

Meeting Object: **WP4 workshop**

Author: **Olof Barring**

Partner: **CERN**

Meeting Date: **19<sup>th</sup> – 20<sup>th</sup> of June 2002**

Meeting Place: **CERN**

Attendees: **David Groep (NIKHEF), Martijn Steenbakkers (NIKHEF), Enrico Ferro (INFN), Marco Serra (INFN), Andrea Chierici (INFN), Paul Anderson (PPARC), Piotr Poznanski (CERN), Jaroslaw Polok (CERN), David Front (LCG/CERN), Markus Schulz (CERN), Mathias Gug (CERN), Sylvain Chapeland (CERN), Michele Michelotto (INFN), Lord Hess (KIP), Julian Blake (CERN), Lionel Cons (CERN), German Cancio (CERN), Bernd Panzer (CERN), Sergey Makarychev (CERN), Juan Pelegrin (CERN), David Foster (CERN), Thorsten Kleinwort (CERN), Jan van Eldik (CERN), Jan Iven (CERN), Marcus Hardt (CrossGrid/Karlsruhe), Thomas Röblitz (ZIB), Maite Barroso (CERN), Ian Nielsen (LCG/CERN)**

---

## 1. DECISIONS

- It was agreed that we should go ahead with developing the global schema. A large fraction of the global schema should be optional and components must declare what parts they need and add additional validation as required. It must be written down what component providers need to do to write the configuration for her/his component and how the components can be included in the global schema.
- It was decided to abandon the CCM daemon and provide a fat direct access library instead. Encryption/access control has to be taken into account.
- It was decided to use the new XML schema proposed by Piotr. It remains to decide who works on what for its implementation.
- Go for approach #4 in German slides on LCFG – HLD integration. EDG will start to work on details outlined under approach #3 and Paul will ask Lex when he is back to work on a flattening convention for default serialization.

## 2. FIRST DAY, 19<sup>TH</sup> OF JUNE, TASK STATUS AND PROGRESS REPORTS

### 2.1. MONITORING, SYLVAIN CHAPELAND

Sylvain's slides: <http://documents.cern.ch/cgi-bin/setlink?base=agenda&categ=a02760&id=a02760s2t2/transparencies>

#### 2.1.1. Discussion

- Olof said that the alarm display had originally been foreseen for release 1.4 (end of July) but it had now been delayed to release 2.
- Jan van Eldik commented that the low rate of lost UDP packages reported by Sylvain depends strongly on the number of metrics per node and the sampling frequency. He

said that 10 metrics on 350 nodes probably works fine while for him with 100 metrics on 1,000 nodes the situation is much worse.

- Lionel wondered why a pull model had been chosen for the transport. Sylvain answered that push or pull transport had not been decided yet. Probably both will have to be supported. The example shown in the presentation was based on pull.
- Michele pointed out that using Oracle for backend database would pose a problem for other sites than CERN. Sylvain and Olof explained that the database has not yet been selected and that both Oracle and MySQL are tested. If Oracle is selected it is clear that at least one Open Source database will be supported as well.
- Lionel wondered how connection overhead could be avoided in the TCP transport. Was the plan to use open connections? Yes. Jan van Eldik then commented that it is important that the transport is easily configurable.
- David Groep was worried about using FTP or GridFTP for the transport. With FTP (and GridFTP) there is a new data connection for every file transfer, which means that there will be a significant connect time overhead, in particular if GSI was to be used for the authentication. David suggested that we should look at https instead.
- David Front wondered what the repository query API would look like? Sylvain answered that only a very simple query interface will be provided. Since most likely the database will natively support SQL there is no reason to try to implement another sophisticated query language on top of it.

## 2.2. LCFG AND EDG MONITORING, MATHIAS GUG

Mathias' slides: <http://documents.cern.ch/cgi-bin/setlink?base=agenda&categ=a02760&id=a02760s2t5/transparenties>

### 2.2.1. Discussion

- German said that there seems to be a huge overlap with what Mathias has been doing and what is provided with the new version of LCFGng. Markus said that the scope of this project was to keep us (the CERN testbed administrators) going while we are waiting for a proper solution.
- German also wondered why the logfiles were sent to the server for parsing rather than being parsed on the client hosts? Mathias said that this was easier to implement.
- Marco pointed out that another problem that needs to be monitored is if the transfer of the profile to the client node fails. In this case the old profile will be used and the monitoring will happily report that everything is OK while the profile is actually wrong. Paul said that the standard timestamps on the status web-page could be used to detect whether or not the update of a profile has succeeded.

## 2.3. FAULT TOLERANCE, LORD HESS

Lord's slides: <http://documents.cern.ch/cgi-bin/setlink?base=agenda&categ=a02760&id=a02760s2t3/transparenties>

### 2.3.1. Discussion

- Olof said that the user API looked very much application monitoring and he wondered how this relates to the GRM work of WP3? Lord agreed that it is a type of application monitoring but the aim is to be able to correlate application errors with other system monitoring data (stored in the monitoring repository) and decide on appropriate actions. The user can provide recovery actions that may be called at certain conditions.

- Olof was still worried how this fit into the originally planned architecture, the focus now seem to have been moved from defining actuator and dispatcher interfaces and define recovery rule language. Lord said that this was still the plan.

## 2.4. GRIDIFICATION REPORT, MARTIJN STEENBAKEN

Martijn's slides: <http://documents.cern.ch/cgi-bin/setlink?base=agenda&categ=a02760&id=a02760s2t4/transparenties>

### 2.4.1. Discussion

- German wondered if the authorization policy language foreseen for release 2 would be based on the results from the IETF AAA group? David Groep said that this will be partly the case depending on how much the AAA group has currently defined.
- German also wondered why they foresee to move from dynamic shared objects to service implementations of the authorization plug-ins? David said that the reason is that some authorization decisions may have to be performed not as root.

## 2.5. RESOURCE MANAGEMENT, THOMAS RÖBLITZ

Thomas' slides:

### 2.5.1. Discussion

- German wondered what are the open issues with Condor? Thomas said that it works quite differently from other batch systems. The RMS can probably easily support the management of user jobs while it may be more difficult to support the administrative interface for scheduling of maintenance operations.

## 2.6. INSTALLATION MANAGEMENT, GERMAN CANCIO

German's slides: <http://documents.cern.ch/cgi-bin/setlink?base=agenda&categ=a02760&id=a02760s2t7/transparenties>

### 2.6.1. Discussion

- Marco was concerned with the mechanism for upgrading from current LCFG to LCFGng. A complete re-install of the LCFG server will not be well appreciated by the testbed site administrators. German said that possibly (hopefully) only the LCFG clients need to be re-installed but this could not be guaranteed yet.
- Marco was also worried over a statement from Fab at the recent INFN meeting: it appeared as if WP4 would support LCFGlight. Olof said that it was clear from the reply to the EU reviewers that WP4 will *not* support LCFGlight. However, what we aim to support for smaller sites is a WP4light consisting of components only from the gridification, resource management and monitoring subsystems packages for manual installation and configuration.
- Paul said in connection to the above point that it is important to carefully design the LCFG objects. For instance, one should take care to put as much as possible of the package control in the *init.d* scripts and simply call the standard methods from the LCFG object. This would facilitate manual configuration of packages at installation not managed by LCFG.

## 2.7. CONFIGURATION MANAGEMENT AND PAN'S COOLEST FEATURES, LIONEL CONS

Lionel's slides:

<http://documents.cern.ch/cgi-bin/setlink?base=agenda&categ=a02760&id=a02760s2t10/transparenties>

<http://documents.cern.ch/cgi-bin/setlink?base=agenda&categ=a02760&id=a02760s2t8/transparenties>

## 2.7.1. Discussion

- Paul wondered if the “types” was referred to as real language types or as strings associated with validation code? The latter
- Paul asked if the compiler support dependencies? Yes, derivation is already supported. Dependencies are currently handled at template level. In XML there will be a list of HLD templates used to generate the information.
- Paul said that this would become important: a typical LCFG client depends on 100 source files
- Paul commented that it is interesting to compare the HLD and LCFG. He said that he is worried about some fundamental issues like who will be programming the templates? Is it the component providers? The language is complex and for LCFG they have taken the rather opposite direction by simplifying their language as far as possible.
- Thorsten wondered how the configuration data would be stored and how the information can be queried? Lionel said that
  - Templates will end up in CVS + some glue to make it transactional
  - Server modules to be plugged to the database will be supported. The query will be done on low-level definition (the real data).
- Jan van Eldik asked if would it be possible to import code from other languages, e.g. perl modules? No, not directly but if there is some cool perl module, which is found to be useful, it could be plugged into the compiler like it was done for PCRE to support Perl Regular Expressions.

## 3. SECOND DAY, 20<sup>TH</sup> OF JUNE

### 3.1. GLOBAL SCHEMA, MAITE BAROSSO

Maite's slides: <http://documents.cern.ch/cgi-bin/setlink?base=agenda&categ=a02760&id=a02760s1t1/transparenties>

#### 3.1.1. Discussion

- Slide 10 (hardware branch):
  - Paul: what about the defaults? Lionel: in the HLD you can define optional or compulsorily fields. If you decide that the information is necessary you make that information compulsorily.
  - Paul: what does the “global” mean here? Do you foresee one schema per site or one for the whole world? Lionel: no it is rather one schema per group (likely to be per site).
  - Jan Iven asked if not all fields could be made optional?
  - Paul liked Jan's idea. It would allow each particular component to specify what information it requires and if you don't provide that information you cannot use that component. German said that this is just another type of validation? Lionel commented that this could be done but that it puts more burdens on the programmer of the components.

- German asked what should be “optional” in the model proposed by Jan and Paul? Is it the HARDDISK or the fields in the HARDDISK definition? Paul said that it must be the fields otherwise it is not practical. He said that if you want to configure your laptop for the first time you only care about the size of the disk but not about the number of sectors.
- Slide 14 (system branch):
  - German commented that the nlist enumeration of the partitions is just one way of doing it. There are other ways.
  - Paul said that HLD provides a nice way to specify abstract information but then you need to translate that into real implementation information. Currently we do it with the profiles. Maybe it is worth performing that translation at the compiler level. This allows more validation at compile time rather than run time.
  - Jan Iven suggested splitting architecture checks in the compiler and leaving the syntax translation to the perl components. The component specifies what it expects to be present in the database.
  - Maite objected that different components might require different information. Jan: that’s right, you make the union of all the information needed by the components you require for your system.
- Slide 17 (system branch):
  - Jan: can one validate that NIS is installed before it is configured? Yes and no. One can add validation that the NIS package is installed in /sw/packages . Jan commented that one would rather search for capabilities than package names.
- Slide 21 (software branch):
  - Paul: are you catering for supporting other packages format than RPM? Yes. Maybe one should then have a more generic PACKAGE type.
  - Olof wondered if it would be possible to include RPM dependencies in the validation. This was proposed by Lionel but rejected by testbed site administrators. They said that RPM dependencies are checked anyway on node. Lionel said that validating RPM dependencies is possible to do but it is risk to become very complex (not simply Boolean).
  - Jan Iven said that maybe one could allow for an ad-hoc tests (e.g. dummy installation in /tmp) to be added as external validation? In all cases one would start from something simple and evolve it from experience of common errors.
- Slide 23 (examples):
  - Marco wondered if it is possible to define the protocol to access the repository? It is currently specified with the “URL” but could be split: “host”; “protocol”; “directory”.
  - After some discussion it was thought to better to break down the URL in host, protocol and directory.
- Slide 24 (examples):
  - Jan Iven asked if the intention was to support different version of the same package with single name in different repositories. Yes. Do we want this? It is optional.
  - Lionel said that this was just a proposal: all tasks should review the global schema and comment.

- Slide 26 (examples):
  - Jan asked if one could specify in node\_profile a list of valid IP addresses? Yes.
  - Jan wondered if the global schema definition also proposes standard naming of templates? Yes it has to be there but it is not included yet.
  - Paul said it should be more than just a convention, e.g. GUIs showing disk information needs to rely on the schema.
- Slide 28 (examples):
  - German said that there is a modified 3<sup>rd</sup> solution where the external file is not specified at all and provided with the RPM instead.
  - Paul preferred this solution.
- General comments after Maite finished her presentation:
  - Enrico said that he is sceptical about the usefulness of the global schema. The language appears very useful. He said that he would really like to soon see an HLD implementation of the LCFG schema otherwise it would be a too difficult migration.
  - Paul said that he liked the fact that you can write PAN code that generates resources for my component from the global schema.
  - Lionel said that the decision he would like to day is to whether or not the global schema is something good for the future while we still can work on immediate solutions for testbed administrators in parallel.
- Wider involvement and organisation of the work in defining the global schema:
  - Maite said that so far configuration task, German, Thorsten and Markus have participated. She know really want like to have other sites and tasks involved.
  - Thorsten said that to populate the schema we need some tools to facilitate for the operators. Lionel answered that GUIs are foreseen for later, beyond release 2.
  - Olof said that working organisation must be put in place for sharing templates. Maite and Lionel proposed to setup a mailing list and to put files and templates in EDG CVS.
  - Paul wondered how this is going to be managed in the longer term? Are we going to have a global schema per site or really global? If we have different schemas between sites there will inevitably be cases where components become incompatible. We have to agree on a core set.

### 3.1.2. Decision:

- We shall go ahead with the global schema. A large fraction of the global schema should be optional and components must declare what parts they need and add additional validation as required. It must be written down what component providers need to do to write the configuration for her/his component and how the components can be included in the global schema.

## 3.2. CCM OR NOT CCM, LIONEL CONS

Lionel's slides: <http://documents.cern.ch/cgi-bin/setlink?base=agenda&categ=a02760&id=a02760s1t2/transparencies>

### 3.2.1. Discussion

- The support for access control in the direct access library was discussed in some detail. Paul suggested that the cache manager could write DBM files with different UNIX access rights.

### 3.2.2. Decision

- It was decided to abandon the CCM daemon and provide a fat direct access library instead. Encryption/access control has to be taken into account.

### 3.3. UPDATED XML, PIOTR POZNANSKI

Piotr's slides: <http://documents.cern.ch/cgi-bin/setlink?base=agenda&categ=a02760&id=a02760s1t3/transparencies>

#### 3.3.1. Discussion

- Slide 8
- Paul wondered what is the XML schema needed for? Validation of the compiler output.
- Paul said that types are properties of the schema and the values are the properties of the profile. Piotr replied that the XML schema is used by the validation. The idea is to have all required information in the document.
- Lionel said that for him the stripping out the types is just a win in a couple of bytes to be transported.
- Olof asked Paul if he thought we could agree on same format between LCFG and WP4 for the longer term? Probably not.

#### 3.3.2. Decision

- It was decided to use the new XML schema proposed by Piotr. It remains to decide who works on what for its implementation.

### 3.4. NVA API, PIOTR POZNANSKI

Piotr's slides: <http://documents.cern.ch/cgi-bin/setlink?base=agenda&categ=a02760&id=a02760s1t3/transparencies>

#### 3.4.1. Discussion

- Piotr pointed out that the configuration locking may need to be removed from the API now when we have decided to not keep the CCM (see section 3.2.2).
- Paul sees an issue in what one expect the end user to see? He said that 90% of components would use a simplified layer of tools on top of the API. This common layer of simpler access appears to be missing.

### 3.5. KICKSTART FILE GENERATION, ANDREA CHIERICI

Andrea's slides:

#### 3.5.1. Discussion

- German wondered when the starting LCFG takes place in the post-install? After the reboot as it is now with a standard floppy installation.
- Paul suggested that if parsing of the XML profile is needed it would be good to start from LCFG parser. Another possibility would be exploiting the CDB plug-in to

generate KickStart file. German said that he would prefer to leave the generation of the KickStart file separate from the CDB.

- Jan asked why using a minimal set of RPMs from a special repository. Couldn't the KickStart file be generated to install the complete desired set of RPMs? With the present proposal one needs to work with two separate RPM repositories. This issue caused some discussion and the conclusion was to defer the decision to whether or not two repositories were needed until after some prototyping had been performed.

### 3.6. LCFG EVOLUTION, PAUL ANDERSON

Paul's slides: <http://documents.cern.ch/cgi-bin/setlink?base=agenda&categ=a02760&id=a02760s1t13/transparencies>

#### 3.6.1. Discussion

- Lionel asked if the distributed development of components could be quantified? Paul said that a couple of people are responsible for most of the components + 6-7 people for the rest. In total about a dozen.
- German commented that there is quite a lot of overlapping between LCFG and EDG WP4 in what we want for developments but the problem is that the timescales and focuses are different. He said that we would probably need to make parallel developments on the client part. The EDG WP4 would aim to change the configuration accessing part while the component model could probably stay the same.

### 3.7. STATUS AND ISSUES OF THE HLD – LCFG INTEGRATION

German's slides: <http://documents.cern.ch/cgi-bin/setlink?base=agenda&categ=a02760&id=a02760s1t6/transparencies>

#### 3.7.1. Discussion

- Paul wondered why approach #3 was not backwards compatible? German said that with this approach you don't have the flat structure anymore (serialization) needed by the LCFG components. Options are
  - To write a serialization component on top of the NVA API
  - Hack the components so that they use a default serialization convention
  - Write *pan* code to generate component specific configuration
- Marco asked what happened to the development and production threads that were decided a couple months ago? Enrico explained that the XML generated by *pan* is not completely compatible with LCFG. Olof added that the last couple of months there have been a lot of efforts invested in trying to understand how LCFG could be integrated with the HLD. To him it now appears as if that approach #4 presented by German appears to be the least expensive even if it breaks what was planned originally for the production thread. The production thread will now instead focus on deploying LCFGng, which is going to be used for release 2.

#### 3.7.2. Decision

- Go for approach #4 in German slides. EDG will start to work on details outlined under approach #3 and Paul will ask Lex when he is back to work on a flattening convention for default serialization.

### 3.8. LCG PROJECT OVERVIEW, DAVID FOSTER

David's slides: <http://documents.cern.ch/cgi-bin/setlink?base=agenda&categ=a02760&id=a02760s1t17/transparencies>



### 3.8.1. Discussion

- Some worries were raised as to whether EDG WP4 is kept out from the GLUE packaging discussions. David assured that this was not the case. The GLUE packaging discussion has not started yet.
- Olof wondered if LCG would provide software certification testbed and people to perform the certification? Yes

### 3.9. ATF REPORT, GERMAN CANCIO

German's slides: <http://documents.cern.ch/cgi-bin/setlink?base=agenda&categ=a02760&id=a02760s1t8>

### 3.10. GLUE SCHEMA, OLOF BÄRRING

Olof's slides: <http://documents.cern.ch/cgi-bin/setlink?base=agenda&categ=a02760&id=a02760s1t9/transparencies>