



Deployment work at CERN: installation and configuration tasks

WP4 workshop

Barcelona project conference 5/03

German Cancio CERN IT/FIO

The logo consists of the word "Data" in orange above the word "GRID" in black. The "G" and "I" in "GRID" are partially overlaid by a blue globe icon with latitude and longitude lines.

Data GRID Agenda

- ◆ Why, what, when
- ◆ CERN enhancements
- ◆ Issues
- ◆ Summary



Why?

- ◆ CERN Computer Centre is being scaled up in capacity for handling LHC capacity requirements
 - *O(10K) nodes expected in 2006/7*
 - *Currently, ~ 1.5K*
 - *Concentrating on Linux on Intel Hardware*
- ◆ Current management software in production was written (many!) years ago, for a reduced number of heterogeneous RISC UNIX nodes
 - *No central configuration mechanism (>20 databases holding configuration information!)*
 - *Insufficient automatization (manual steps required for node installation and configuration)*
 - *Scalability issues: eg. software distribution*
 - *Outdated technology obsoleted by upcoming de-facto standards (eg. RPM)*
- ◆ *CERN's aim participating in WP4 was to find solutions to this problem!*



What?

◆ Central Configuration Database

- CDB set up for hosting configuration profiles of the main CERN interactive and batch services (LXPLUS and LXBATCH)
- Templates created for hardware information (currently: new SEIL nodes):
- System information including partition layout
- Software packages configuration structured in templates (RH distribution, CERN CC packages, EDG packages)

◆ Software Distribution

- SWRep: hosting RPM packages
- SPMA: Replaces legacy software update mechanisms (ASIS, rpmupdate, SUE security)

◆ Node configuration

- CCM: access to CDB information via NVA API
- NCM: will replace the CERN 'SUE' node management framework



When?

◆ SPM (SPMA and SWRep)

- Started pilot for phasing out the CERN SW distribution systems on batch worker nodes for LCG-1
- Using HTTP as package access protocol (scalability)
- > 100 nodes currently running it in production
- Deployment page: <http://cern.ch/wp4-install/CERN/deploy>

◆ CDB

- Database in production status, holding site-wide, cluster and node specific configurations for ~ 350 clients
- Using WP4 Global schema with site specific extensions
- Enhancements: GUI (next slide)

◆ CCM

- About to be deployed

◆ NCM

- As soon as it gets available with components



CERN Enhancements (I)

- ◆ GUI for CDB
 - **PanGUIIn**: graphical front end for browsing, modifying, validating and committing CDB templates
 - *Generic development (no site dependencies): Can be feed back into EDG*
- ◆ Configuration access shortcuts and migration
 - *CCM provides a 'low-level' API for traversing the node configuration profile tree*
 - *CCConfig: 'high-level' API giving access for recurring information chunks (partitions, package lists, network information, cluster name, etc)*
 - *CCConfig allows to access CDB info from the current node configuration system (SUE)*
- ◆ Server clustering solution
 - *For CDB (XML profiles) and SWRep (RPM's over HTTP)*
 - *Replication done with rsync*
 - *Load balancing done with simple DNS round-robin*
 - *Currently, 3 servers in production (800 MHz, 500MB RAM, FastEthernet)*

Location Configuration Edit Template Tools Help

Status

profile

- profile_lxb0201
- profile_lxb0202
- profile_lxb0203
- profile_lxb0204
- profile_lxb0205
- profile_lxb0206
- profile_lxb0207
- profile_lxb0208
- profile_lxb0209
- profile_lxb0210
- profile_lxb0211
- profile_lxb0212
- profile_lxb0213
- profile_lxb0214
- profile_lxb0215
- profile_lxb0216
- profile_lxb0217
- profile_lxb0218
- profile_lxb0219
- profile_lxb0220
- profile_lxb0221
- profile_lxb0222
- profile_lxb0223
- profile_lxb0224
- profile_lxb0225
- profile_lxb0226
- profile_lxb0227
- profile_lxb0228
- profile_lxb0229
- profile_lxb0230
- profile_lxb0231

```

# CERN RH7.3 BASE packages with security updates
#
# Based on KickStart profile:
# %packages
# @ CERN Server
# @ Software Development
# minicom
# rxvt
# enscript
#
# and removing
#   tetex, tetex-latex, linuxdoc-tools,
#   anacron, logwatch, tmpwatch
#
# (Does not include @GNOME nor @KDE)
# RESPONSIBLE: German Cancio <German.Cancio@cern.ch>
#
template pro_software_packages_cern_redhat7_3_release;

"/software/packages"=pkg_add("4Suite", "0.11.1-8", "1386")
"/software/packages"=pkg_add("CERN-HEPiX-scripts-X11-com")
"/software/packages"=pkg_add("CERN-SUE", "7.2-1.50", "1386")
"/software/packages"=pkg_add("CERN-SUE-starter", "7.3.cern")
"/software/packages"=pkg_add("CERN-compat-locallinks", "1")
"/software/packages"=pkg_add("CERN-images-X11", "1-5.cern")
"/software/packages"=pkg_add("CERN-images-base", "1-5.cern")

```

Welcome | pro_software_packages_cern_redhat7_3_inte

pro_software_packages_cern_redhat7_3_release | profile_b

Logging History:

- Started Panguin

Logger

Profile Viewer - [profile_lxb0219.xml]

https://lxserv.cern.ch/xml/profile_lxb0219.xml

- profile
 - software
 - hardware
 - cpus
 - harddisks
 - ram
 - cards
 - location = 513V 19-M504
 - system
 - network
 - domainname = cern.ch
 - hostname = lxb0219
 - interfaces
 - cluster
 - filesystems
 - ccdbname = PLUS

Info:



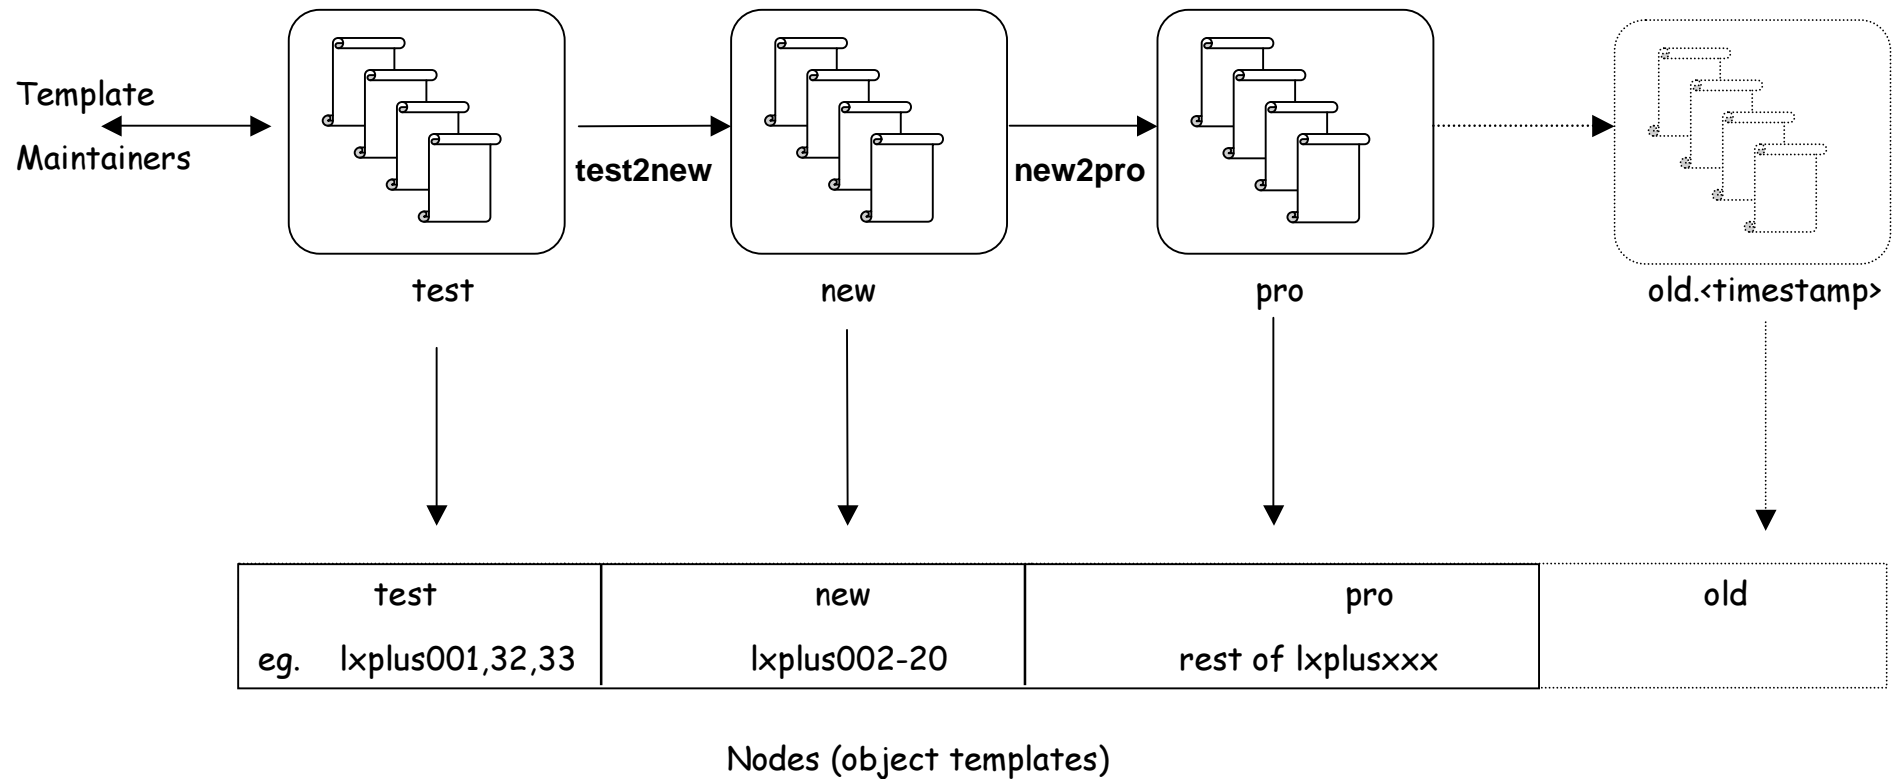
CERN enhancements (II)

◆ Release life cycle in CDB

- A high-level configuration change may work on 99.9% of the nodes but break on some
- Even if CDB keeps history information and is able to roll back, it cannot easily keep multiple configurations in parallel
- Solution: keep 'test', 'new', 'pro' and 'old' templates **simultaneously** in CDB; and migration scripts which perform transitions from 'test'->'new'->'pro'->'old'.
- Most node configs point to 'pro'; test nodes use 'test', certification nodes use 'new'
- If a node fails with a new 'pro' configuration, it is pointed to use 'old' while understanding the problem, instead of rolling back the complete CDB



CERN enhancements (III)





Issues (I)

- ◆ Understand the current system
 - Identify current dependencies (on shared file systems, legacy services)
 - We cannot start over from scratch!
- ◆ Migration considerations
 - Production environment -> minimal downtimes!
 - Large computer centre
 - WP4 not a drop-in, turn-key solution, needs integration with existing environment
 - Produce 'glue' procedures and scripts
- ◆ Scalability and Performance on large clusters
 - SWRep: How many servers do we need for n nodes?
 - CDB: how long does it take to generate $O(100)$, $O(1000)$ configuration profiles? What are the CPU requirements?



Issues (II)

- ◆ Careful planning and execution required
 - Find appropriate time slots for phase-ins and phase-outs (one at a time!)
 - Announcements to users when non-transparent changes required
 - update operation manuals for site administrators and technicians
 - Problems and failures have much higher impact than on EDG testbeds!
- ◆ Not all R3 functionality available yet
 - ... some new procedures are harder to follow than old ones!
- ◆ Site wide deployment (outside computer centre)
 - Until now, same tools used for farms, servers and desktops
 - Desktop support kept in mind during WP4 design, but needs additional work
- ◆ WP4 deployment timescale will exceed the lifetime of EDG
 - Support and bugfixes
 - Future developments
 - Port to new platforms (RH9, IA64, Solaris...)



Summary

- ◆ WP4 software is becoming a key part of the CERN computer centre management
- ◆ CERN CC has very high requirements on stability and reliability
 - *WP4 is coping with it! (at least up to now..)*
- ◆ Most valuable feedback for EDG, both in terms of testing and functionality enhancements