

# Controls issues experienced during the LHC start-up: status report

Wojciech Sliwinski for BE-CO group

Special thanks to: E.Hatziangeli, K.Sigerud, P.Charrue,  
V.Baggiolini, M.Sobczak, M.Arruat, F.Ehm

LHC Beam Commissioning workshop, Evian 19-20 January 2010

# Abstract

- During 2009 the LHC control system became operational through several dry runs and injection tests and contributed to the successful LHC start-up
- Today is a good opportunity to assess the quality and performance of the various components of the controls infrastructure both software and hardware
- The presentation will give a summary of the controls issues experienced during the LHC start-up and will outline applied and planned actions needed to resolve them

# Emphasis on few critical systems

- Infrastructure – disk space, consoles, ...
- CMW – proxies, subscriptions, ...
- FESA – front-end instabilities
- RBAC – introduction of STRICT policy
- JMS – data publishing
- BE–CO Development Process

# Infrastructure – NFS disk space

- **Massive increase of the total amount of controls operational data**

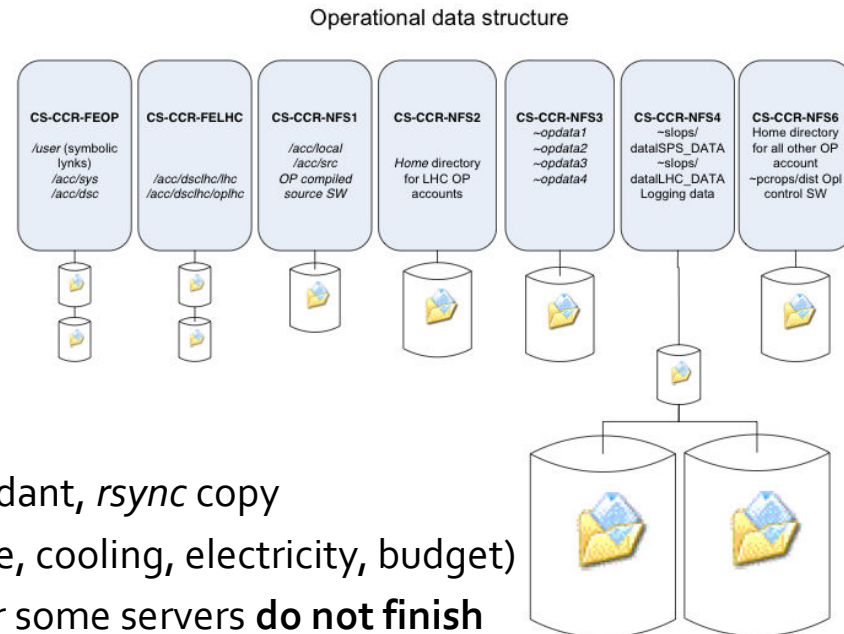
- 2005 => ~400 GB, 2010 => 4Tb
- The last 2Tb were filled in only 2 months by users and application data

- **Problem description**

- Data is duplicated **three times**: source, redundant, *rsync* copy
- Reaching the **physical limits** of the CCR (space, cooling, electricity, budget)
- The tape backups take **too much time** and for some servers **do not finish** during the night!

- **Planned action**

- CO-IN **analyses** the disk space usage and **will propose** a long-term solution (February 2010) based on **new storage** technology from HP



# Infrastructure – Operational Consoles


## ■ High load on consoles

- **Problem:** Consoles became **slow** and **not responsive**, even with only few apps. running
- **Performed actions**
  - Introduced visual **indicator** in CCM
  - Identified an issue (Vistar zombie processes) and **fixed it promptly** before the start-up
  - Memory of all LHC consoles was **doubled** (4GB)
  - EIC consoles were upgraded to **the latest HW** configuration

*Since then ... green indicators and .... happy people*

## ■ Java cache for applications on consoles

- **Problem:** Applications started via CCM may use older jar versions from the local cache
- **Performed action:** Provided context menu to **completely clean** the local Java cache
- **Planned action:** Validate with **the latest Java** version installed in January 2010



# CMW – Unstable Proxies

## ■ Problem description

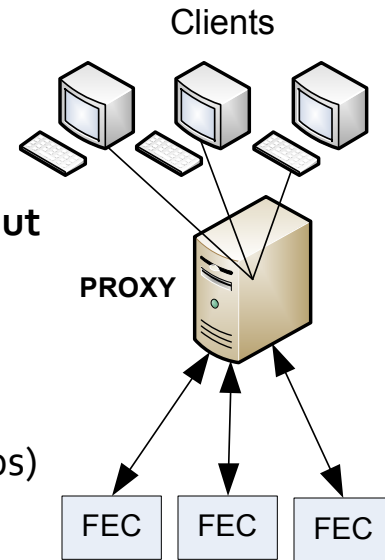
- Under high load Proxy **did not respond** promptly to a calling client
- Client used to interpret it as Proxy being **down** or **unreachable** or **timeout**
- Observed particularly when setting up new subscriptions
  - E.g. in Fixed Displays applications (e.g. LHC Page1)

## ■ Performed actions

- Performed several iterations to solve the problem (with OP & eqp. groups)
- Improved significantly level of **concurrent processing** in Proxy
- Increased the **client timeout** for RDA calls
- Added dedicated **diagnostics** to monitor current client connections
- **Validated new Proxy** during the dry-run of 13th Jan. on selected BI front-ends (BTV, BCT)

## ■ Planned actions

- Provide **statistics of traffic** in order to detect early the overloaded Proxies
- Integration of statistics/diagnostics info with the **DIAMON** facilities



# CMW – Proxy Redirection

## ■ Problem description

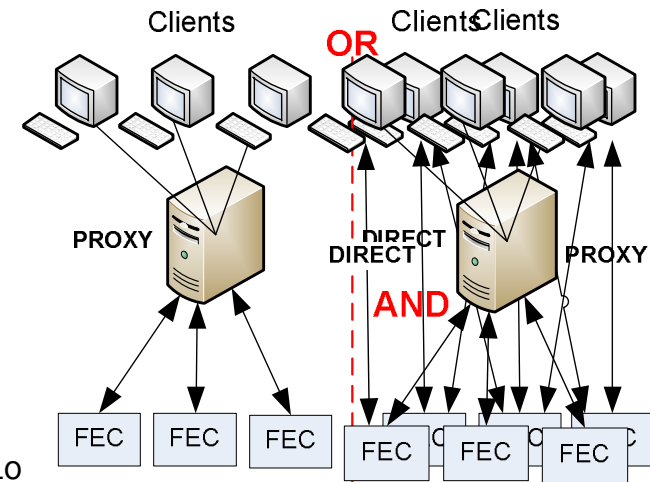
- It was **not possible to mix** direct & Proxy connections from the same client application
- Proxy redirection was **"all or nothing"**
  - Either all via **the same Proxy** or all connect **directly**

## ■ Performed actions

- **New centralized** architecture for Proxy resolution
  - New CMW-RDA (C++ 2.8.9 & Java 2.8.5) released on 07.01.2010
- Now **possible to mix** direct & Proxy connections from the same client application
- Redirection granularity set on **front-end server** level
- Proxy configuration is **database driven** (Configuration DB)
  - If proxy defined for a given front-end server, by default all client connections go via this Proxy
- Experts **can force** old behavior via system properties

## ■ Planned actions

- New **expert** feature: **disable Proxy handling** for selected front-end servers directly from the expert application (requested by BI)
- Contact eqp. groups and **configure Proxy redirection** for their front-end servers



# CMW – Unreliable Subscriptions

- **Problem description**
  - **Disconnected** subscriptions caused **loss of important data**
  - Observed in several **critical systems**: BLM concentrators, BCT, RF
  - Problem was also impacted by some **faulty** network switches
- **Performed actions**
  - **Significant troubleshooting** and investigation to locate the cause of the problem
  - Identified faulty network switches **were replaced**
  - Identified CMW timeout was **too short** for the front-end servers to process the notifications
    - Increased CMW timeout to 4 seconds to give **more time** to server **to push the notifications**
    - **Reduced** the subscription disconnections & drastically improved the situation
  - Identified the **configuration parameters** (CORBA) which were impacting the performance
    - New set of CORBA configuration parameters **were deployed** on the front-end side
  - The major clients (BLM,BCT, ...) **highly appreciated** the efforts & final outcome



# FESA – Crashing Front-End Servers

## ■ Problem description


- Front-end server was occasionally **crashing** in several situations
  - **Many clients** were **directly subscribed** to a front-end publishing data with a high rate
  - At the front-end **start-up** time when many clients were **re-establishing** the subscriptions
- Observed several times for **RF FESA** servers

## ■ Performed actions

- FESA team investigated the issue and managed to locate the cause of the problem
- The bug was fixed and released as a patch to FESA 2.10 on 14.01.2010
- The **new fixed** FESA version was validated together with the RF team
  - RF front-end server was stressed with many clients connected and publishing rate 10 times higher than in operational mode
  - Front-end was restarted to test the behavior at the **start-up time**
- **Stability** of the FESA servers **was improved**

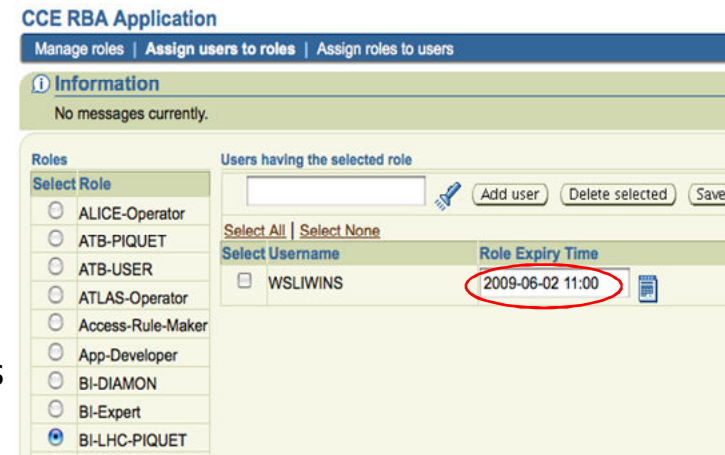
# RBAC – introduction of STRICT policy

- From March 2009 ongoing campaign to introduce STRICT policy in LHC
  - Several dry runs and tests in collaboration with OP & equipment groups
  - End July 2009: default RBAC policy for LHC => STRICT
    - RBAC token mandatory
    - Protection of SET operations on properties
  - Main goal for the start-up => STRICT mode everywhere
  - All LHC applications and majority of equipment became RBAC enabled
  - Already in STRICT: PO, BI (BLM, BPM, BCT, BTV, ...), BT, Collimators, BIS, QPS, most of RF
  - TODO: Rest of RF (~55 FECs), BI(2 FECs), PVSS-Cryo (~69 FECs), PVSS-Survey (10 FECs), IPOC (8 FECs)



moreover ...

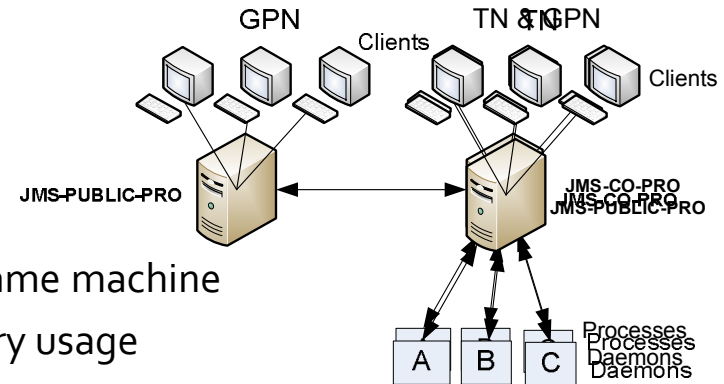
- New RBAC piquet ROLE with Expiry Time
- Access maps dependent on the Operational Mode
- New Central CMW Configuration server
  - Distribution of the Operational Mode and RBAC policy
- **NEW!** Proxy extension: RBAC token for get/set comes from a client but for monitor Proxy token is used



# JMS – data publishing

## ■ Problem description

- Overloaded brokers **stopped publishing** the data
  - It was seen by BLM and Logging
- Both brokers **CO (TN) & Public (GPN)** were on the same machine
- **Too many** client connections => high CPU & memory usage
- Machine (hardware) **was not able** to process all requests in a timely manner



## ■ Performed actions

- **Physical separation** => **CO** broker moved to a new 16 core machine, **Public** broker stayed on old machine
- Brokers are constantly **monitored** for performance

## ■ Possible improvements

- Need to evaluate different broker configurations (Load Balancing & Hot Backup)
- Usage of **dedicated brokers** for „*special*“ clients e.g. for Concentrators

# Improving BE-CO Development Process

- Identified issues
  - Poorly **coordinated** and **communicated** upgrades
  - Upgrades **break existing** controls software (side-effects, bugs)
- We cannot **freeze** the controls system and **stop** all our developments
  - Operations request new /enhanced functionality
  - Bug fixes
- Better development and deployment procedures
  - **Coordinated** development, testing and releases for the controls system core (FESA, CMW, RBAC, JAPC, LSA)
  - Changes must be **motivated, prioritized** and **traced** in **JIRA**
    - Cooperation of the clients (OP, eqp. specialists) is essential for this
  - **Testbed** to validate Control System **core**
  - Rules and tools to manage **dependencies** and guarantee **backward compatibility**
- To be presented soon in the BE-CO Technical Committee

# Summary



- Overall, the Controls system was stable and performed reliably during the LHC start-up and commissioning
- During operations certain shortcomings were identified in the Controls infrastructure which were thoroughly followed-up and solutions were rapidly applied
- In 2010 the emphasis will be on:
  - Consolidate the communications infrastructure to withstand reliably the load
  - Push for the remaining equipment to go to the RBAC STRICT policy
  - Develop and deploy a solid development and release process across the whole Controls infrastructure



# DIAMON – inconsistent host status info.

## ■ Problem description

- State of a host displayed as **correct** (green) but sometimes one of the servers (running on that host) is in an **error state** (red)

## ■ Planned actions

- **Known bug:** after searching for a host, host state **does not change** even if one of the server processes goes down (host status should become also red)

The screenshot shows the DIAMON console interface. On the left, a tree view shows the hierarchy of hosts, with 'Leir (LEI)' highlighted in red. The main area displays a table of host status information. The table has columns for host names and their states. The 'dleiect' host is highlighted in red, indicating an error state. A dialog box titled 'Question' is open, asking 'Do you really want to send a 'Repair' command to 'Host dleiect'?' with 'Yes' and 'No' buttons. Below the dialog, the 'dleiect (LEIR ELECTRON COOLING CONTROL)' host details are shown, including a 'Repair' button and a 'Properties' tab. The console also shows a message log at the bottom.

Host Name	Status
cfv-363-ctrl1	Green
cfv-354-ctmccm	Green
cfv-363-hjlei	Green
cfv-ccr-ctmcc	Green
cfv-ccr-ctmemm	Green
cfv-ccr-ctmmaina	Green
cfv-ccr-ctmmainb	Green
cs-ccr-cmw1	Green
cs-ccr-cmw2	Green
cs-ccr-cmw3	Green
cs-ccr-colsa	Green
cs-ccr-ctm	Green
cs-ccr-feop	Green
cs-ccr-isa2	Green
cs-ccr-isa3	Green
cs-ccr-laser1	Green
cs-ccr-laser2	Green
cs-ccr-logging0	Green
cs-ccr-isa2	Green
cs-ccr-oasis	Green
cs-ccr-tomo1	Green
cs-ccr-veir	Yellow
dleiaos1	Green
dleiaos2	Green
dleiaos3	Green
dleibgen	Green
dleibump	Green
dleieec	Red
dleigrd	Green
dleim1	Green
dleim2	Green
dleikick	Green
dleipw	Green
dleipow1	Green
dleipow2	Green
dleipow3	Green
dleirf	Green
dleirf2	Green
dleirfd	Green
dleirfu	Green
dleitm	Green
dleitraf	Red
dln3aos1	Green
dmcsvos	Green

# LHC ALARMS – Known Issues

- **Problem description**
  - Flooding of the Alarm screen with **hundreds** of alarms (many not relevant to Operations)
  - OP are **not looking** to the Alarm screen when there is a **problem**
- **Planned Actions**
  - A **new LASER admin tool** to aid OP to manage the **configuration** of alarm categories (PC alarms) is under development by CO for 2010
  - **Clean-up** unnecessary alarms