2011 DESY Grid Center Review

May 24, 2011

Abstract

The DESY Grid center includes the Grid Sites in Hamburg and Zeuthen and the National Analysis Facility (NAF), an analysis facility for the German LHC community, integrated in and complementing the DESY-HH and DESY-ZN Grid infrastructure.

A review of the services provided and functionality, capacity and performance upgrades required was performed end of 2010. This report summarizes the findings and recommendations of this review.

Members of the review team:

Gerhard Brandt (ATLAS), Ulrich Husemann (ATLAS), Matthias Kasemann (CMS, chair), Hartmut Stadie (CMS), Kai Leffhalm (DV), Yves Kemp (IT) and Christoph Wissing (IT).

Table of contents:

Introduction and executive summary

The DESY Grid center includes the Grid Sites in Hamburg and Zeuthen and the National Analysis Facility (NAF), an analysis facility for the German LHC community, integrated in and complementing the DESY-HH and DESY-ZN Grid infrastructure. Major users communities of the Grid center are the LHC experiment groups ATLAS, CMS and, to a smaller extend, LHCb. The HERA experiments H1 and ZEUS and the ILC and Calice groups are also significant users of the DESY Grid Center. Together with several smaller communities they use up to 15% of the total capacity provided.

The Grid Center consists of Grid computing infrastructure running the gLite middleware and resources for the NAF. For the Virtual Organizations (VO) from the HERA experiments and the ILC activities the DESY Grid Center provides in addition to CPU and storage resources central Grid services like VO registration services, central file catalogs and work load management systems.

The DESY Grid Centre is part of the World-Wide LHC Computing Grid and operates very reliable. On the international level it is among the most heavy used Tier-2 centers for LHC data analysis. The NAF is augmenting the Tier-2 center functionality for end-user analysis; it is well adopted by German groups. A significant role in the success of the DESY Grid center plays the long-standing and successful participation in the developments and operations of global HEP Grid projects, where DESY is a major player worldwide. In the year 2010 the DESY Grid infrastructure was heavily used and contributed significantly to produce first results based on LHC data. Wide area network connections and bandwidth available has proven adequate for now to allow fast and timely data transfers and access to remote data.

Against this background a review was instantiated in fall 2010 to examine the usage profile of the existing Grid Center resources, to review existing monitoring tools and recommend performance metrics and monitoring values and to examine user community needs and recommend investments for the next 5 years. The results are documented in a report, addressed to DESY FH management.

In Chapter 1 the existing resources are summarized for fall 2010 and after an upgrade in early 2011; chapter 2 summarizes the findings and recommendations for metrics and monitoring of the Grid Center performance and usage. Required capacity upgrades are recommended in chapter 3 separated for the Tier-2 center, where the upgrade should follow the approved plans of the Worldwide LHC Computing Grid project (WLCG) and for the NAF. For the NAF a list of strategic and operational recommendations is given in section 3.3. They are based on the experience of intensively using the infrastructure for LHC Data analysis and should help to increase reliability and performance and to shape and configure the infrastructure to meet the requirements of high luminosity data analysis in the coming years. It is also recommended to increase the teams across the experiments and IT to strengthen operation, support and active development of the infrastructure. The capacity upgrades recommended in chapter 3.4 complete the recommendations of this report, for 2011 a NAF capacity increase of about a factor of two is recommended. For the following years the upgrades should be similar to the Tier-2 upgrades.

The mandate of this review is presented in chapter 4, together with the members of the review team and a description of the execution of the review. A section with links to further information and an overview of German Tier-2 center capacities in Appendix A complete the report.

1. Review of existing DESY Grid Center resources

The existing DESY Grid Center resources installed in Hamburg and Zeuthen in November 2010 are listed in Table 1. The CPU capacity is approximate, since the computer models differ from different purchases. The Tier-2 capacity is pledged towards the WLCG. CPU capacity of the sites is measured in HEPSPEC06 (HS06), based on performance benchmarks carefully checked against typical HEP MC and reconstruction programs. Storage capacity is measured in units of TeraBytes (TB) or PetaBytes (PB). Not included are non-Grid resources for Theory and Photon science groups, local farms for HERA experiments, work group servers, and desktops.

	November 2010			April 2011		
	DESY-HH	DESY-ZN	NAF	DESY-HH	DESY-ZN	NAF
CPU [#slots]	4224	784	1620	4224	784	1620
CPU total [kHS06]	33.9 kHS06	8.6 kHS06	22 kHS06*	33.9 kHS06	8.6 kHS06	22 kHS06*
ATLAS T2 pledged CPU	4.8 kHS06			6.2 kHS06		
CMS T2 pledged CPU	8.0 kHS06			11.8 kHS06		
LHCb T2 pledged CPU	3.2 kHS06			3.2 kHS06		
ATLAS disk total (dCache/Lustre)	500 TB	518 TB	78 TB	890 TB	750 TB	135 TB
ATLAS T2 disk pldeged (dCache)	740 TB			1050 TB		
CMS DISK total (dCache/Lustre)	660 TB*		58 TB	1057 TB*		58 TB
CMS T2 DISK pledged (dCache)	400 TB			640 TB		
LHCb Disk total (dCache/Lustre)		163 TB	47 TB		180 TB	47 TB
LHCb T2 Disk pledged (dCache)	2 TB			2 TB		
Disk (other VOs, dCache/Lustre)	140 TB	357 TB	33 TB	348 TB	500 TB	33 TB

***: includes UniHH resources of 240 TB of dCache disk and 7.4 kHS04 of NAF CPU**

Table 1: Existing DESY Grid Center resources in November 2010 and April 2011.

The DESY Grid center resources are used by the LHC experiments (ATLAS, CMS and LHCb), but also by the general ILC, the CALICE, the H1 and ZEUS groups. For the NAF the resource allocation is about 4:2:1:1 for ATLAS : CMS : LHCb : ILC. The Astroparticle groups (IceCube and CTA) use non-Grid computing resources to a large extend not included in the table.

The NAF local batch CPU resources and the dCache capacities were upgraded in January 2011 after the initial information was collected for the review; the up-to-date status was shown at the DESY PRC in April 2011, see rightmost columns in Table 1.

2. Grid Center performance metrics and monitoring values

A large amount of information is regularly collected and made available by DESY and as part of the EGEE and WLCG grid projects for the Tier-2 part of the Grid Center. Information is provided at different levels of details for VO administrators, service providers and end users. Some relevant links to available monitoring information are given in the links section; some parts are only available in the DESY domain. For several of these values a more widespread availability would help users to understand the behavior and performance of the systems, especially when faced with performance bottlenecks.

For the review the team evaluated the existing monitoring information for the services, the usage pattern and the resource loading of the DESY Grid Center infrastructure. The goal was to recommend metrics and graphs for summary information showing usage and resource load relevant for the users performing analysis, for reviews and to prepare decisions for major upgrades. Since the user community is very heterogeneous with different external milestones, schedules and sometimes funding lines, it is required to provide most of these information for each VO separately as well as globally for the DESY Grid Center.

The review team recommends providing the following plots in addition to the existing information on a regular and continuous basis.

Monitoring CPU capacity:

For the end user the total time to successful job completion is the relevant value, i.e. the sum of waiting time in the queue plus processing time. This time is impacted by overall system load, free CPU capacity, I/O bandwidth for input and output as well as system services failures and downtimes.

It is recommended to make available on a continuous basis actual measurements and history plots of the ratio of <job waiting time>/<job running time> for each of the VOs and summed up for the NAF.

To evaluate the total system CPU capacity the number of job slots available and used by different VOs is a good measure. For an analysis infrastructure like the NAF, it is recommended that the load does not exceed 75% regularly for more than few days with peaks up to 90% in order to guarantee the required response for a service used interactively to a large extend, Figure 2 below shows this plot for the period October 2010 until March 2011.

It is recommended to make the plot in Figure 2 available to all users on a continuous basis.

Monitoring and managing disk usage and data transfer bandwidth

Storage capacity is allocated to VOs configured in disk pools according the experiments needs. Monitoring of the pool occupancy is provided by the dCache service; see links in the section below. Disk space is a precious resource and should be used and actively managed according to the priority of the data sets for the analysis as identified in the corresponding VO and user community. Tools, which regularly collect access statistics, should help to decide which data set to keep on disk.

A detailed monitoring of the bandwidth used to transfer data from the storage pools to the processing nodes helps to identify possible limitations in the LAN infrastructure. A fine-grained monitoring of the used bandwidth of the WAN connections is needed to compare the actual usage with the expectations from the plans from the experiments.

It is recommended to provide access statistics and detailed network usage reports regularly to the data managers in the experiments.

Monitoring usage shares

The NAF services are provided for the whole German HEP community. The usage by institute can be monitored and plots produced like Figure 1, which shows, that about 60% of resources are provided to users and groups from German institutes outside DESY.

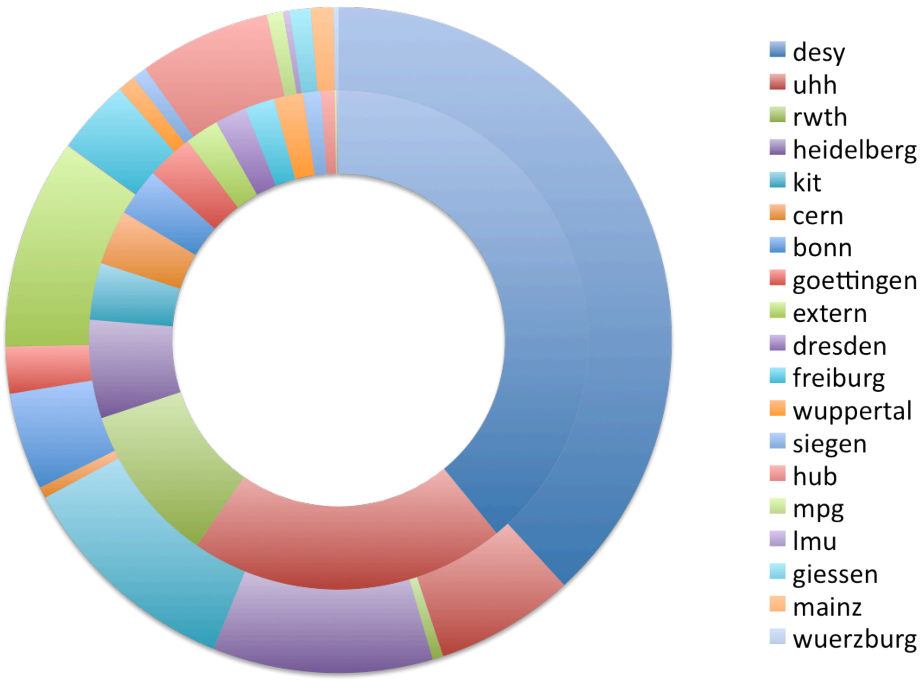

Figure 1: NAF usage by German institute, inner circle: wall clock time for jobs, outer circle: Lustre storage. Institutes are listed clock-wise, starting at 12:00h. (V.Gülzow, DESY PRC, April 28, 2011)

3. Proposed plans for 2011 – 2014

Every year the LHC experiments calculate the resources required to perform MC production and data analysis, based on the LHC running schedule and their computing models for the following two years (N+1 and N+2). These requests are reviewed by the C-RSG (Computing Resource Scrutiny Group) review group, which reports their findings and recommendation to the CERN Computing Resource Review Board (C-RRB), which subsequently approves the requirements. The centers involved pledge capacity for the fall C-RRB every year, at which time the expected total capacity for the following year (N+1) is known to the experiments. For the years further into the future it is obvious that data analysis will continue, but the schedule for new data recording is not known in detail, since the LHC schedule is only known tentatively. No detailed resource planning is attempted, only rough estimates based on experience and knowledge of hardware replacements is known. For 2011 the request and pledge process finished in October 2010, currently the revised and approved 2012 requests are awaiting pledges from the sites. For 2013 the experiments estimated their probable needs, the approval and pledge process will follow in 2012.

The HERA experiments H1 and ZEUS and the ILC and Calice groups are also significant users of the DESY Grid Center. Together with several smaller communities they use up to 15% of the total capacity provided. The expected computing and storage capacity requirements for these groups are stable or require only a very moderate upgrade of the allocated routed resource, included in the proposed NAF plans below.

3.1 Tier-2 plans for 2011

The DESY pledges for 2011 and 2012 on the basis of the WLCG 2010 requests are listed in Table 3. The DESY Tier-2 capacity for 2011 was already provided by April 2011, as foreseen in the annual WLCG upgrade process.

	2011 (CPU)	2011 (disk)	2012 (CPU)	2012 (disk)
ATLAS DESY T2 (pledged)	6.2 kHS06 (+29%)	1050 TB (+42%)	6.6 kHS06	1350 TB
CMS DESY T2 (pledged)	11.8 kHS06 (+47%)	640 TB (+60%)	12.9 kHS06	900 TB
LHCb DESY T2 (pledged)	3.2 kHS06 (+0%)	2 TB (+0%)	3.2 kHS06	2 TB
Total pledged	21.2 kHS06 (+33%)	1692 TB (+48%)	22.7 kHS06	2252 TB

Table 3: DESY Tier-2 center pledges in response to the 2010 WLCG requirements estimates. For 2011 the relative increase compared to the 2010 capacity is indicated. The 2012 pledges listed are outdated; they were revised based on updated requirements presented at the C-RRB in April 2011 (see Table 4).

3.2 Proposed Tier-2 plans for 2012 – 2014

The LHC running schedule has substantially changed since early 2010. It is now planned to extend the running at center of mass energy of 7 TeV and luminosities of the order of a few times 10^{33} cm^{-2} s^{-1} with 50 nsec bunch spacing until at least end of 2012, with the possibility to run until middle of 2013. A long shutdown of 18 months will follow to prepare the machine for 14 TeV running and to upgrade various detector components in the experiments. Including expected commissioning time, the time from physics data recording before the shutdown until after the shutdown may be as long as 24 months.

The experiments revised their resource requirements for the 2011 review and pledge process taking these plans into account. Table 4 lists the relative annual DESY Tier-2 capacity increase request, based on the updated WLCG requirements as presented at the C-RRB in April 2011 (see links section for further documentation). LHCb is not requesting any upgrade for the global Tier-2 center capacity. The total DESY increase requested is the weighted sum of ATLAS, CMS and LHCb increase. DESY will have to pledge resources by summer 2011 for the C-RRB in October 2011.

	2012 (CPU)	2012 (disk)	2013 (CPU)	2013 (disk)
global ATLAS T2 increase	6%	25%	7%	13%
global CMS T2 increase	10%	31%	0%	0%
global LHCb T2 increase	33%	0%	0%	0%
Total DESY T2 increase	22%	27%	1%	4%

Table 4: CRSG-recommended increase of WLCG Tier-2 center capacity for 2012 and 2013 (preliminary). The total DESY increase requested is the weighted sum of ATLAS, CMS and LHCb increase. Pledges will by collected for the C-RRB in October 2011.

For 2013 the CMS experiment plans require a Tier-2 center capacity increase of about 30% for CPU and 20% for disk storage. It is planned to use large parts of the T0 and CAF resources at CERN (up to 120 kHS06 and up to 5.5PB of disk) for analysis during the LHC shutdown period in 2013-14 when they are not needed for data collection. Correspondingly there is no capacity increase required for CMS Tier-2 centers in 2013.

Maintaining low trigger thresholds for efficiently detecting decays of a low mass Higgs presents a challenge for the LHC experiments at high luminosity and high pile-up conditions, if the rest of the LHC physics program is not to be severely condensed. In particular in 2012 the experiments have to be prepared to face luminosities of up to 5^{33} $cm^{-2}s^{-1}$ and pile-up approaching up to 16 pp collisions per bunch crossing. CMS presented a plan to the LHC Committee in March 2011 to increase the data-taking rate from nominally 300 Hz to 400 Hz to extend the physics reach. This plan will be reviewed at subsequent LHCC reviews based on experience gained running and triggering with higher luminosities. Would this be approved the corresponding Tier-2 requirements for 2012 are bigger than listed in Table 3 by 12% for CPU and 30% for disk storage, for 2013 the requirements grow by 12% for CPU and 20% for disk. The decision and a more detailed CMS request are expected in time for the C-RRB in October 2011. Concrete plans from ATLAS are not known the to review team, but similar measures and consequences for the required resources for analysis are to be expected.

In Appendix A the capacities of the German ATLAS and CMS Tier-2 centers are shown as pledged to WLCG for 2009-2011. The total number of nominal Tier-2 centers in Germany is 3.5 for the ATLAS experiment and 1.5 for CMS. Resources are adding up to about the relative participation of Germany in the international experiments, which is 11.5% for ATLAS and 7.8% for CMS. DESY provides the capacity of one nominal Tier-2 center each for ATLAS and CMS. Based on C-RRB documents April 2011 the necessary pledges for the years 2012-13 are indicated which would keep the relative German strength. Currently DESY contributes 35-40% to the total German Tier-2 center capacity.

Recommendation:
The DESY Grid Center review team recommends increasing the DESY Tier-2 capacities for the years 2012 and 2013 according to the plans presented at the C-RRB in 2011. The plans for 2013 will become final and preliminary 2014 requirements will be presented to the C-RRB in April 2012.

Due to changes in the LHC running schedule and the high luminosity provided in 2011 it is expected that the Tier-2 resources pledged for 2011 might fall short. This needs to be taken into account when planning the 2011 NAF resources, see below.

3.3 NAF strategic and operational recommendations

The NAF provides computing resources to the German groups to play a leading role in the data analysis. In 2008 the NAF capacity was designed to be roughly equivalent to 1.5 average Tier-2 centers, with a special focus on data storage.

End-user analysis, in contrast to data reconstruction and simulation, uses high statistics data samples of relative small event sizes (20-150kb). This requires a high I/O rate capacity of the storage infrastructure. At the NAF storage is provided by the dCache storage element shared with the DESY Tier-2 instances. This provides Grid-enabled access to all Tier-2 data sets. To support interactive and local batch based analysis a high-performance Lustre storage infrastructure is provided. In addition AFS is used to store user results.

The DESY ATLAS and CMS groups collected feedback on the experience performing analysis at the Tier-2 and the NAF. The results of these investigations will be discussed with IT in detail. The following list identifies the major observations and recommendations of the experiments groups and the review team:

1. The I/O performance observed on the Tier-2 and NAF is not always stable and does not match the expected needs for 2011 data analysis. Reasons could be manifold, from poor object alignment in files, too many parallel accesses to single files, too few replicas of frequently

accessed files, hardware and configuration bottlenecks and others.
Recommendation: Investigate and improve storage I/O performance in a team effort of experienced users performing analysis and IT.

2. Several storage systems are used for analysis (dCache, Lustre, afs). Frequently analysis use cases depend on more than one file system. The reliability of the afs file system was not at a satisfying level. This is absolutely essential to produce results quickly, reliably and predicable in the competitive environment of LHC experiments.
Recommendation: The stability of storage systems involved needs to be monitored closely and improved with high priority.
3. Besides stability and performance issues for the Lustre file system the long-term support of this product is not assured after the product was purchased by Oracle.
Recommendation: Perform prototype studies and systematic analysis case studies to prepare the decision on future file systems matching the performance and capacity needs for the NAF.
4. A sizeable fraction of the analysis steps are performed accessing the resources interactively. Code development and debugging tools, as well as desktop sharing, document browsing and paper writing tools need to be made available in the NAF to provide a direct coupling from the desktop to the analysis environment used.
Recommendation: Provide the necessary workgroup server like, interactive services as identified by the users.
5. Some experiment specific, dedicated analysis and processing steps require specialized computers, like the precision tracker calibration or a dedicated machine for analysis job submission and report logging. Several of these machines were setup and provided as part of the NAF. Additional needs will arise in future with further development of more advanced analysis steps.
Recommendation: Continue to support specific, dedicated computing for experiments analysis steps, either as part of the NAF or elsewhere at DESY.
6. The NAF is used in a highly collaborative mode; storage needs to be allocated to groups and individuals. Users expressed the need for a quota system with several TB of storage capacity on the high performance analysis file system (Lustre for now).
Recommendation: Together with the users identify a solution and provide adequate functionality and storage space.
7. Questions related to operations of the NAF are regularly handled at the NAF Users Committee (NUC). It is essential that the NAF concept is actively developed and extended and the performance optimized in close collaboration with its users. In some cases systematic studies need to be performed, like for the high performance file system (see above). The performance of the NAF is best tested and optimized with an instrumented benchmark tool, based on a real physics analysis. These studies need to be performed and steered in a coordinated way, either through extending the scope of the NUC or by a technical investigation group.
The GridLab at DESY, which was presented at he PRC in April 2011, can be the infrastructure to perform many of the proposed systematic studies, some others needs to be performed under realistic load conditions on the production sites.
Recommendation: Coordinate and build teams with experiment and IT members to perform technical studies to prepare NAF upgrade and extensions decisions with schedule and work plans, reporting results publicly. An analysis benchmark tool should be used to perform system optimizations.

Highest priority for the Grid center infrastructure is to maintain and consolidate the services at a high level, to increase reliability and availability to the required level and allow for faster troubleshooting. Additional work needs to be done both in developing system health monitoring and failure recovery tools as well as in direct user support and debugging of code and system failures. To implement the recommendations listed above and invest into the improvements the currently existing manpower is deemed not sufficient. Additional manpower should strengthen the teams across the experiments and IT to use the investment in the infrastructure most efficiently.

Recommendation of the review team:
The NAF support team should be increased by 2 FTE and work closely with the experiments groups using the Grid Center for analysis.

3.4 NAF capacity recommendations

As evaluated in November 2010 users from German universities performing individual and physics group directed analysis consume about 60% of the delivered analysis capacity, see Figure 1 above.
It is expected that the user base for the NAF may grow slowly in numbers for the coming years when more scientists shift their focus towards analysis.

The NAF capacity in November 2010 is shown in Table 1, see above. The CPU occupancy for the last 6 months is shown in Figure 2. An increase of CPU capacity is visible for 2011 where older hardware was replaced.

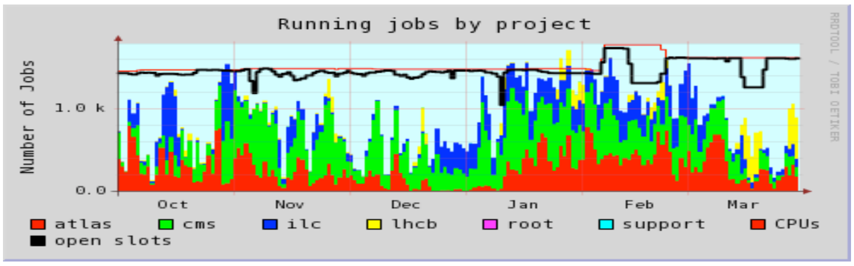

Figure 2: NAF CPU occupancy, the different colors indicate usage by experiment

Clearly visible is the time dependant occupancy when results for conferences need to be finished, which indicates the need for an immediate upgrade of the NAF resources in order to provide sufficient resources for the 2011 luminosity increase, expected to be at least a factor of 30.

Recommendation of the review team:
The recommended occupancy level should not go regularly above 75%. An analysis facility, which is significantly used in interactive mode, looses the required responsiveness at load levels above 75%.

The capacity scaling factor for LHC data analysis depends on the data volume to be analyzed and is very similar to the specified increase for the Tier-2 centers. For the year 2011 this would require a total increase of at least 33% for CPU and 50% of disk capacity. For the year 2011 such upgrades are considered too small, see below.

The experiments plans for 2011 were made early 2010 based on a different LHC schedule. When the LHC schedule was revised the requirements setting and approval process for 2011 was finished already. The LHC continues to perform well; the experiments will collect significantly more data at higher instantaneous luminosity than planned in 2010. It is therefore not unexpected that the Tier-2 center capacity pledged for 2011 will fall short. For 2011 it is therefore not sufficient to align the NAF upgrade to the relative 2011-increase of the Tier-2 center. The German user community would benefit substantially from a substantially bigger upgrade. Only this will provide sufficient capacity for analysis of a larger data set data recorded under high specific luminosity conditions in 2011, i.e. with high pile-up and will provide additional analysis capacity, which will be lacking at Tier-2 centers.

Based on the experience from the analyses of the 2010 data the ATLAS and CMS storage needs are reevaluated for the NAF: The total size of data needed by ATLAS and CMS for analysis is expected to be 1 PB for each experiment for 2011. For fast data access of commonly used data sets 10 % of the data should be on a High-I/O storage, this results in 200 TB of

additional High-I/O storage. Part of the interesting data set will already be stored at the DESY Tier-2 sites. As the sites have only limited influence on the data placement on the pledged resources, we assume an overlap of 200 TB for ATLAS and 300 TB for CMS, respectively. Furthermore, additional user space is needed for the output files of the individual analyses. This is estimated as 300 TB for ATLAS and 150 TB for CMS: Hence, 1.95 PB of dCache space and 200 TB additional High-I/O storage are requested in total for the ATLAS and CMS experiments in 2011. CPU should grow accordingly and should double in 2011.
For ILC and LHCb we expect only moderate growth rates. Together with the storage the I/O bandwidth capacity must grow, to keep the turn around time for analysis about constant. Once, the NAF resources have been adapted to the new requirements, we expect an average growth rate similar to the overall experiment's needs (e.g. Tier-2 growth). Table 5 lists these recommended upgrades.

	2011 CPU	2011 storage dCache / High-I/O	2012 CPU	2012 storage	2013 CPU	2013 storage
ATLAS	-	1100 / 200 TB	+25%	+25%		
CMS	-	850 / 150 TB	+31%	+31%		
ILC, LHCb	-	200 / 50 TB	+10%	+10%		
Total	+100%	2150 / 400 TB	+25%	+25%	+25%	+25%

Table 5: Recommended NAF capacity upgrades based on experiments requirements. 2013 values will be revised in 2012 when the experiments planning process finishes. For 2011 only the total CPU increase is given. For 2013 details are not known yet and will depend on the LHC running performance and schedule in 2012 and 2013, the rough expected total growth is indicated.

Recommendation:
The DESY Grid Center review team recommends for 2011 increasing the NAF capacities by about a factor of two. This is substantially more than the requested relative Tier-2 center upgrade to compensate for the expected shortage of analysis capacity at the Tier-2.

For the following years the annual NAF capacity increase should be aligned with the total annual Tier-2 center capacity upgrade of the experiments, including Tier-2 capacity at CERN.

4. Mandate and execution of the review

The scope of the review was to

a. examine the usage profile of the existing Grid Center resources
b. review existing monitoring tools and recommend performance metrics and monitoring values as well for generic utilization and for more in-depth operations control
c. examine user community needs and recommend investments for the next 5 years.

The results should be documented in a report, addressed to DESY FH management.

Members of the review team were Gerhard Brandt (ATLAS), Ulrich Husemann (ATLAS), Matthias Kasemann (CMS, chair), Hartmut Stadie (CMS), Kai Leffhalm (DV), Yves Kemp (IT) and Christoph Wissing (IT).

The team met several times in the months September-December 2010. Detailed feedback on the experience and needs for LHC data analysis was collected from ATLAS and CMS groups. Information was collected by email from the ILC, Calice and HERA groups. Information was collected and documented on an internal web page. Draft recommendations were formulated towards the end of 2010. Included in the report are the updated computing resource requests from ATAS and CMS based on the modified LHC running conditions for the next years, which became available in March/April 2011.

Links to further information

- DESY Grid Center web page: http://gridcenter.desy.de/
- NAF website: http://www.naf.desy.de/
- WLCG accounting portal: http://www3.egee.cesga.es/gridsite/accounting/CESGA/egee_view.html
- EGEE Grid health map: http://gridmap.cern.ch/gm/#topo=regions&layout=tc&vo=OPS&serv=Site&sitenames
- DESY Hamburg dCache disk services and monitoring, visible on the DESY intranet:
    - ATLAS: http://dcache-ses-atlas.desy.de:2288/
    - CMS: http://dcache-ses-cms.desy.de:2288/
    - General: http://dcache-ses-desy.desy.de:2288/
- WLCG utilization monitoring and resources planning: http://lcg.web.cern.ch/LCG/resources.htm
- April 2011 C-RRB documents: https://indico.cern.ch/conferenceDisplay.py?confId=128046

Appendix A: Overview of German Tier-2 center capacities

	2009-pld	2010-pld	2011-pld	2012-req	2013-req	2012-400Hz	2013-400Hz
C int'l [HS06]	116000	196221	315202	350000 (+12% p.a.)	350000 (+0% p.a.)	410000 (+30% p.a.)	410000 (+0% p.a.)
C int'l [TB]	8400	13627	20219	26000 (+30% p.a.)	26000 (+0% p.a.)	32000 (+58% p.a.)	32000 (+0% p.a.)
C: Desy [HS06]	4400	8000	11800	13216	13216	15340	15340
C: Desy [TB]	230	400	640	832	832	1011	1011
RWTH [HS06]	2200	5000	6600	7392	7392	8580	8580
RWTH [TB]	115	250	330	429	429	521	521
D-CMS [HS06]	6600	13000	18400	20608	20608	23920	23920
D-CMS [TB]	345	650	970	1261	1261	1533	1533
C: Desy/d-tot [HS06]	66,7 %	61,5 %	64,1 %	64,1 %	64,1 %	64,1 %	64,1 %
C: Desy/d-tot [TB]	66,7 %	61,5 %	66,0 %	66,0 %	66,0 %	66,0 %	66,0 %
C: D/int'l [HS06]	5,7 %	6,6 %	5,8 %	5,9 %	5,9 %	5,8 %	5,8 %
C: D/int'l [TB]	4,1 %	4,8 %	4,8 %	4,9 %	4,9 %	4,8 %	4,8 %
A int'l [HS06]	114000	214976	281228	295000 (+5% p.a.)	321000 (+9% p.a.)		
A int'l [TB]	11200	21238	34203	49000 (+43% p.a.)	56000 (+14% p.a.)		
A: Desy [HS06]	2400	4800	6200	6510	7095,9		
A: Desy [TB]	290	740	1050	1502	1712		
Wu [HS06]	1800	3436	4633	4865	5302		
Wu [TB]	219	369	633	905	1032		
Fr [HS06]	1800	3436	4610	4841	5276		
Fr [TB]	219	369	518	741	844		
LMU [HS06]	1800	3440	9200	9660	10529		
LMU [TB]	220	370	1040	1487	1695		
Goe [HS06]	1200	2280	3800	3990	4349		
Goe [TB]	150	250	400	572	652		
MPI [HS06]	2116	3440	0	0	0		
MPI [TB]	265	370	0	0	0		
D-ATLAS [HS06]	11116	20832	28443	29865	32553		
D-ATLAS [TB]	1363	2468	3641	5207	5936		
A: Desy/d-tot [HS06]	21,6 %	23,0 %	21,8 %	21,8 %	21,8 %		
A: Desy/d-tot [TB]	21,3 %	30,0 %	28,8 %	28,8 %	28,8 %		
A: D/int'l [HS06]	9,8 %	9,7 %	10,1 %	10,1 %	10,1 %		
A: D/int'l [TB]	12,2 %	11,6 %	10,6 %	10,6 %	10,6 %		
D: Desy/d-tot [HS06]	38,4 %	37,8 %	38,4 %	39,1 %	38,2 %		
D: Desy/d-tot [TB]	30,4 %	36,6 %	36,7 %	36,1 %	35,3 %		

German T2 centers for ATLAS, CMS, pledged capacity for 2009-2011, required pledge to keep relative strength, based on C-RRB documents April 2011.
Green: CMS experiment, Blue: ATLAS experiment, Gray: DESY Tier-2 capacity in total.