Reference installation for the German grid initiative D-Grid

View the article online for updates and enhancements.

Reference installation for the German grid initiative D-Grid

W Buehler, O Dulov, A Garcia, T Jejkal, F Jrad, H Marten, X Mol D Nilsen and O Schneider

Steinbuch Centre for Computing, Forschungszentrum Karlsruhe

E-mail: `Oleg.Dulov@iwr.fzk.de`

Abstract.   The D-Grid reference installation is a test platform for the German grid initiative. The main task is to create the grid prototype for software and hardware components needed in the D-Grid community. For each grid-related task field different alternative middleware is included. With respect to changing demands from the community, new versions of the reference installation are released every six months.

Keywords
Grid Computing, D-Grid, job submission middleware, data management middleware, reference installation, grid prototype.

1. Introduction
As a new service in the German grid initiative D-Grid [1], the reference installation [2] is conceptualised, set up and administered at the Steinbuch Centre for Computing (SCC). It serves as reference system and test platform for various grid middleware solutions oriented towards the demands of the German grid community. Its purpose is to demonstrate and document how to build a grid infrastructure as a composition of the most favorable middleware. The reference installation includes:

- A Mini-cluster for software installation. It consists of hardware devices (physical and virtual) to install the grid-related software for tests.
- A sequence for maintaining software as follows:
  - Install the software.
  - Identify errors and deficiencies compared with the requirements from the customers through tests.
  - Carry out changes according to tickets from helpdesk if necessary. These changes may entail new tests or even new installation. The current reference installation does not include "standard" procedures or frameworks for testing. However, there is the Nagios [3] based initiative in D-Grid to work on the introduction of such a framework.
- Technical documentation[1] as the end product of the reference installation based on the MediaWiki [4] open source wiki engine. The technical documentation is structured in two

---

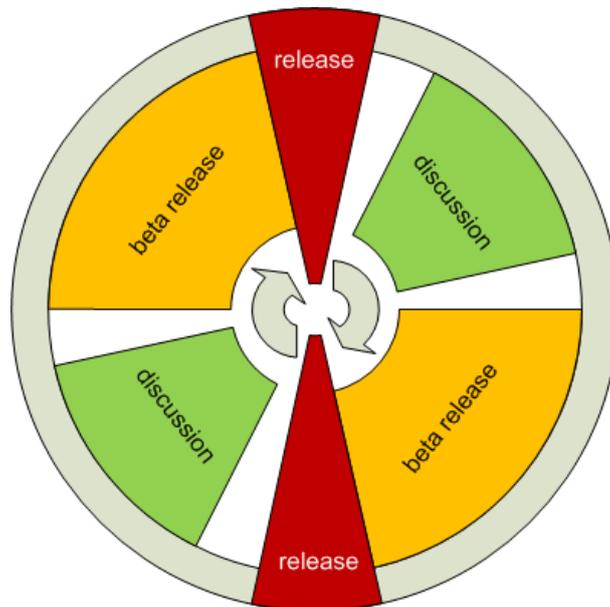[1] used for D-Grid at `http://dgiref.d-grid.de/wiki/Introduction`

Figure 1. D-Grid reference installation release cycle. Each new release for the reference installation will start with *discussion*, where all D-Grid users are invited to contribute via e-mail. With approval from the D-Grid TAB the *discussion* phase is closed and *beta-release* phase is entered. Decisions from discussion and TAB are incorporated into the installation and system functional tests are performed. When the release package proofs to be functional, the procedure passes over to the terminal *release* phase. Now the software is installed on the reference installation's hardware and all documentation is published.

> internal logical parts. Content is separated as "pure technical content", like bash scripts with comments, and "common sense content", which includes configuration notes, links to other content, common words, descriptions and so on.

Note, that the term *change* has a very important role in the reference installation and is coupled with another term — *configuration*. Every support team or end user of the reference system can initiate changes by writing a ticket in the DGUS [5] support portal. When necessary these will lead to changes in the system configuration, or in the technical documentation, or both.

2. Release cycle

The reference installation is released in a new version on well-defined dates — normally every six months. Each version has to go through three phases: *discussion*, *beta-release* and *release* (see the Figure 1). Throughout the whole process the system maintainer of the reference installation is moderating discussions or performing actions on hard- and software respectively.

During *discussion phase* the D-Grid community is invited to consult the software and hardware components for the upcoming new release. The main conversation is based on exchange of e-mails via a dedicated mailinglist[2], important statements are collected in the D-Grid reference installation's Wiki. After the discussion is closed, all suggestions are summarized in a wish list and submitted to the D-Grid Technical Advisory Board (TAB) [6]. The TAB will come to the decision on which software stacks are to be included in the upcoming version of the reference installation and whether existing components have to be excluded or updated/upgraded. Finally

---

[2] `dgrid-ref-support@d-grid.de`, for registration on it, please contact `oleg.dulov@kit.edu`

the reference system maintainer sends the request for software packages to appropriate support teams.

The *beta-release phase* includes the installation of the new version for the reference installation on a pre-production system (often called certification system) and performing checks for conflicts between different components. This is done by the maintainer of the reference system with support by the middleware supporters. It is still possible to change the version for a software package in case of issues that turn out irresolvable otherwise.

When *release* phase has started the beta-release installation is repeated on the reference system. From that point onward, versions and software changes are not possible anymore (except for fixing critical bugs). After final migration tests and operability checks the technical documentation is published on the reference installation's Wiki and the D-Grid resource providers are demanded to update their systems.

3. Hardware architecture

The D-Grid reference installation includes two types of systems — *certification* and *reference systems*. The purpose of certification system is to serve as "pre-production" system and is associated with the beta-release phase.

The current certification system consists of a set of 26 identical virtual machines (VMware ESX [7] based) with minimal standard operating system. Thanks to the virtual environment, machine parameters, such as 32 bit or 64 bit architecture, can be easily adapted. Whereas the processor family (e.g. Intel Xeon, AMD Opteron, ect.) is predetermined by the underlying hardware.

On the other hand, the reference system is set up on 22 physical machines (a mini-cluster). All of them are 64 bit architecture, two of them possess very large hard drives and small memory. Just the opposite for the remaining 20 machines that have build in high-end memory and slender disk space.

The advantage of using the certification system is that the working reference installation is not touched until the software has been tested (certified). But it is still in question, whether we can handle the inflexibility to change the processor family.

4. Software architecture

Figure 2 shows a structural overview of the setup for the reference installation released on 15th of January, 2009. A unique feature of D-Grid infratructure is that all functional aspects are covered not by just one single middleware, but several middleware stacks are available [11].

The **cluster** includes the software to make a working compute cluster (e.g. network file server, batch system, automated configuration tools, worker nodes with configured software) and a special type of node, named interactive node. This node is configured as a access point for the grid administrators and programmers.

Job submission frontends manage the users' tasks (jobs) entering from user interface and assign them to the cluster for execution. The user interface (UI) provides the protocols for communication. The D-Grid reference installation features gLite [8], Globus Toolkit [9] and UNICORE [10] software products in its first release.

Data management frontends are software that manage a potentially huge amount of storage space — in the magnitude of Petabytes — in grid environments. Additionally they may serve as recipient for data output from running jobs on linked clusters. In a similar way to the job submission front ends this type of middleware can also include an own user interface to access and retrieve stored data with appropriate tools for interaction.
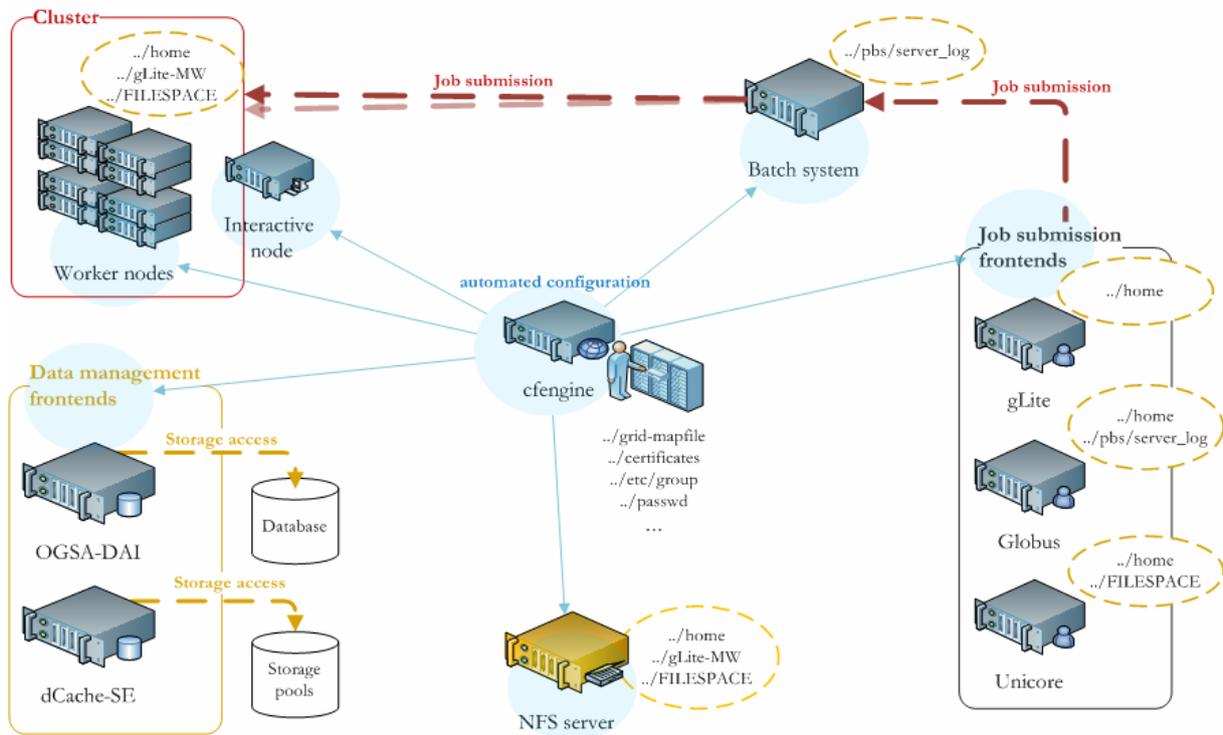
Figure 2. D-Grid reference installation infrastructure. The D-Grid reference installation released on 15th of January, 2009, includes one NFS server for files distribution between other components, one cfengine server for automated configuration of other components, one server for the batch system, one cluster with worker nodes, three frontends for job submission (UNICORE, Globus Toolkit, gLite), two frontends for data management (OGSA-DAI, dCache-SE), one interactive node as an entry point to the cluster for users.

5. Proceeding

Software packages with their technical documentation are received by the system maintainer and the installation into the certification system is started. The working procedure is the following: during the beta phase, the system maintainer (or support team) is writing the bash scripts to automate the installation and configuration procedures for every included software package. The package can be excluded from release if there is a serious problem to maintain it into the system.

The results of the beta-phase are:

a) bash scripts used to install and document the system are created

b) reference system is installed into the certification system

The bash scripts from beta-phase are the part of the technical documentation. There is the procedure which uses the MediaWiki API (Application Programming Interface) to copy this to the wiki-based technical documentation system. Such a procedure avoids the data duplication or some misunderstandings between installation procedure itself and their description.

During the release phase, the system maintainer makes a migration from certification into the reference system. This can be done manually (by using written bash scripts during beta-phase) or automatically (for example by cfengine [12]). The technical documentation is updated with some important data and published in the end of this phase.

6. Conclusion

The D-Grid reference installation is used for different purposes. The most important is to test new grid components (or their new versions) and detect if there are compatibility issues or inconsistencies. Additionally the reference installation system can be used for reconstruction of errors as there is a well-defined environment defined. At the moment around 20 Virtual Organizations (VO) are supported by the reference installation. Any member of these VOs can access the reference system by using their user interface and check out the setup. However, the main target group for the reference system consists of the software developers of VOs and system administrators of the resource providers.

References

[1] D-Grid official website: `http://www.d-grid.de/`
[2] D-Grid reference installation official website: `http://dgiref.d-grid.de/`
[3] Open Source monitoring project Nagios: `http://www.nagios.org/`
[4] Open Source Wiki project MediaWiki: `http://www.mediawiki.org/`
[5] Official web service portal for support requests in the German grid environment `https://iwrdgus.fzk.de/pages/home.php`
[6] Information about D-Grid TAB: `http://www.d-grid.de/index.php?id=464&L=1`
[7] Virtualization software VMware ESX `http://www.vmware.com/`
[8] Grid computing middleware gLite `http://glite.web.cern.ch/glite/`
[9] Open source software toolkit used for building grids `http://www.globus.org/toolkit/`
[10] A ready-to-run Grid system including client and server software UNICORE `http://www.unicore.eu/`
[11] Alef M, Fieseler T, Freitag S, Garcia A, Grimm C, Gürich W, Mehammed H, Schley L, Schneider O and Volpato G L 2009 Integration of multiple middlewares on a single computing resource *Future Gener. Comput. Syst.* **25** 268-74
[12] Open source automated configuration tool cfengine: `http://www.cfengine.org/`