



Linux in the Computer Center at CERN

Zeuthen
21.10.2002
Thorsten Kleinwort
CERN-IT

1



Overview

- Linux at CERN:
 - Past
 - Present
 - Future
- Linux at CERN: Some details:
 - "Legacy stuff"
 - Configuration
 - Installation
 - AOB
 - Outlook



(Pre-) Linux at CERN: Past

- Several private Clusters
- Few machines in each Cluster
- All types of hardware: HP, AIX, SGI, Sun,...
- Proprietary base installation (CD, tape...)
- OS independent post installation (SUE)
- OS independent software distribution (ASIS)

29 January 2007

Thorsten Kleinwort IT/FIO/IS

3



Linux at CERN: present I

- Decommissioning of RISC hardware:
 - AIX, HP, DEC, and SGI have all gone and are not supported in the Computer Center any more
- Focus on Linux (Intel) & Solaris (Sun)
 - Linux for login and batch, as desktop machines, and for disk and tape servers
 - Solaris for servers & cross check platform but no general login and batch any more
- Now:
 - LXPLUS ~ 75 nodes, LXBATCH ~600 nodes

29 January 2007

Thorsten Kleinwort IT/FIO/IS

4



Linux at CERN: present II

- **Installation & Maintenance outsourced:**
 - Done by at company (Serco)
 - Using our (old) tools
 - Big problems describing the "Service" they have to provide
 - Installation became incomprehensible
 - Administration far from being automated
- For the new Linux RedHat 7, we took back the responsibilities

29 January 2007

Thorsten Kleinwort IT/FIO/IS

5



Linux at CERN: present III

- Current installed version still RedHat 6.1
- Certification for RedHat 7.3 is ongoing, we (CC) are ready
- We (CC) have a complete new, automated installation for Linux 7.3
- We have started to use a configuration database (CCConfig)
- We have redone the monitoring

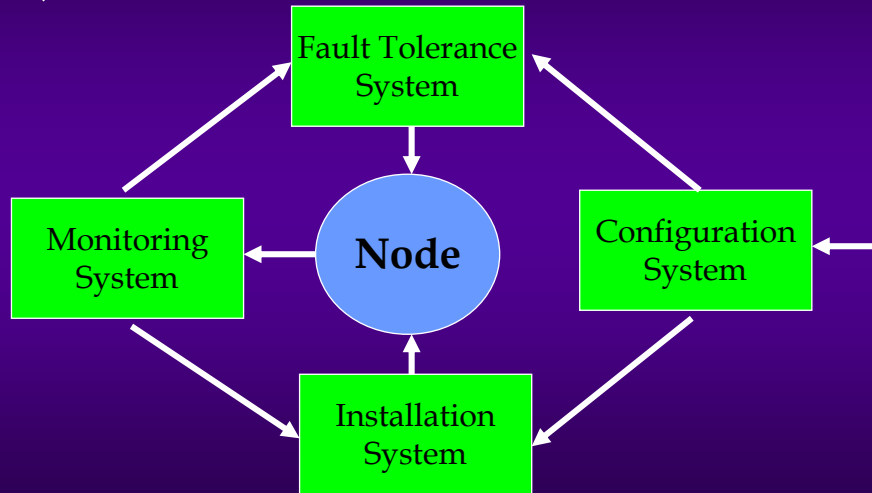
29 January 2007

Thorsten Kleinwort IT/FIO/IS

6



Linux at CERN:



29 January 2007

Thorsten Kleinwort IT/FIO/IS

7



Linux at CERN: future

- **LCG:**
 - ~10000 Linux nodes
 - Computing Grid
- **EDG:**
Split up in Tasks: WP4, Fabric Mgmt:
 - Installation
 - Configuration
 - Monitoring
 - Fault Tolerance

29 January 2007

Thorsten Kleinwort IT/FIO/IS

8



Linux: Details

- **ASIS:**
 - Was a tool for a platform independent software distribution
 - Now (on Linux): RPM based
 - Uses now system data base
- **SUE:**
 - Was a common tool for all platforms
 - Still in use, but deficiencies are apparent
 - Now only used for configuration
- **Kickstart:**
 - The RedHat tool for automatic installation

29 January 2007

Thorsten Kleinwort IT/FIO/IS

9



Linux: Details II

- **"BIS" installation:**
 - Reflected Cluster & OS dependencies
 - Was 'junked' for a better and cleaner installation
- **Configuration:**
 - New Configuration interface, *CCConfig()*
 - Conform with WP4 Configuration Task
 - Already available now as a PERL module on the node:
 - Can be used, e.g. within SUE to provide node information

29 January 2007

Thorsten Kleinwort IT/FIO/IS

10



Linux Details: III

- **Linux 7.3 in the Computer Center:**
 - Automatic generation of kickstart files
 - Boot machines with netboot or floppy
 - Install base installation with Kickstart
 - Install CERN and CC stuff afterwards with RPM
 - Configure with SUE, configuration from CCConfig
 - SUE may be replaced by the WP4 installation tool
 - Maintenance of the machine: rpmupdate and SUE run on demand, triggered by notification (No regular run)
 - Monitoring done with a prototype of WP4 monitoring

29 January 2007

Thorsten Kleinwort IT/FIO/IS

11



Linux Details: IV

- **Next steps:**
 - Upgrade our whole farms (~700) and install some new arrivals (~300) until mid next year to RedHat 7.3
 - Make installation AFS independent
 - Collaborate with WP4:
 - Provide replacement for SUE
 - Use Configuration Management for hardware and software database
 - Enhance Monitoring (Correlation engine, ...)

29 January 2007

Thorsten Kleinwort IT/FIO/IS

12



Linux Details: V

- The batch System LSF
- Some Security issues
- Handling of /etc/passwd, /etc/group
- Configuration



The batch system

- LSF: Introduced in 1997, Version 3.2:
 - Multicluster: several submission & execution cluster, due to the large number of Clusters
 - Using fixed partitions per group/experiment
- Current Version 4.2:
 - Back to one cluster (submission & execution)
 - Using "fairshare": Better utilisation
 - But slow reconfiguration times: around 15 min
- Good cooperation with Platform (Canada, UK, Germany {Munich})



Secure host information

- **Problem: How to get a private key on a new installed host:**
 - **Floppy boot: Put a key on the floppy**
 - **Network boot: Trust your network (Bootp)**
 - **Put a private/public key (gpg) on the host in an early stage of the installation**
 - **Use this key to encrypt secure information:**
 - SSH host keys
 - The header of /etc/passwd (contains crypted pwds)
 - **Framework to generate and manage keys and secure information**

29 January 2007

Thorsten Kleinwort IT/FIO/IS

15



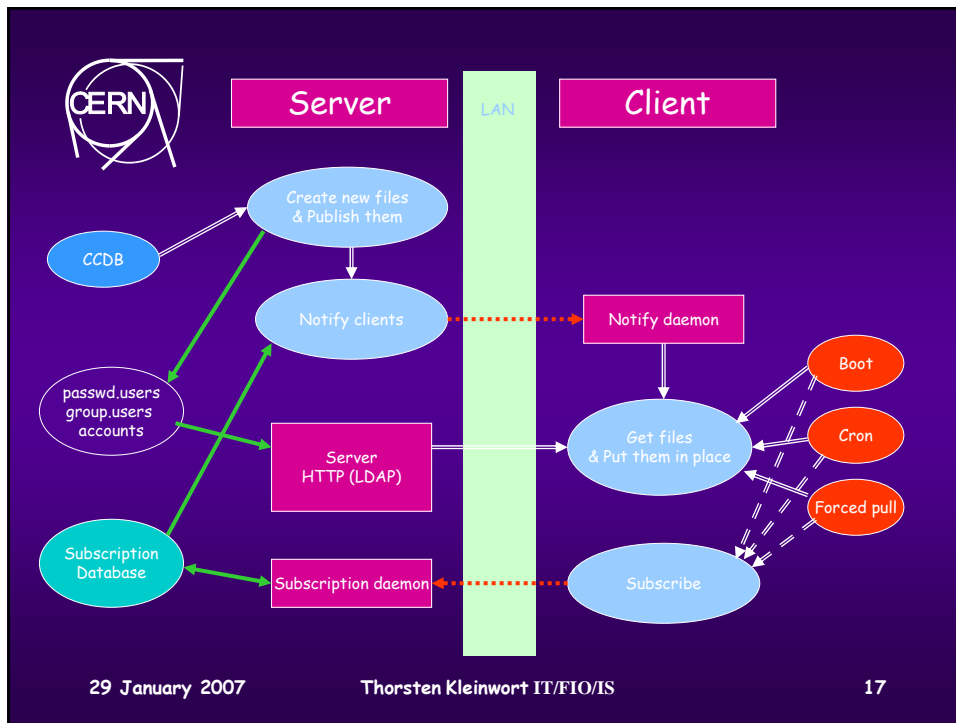
User Account Management

- **Problem: Big amount of data (~1MB), changes irregularly:
/etc/passwd, /etc/group**
- **We keep this information local**
- **Using "client poll", together with a notification mechanism, for updates**

29 January 2007

Thorsten Kleinwort IT/FIO/IS

16



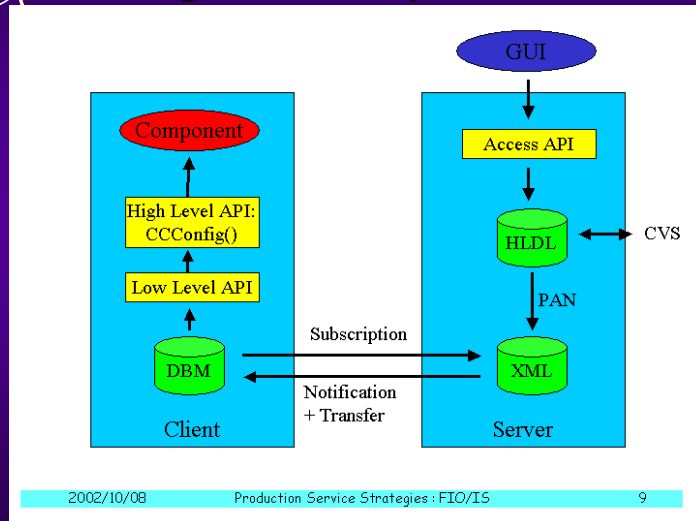
CERN Configuration

- **Problem: SUE does not have a configuration information interface**
- **Invented CConfig():**
 - **High level API for getting host information**
 - **In collaboration with WP4 configuration task**
 - **plans:**
 - Use a HLDL (High Level Description Language) for describing a host
 - Use a compiler to create XML
 - Download the XML file into local cache
 - Use low level API for config info and for CConfig

At the bottom, the date **29 January 2007**, the name **Thorsten Kleinwort IT/FIO/IS**, and the page number **18** are displayed.



Configuration: plans



29 January 2007

Thorsten Kleinwort IT/FIO/IS

19



Summary

- Due to increasing number of Linux hosts and out-of-time tools, we have redone our Linux installation
- We are a little bit ahead of the EDG project: Needed our own solutions
- But in collaboration with them
- We are now preparing for LHC (~10000)
- All tools have to be re-evaluated

29 January 2007

Thorsten Kleinwort IT/FIO/IS

20