



Enabling Grids for E-science

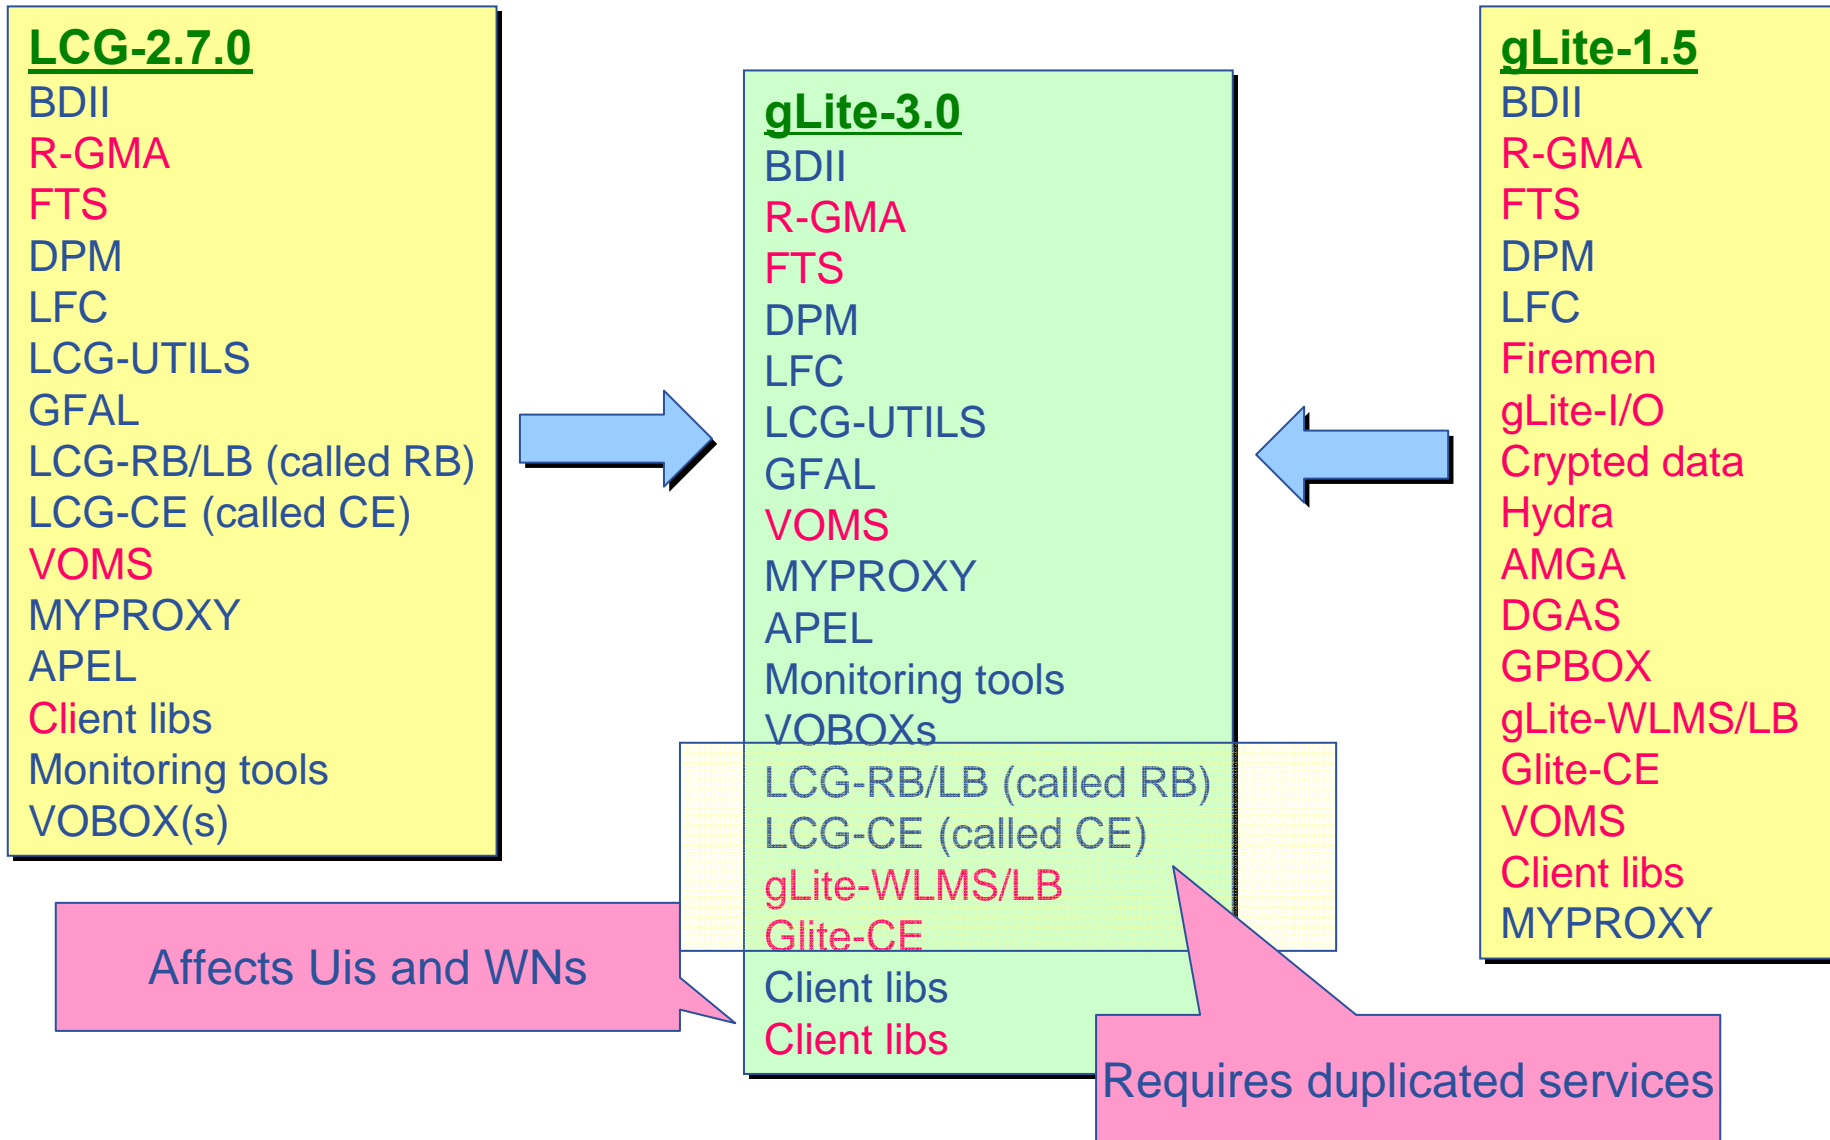
gLite-3.0.0 Release Process and Distribution

Markus Schulz
CERN-IT-GD-ITR
EGEE-II SA3

www.eu-egee.org



What is gLite-3.0.0?



LCG-2.7.0

LCG-build system
 LCG Documentation
 Configuration management
 Dependency management
 Deployment oriented
 Test&Certification
 Certification testbed
 Bug tracking
 Site manager expertise
 Procedures
 Naming conventions
 Compact local team

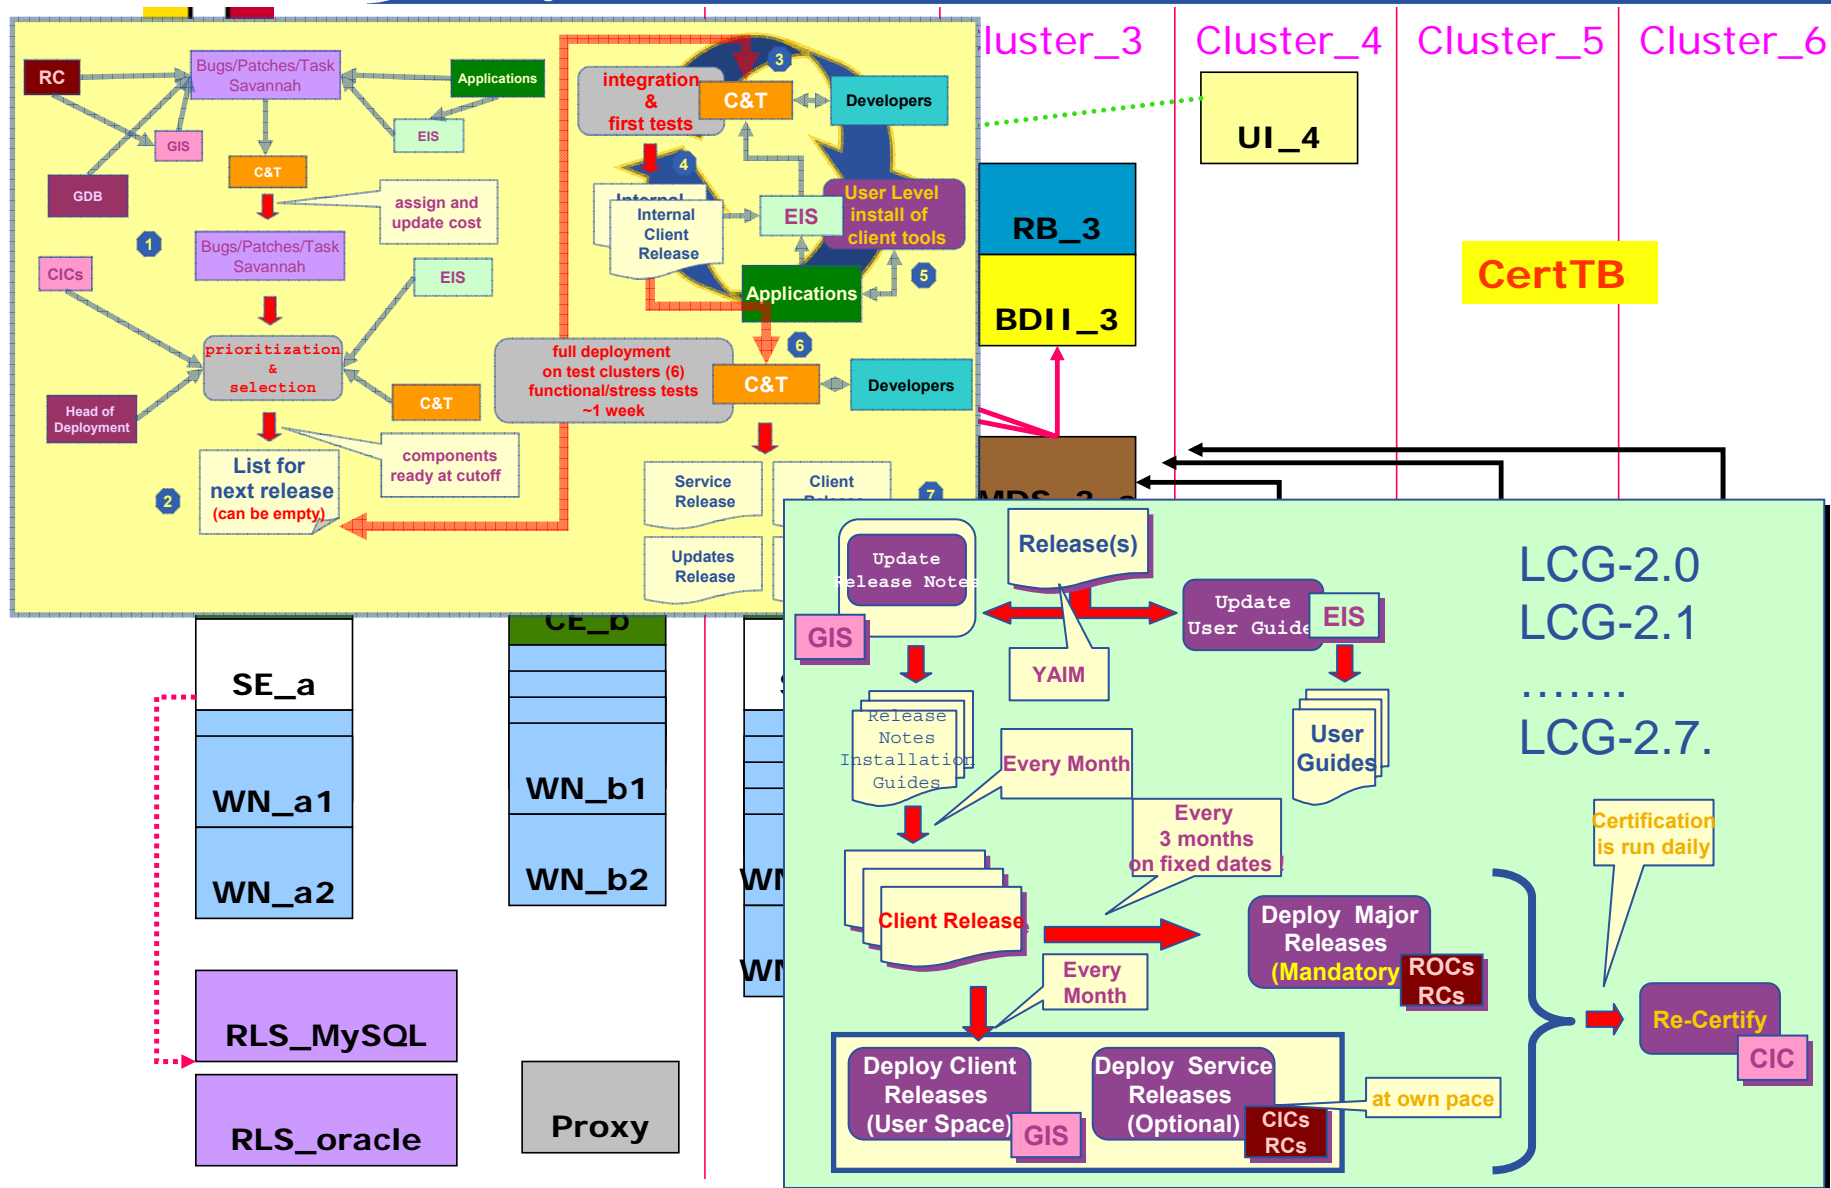
gLite-1.5

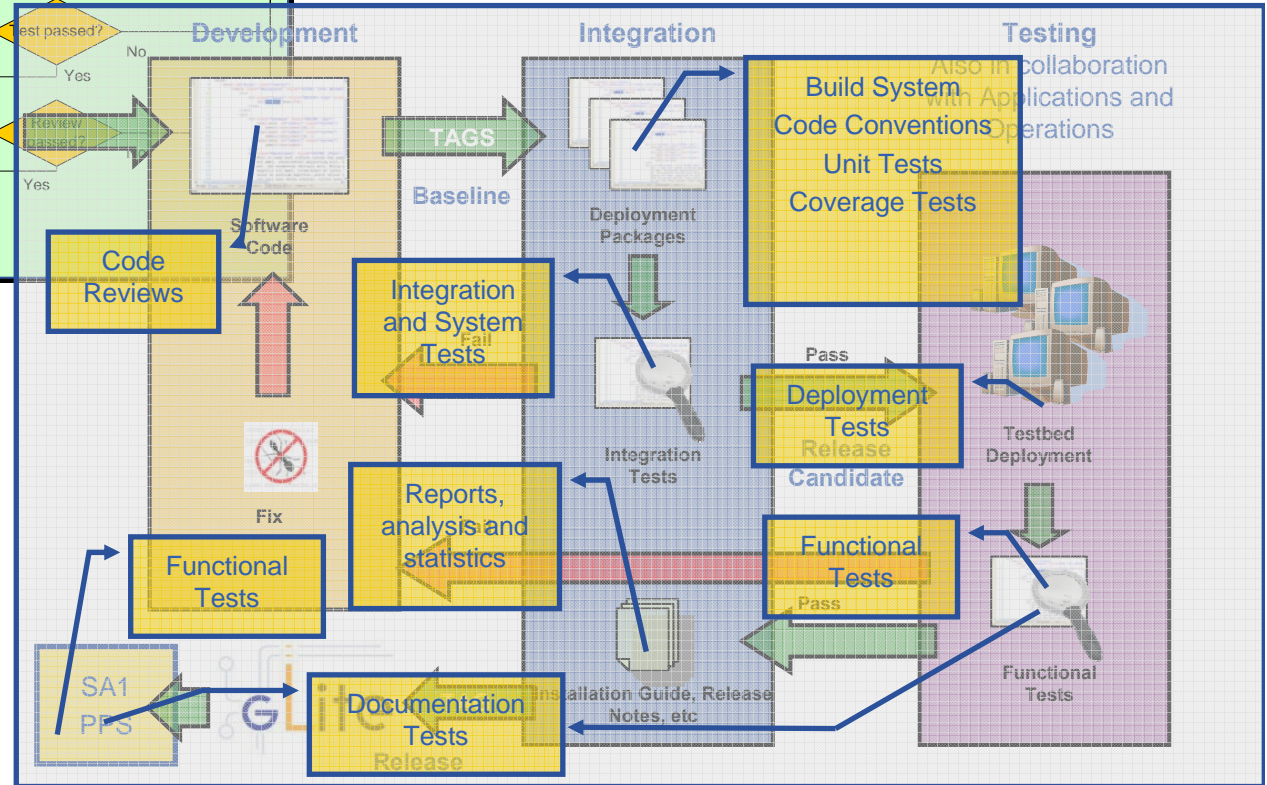
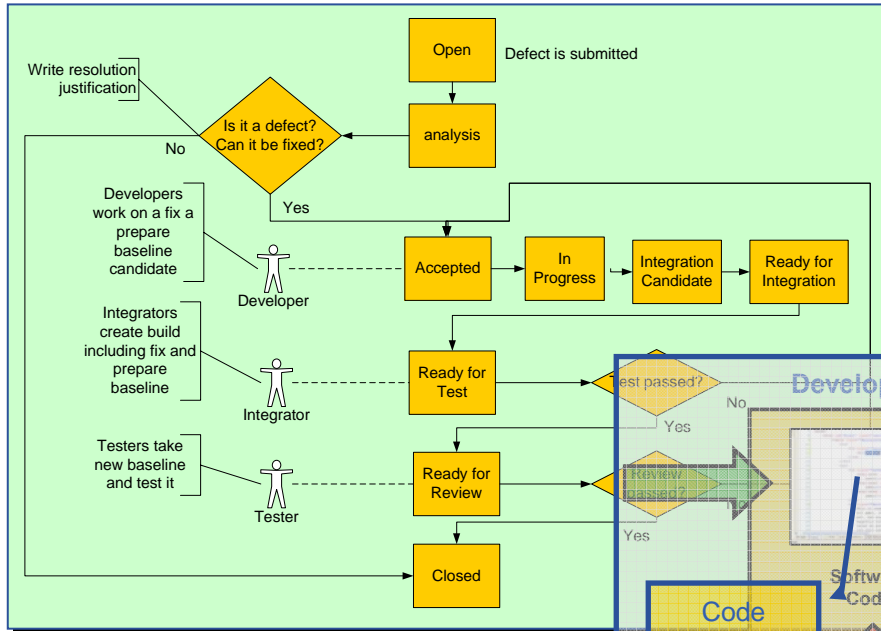
gLite build system
 gLite documentation
 Procedures
 Certification testbed
 Bug Tracking
 Testing and certification
 Configuration management
 Dependency management
 Developer oriented
 Certification testbed
 Naming Conventions
 Distributed Team

gLite-3.0

LCG-build system
 gLite build system
 LCG Documentation
 gLite documentation
 Configuration management
 Configuration management
 Dependency management
 Dependency management
 Deployment oriented
 Test&Certification
 Testing and certification
 Certification testbed
 Bug tracking
 Bug Tracking
 Procedures
 Procedures
 Naming conventions
 Naming Conventions
 New team+ external developers

- Two processes
 - With slightly different focus had to be integrated
- Teams worked independently until late January
 - gLite-1.5
 - LCG-2.7.0
- No time for in depth integration of process





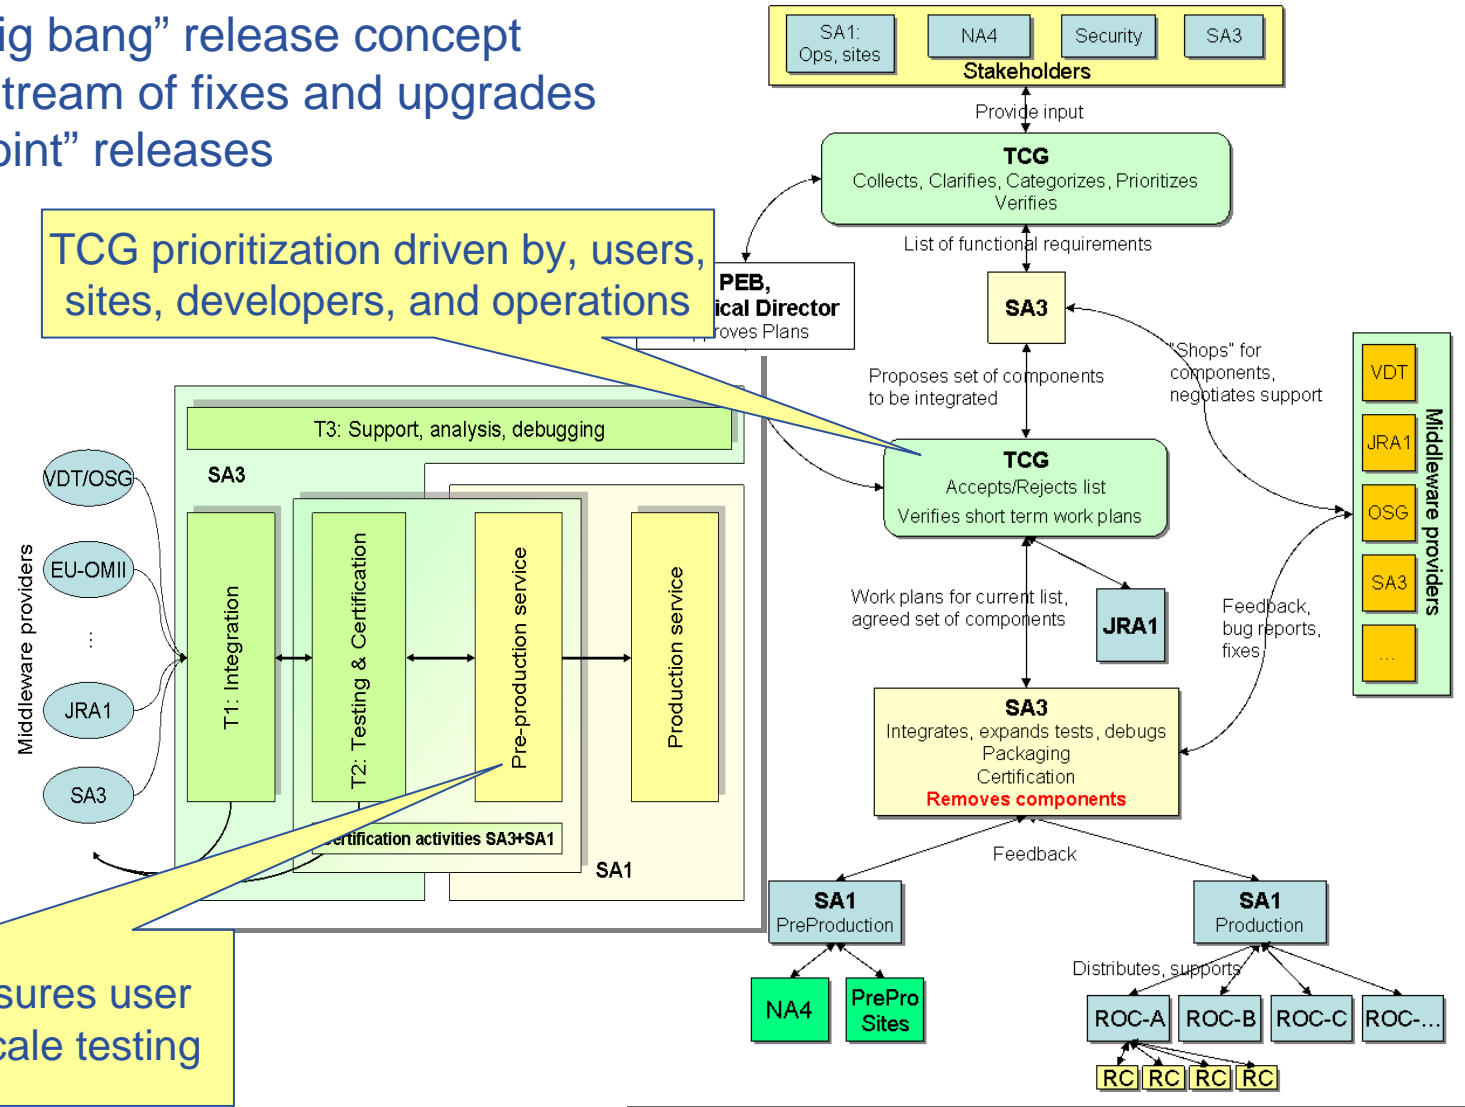
Based on standard methods and frameworks

Very complete (tests, tracking, QA, builds, documentation generation)

End of the “big bang” release concept
 Continuous stream of fixes and upgrades
 1-2 “check point” releases

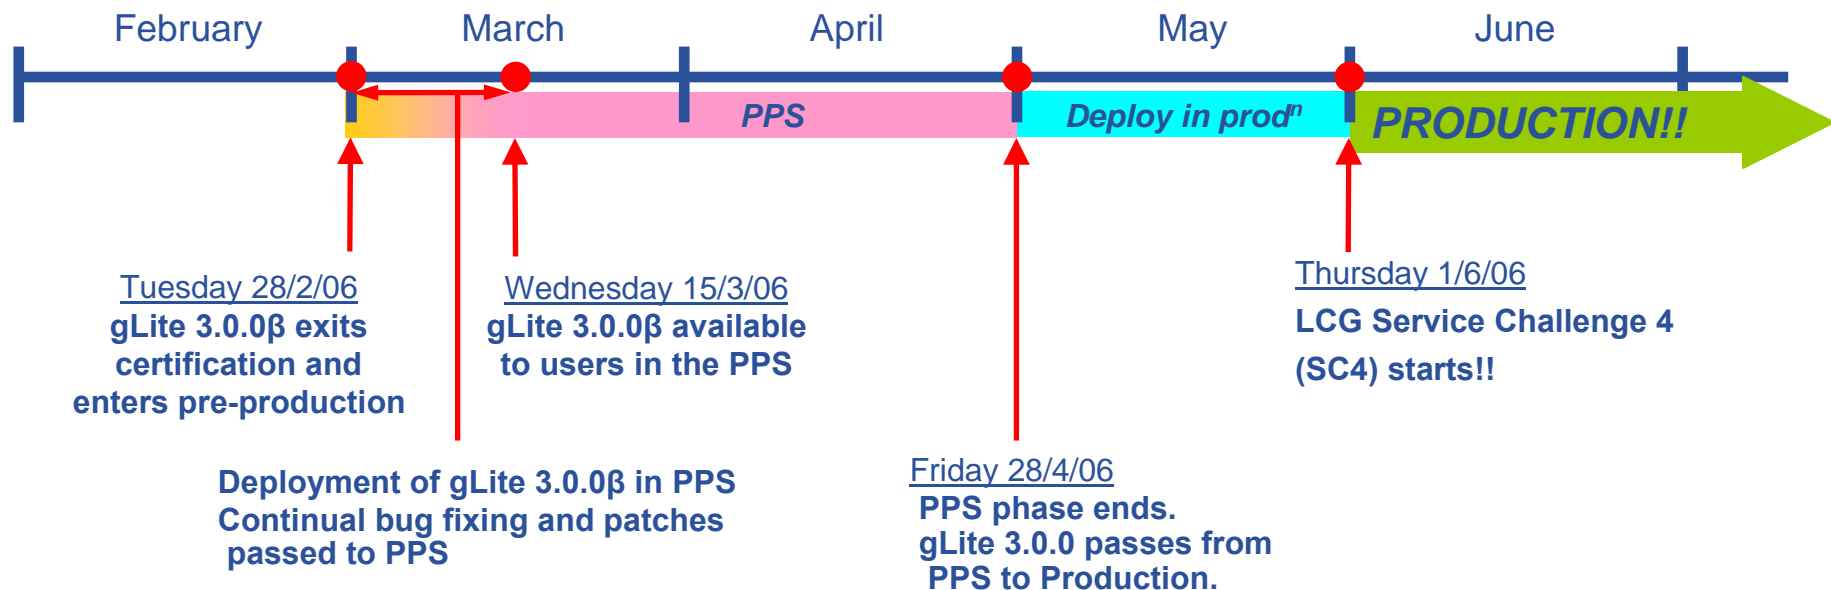
TCG prioritization driven by, users, sites, developers, and operations

Preproduction ensures user input and large scale testing



History of gLite-3.0

- Detailed planning began end Jan, with clear Jun 1st deadline.
 - 20th Feb, freeze, patches only thereafter
 - 28th Feb, beta release to PPS
 - **Mid March – gLite 3.0 to be available on PPS for user inspection**
 - 1st May, gLite 3.0 becomes production ready (patches, docs, installation/config...)
 - 1st May, gLite 3.0 released to production
 - 1st June, SC4 sites on gLite 3.0



– What happened

- 20th Feb; freeze
- 3rd March, beta released to PPS
- All of March, deployment on PPS, close monitoring and creation of RC2
 - *PPS not available for users*
- 11th Apr – RC2 hits the PPS
 - *too late, but what's the average bug lifetime until integration?*
- Apr – updates and patches
 - *PPS sites are trying to run a stable service;*
 - Less than ideal conditions for testers
- Before easter: ROC deployment testing (5 ROCs volunteered)
 - CE ROC, IT ROC, UK ROC, EGEE-SEE.
 - #16388 #16355 submitted
- 4th May gLite 3.0 released to production with staged deployment plan
- 5th June 8 'blah' CEs, >50 sites have installed client lib.

- **2 weeks after release**
 - 1st upgrade (YAIM only)
 - Fixes for relocatable UI/WNs
 - Many documentation glitches fixed

- **2nd June: Full day postmortem (follow link for more details)**
- **6th June**
 - 1st upgrade ready
 - gLite WLMs, CE, UI, WN
 - Some desired WLMS bugfixes didn't make it (2 weeks later)

- **Plan:**
- **Aggregate bug fixes and release upgrades every 2-3 weeks**
- **Major changes have to pass through the PPS**
 - No clear set of criteria
 - Needs to be discussed for each upgrade
 - EMT/GD?

Problems

- (Fixed release date + Fixed set of components) := Problems
- Release has still significant deployment issues
- The release for gLite-3.0 hit most sites unprepared
 - Release had been announced already in January
 - Weekly status updates (ops meeting, wiki page)
 - The link to the SC4 schedule was not clear (developers, sites)
- The release and upgrade notes confused sites
 - Who should deploy what when?
- The staged rollout made everyone wait
 - 2 waves of T1s followed by the SC4 T2s
 - No trailblazers (small to medium sites that upgrade rapidly)
- There was not enough time left for localization
 - Pre-release for deployment tests was released **before easter**
 - One month not enough for large sites to integrate new services
 - Planning, Resource allocation
 - Adaptation to local fabric management and batch systems
 - Internal testing of new services/releases
 - **Resolving coexistences problems with other grid stacks**

- **Release Notes 31.01.2006**

This release of LCG-2_7_0 should be understood in the context of this year's timetable. As a ROC manager or site admin you may have heard about the planned merger of the LCG-2.x middleware stack and the gLite middleware fusing into gLite-3.0.

While we can't give you an accurate release date for this important change we think that we should share with you what drives the process that will lead to gLite-3.0. The main event on the production system this year will be Service Challenge 4. **For this the start of the production phase is June 1st. At this time gLite-3.0 has to be up and running at all sites that participate.** To allow for some local adaptation, verification, and the like, the release has to be given to the sites **at least one month before the start.** This makes it very likely that the next release will be there before April 30.

- **Release Notes 4.05.2006**

Who can upgrade?

Upgrades are supported from LCG-2.7.0. If you upgrade from gLite-1.5.0 upgrades might work for the components that are in common, however this has not been tested and it might be advisable to reinstall instead.

Why two workload management systems?

Since many applications haven't migrated all their production systems to the gLite WLM we have to keep the LCG RBs and CEs operating.

- **Release Notes 4.05.2006**

gLite CE or LCG CE?

All larger sites should deploy both CEs to ensure that the majority of resources is available from both worlds. If you are running two CEs please take care to avoid collisions of pool account mapping. This is typically achieved either by allocating separate pool account ranges to each CE or by allowing them to share a gridmapdir. Since the gLite WLMS can utilize LCG-CEs, smaller sites should stay with the LCG CE to allow access to their resources and data that is stored at their site. The project will make an announcement when this preference will change.

gLiteWMS/LB or LCG-RB?

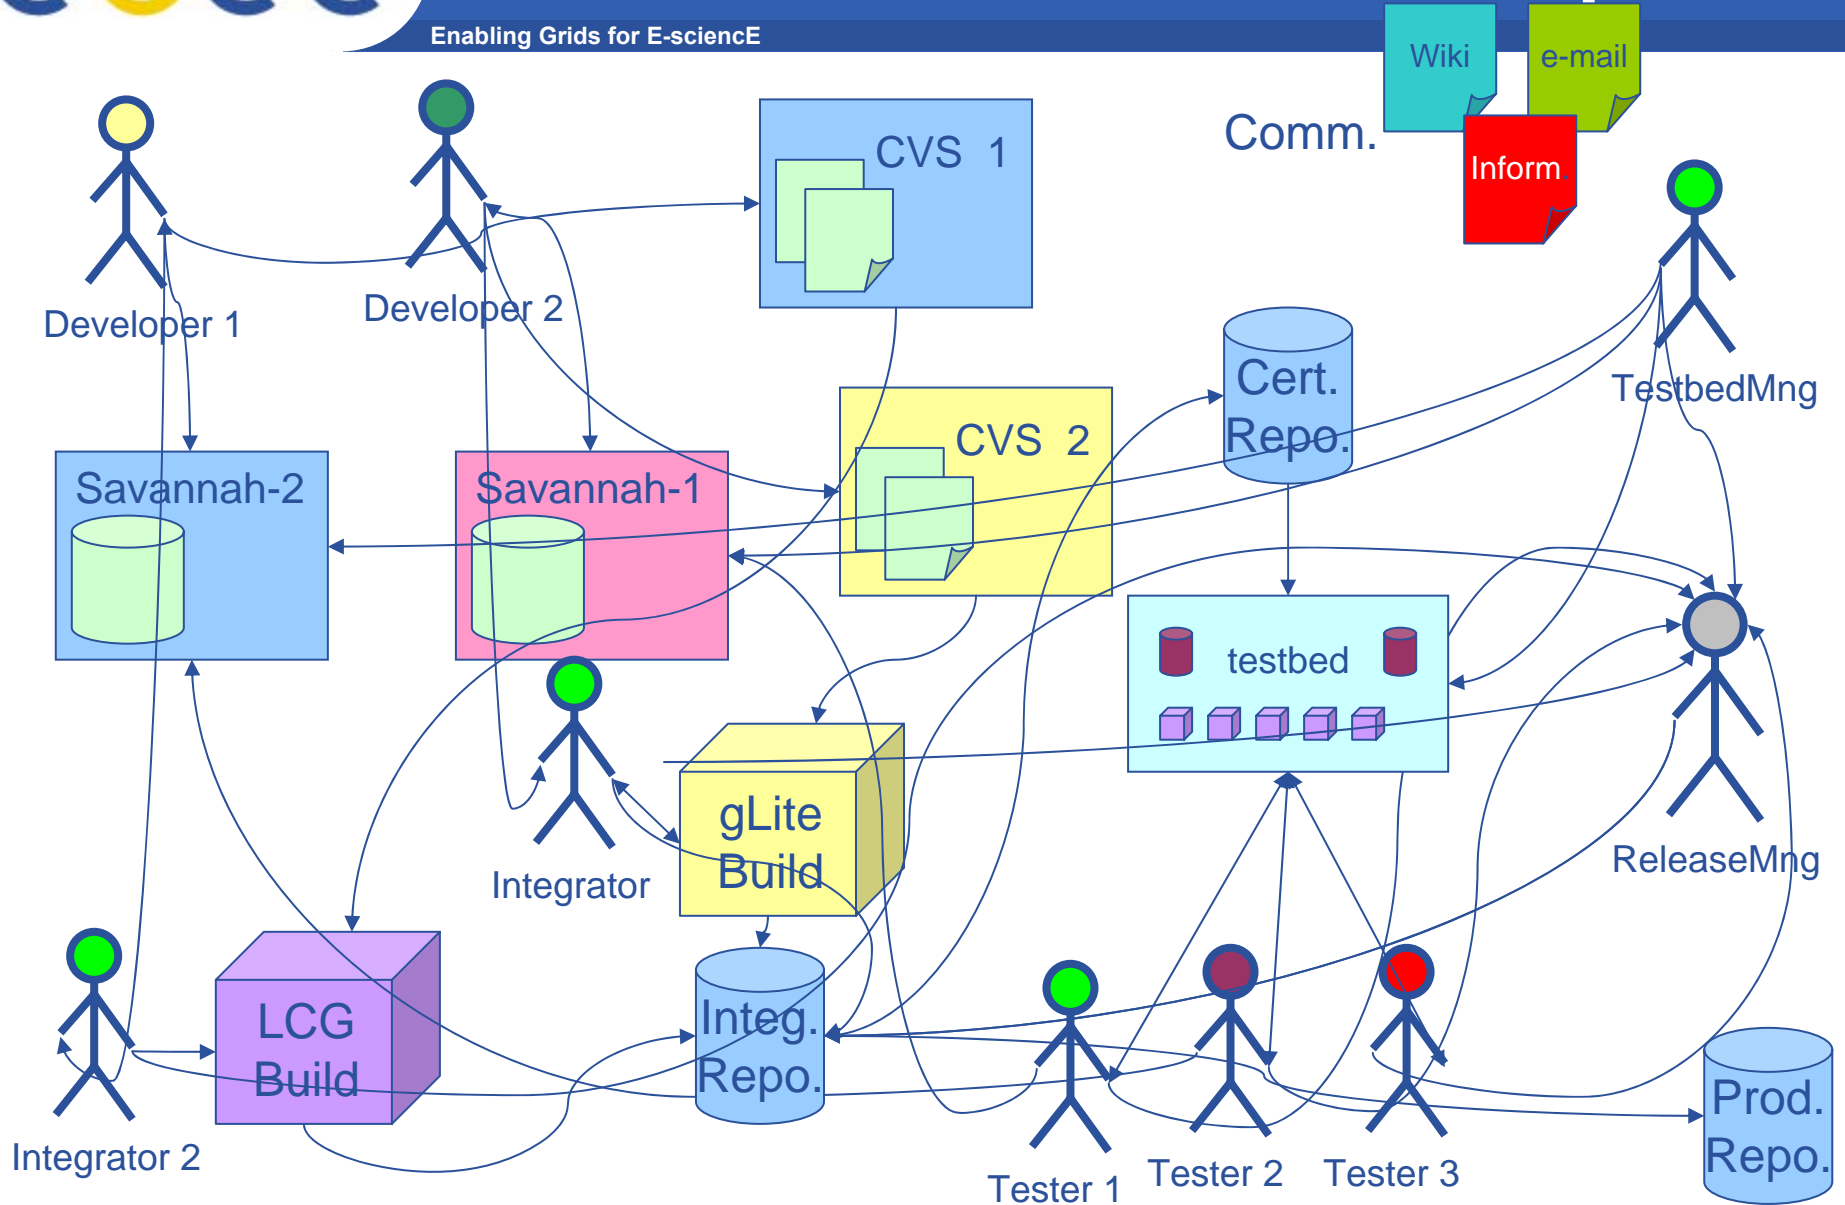
Since the work load manager end user APIs differ for these services you have to get an agreement with you local user community which node type they prefer. Large sites should try to add at least on gLite WMS/LB node to the set of RBs that they are operating.

- **Communication**
 - Within the team
 - Integration, cert testbed, testers, release manger
 - Synchronization via Wiki page (temporary solution)
 - Between teams and developers
 - Between team and sites (ROCs)
 - Ad hoc synchronization via Wiki page frequent meetings
 - For full cycle of workflow we needed a multi step communication

- **Process (two are worse than none!!!)**
 - Non uniform tracking
 - 2 * Savannah Projects + GGUS
 - Temp. adhoc process (gLite/LCG) (Improved rapidly)
 - Relied on informal comm.

 - Tracking between bugs, patches and release candidates unreliable

- **Testing**
 - “Scattered” tests (linked via Wiki)
 - Tests not complete (info sys) (Andreas Unterkircher maintains a gap analysis page)
 - Test results not always conclusive
 - Test require full setup
 - Old sins slipped through (tarball WN, setting environment not via profile.d.....)
- **Integration**
 - Two build systems (hard to manage dependencies)
 - Conflicting dependencies
 - New structure and management of repositories created many minor problems
- **Time**
 - **One full cycle (patch, build, repository, cert install, test) took too long**



- **Configuration management**
 - Goal: uniform interface for site admins
 - YAIM *wrapper* around gLite configuration tools
 - Main focus on “translation”
 - Functional overlap (like gridmapfile conf.) had to be managed
- **Wrapper needed much testing**
 - Configuration management affected certification (slowed it down)
 - At the end we missed some bugs.
- **gLite scripts needed patches and functional enhancements**
 - Slow turnaround
 - Missed some functionality gaps of the underlying tools (glite-CE)
- **“Double headed” site layout (gliteCE + glite-LCG-CE)**
 - More complex than in previous releases
 - Created confusion



Steps to improve

- **Speed up the process** <----- **Key!!!!**
 - We don't get anything right the first time
 - Simplify the build/integration process
 - Define acceptance criteria for components
 - Readiness for integration, testing, pps release, production
- **Make the process more robust**
 - Formal tracking or workflow
 - Status of release generated from workflow traces
- **Ensure reliable link between bugs-patches-releases**
- **Avoid communication channels with less than clear scope**
 - Move away from rollout list
 - More formal communication with ROCs and sites
 - A bug is not a bug until it has a bug number
- **Configuration management**
 - Less wrappers around wrappers around wrappers
 - But how?
- **Uniform Documentation**

- **Process**

- Draft for new process

- <https://edms.cern.ch/file/724371/1/EGEE-MSA3.2-724371-v0.0.doc>
 - Will be presented at the SA3 all hands meeting (June 21st, 22nd)
 - Jointly developed by all stakeholders:
 - Next meeting with JRA1 on June 21st (closed session)
 - Proposal for changes in bug tracking (Joachim Flammer drives this)
 - *one tracker*
 - *back link to GGUS tickets*
 - *additional fields to ensure bug/patch/rpm version links*
 - *access to Savannah data base??*



- **Testing**
 - Inventory and list of missing and existing tests (available in wiki)
 - Will be presented at the SA3 all hands meeting (June 21st, 22nd)
 - To implement the missing tests has to be distributed activity
 - SA3 at CERN can integrate, coordinate, and run tests

 - Guideline for writing tests
 - Draft exists

 - Splitting tests to support multilevel approach
 - Ping tests
 - System wide tests
 - Stress tests
 - Deployment tests (ROCs)
 - Performance tests ----> outside participation
 - Interoperations tests ----> OSG soon, ARC and Unicore later

 - Integration of tests into ETICS
 - First 30 test cases have been ported
 - Feedback has been given to ETICS
 - More after ETICS increases functionality

- **Integration**

- moving all components with a long term future into the ETICS build system
 - LFC, DPM, BDII already ported
 - Moved to one CVS system
- Minor adjustments on how repositories are populated to speed up process
- Scripting for creation of material that is needed after certification
 - RPM-lists, tarballs.....

- **Cert testbeds**

- We research ways to speed up the upgrade/ re-install process
 - And detect problems as early as possible (upgrade/test/upgrade/test..)
- Virtualization is a good candidate
 - Prototype by Andreas, Di, and Dimitar is now available(based on XEN)
- External testbeds for interoperation and platform support needed
 - SA3 all hands and OSG/LCG operations meeting

- **Documentation**

- Split between tool description, site layout guide, and upgrade/install documents
- Move more to wiki based documentation
- Remove outdated documentation
- Uniform format/naming convention
- Work with the PPS on documenting trouble shooting guides
- Re-organize our web pages

- Not clear that we can cover all of this, need help.

- **Communication**

- gLite-announce + news page on gLite and CIC on duty page
 - Exists, policy decision
- Rollout -----> discussion forum
- More effective meetings
 - gLite-3.0 release prep. meetings has been merged with the EMT
 - *handling technical aspects of upgrades, fixes etc.*

- **gLite-3.0 was in time ready for release**
 - Despite CHEP, and kick offs (ETICS, EGEE-II, egee- extension projects...)
- **gLite-3.0 contained agreed components**
 - With almost full functionality
- **Merging 2 stacks and processes was as hard as expected**
 - Will take some more time to be transparent
- **New process definition and tool integration are progressing well**
 - CVS, tracker, build....
 - Addresses communication problems within release prep and with developers
- **Testing is still not where it should be in terms of coverage and ease of use**
 - Structural progress has started
 - Outside contribution needed (enormous amount of work)
- **Communication before and during rollout & site support has to be improved**
 - SA1/3 have to come up with better concepts
- **Documentation**
 - Frightening amount of work
 - Rolling out bug fixes and critical enhancements still have priority
 - Consolidation effort needed (1-2 members of the team 3 months)
- **Configuration management**
 - The “Russian Doll” approach proved to be less than ideal.



Enabling Grids for E-scienceE

Move to SL4

www.eu-egee.org



- **SL4 32 bit**
 - can be used with VDT.1.2
 - can be used with VDT 1.3
 - The move should be simple for C code, for C++ some difficulties are expected because the compiler is more strict

- **SL4 64 bit**
 - VDT.1.2
 - IA 64 no known issues
 - IA86 64 (opteron) and 64 bit Pentium style CPUs
 - *can't be moved with VDT1.2 to 64 bit.*
 - *open-SSL (v0.96) has problems with the underlying libs.*
 - VDT1.3
 - Can be made to work with all 64 bit versions.

- **VDT 1.3 has currently two known problems**
 - open-ssl version
 - affects the mail address extraction from the grid-map file
 - *perl script that creates the file (under out control).*
 - *order of a days, including testing*
 - Perl base classes used by the job manager scripts have changed
 - job manager scripts need to be adjusted.
 - Expert will need 2 week
- **64 bit move depends on work needed on the gLite middleware**

- **Builds of gLite components on SLC4 are regularly done on ia32 and x86_64 architectures**
- **Build logs are available at:**
 - <http://lxb2072.cern.ch:8080/etics>
 - <http://lxb5588.cern.ch:8080/etics>
- **Not all components are currently building**
 - bug reports are sent to the developers
- **Example issues include:**
 - Need new version VDT/Globus. Currently used version (1.2.2) doesn't work on 64bit platforms.
 - Pre-release version VDT 1.3.10 is under evaluation
 - A number of external dependencies have to be compiled/ported first (classad, fcgi, xerces, xalan, etc)
 - Incompatible changes in some system libraries (libxml? TBV)

- **Continue to build for 32 and 64 bit**
- **Sequence**
 - 1: SL4 32bit with VDT 1.2
 - 2: SL4 32bit with VDT 1.3
 - 3: SL4 64bit (opteron) VDT 1.3 (affect more sites)
 - 4: SL4 64bit IA64 VDT 1.3
- **LAL works on port of LCG-2-7-0 to SL4**



Enabling Grids for E-scienceE

Some slides from the postmortem

www.eu-egee.org





Enabling Grids for E-science

Bug-Issue-State-Progress Tracking (before and after the release)

www.eu-egee.org



- **During release preparation:**
 - Savannah
 - LCG components
 - *Release relevant and general problems mixed*
 - gLite components
 - *Release relevant and general problems mixed*
 - Pre-production mailing list
 - Open issues via Wiki Page
 - Here issues are jointly **listed**
 - bug numbers **referenced**
 - Status changes **manually updated** -----> **consistency with bug trackers?????**
 - Test status
 - Some components via Web result page
 - Some components via e-mail messages (works/doesn't)
 - Impossible to link test results with status of release
 - **Situation improved with additional Wiki table for test tracking**
 - *Manually updated (trust?)*

- **During Rollout:**
 - Savannah
 - LCG components
 - *Release relevant and general problems mixed*
 - gLite components
 - *Release relevant and general problems mixed*
 - GGUS ticketing system
 - Install problems
 - User problems
 - Software problems
 - Configuration problems
 - **Answered or transferred to Savannah**
 - *Who closes when what?*
 - Mailing lists & private communication
 - We did not manage to stick to the rule
 - *“If it is not in a tracker it is not a problem”*
- **Never a clear, visible status of the problems -----> Doubt, Rumours,**

- **In depth research needed to answer questions like this:**
 - Which bugs are fixed in gLite-3.0.0-RC2?
 - Is bug #1216 fixed in gLite-3.0.0-RC2?
 - Which tag version/RPM version fixes bug #1786?
 - New bug opened during certification outside the cert process
 - Does this affect the version on the certification testbed?
- **Some changes underway:**
- **Merging all Savannah tracker**
 - With additional information
 - Need for Savannah “data mining tools”
 - Security issues (all DB tables of all project owned by one user)
- **New integrated process (in progress)**
- **GGUS ticketing not resolved**

Communication (?)

- **Project management -> release team**
 - Release team was aware of the target release date (May 1st) and derived dates
 - -6 weeks pre-pro (was missed!!!)
 - End user documentation team assumed 1st of June
- **Project management -> experiments**
 - Different views in the experiments about what will be in the release and how long this can be negotiated
 - ATLAS (gLite-CE not needed)
 - Release date clear
- **Project management/ release team -> developers**
 - **Release date took several of the developers by surprise**
 - Assumed 1st of June as closing date
 - Totally unaware of the date
- **Release team -> ROCs**
 - Pre-release before eastern was not seen as an indication that the release will come on time
- **Release team/ ROCs -> sites**
 - Most of the sites are not aware of the release schedule

- System Elements:

- Release Coordinator
- Developers
- Integrators
- CertTestBed Manager(s)
- Testers

This has to be more formalized
(+ common sense)

- Problem statement:

- All have to *synchronize some* of their activities with each other and some need to have an *aggregated*, up to date, *status* view
- Most problems need more than one iteration on each comm. channel
- “some of their activities” -----> 100% are 100% active
- -----> communication handshakes fail without being noticed for hours
- **This was extremely expensive**
 - Time
 - Source of conflicts

- **Spreading the gospel (Announcing releases)**
 - Shotgun approach
 - Rollout, glite-announce, ROCs, forward to all management lists ...
 - Rollout list
 - a discussion forum where all kinds of interesting ideas are floated
 - new releases are treated like this
 - frequently ignored/ seen as an recreational activity
 - Release pages
 - Static, only visited when people got the signal
 - ROCs
 - Don't reach the extension projects
 - Different speed limits
- **Need to use gLite-announce as a one way channel**
- **Convert rollout to a discussion list/ newsgroup**
- **Maintain an interesting news page**

- **Rumors and scaremongering**
 - A lot of time and energy is devoted to verify or falsify rumors
 - Constant review mode for SA3/SA1
 - Samples:
 - No gLite-CE ever worked
 - gLite-3-0 was release without testing
 -
 - This had a severe effect for the “morale” in the trenches
 - People try to avoid making mistakes by all means -----> productivity meltdown

- **We have to keep a daily, public release log**



Enabling Grids for E-scienceE

Documentation

www.eu-egee.org



- **Problems**
 - Too much
 - Not enough
 - New and old
 - Quality
 - Hard to find
 - Wrong/outdated
 - Too aggregated
 - No “one stop” shop
 - Branding
 - Hard to maintain
 - Documentation from partners not linked

 - Not read

- **Too much**
 - We basically provided LCG-2.7 + gLite-1.5 + gLite-3.0 specific
 - Separated already information on non included components
- **Not enough**
 - Material on boot strapping a site has to be added again
 - Was there in early LCG-2 releases
 - What is a site, central services etc.
 - Which notes, hardware requirements, how to start
- **New and old**
 - Site managers get confused by finding different instructions for the same components (tarball UI)
- **Quality**
 - Need to proof read documentation and check links before release
 - Wiki material has to be reviewed and removed/updated

- **Hard to find**
 - Documentation is scattered between several pages
 - Several Wiki sites with information
 - Guide to documentation needed
 - short summary for all documents
- **Wrong/outdated**
 - See “quality”
 - Good example “Testing your site”
 - Review!!!
- **Too aggregated**
 - YAIM description and install guide is mixed
- **No “one stop” shop**
 - We need a high level intro
- **Branding**
 - No uniform naming of documents/ services
 - No classification of documents
 - No common appearance

- **Hard to maintain**
 - Some of the documentation can be only build by an elite
 - Can have site effects
- **Documentation from partners not linked**
 - Example:
 - Material that covers other batch systems
 - Special setups via NAT
- **Not read**
 - Especially the “Release Notes” are not read by half of the sites



Enabling Grids for E-scienceE

Merging of the two stacks

www.eu-egee.org



- **Harmonization of dependencies**
 - Internal dependencies
 - Old/new version of voms libraries
 - Fetch-crl vs. edg-system-utils
 - External dependencies
 - E.g. different MySQL version
 - Especially User Interface merge took some time
 - Stops still co-installation of some LCG and gLite modules on one node
 - **Status: Partially solved (will need more iteration)**
- **Harmonization of naming convention and creation of metapackages**
 - e.g. glite-VOMS_mysql
 - caused old version still being distributed
 - **Status: Solved (some still might to be renamed)**

- **Handling of non-freely distributable dependencies**
 - J2SDK
 - gLite stack uses j2re – can LCG stack be moved as well?
 - J2sdk freely distributable in future?
 - oracle-instantclient
 - gLite stack: no rpm dependency, checking during configuration
 - LCG stack: rpm dependency
 - Does oracle always provide rpms?

→ **Status: Solved (oracle instant-client dependencies introduced)**

- **Naming convention for RPMs**
 - Both stack use different naming conventions
 - Should we move to glite-xxx in general?
 - Platform specification should be homogeneous

→ **Status: Not solved (but not critical)**

- **Conflicting dependencies with SLC3**
 - Conflict of lam/mpi requires un-installation of RPMs from standard SLC3 installation
 - **Status: Partially solved**

- **Problems due to handling of two software stacks**
 - duplication of rpms ...
 - See repository presentation
 - **Status: Partially solved (Merging in ETICS)**

Joachim Flammer

- **Different ways of working**
 - **Different names for the same things causing misunderstanding**
 - Deployment modules ⇔ Metapackages
 - **Stress situation**
 - Sometimes unnecessary aggressive communication
 - Communication could be improved
- **Status: Partially solved ...**

Joachim Flammer

- **Procedure has been shaped during the release**
 - Several changes in the way how to create combined release
 - Unclear mechanism of update strategy
 - Several changes in the way how to request/create update

→ **Status: Solved (for the moment ...)**
- **Savannah**
 - Does not completely represent our workflow
 - Missing fields makes tracking of bugs difficult
 - Developers often not updating bug status / providing necessary information

→ **Status: Partially solved (Savannah update)**

- **Configuration**
 - Harmonize on one configuration
- **Documentation**
 - Harmonize to one stack of documentation
- **GGUS**
 - Huge amount of mails
 - How to handle support