



The CERN Computer Centre Supervision Project

Helge Meinhard / CERN-IT

HEPiX Catania

2002/04/18

<http://cern.ch/it-proj-ccs>

Outline

- Motivation
 - Strategy, people, history
 - Objectives for 2002
 - Prototype 0
 - Data gathering, internal aspects, visual aspects
 - Live demonstration
 - Next steps
 - Communication
 - Relationship with WP4
 - Conclusions
-

Motivation (1)

- SURE system used successfully for operators and (partly) service managers
 - But very close to its limits in terms of extending functionality, and scaling
 - No way to integrate another few hundred boxes
- PEM started with high ambitions, now defunct
 - Prototype with web page presenting services (JvE) very useful, maintained until March 2002
- EDG WP4 again very ambitious
 - Not clear that functionality required at CERN will be provided in time
 - Prototype scheduled for September 2002
- Summer 2001: Clara Gaspar starts investigating PVSS for LHCb EF control

Motivation (2)

- Something needed that offers required functionality asap
 - CCS project launched late in 2001
 - Mandate: develop a flexible, scalable system for supervising (i.e. monitoring and controlling) all computer centre equipment and services
 - Requirements for framework functionality no different from those by detector control systems
 - Tool of choice (JCOP) for DCS systems of experiments: PVSS from Austrian company ETM
 - Try and use this to start with
-

Strategy

- Requirements collection by PEM and WP4 still very valid
 - No need for another time-consuming round
 - New implementation taking into account functionality and constraints of PVSS
 - Deliver series of prototypes with increasing functionality
 - Early user feedback facilitated
-

People

- Core team

- Pete Jones (50%)
- Jan van Eldik (50%)
- Fabio Trevisani (50%)
- Hugo Cacote (100%, started in 02/2002)
- Helge Meinhard (70%)

- Advisors

- Olof Barring (link to WP4)
 - Sylvain Chapeland
 - Rod McKenzie
-

History

- 2001/11/01 Informal launch of FIO group and FIO-CCS section
 - 2001/11/15 Brainstorming meeting about CCS project
 - 2001/11/27 One-day course “Overview of PVSS II”
 - 2001/12/10 – 2001/12/14 Five-day course on PVSS II
 - 2002/01 Start work on Prototype 0
 - 2002/03 End work on Prototype 0
-

Objectives for 2002 (1)

- Prototype 0 – March 2002
Learning experience with PVSS
 - Re-implement functionality of PEM prototype with PVSS
 - Lxbatch, lxplus, lxshare
 - No changes on node side (i.e. UDP-based PEM transport)
 - Proof-of-concept for automatically generating alerts for the operators
-

Objectives for 2002 (2)

- Prototype 1 – June 2002
 - Improved protocol between nodes and PVSS system
 - Improved handling of configuration info
 - Added functionality for Linux clusters
 - Including perhaps triggering automatic actions
 - Person-power permitting, add more services
 - Decision on PVSS suitability
-

Objectives for 2002 (3)

- Prototype 2 – December 2002
 - subject to PVSS suitability
 - Add more services
 - Iterate on objectives for P0 and P1 if needed
-

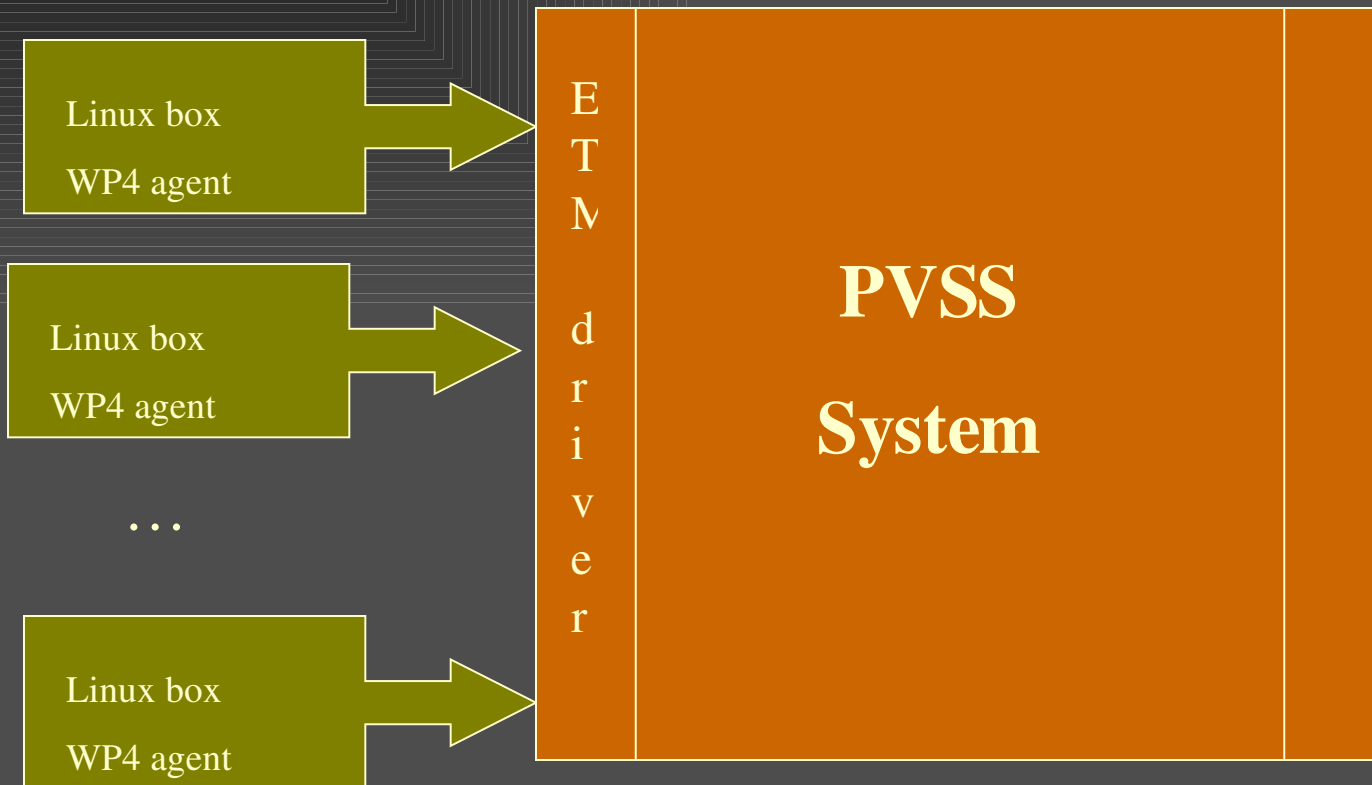
Work done: Prototype 0

- Data gathering
 - Internal aspects
 - Visual aspects
-

Data gathering (1)

- PVSS does not supply monitoring clients for computers
 - Re-use of the WP4 agents (client side)
 - ETM supply a driver prototype (server side)
 - Linux PCs (lxbatch, lxplus, lxshare, disk server clusters, ...)
-

Data gathering (2)



Internal aspects (1)

- Definition of the data structure
 - Basic datapoint types: FileSystem, File, Daemon, ...
 - Derived datapoint types: LSF, Registration, MyPC, ...
 - Instantiations of MyPC look like
 - `lxplus001.System.AFSversion`
 - `lxplus001.LSF.sbatchd`
 - `lxplus001.Registration.passwd_header`

Internal aspects (2)

- Configuration generated by an external program
 - Machine info extracted from existing databases (CS, BIS, ...)
 - Uploaded using the PVSS ASCII manager
 - 1000 machines with up to 100 metrics each
 - Datapoint elements configured for.
 - Driver specifics (the source of the data)
 - Alert handling (for creating alarms – *numeric values only*)
 - Archiving (useful for trending – *does not work yet*)
-

Visual aspects

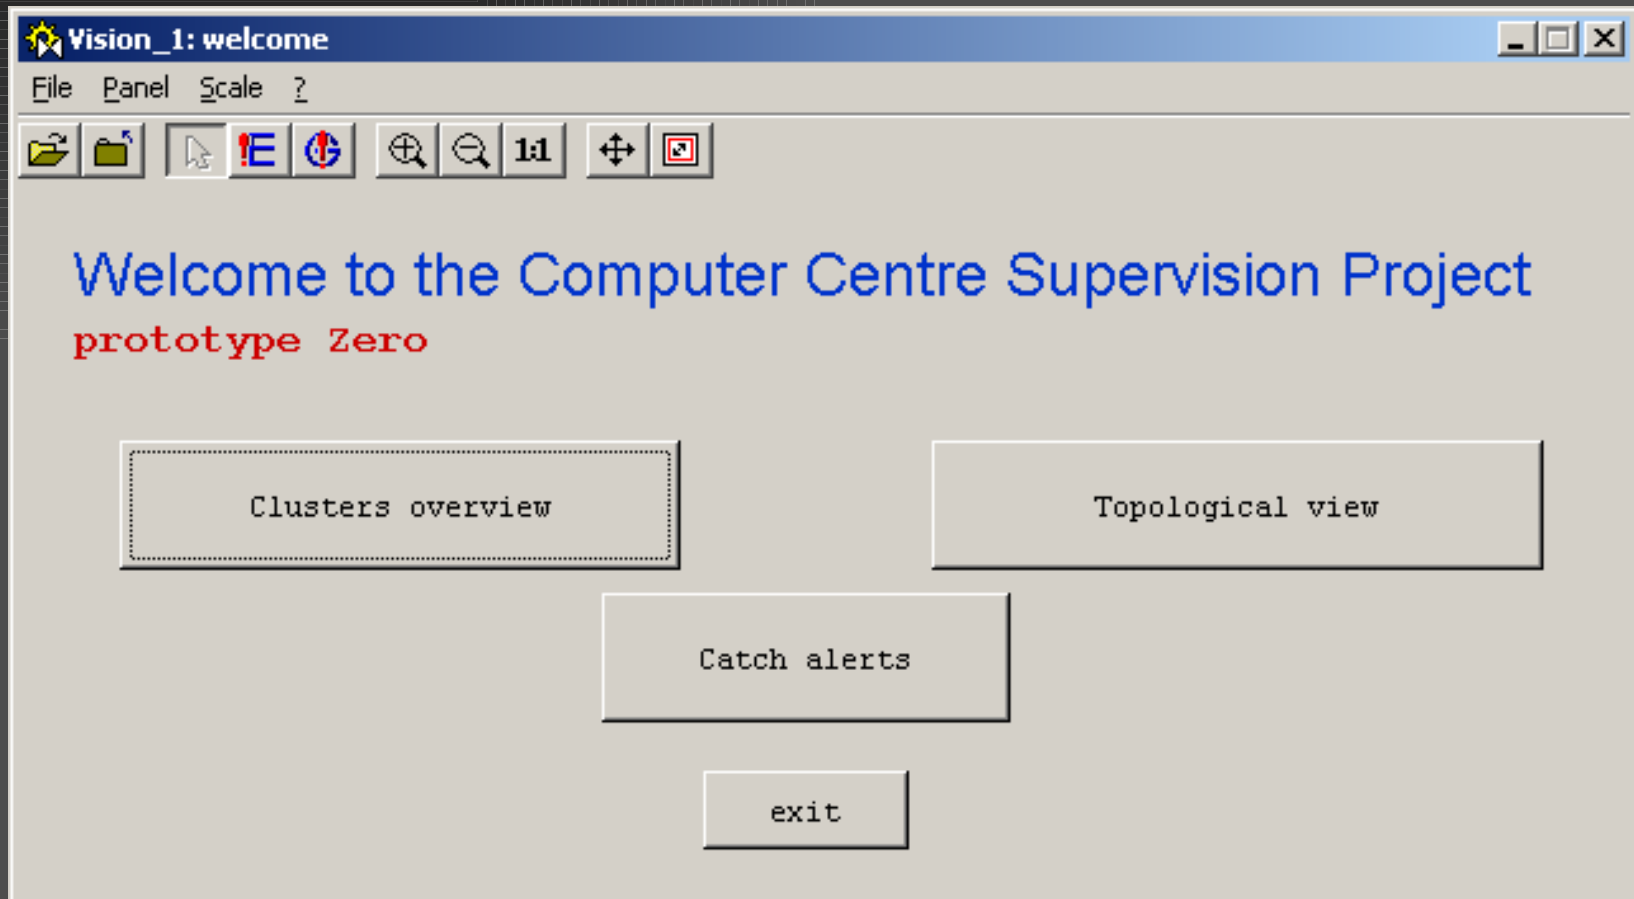
- PVSS uses panels to interact with users
- PVSS provides a toolkit to build them
 - Graphical objects (windows, buttons, circles, ...)
 - C-like scripting language
- Three views were implemented:
 - Clusters a la PEM
 - Topological
 - Alert catching
- Static “PEM” Web page now generated from PVSS data
(<http://cern.ch/it-proj-ccs2/PrototypeZero/ClusterStatus>)



Demonstration of Prototype 0



Welcome panel



Clusters overview

The screenshot shows a window titled 'clusters' with a blue title bar. The main content area is titled 'Computer Centre Topology Clusters'. On the left, there is a vertical list titled 'Alphabetical list' with a scroll bar, containing the names of the clusters. The main area on the right is titled 'Available Clusters' and contains a 4x12 grid of buttons, each representing a cluster name.

Alphabetical list

- afs
- atlas
- atljprgd
- castor
- cms
- cnwgs
- cvs
- dxplus
- eff
- geant
- globus
- hpplus
- hps
- infrastructure
- inst_hp
- iss
- l3c
- licman
- linux_inst
- lock_servers
- lsfmaster
- lxbatch
- lxplus
- lxshare
- mail_servers
- misc_sgi
- monarc
- mta

Available Clusters

afs	atlas	atljprgd	castor
cms	cnwgs	cvs	dxplus
eff	geant	globus	hpplus
hps	infrastructure	inst_hp	iss
l3c	licman	linux_inst	lock_servers
lsfmaster	lxbatch	lxplus	lxshare
mail_servers	misc_sgi	monarc	mta
na48	na49i	nsr	parc
pdpdev	pdptst	ps_servers	pttools
public_servers	registration	remedy	shiftableph
shiftatlas	shiftccf	shiftcms	shiftdelphi
shiftl3	shiftlhcb	shiftna45	shiftna48
shiftna49	shiftnomad	shiftopal	shiftslap
sundev	sure	tape_servers	thwgs
tomog	www		

close

Topological view

ComputerCentre

MR8

MR7

MR17

MR6

MR16

MR5

MR15

MR4

MR14

MR3

MR13

MR2

MR12

MR1

MR11

MR0

MR10

MR32

MR22

IT Computer Centre B513

14:03:36 15-Apr-02

BIS services

Close

MR7:

- atlasd01
- lxbatch112
- lxbatch121
- lxbatch122
- lxbatch133
- lxbatch247
- lxbatch248
- lxbatch249
- lxbatch250

One rack


rackMR7							
atlasd01	lxbatch112	lxbatch121	lxbatch122	lxbatch133	lxbatch247	lxbatch248	lxbatch249
	0.08	2.05	2.01	0.08	0.11	1.14	0.07
lxbatch250	lxbatch251	lxbatch252	lxbatch253	lxbatch254	lxbatch255	lxbatch256	lxbatch257
1.07	0	0	0.07	1	0	0.05	0
lxbatch258	lxbatch259	lxbatch260	lxbatch261	lxbatch262	lxbatch263	lxbatch264	lxbatch265
0.26	0	0	0.1	0.01	0.04	0.07	0.12
lxbatch266	lxbatch267	lxbatch268	lxbatch269	lxbatch270	lxplus016	lxplus017	lxplus018
0.02	1.08	0.07	0	0.01	0.06	0.21	1.74
lxplus019	lxplus020	lxplus021	lxplus022	lxplus023	lxplus024	lxplus025	lxplus026
2.8	0.36	0.33	0.22	59.65	0.21	1.7	1.63
lxplus027							
0.26							


One machine

PClxplus023























SYSTEM REGISTRATION LSF MONITORING AFS

--lxplus023 @ 15-Apr-02 14:12:27 --

LOAD  59.65 --> High Load

ZOMBIE  1 --> small number of zombies

daemons

inetd->			klogd->			sendmail->		
crond->			sshd->			portmapd->		
rfiod->			atd->			syslogd->		
named->			rpcd->					

ALERT STATUS


Numeric Readings


N Sockets --> 217

N Proc --> 227

UpTime --> 2 d 21 h 40 m 40 s

BootTime --> 11789 d 13 h 55 m 5 s

var  10.00

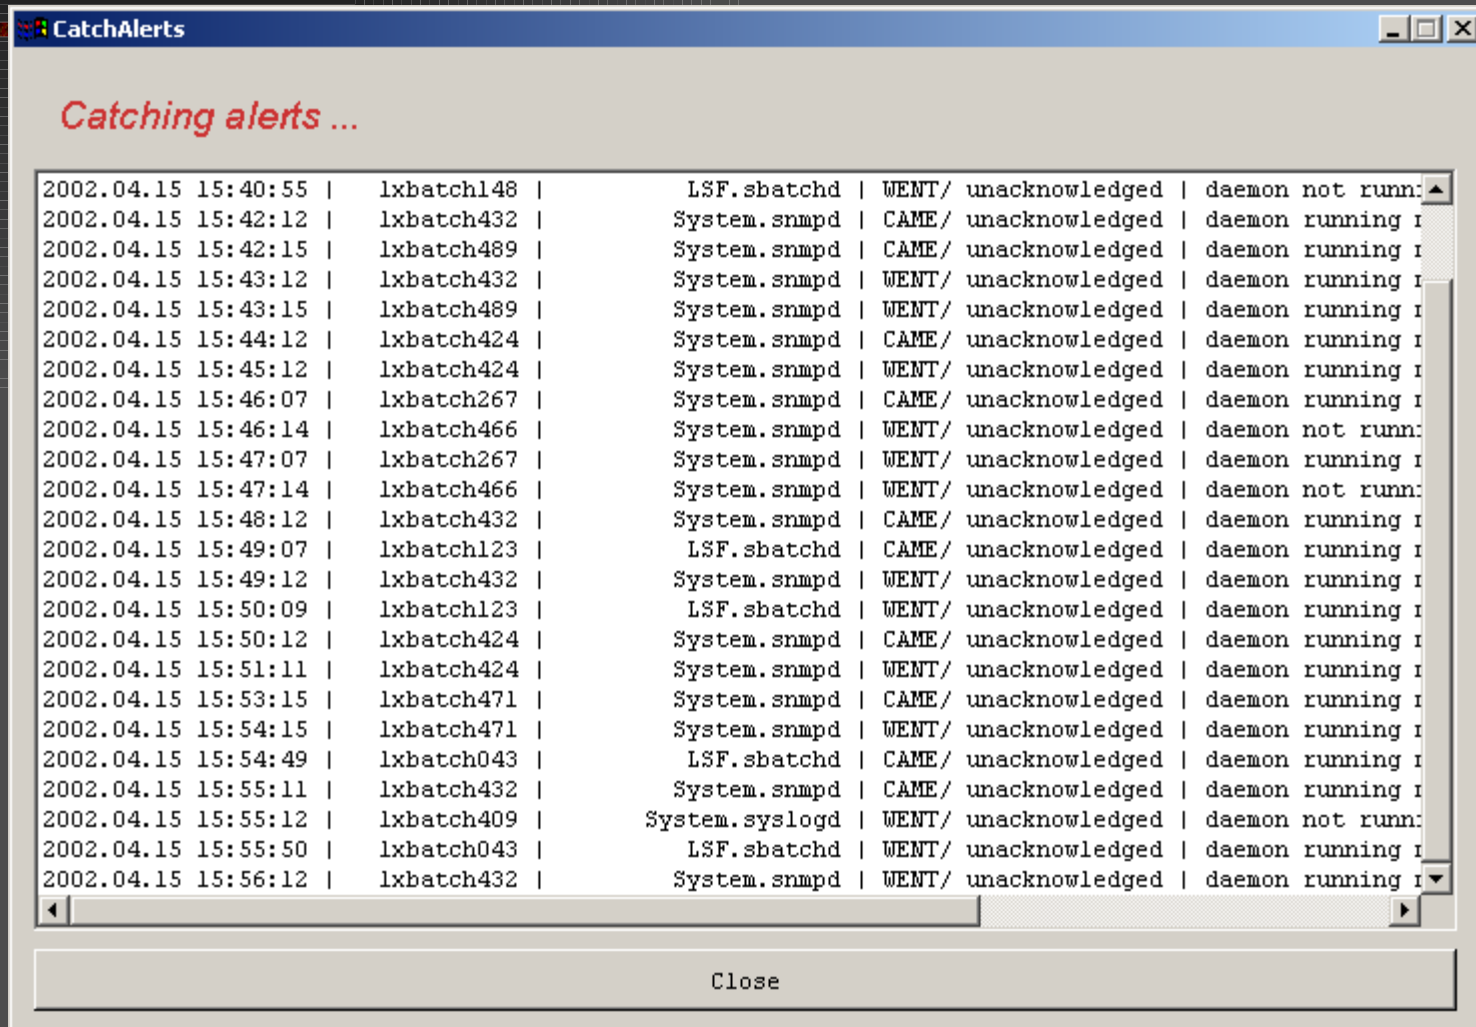
tmp  4.00

50 % 100%

OPERATING SYSTEM

2.2.19-6.2.1.1smp

Alert catcher



Catching alerts ...

2002.04.15	15:40:55	lxbatch148	LSF.sbatchd	WENT/	unacknowledged	daemon not runn:
2002.04.15	15:42:12	lxbatch432	System.snmpd	CAME/	unacknowledged	daemon running r
2002.04.15	15:42:15	lxbatch489	System.snmpd	CAME/	unacknowledged	daemon running r
2002.04.15	15:43:12	lxbatch432	System.snmpd	WENT/	unacknowledged	daemon running r
2002.04.15	15:43:15	lxbatch489	System.snmpd	WENT/	unacknowledged	daemon running r
2002.04.15	15:44:12	lxbatch424	System.snmpd	CAME/	unacknowledged	daemon running r
2002.04.15	15:45:12	lxbatch424	System.snmpd	WENT/	unacknowledged	daemon running r
2002.04.15	15:46:07	lxbatch267	System.snmpd	CAME/	unacknowledged	daemon running r
2002.04.15	15:46:14	lxbatch466	System.snmpd	WENT/	unacknowledged	daemon not runn:
2002.04.15	15:47:07	lxbatch267	System.snmpd	WENT/	unacknowledged	daemon running r
2002.04.15	15:47:14	lxbatch466	System.snmpd	WENT/	unacknowledged	daemon not runn:
2002.04.15	15:48:12	lxbatch432	System.snmpd	CAME/	unacknowledged	daemon running r
2002.04.15	15:49:07	lxbatch123	LSF.sbatchd	CAME/	unacknowledged	daemon running r
2002.04.15	15:49:12	lxbatch432	System.snmpd	WENT/	unacknowledged	daemon running r
2002.04.15	15:50:09	lxbatch123	LSF.sbatchd	WENT/	unacknowledged	daemon running r
2002.04.15	15:50:12	lxbatch424	System.snmpd	CAME/	unacknowledged	daemon running r
2002.04.15	15:51:11	lxbatch424	System.snmpd	WENT/	unacknowledged	daemon running r
2002.04.15	15:53:15	lxbatch471	System.snmpd	CAME/	unacknowledged	daemon running r
2002.04.15	15:54:15	lxbatch471	System.snmpd	WENT/	unacknowledged	daemon running r
2002.04.15	15:54:49	lxbatch043	LSF.sbatchd	CAME/	unacknowledged	daemon running r
2002.04.15	15:55:11	lxbatch432	System.snmpd	CAME/	unacknowledged	daemon running r
2002.04.15	15:55:12	lxbatch409	System.syslogd	WENT/	unacknowledged	daemon not runn:
2002.04.15	15:55:50	lxbatch043	LSF.sbatchd	WENT/	unacknowledged	daemon running r
2002.04.15	15:56:12	lxbatch432	System.snmpd	WENT/	unacknowledged	daemon running r

Close

Next steps (1)

- Prototype 1, due June 2002
 - Deployed by
 - Operators (partly phase out Sure)
 - Service managers concerned
 - Additional objectives:
 - Scalability tests (criteria to be discussed and defined)
 - Re-design of data points
 - Web interface
-

Next steps (2)

- Decision about suitability of PVSS
 - Develop list of criteria – user involvement
 - Functionality
 - Performance
 - Ease-of-use for users and developers
 - Cost
 - Experience with PVSS from Prototype 0:
 - Very useful to use control framework functionality provided
 - Complex toolkit - need expert advice from ETM
-

Next steps (3)

- Prototype 2, due December 2002
 - subject to positive result of suitability decision
 - Deployed by
 - IT managers
 - End users
 - Service managers of additional services
 - Additional objectives:
 - Correlation engine on PVSS system
 - Local correlation engine (?)
-

Communication

- Moving from learning experience into an effort eventually leading to deployment
 - Wider communication required
 - Receive requirements/use cases from users (IT service providers and end users alike)
 - Discuss acceptance criteria
 - Present regularly what we have done
 - Monthly open project meeting planned with users at CERN
-

Relationship with EDG WP4

- Short-term: Plan to use their configuration database services
 - Long-term: Make sure the two projects converge
 - Different approaches towards same goal
 - Use same solutions wherever possible
 - Factorise and modularise as much as possible to support different solutions where necessary
 - Regular contacts with WP4 through WP4 manager
-

Conclusion

- Good start
 - But P0 still has a lot of rough edges...
 - Core team very motivated and committed
 - Potential for further successful work
 - Will need help from users
 - Make their needs known
 - Help implement their service
 - Use the system
 - A very long way ahead...
-