



HappyFace Meta-Monitoring for DESY Hamburg Grid Resources

Christian Georg Wehrberger
Georg-August University Göttingen
Germany

September 8th, 2011

Supervisor: Dr. Yves Kemp

Abstract

The stable and reliable operation of all kind of computing resources relies on close and efficient monitoring. The deployment of meta-monitoring systems like *HappyFace* is an approach to improve the monitoring of resource centres in terms of simplicity, effectiveness, and reliability. This report deals with the setup of HappyFace for the DESY Hamburg resource centre and the development of new modules in course of the *DESY Summer Students Programme 2011*.

Contents

1	Introduction	3
2	The HappyFace Project	4
2.1	Meta-Monitoring	4
2.2	The HappyFace Project	4
2.2.1	HappyFace Requirements	4
2.2.2	HappyFace Data Flow	5
2.2.3	HappyFace Configuration	6
2.2.4	Monitored Resources for DESY Hamburg	6
2.2.5	DESY Hamburg HappyFace Setup	7
2.2.6	HappyFace Module Development	7
3	Conclusion and Outlook	10

1 Introduction

After being initiated by the approval of the CERN [1] council in 1994, the LHC [2], a proton-proton collider with a centre of mass energy up to 14 TeV, started to take data in 2009. Since then, all simulation and all experiment data gathered by the four main experiments ALICE [3], ATLAS [4], CMS [5] and LHCb [6] has to be processed, analysed and stored. The process of data taking will be discussed using the example of the ATLAS detector. It consists of the Inner Detector (Pixel Detector [7], Silicon Microstrip Tracker [8], Transition Radiation Tracker [9]), the calorimetry [10], and a muon system [11–14] as outermost layer. At a nominal bunch crossing rate of 40 MHz, the expected number of collisions per second is approximately 23, taking into account a luminosity of $\mathcal{L} = 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ [15]. Three trigger levels in the ATLAS trigger system reduce the unfiltered output rate of about 10^{12} MByte per day by a factor of about $5 \cdot 10^5$ [16], corresponding to an event rate of 100 Hz and 200 MByte/s. In 2010, the overall amount of data being recorded by all four main experiments is was 13 PByte.

In order to store and analyse all data, a distributed computing infrastructure, the *World-wide LHC Computing Grid* [17] (WLCG) had to be set up, making use of over one-hundred computer centres all over the world. One of these computer centres is provided, administered and maintained by *Deutsches Elektron-Synchrotron* [18] (DESY) Hamburg. In the framework of the *Helmholtz Alliance Physics at the TeraScale* [19], partially interacting with grid infrastructures, DESY Hamburg and DESY Zeuthen also provide the *National Analysis Facility* [20] (NAF), providing its resources to German institutions involved in the LHC experiments ATLAS, CMS, and LHCb and also the ILC [21] experiment.

2 The HappyFace Project

2.1 Meta-Monitoring

Meta-monitoring is the monitoring of monitoring resources by a single system. As described more detailed in [22], a meta-monitoring system preferentially fulfils the following requirements:

- *flexibility and genericness*, being adaptable to all sites
- *single point of access*, showing all relevant monitoring information e.g. on one website
- *up-to-date monitoring information*, in the best case real-time
- *history functionality*, giving its users to possibility to review previous monitoring results
- *fast accessibility*, letting the users access the monitoring information quickly
- *accommodativeness*, giving its users easy access to the monitoring outputs
- *simple warning system*, notifying immediately the responsible persons and displaying the status information unmistakeably and simple
- *customisability and extendibility*, matching perfectly the site's monitoring requirements

That is why a meta-monitoring system provides the advantage to easily correlate monitoring information of different sources in order to identify problems earlier and faster in comparison to regular monitoring systems. An intuitive interface ensures the proper usage even for inexperienced users. The single point of access and the displaying of a monitoring summary enable monitoring at a glance. No time and effort has to be put in the browsing of various monitoring sites and systems any more.

2.2 The HappyFace Project

Such a meta-monitoring tool is HappyFace, designed for generic monitoring of computing resources. The core of the project is currently developed at Karlsruhe Institute of Technology [23] (KIT). Modules are developed at DESY Hamburg, KIT, University of Aachen, Georg-August University Göttingen, and University of Hamburg. The current version 2.0 of HappyFace is a stable release, published in 2009.

2.2.1 HappyFace Requirements

HappyFace runs on *Linux/Unix* operating systems. For the execution of the HappyFace framework, *Python 2.5* (or higher), *SQLite* and cronjob execution is needed. The webserver displaying the HappyFace website requires *PHP 5* (or higher).

2.2.2 HappyFace Data Flow

The HappyFace framework providing its main functionalities is called *HappyCore*. Also a lot of basic modules are included in this framework, suitable for all kind of information collection. HappyCore collects the output of all modules, categorises and displays it. Furthermore, it generates several website navigation bars. In figure 1 the GridKa HappyFace website is shown.

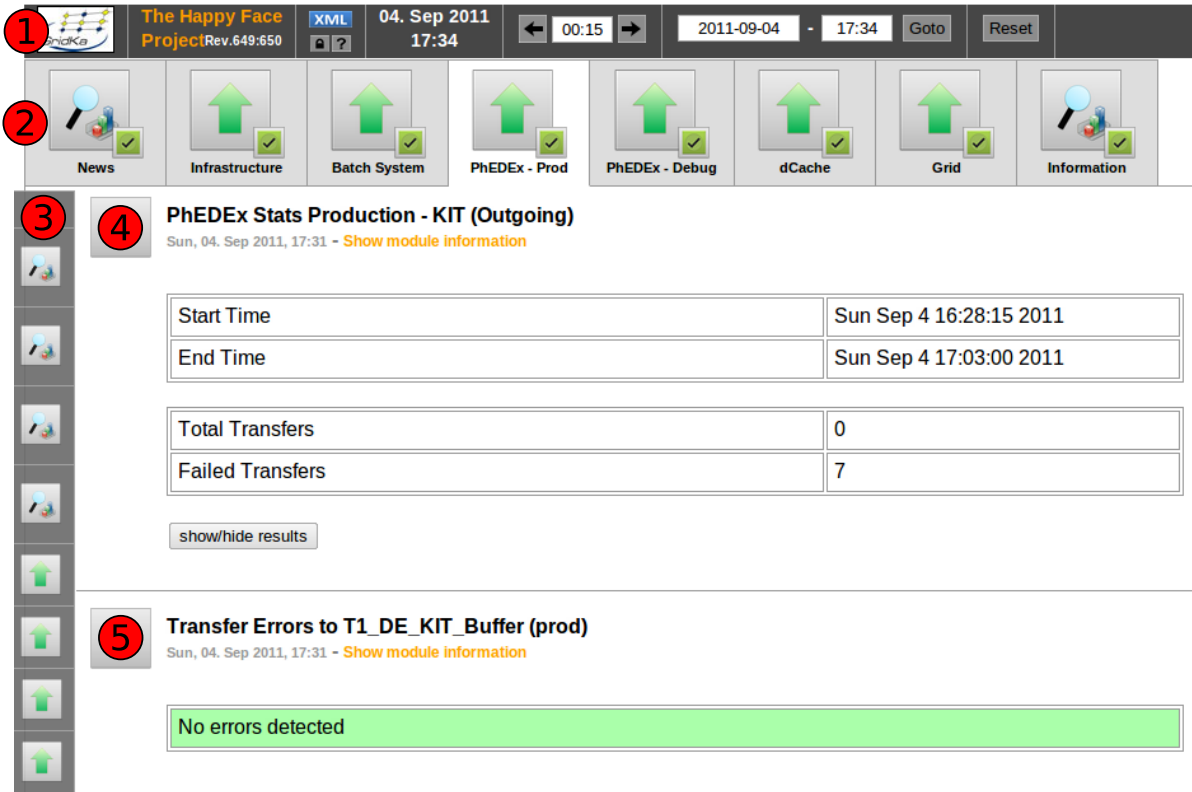


Figure 1: GridKa HappyFace webpage: title bar (1), category navigation bar (2), fast navigation bar (3), specific module (4), specific module (5)

A detailed description of the HappyFace data flow can be obtained from figure 2. Periodically called typically every fifteen minutes, the HappyFace framework is launched by the Python script `run.py`. The monitoring information itself is being collected, processed and written to a PHP output from the various modules written in Python themselves, running in parallel threads. All relevant gathered information is written into a *SQLite* database. Beside the PHP code, the modules also return a status, indicating the success or failure of a monitoring test. The HappyFace webpage, interpreted by a webserver, queries this database for the use of its history functionality.

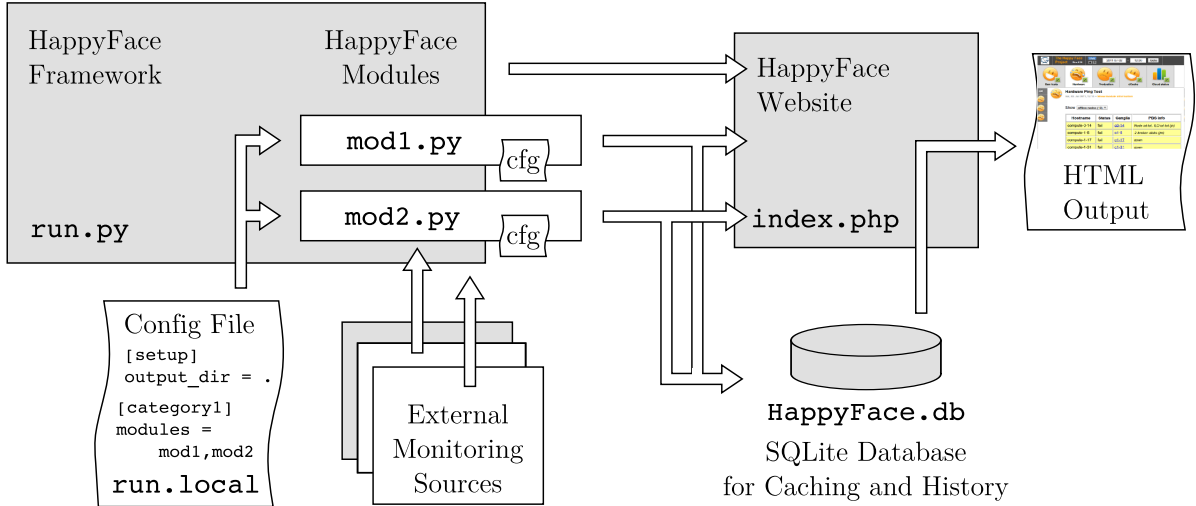


Figure 2: HappyFace data flow

2.2.3 HappyFace Configuration

HappyFace clearly distinguishes between local and global configurations. Besides that, the configuration for `run.py` is separated from the modules' configurations.

All global configuration files are to be found in `HappyFace/config`, whereas module configuration files are located in `HappyFace/local/modules/`. They are named exactly like their corresponding module except their file extension.

In a configuration file, different sections can be specified, indicated by square brackets around the section name (`[section]`). Within those sections, parameter values are assigned using a `name = value` syntax.

2.2.4 Monitored Resources for DESY Hamburg

A central hub for monitoring results at DESY is `grid-mon0.desy.de`. `grid-mon0` is running a webserver, displaying contents in plots and tables. It monitors e.g.

- Berkeley Database Information Index [24] (BDII),
- Computing Elements [25] (CE),
- CreamCE [25],
- Grid Index Information Servers [26] (GIIS),
- LCG File Catalogues [27] (LFC),
- Storage Elements (SE),

- Virtual Organisation Membership Service [28] (VOMS)

and provides various monitoring links. Other information is allocated from `tchobbit.naf.desy.de` and online resources (e.g. <http://lcg-sam.cern.ch>, <http://goc-accounting.grid-support.ac.uk>, and <http://panda.cern.ch>).

2.2.5 DESY Hamburg HappyFace Setup

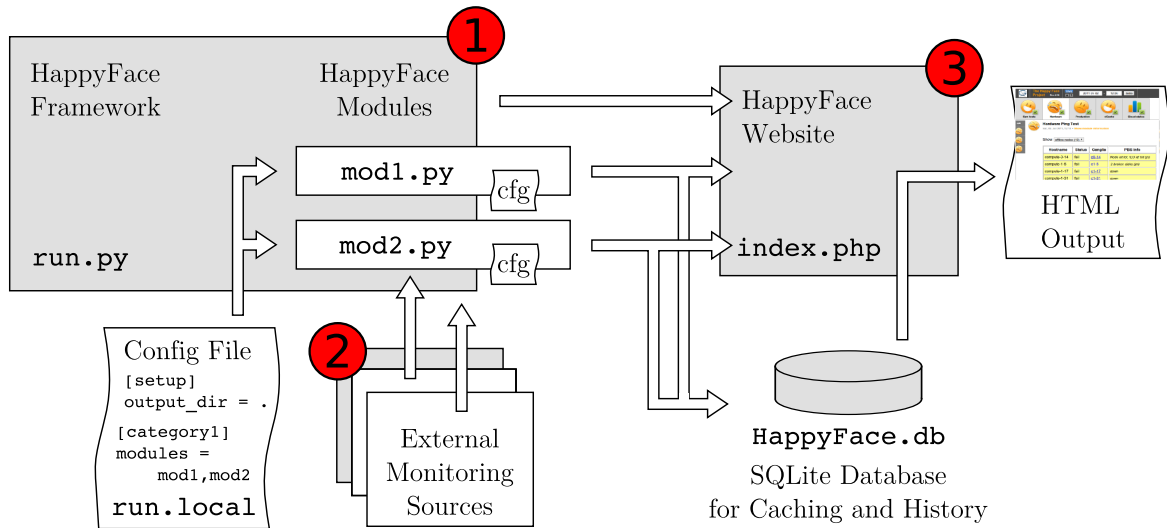


Figure 3: HappyFace data flow at DESY Grid Resources Centre: `grid-happyface.desy.de` (1), `grid-mon0.desy.de`, `tchobbit.naf.desy.de`, online resources (2), `gridcenter.desy.de` (3)

HappyFace features the separation of information collection, processing and displaying. While the relevant monitoring sources is fetched via the DESY *afs* from `grid-mon0.desy.de`, `tchobbit.naf.desy.de`, online resources, etc., the HappyFace framework is executed on `grid-happyface.desy.de`. It creates the HappyFace website, which is interpreted and visualised by `gridcenter.desy.de`. The HappyFace setup is shown in figure 3.

2.2.6 HappyFace Module Development

For the monitoring of DESY grid resources, various modules for different monitoring purposes have been developed.

GIIS Tests This module evaluates the error codes of recent GIIS tests. The alias and both GIIS servers possibly being linked to the alias are being tested. Views for regular grid users and for administrators are provided. The corresponding log file is linked.

Furthermore, the timely duration of the test is taken into consideration for the module status.

SE Tests Testing file accesses via different protocols (*GSIFTP*, *LCG-CR*, and *GSID-CAP*) after a preceding initialisation, a return code is evaluated. Different VOs are taken into account. Log file links are given on the HappyFace webpage.

BDII Tests Evaluating the error codes of recent BDII tests, the alias and all BDII servers possibly being referred to by the alias are being tested. Views both for regular grid users and for administrators are provided. The time duration of the test is taken into consideration for the module status and the test log files are directly linked.

LFC Tests These tests analyse the errors of LFC server tests. The name alias and all individual servers are being tested and test results are displayed separately for regular users and administrators. Different methods and VOs are taken into account.

VOMS Tests For various VOs, several VOMS systems are being tested. The tests' error codes are being rated and define the module's status. Due to the large amount of tests performed for a small time interval, only unsuccessful tests are displayed in the default view.

CE Tests In order to identify problems with CE servers, all DESY LCG and Cream CE hosts are tested. Several subtests provide a more detailed analysis. The corresponding log files are linked.

SE2 Tests These tests monitor SE hosts for VOs *ilc* and *dteam* using two different methods for reading and writing. The corresponding log files and monitoring graphs are directly linked.

SAM Tests The service availability of all tests and for critical tests only are monitored by this modules. The VOs Atlas, CMS and OPS are considered. In case of failure of one of the tests, a link to a more detailed description is given.

Infrastructure Tests In order to monitor the grid infrastructure, several *glite* tests are performed, including proxy initialisation, user interface setup, WMS checks, SE checks, CE checks, and proxy unbinding. Similar tests are performed for the NAF infrastructure.

NAF Infrastructure Tests In this HappyFace module, the NAF infrastructure tests are broken down into their subtests. These tests check critical NAF components and give an overall view of the NAF functionality.

Black Hole Analysis In order to prevent the phenomenon that single compute nodes may consume a large amount of submitted jobs without processing them, a *black hole* analysis is performed. This analysis identifies these compute nodes by comparing the job consumption of each compute node with mean and sigma of an expected distribution. Critical nodes are displayed on the HappyFace webpage.

Plot Grabbing The developed *plot grabber* provides a generic functionality to display all kind of external image content. Defined in the configuration, the user selects post html parameters via drop down menus; The number and kind of parameters is variable.

3 Conclusion and Outlook

Beside recently developed modules, a large variety of already existing modules can be set up easily. Available modules cover e.g. hardware monitoring, dCache [29] monitoring, and PhEDEx [30] for CMS data transfers.

Another very useful tool are *heatmaps*, offering a quick view over the time dependency of a test result. The development of a generic module could cover a large range of modules used at the DESY Hamburg resource centre. Figure 4 displays an exemplary heatmap design.

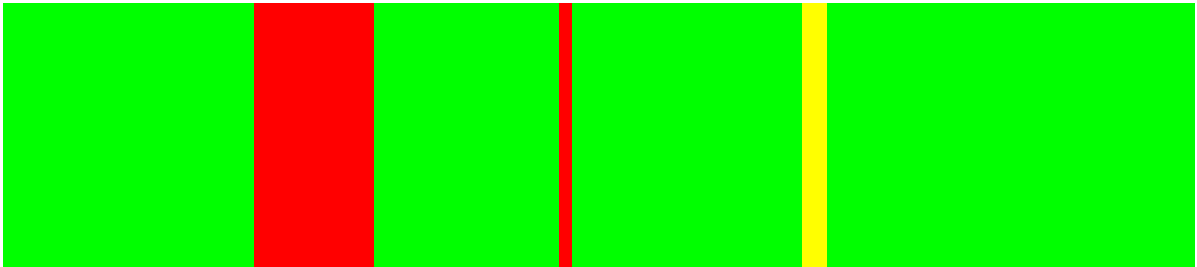


Figure 4: Heatmap exemplary design

As always online tablet computers and smartphones become more and more common, a HappyFace application for such devices is reasonable. It would compensate for the inability of HappyFace to trigger alarms in case of critical situations and promises great benefits in terms of real-time monitoring.

Acknowledgements

Firstly, I like to take this opportunity to thank Dr. Yves Kemp who supervised me during my stay in Hamburg and offered me a warm welcome. He has made available his support in a number of ways. I would also like to thank the whole IT department for their help and advice.

My sincere thanks go to Dr. Olaf Behnke who managed to organise the DESY Summer Students Programme 2011 in such a great way.

Last but not least I want to overstate my gratitude to all summer students for the experience, joy and fun they offered. Thank you!

List of Figures

1	GridKa HappyFace webpage http://ekphappyface.physik.uni-karlsruhe.de	5
2	HappyFace data flow [31]	6
3	HappyFace data flow [31] at DESY Grid Resources Centre	7
4	Heatmap exemplary design	10

References

- [1] *European Organization for Nuclear Research (CERN)*, URL <http://public.web.cern.ch/public/>
- [2] L. Evans, *The large hadron collider project*, CERN Reports pages 275–286 (1997)
- [3] ALICE Collaboration, *ALICE Technical Design Report*, CERN/LHCC 2001-021 (2001)
- [4] ATLAS Collaboration, *ATLAS: technical proposanal for a general-purpose pp experiment at the large hadron collider at CERN*, CERN/LHCC pages 171–173 (1994)
- [5] CMS Collaboration, *CMS - Technical Proposal*, CERN/LHCC/94-38 (1994)
- [6] LHCb Collaboration, *LHCb Technical Proposal*, CERN/LHCC 98-4 (1998)
- [7] G. Aad, M. Ackers, F. Alberti, M. Aleppo, G. Alimonti, J. Alonso, E. Anderssen, A. Andreani, A. Andreazza, J. Arguin, et al., *ATLAS pixel detector electronics and sensors*, Journal of Instrumentation **3**, P07007 (2008)
- [8] A. Ahmad, Z. Albrechtskirchinger, P. Allport, et al., *The silicon microstrip sensors of the ATLAS semiconductor tracker*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **578(1)**, 98 (2007)
- [9] T. Åkesson, F. Anghinolfi, E. Arik, O. Baker, S. Baron, D. Benjamin, H. Bertelsen, V. Bondarenko, V. Bytchkov, J. Callahan, et al., *Status of design and construction of the Transition Radiation Tracker (TRT) for the ATLAS experiment at the LHC*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **522(1-2)**, 131 (2004)
- [10] The Tile Calorimeter Group of the ATLAS Collaboration, *The Production and Qualification of Scintillator Tiles for the ATLAS Hadronic Calorimeter*, Technical Report ATL-TILECAL-PUB-2007-010. ATL-COM-TILECAL-2007-026, CERN (2007)
- [11] K. Nagai, *Thin gap chambers in ATLAS*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **384(1)**, 219 (1996)
- [12] A. D. Simone, *The calibration of the resistive plate chambers of ATLAS*, Journal of Physics: Conference Series **219(3)**, 032036 (2010), URL <http://stacks.iop.org/1742-6596/219/i=3/a=032036>
- [13] M. Livan, *Monitored drift tubes in ATLAS*, Nuclear Instruments and Methods in Physics Research-Section A Only **384(1)**, 214 (1997)

- [14] P. O'Connor, V. Gratchev, A. Kandasamy, V. Polychronakos, V. Tcherniatine, J. Parsons, W. Sippach, *Readout electronics for a High-Rate CSC detector*, in *Fifth Workshop on Electronics for LHC Experiments*, Snowmass Colorado (1999)
- [15] C. Lefevre, *LHC: the guide (english version)* (2008), <http://cdsweb.cern.ch/record/1092437/files/CERN-Brochure-2008-001-Eng.pdf>
- [16] *ATLAS fact sheet*, URL http://atlas.ch/pdf/atlas_factsheet_all.pdf
- [17] M. Delfino, L. Robertson, *Solving the LHC Computing Challenge: A Leading Application of High Throughput Computing Fabrics combined with Computational Grids*, CERN-IT-DLO-2001-003 (2001)
- [18] *Deutsches Elektron-Synchrotron (DESY)*, URL <http://www.desy.de/>
- [19] *Helmholtz Alliance Physics at the TeraScale*, URL <http://www.terascale.de/>
- [20] A. Haupt, Y. Kemp, *Journal of Physics: Conference Series* **219**, 052007 (2010)
- [21] A. Djouadi, J. Lykken, K. Mnig, Y. Okada, M. Oreglia, S. Yamashita, *International Linear Collider Reference Design Report Volume 2: PHYSICS AT THE ILC*, Phys. Rev. D 71 page 015004 (2005)
- [22] V. Büge, V. Mauch, G. Quast, A. S. and Artem Trunov, *Site specific monitoring of multiple information systems - the HappyFace Project*, *Journal of Physics Conference Series* **219**, 062057 (2010)
- [23] *Karlsruhe Institute of Technology (KIT)*, URL <http://www.kit.edu/index.php>
- [24] *Berkeley Database Information Index (BDII)*, URL <https://twiki.cern.ch/twiki/bin/view/EGEE/BDII>
- [25] C. Aiftimiei, P. Andreetto, S. Bertocco, S. Fina, S. Ronco, A. Dorigo, A. Gianelle, M. Marzolla, M. Mazzucato, M. Sgaravatto, *Job submission and management through web services: the experience with the CREAM service*, *Journal of Physics: Conference Series* **119**, 062004 (2008)
- [26] K. Czajkowski, S. Fitzgerald, I. Foster, C. Kesselman, *Grid information services for distributed resource sharing*, in *High Performance Distributed Computing, 2001. Proceedings. 10th IEEE International Symposium on*, pages 181–194, IEEE (2001)
- [27] J. Baud, J. Casey, S. Lemaitre, C. Nicholson, *Performance analysis of a file catalog for the LHC computing grid*, in *hpdc*, pages 91–99, IEEE (2005)
- [28] R. Alfieri, R. Cecchini, V. Ciaschini, L. DellAgnello, A. Frohner, A. Gianoli, K. Lorentey, F. Spataro, *VOMS, an authorization system for virtual organizations*, in *Grid Computing*, pages 33–40, Springer (2004)

- [29] M. Ernst, P. Fuhrmann, M. Gasthuber, T. Mkrtchyan, C. Waldman, *dCache, a distributed storage data caching system*, CHEP2001 Conference Notes pages 152–156 (2001)
- [30] J. Rehn, T. Barrass, D. Bonacorsi, J. Hernandez, I. Semeniouk, L. Tuura, Y. Wu, *PhEDEx high-throughput data transfer management system*, Computing in High Energy and Nuclear Physics (CHEP) 2006 (2006)
- [31] G. Jahn, *Meta-Monitoring of the Grid Resource Centre GoeGrid with HappyFace*, Bachelor's Thesis, Georg-August University Göttingen (2011)