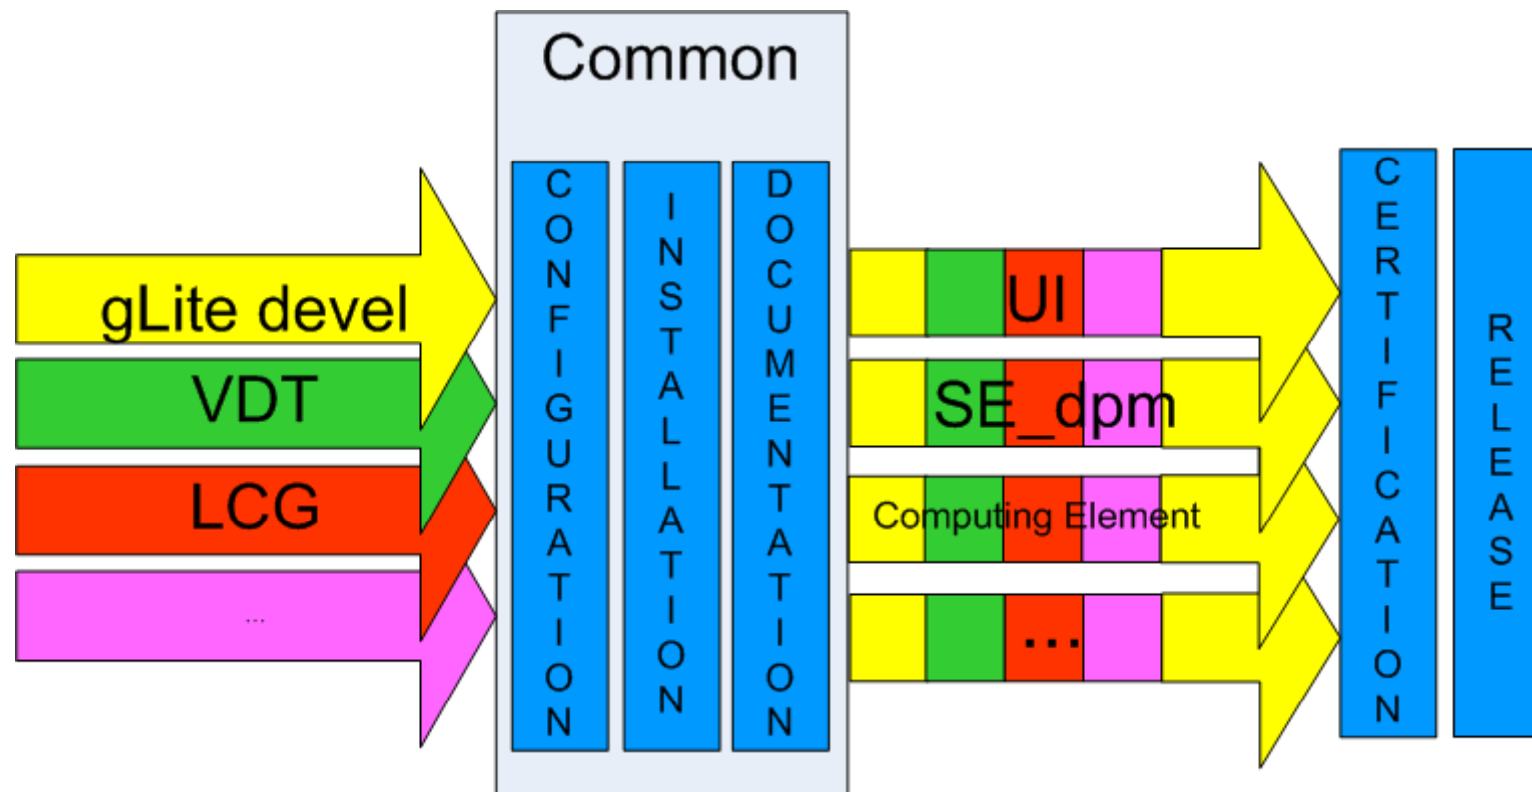
- Process should run concurrently and continuously
- Reality:
 - Several integration points (trivial and critical become intertwined)
 - Several windows of opportunities (certification -> PPS -> Prod)
 - You can miss the bus twice for trivial reasons
 - We have more components than certification “slots”
 - *Some Patches have to wait*



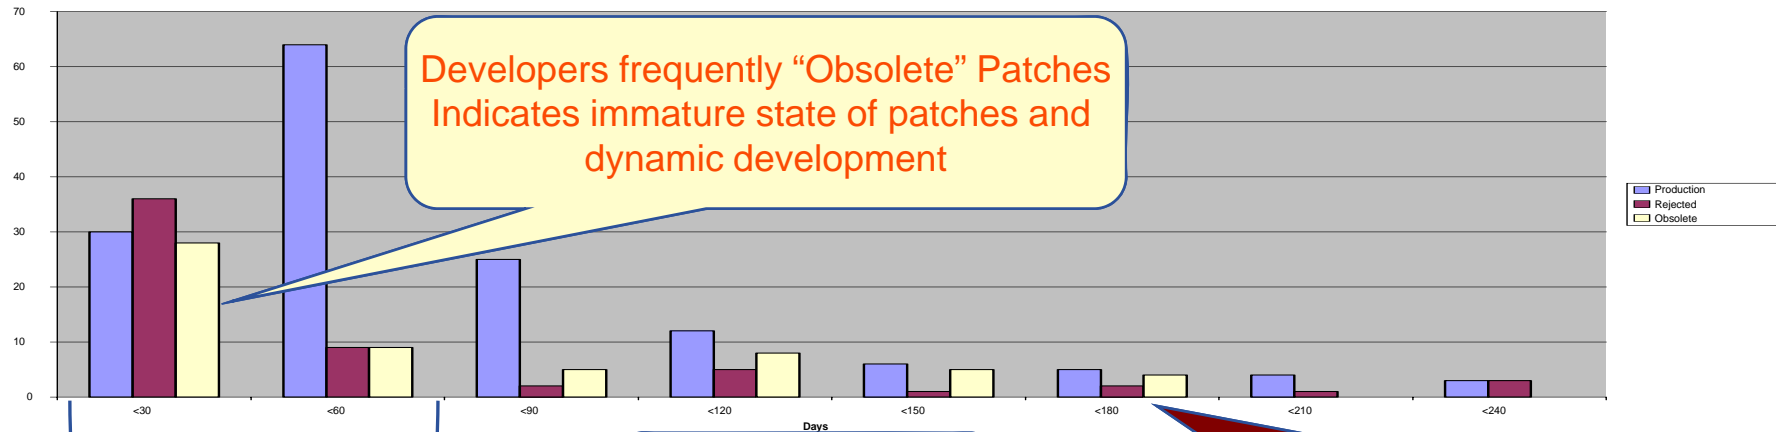
The 3 'Integration Points'

•Process

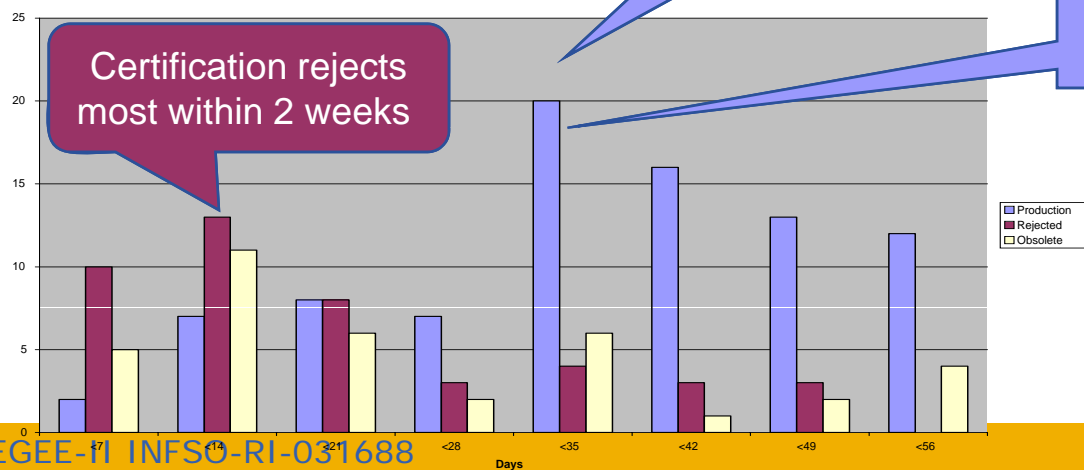
- 'Integration points' introduce delays and extra dependencies
- Important changes can be delayed by trivial ones
- Doesn't change throughput, but affects latency



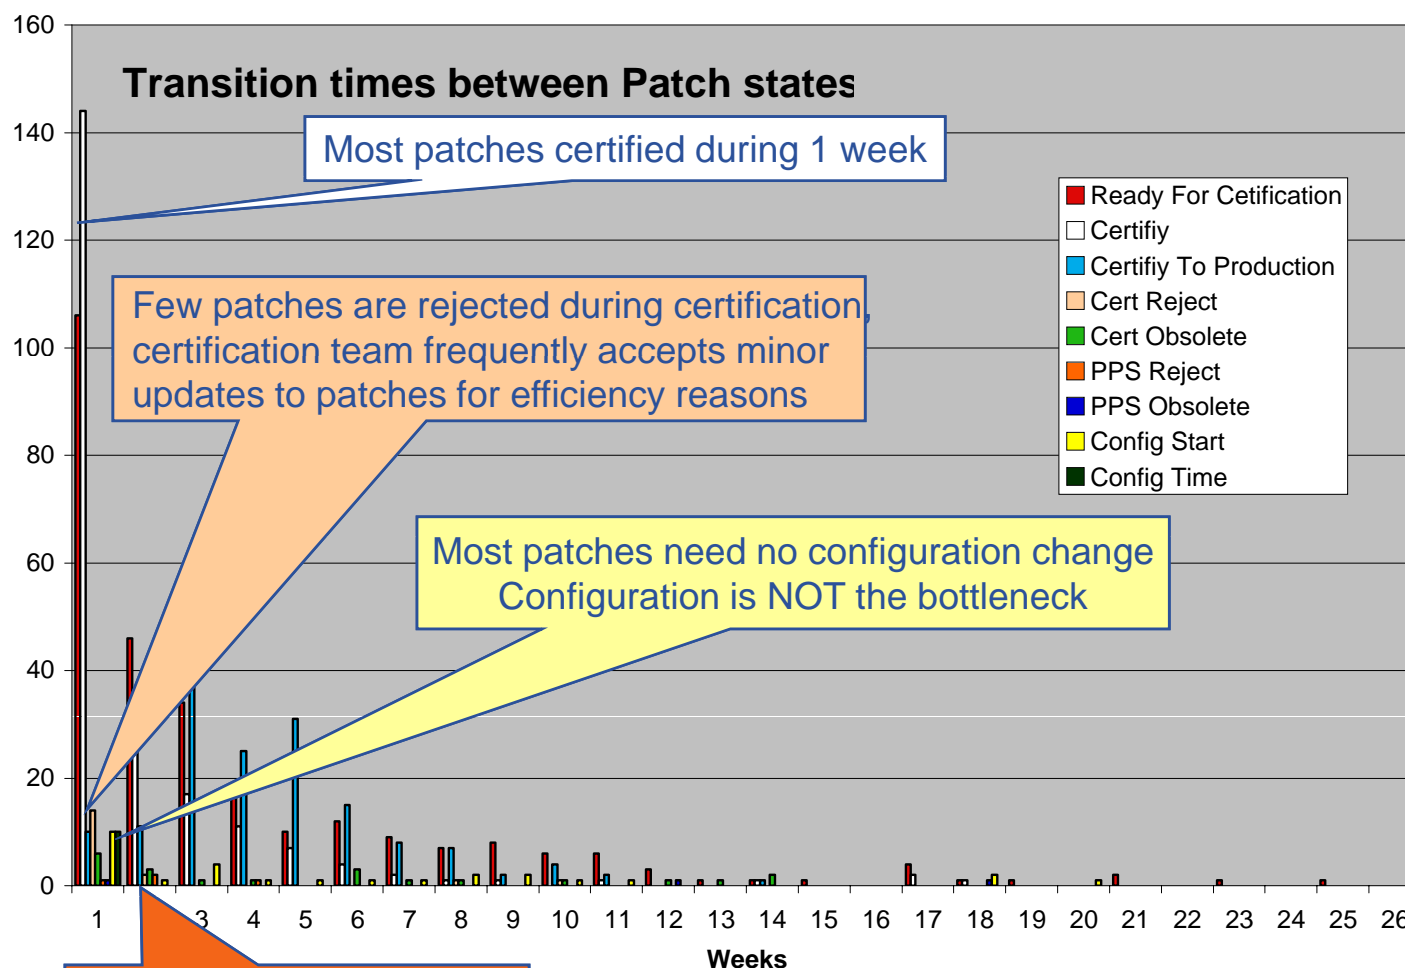
Patch Latency



Patch Latency

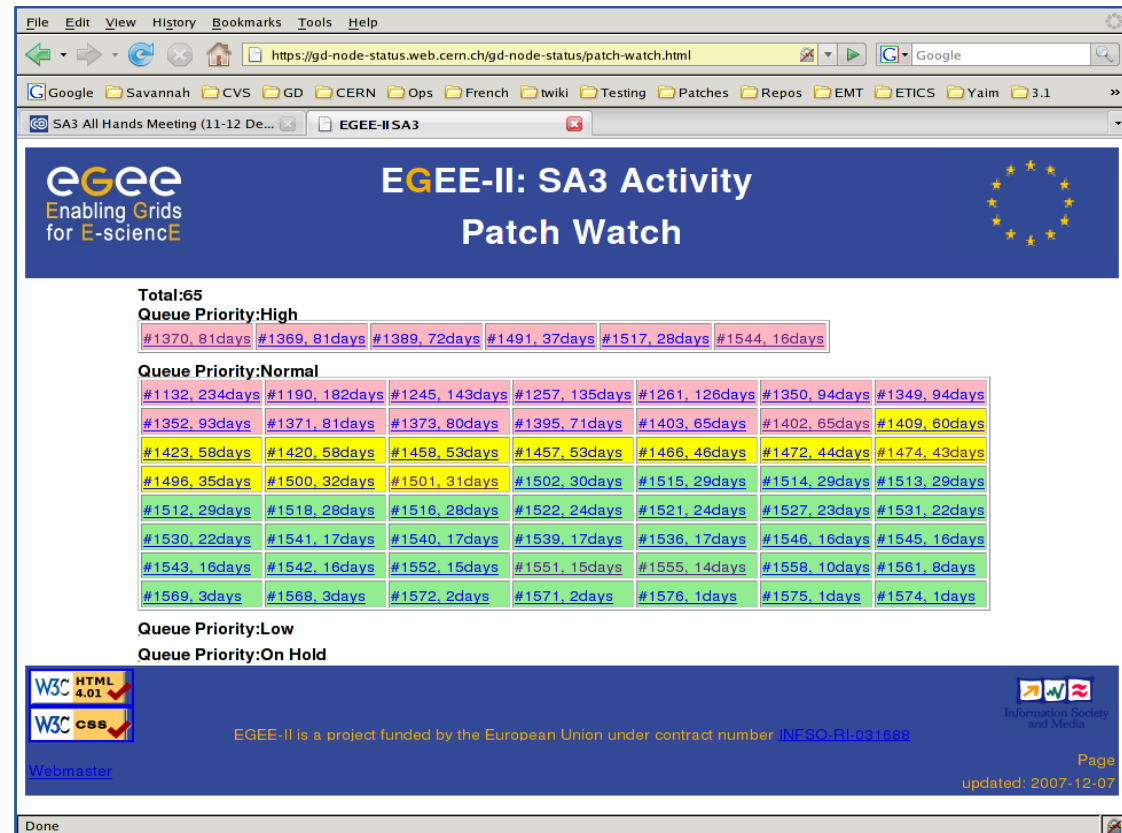


Many patches cycle several times through “ready for cert” and “certify”
The process is responsive, costly iterations could be avoided by creating more complete patches.



How to control the tails?

- Several reasons for long tails:
 - Passive Patch (stays in early state)
 - Priority starvation (there is always something more urgent)
- Addressed by
 - “Patch Watch”



The screenshot shows a web browser window displaying the "EGEE-II: SA3 Activity Patch Watch" page. The page header includes the eGEE logo and the European Union flag. The main content area lists patches categorized by priority: High, Normal, and Low. Each patch entry shows its ID and the number of days it has been in the queue. The page also includes a footer with logos for W3C HTML 4.01 and W3C CSS, and a note about the project being funded by the European Union under contract number INFISO-RI-031688. The page was updated on 2007-12-07.

EGEE-II: SA3 Activity Patch Watch

Total:65

Queue Priority:High

#1370, 81days	#1369, 81days	#1389, 72days	#1491, 37days	#1517, 28days	#1544, 16days
---------------	---------------	---------------	---------------	---------------	---------------

Queue Priority:Normal

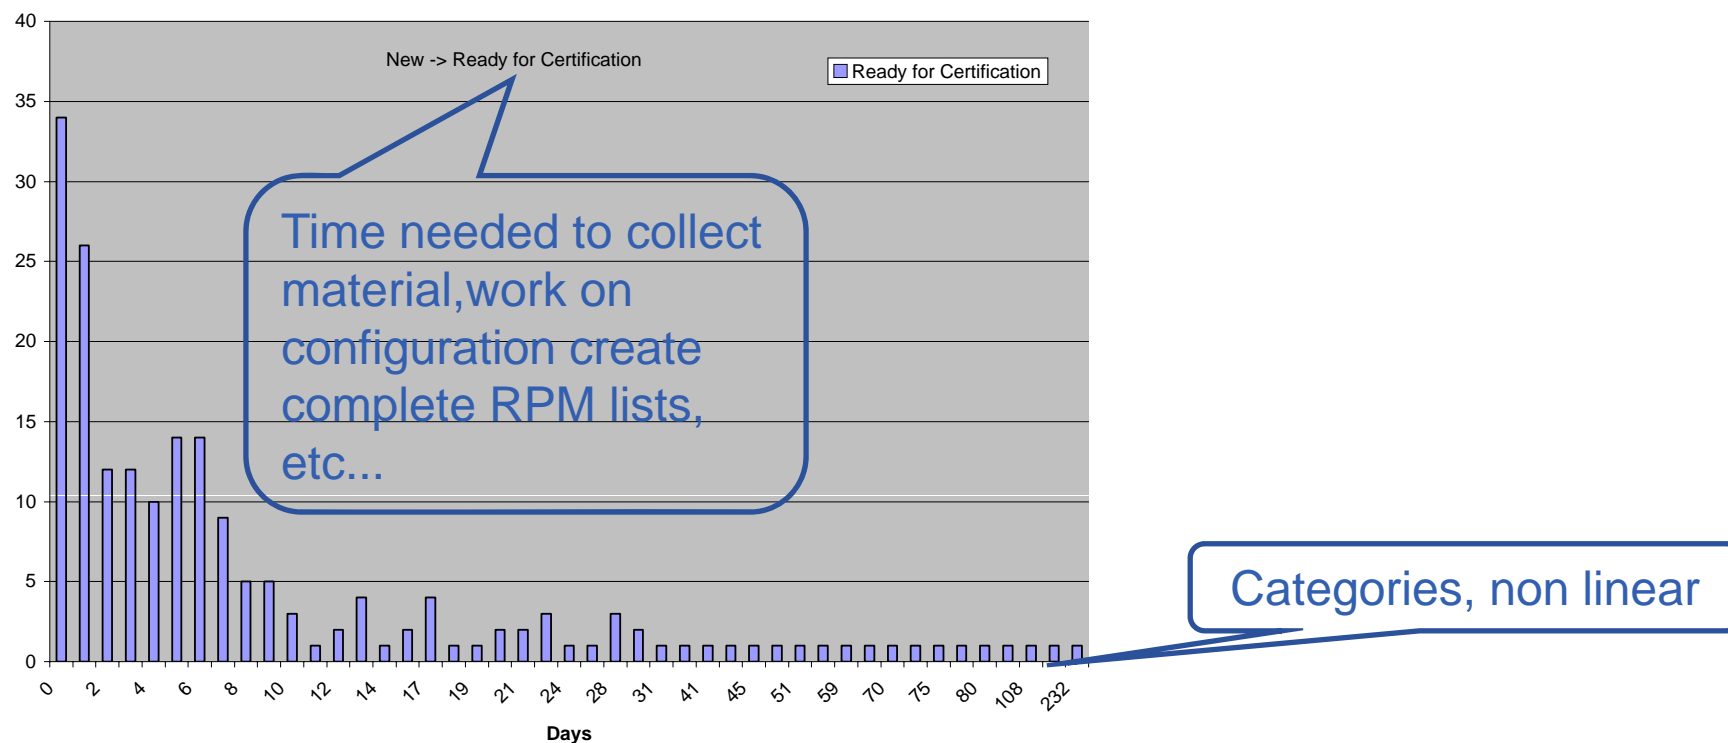
#1132, 234days	#1190, 182days	#1245, 143days	#1257, 135days	#1261, 126days	#1350, 94days	#1349, 94days
#1352, 93days	#1371, 81days	#1373, 80days	#1395, 71days	#1403, 65days	#1402, 65days	#1409, 60days
#1423, 58days	#1420, 58days	#1458, 53days	#1457, 53days	#1466, 46days	#1472, 44days	#1474, 43days
#1496, 35days	#1500, 32days	#1501, 31days	#1502, 30days	#1515, 29days	#1514, 29days	#1513, 29days
#1512, 29days	#1518, 28days	#1516, 28days	#1522, 24days	#1521, 24days	#1527, 23days	#1531, 22days
#1530, 22days	#1541, 17days	#1540, 17days	#1539, 17days	#1536, 17days	#1546, 16days	#1545, 16days
#1543, 16days	#1542, 16days	#1552, 15days	#1551, 15days	#1555, 14days	#1558, 10days	#1561, 8days
#1569, 3days	#1568, 3days	#1572, 2days	#1571, 2days	#1576, 1days	#1575, 1days	#1574, 1days

Queue Priority:Low

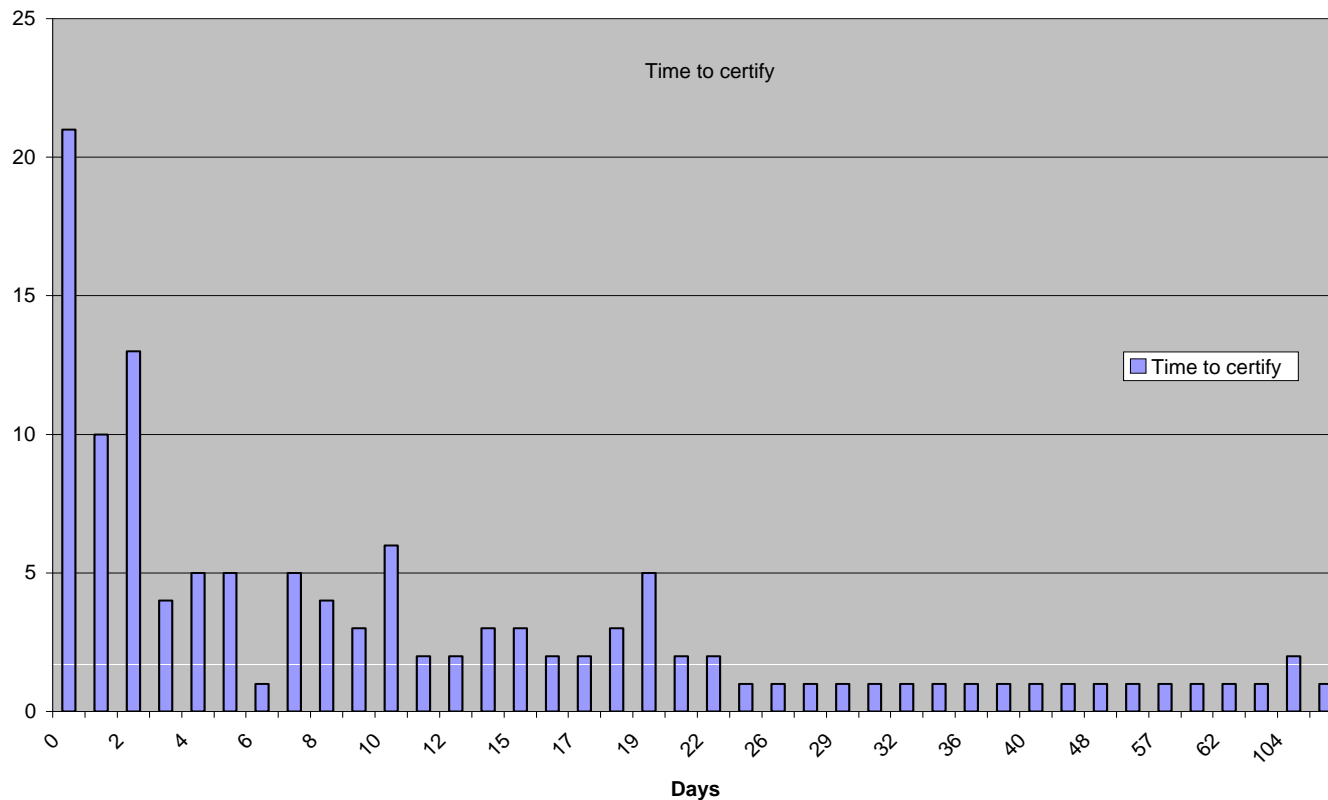
Queue Priority:On Hold

EGEE-II is a project funded by the European Union under contract number [INFISO-RI-031688](#)

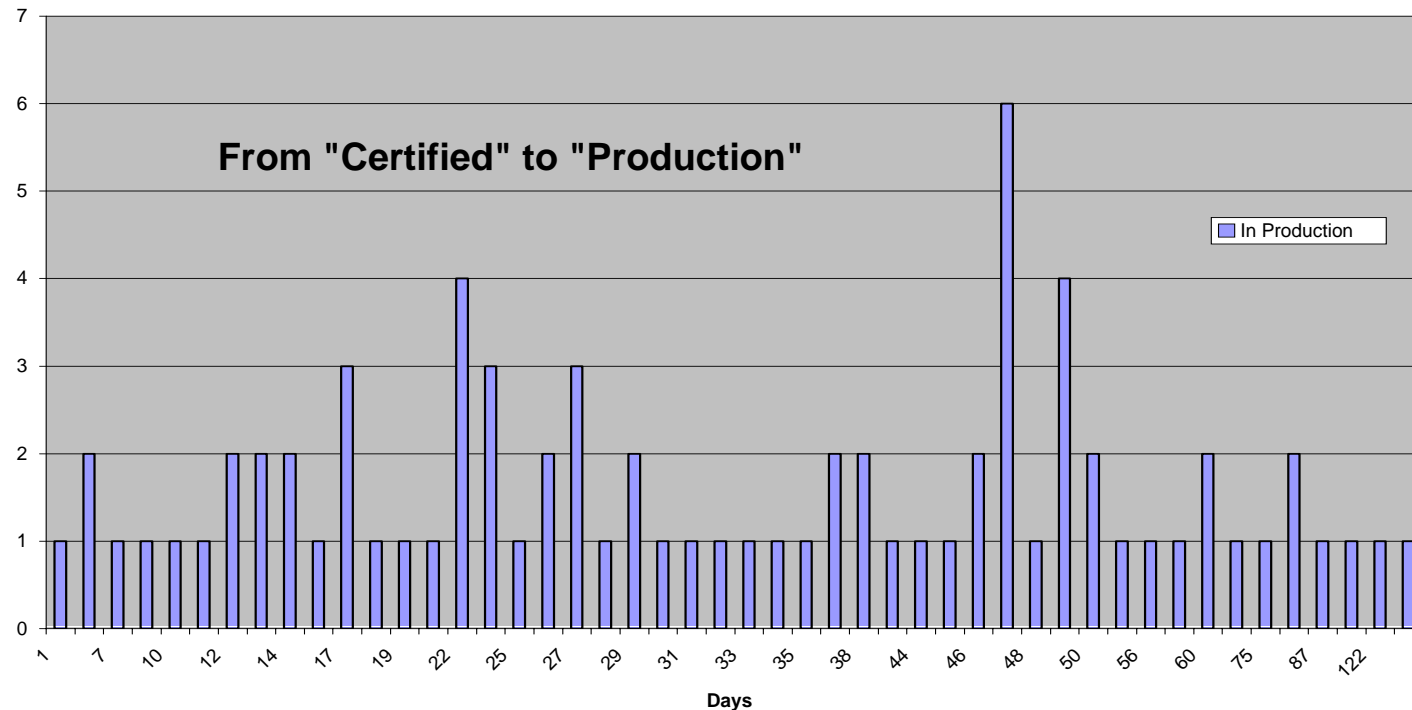
Page updated: 2007-12-07



- Most work can be done before the Patch is opened



- Simple Patches pass through quickly
- Some need development work for tests
- Some are “reprioritized”



- **Possible delays:**
 - Certified to PPS (2 weeks),
 - PPS (> 2 weeks), Release (2 weeks)
- **Some Patches wait for user feedback**

- **Why don't we reject more incomplete or bad Patches**
 - Some Patches have more than 20 comments (average 10)
- **Patches contain a lot of information and history**
 - What is changed?
 - RPMs, configuration, RPM lists, dependencies....
 - Savannah doesn't allow to clone patches
 - And has only GUI interface (no scripting)
 - Patches cannot be locked
- **What to do?**
 - Make these iterations explicit by introducing extra state
 - To get a realistic start date for a patch
 - Maybe ETICS can enhance Savannah

- **Releasing services and clients differently can help**
 - Clients can be “pushed” to sites
 - Multiple versions per site
 - Users can select version via JDL
 - Can replace in many cases the PPS
- **Releasing configuration changes independently**
 - YAIM configuration allows now independent versioning for each component
- **New services have to be run first as pilot services**
 - Cooperation with user communities is a must
- **Deployment tests can be run in parallel with certification tests (one future role of PPS?)**

- **With minor changes to the repositories Patches can be handled very independently (including configuration)**
 - Clear criteria for certification need to be defined
 - Active tracing of Patch status to avoid stale Patches
 - SA3 partners co-located with JRA1 have to ensure quality
 - Less iterations
 - Use of standardized checklists
- **We need an incentive to participate in Patch processing**
 - 36 Patches have been handled by Partners
 - An accounting system and quantitative commitments
 - Credits for different Patch categories certified within time limits
 - *Trivial (1) , Average (2) , Challenging (5 -10)*

Switching on the auto pilot

- **Many steps in the release process can be automated**
 - Almost all data is available in Savannah, ETICS, or test traces
 - Laurence offers daily updated Savannah DB snapshots
- **Automatic deployment and regression testing via ETICS**
 - Many problems that are found in certification are trivial
 - Better handled asynchronously in parallel by development teams
- **Using virtualized testbeds for certification**

- **All work on new services has to be documented via electronic logbooks**
- **All tests and their results have to be logged and archived**
 - And have to be traceable to specific configuration and software versions
- **The goal is reproducibility**
 - This is plain laboratory work!
- **All this sounds trivial, but it would really help**

- **Some fraction of the extended certification testbed and the PPS have overlapping functionality**
- **Not all PPS sites need to move in lockstep**
 - Deployment tests can be carried out by subsets
 - PPS shifts to ensure quick turnaround
 - New services can be introduced by a small subset of partners
 - Either SA3 or SA1
 - Sites will need to change less frequently
 - Workload for each site would be more blocked
 - Site handles interaction with users

- **Ideas, Suggestions**
- **SA3 @ JRA1 sites**
- **Focus on preparation phase:**
 - Identify complete rpm list and define the required build
 - Perform the build against the project definition
 - Run, as part of the build, deployment tests and regression tests
 - Update configuration and produce any necessary yaim packages
 - Perform smoke test on service and clients, communicate results to dev cluster
 - Write the patch release notes
 - Submit a patch, referencing packages in the ETICS repository
 - Offer support as necessary to the assigned patch certifier
 - Regression tests maintenance
 - Add new tests to ETICS