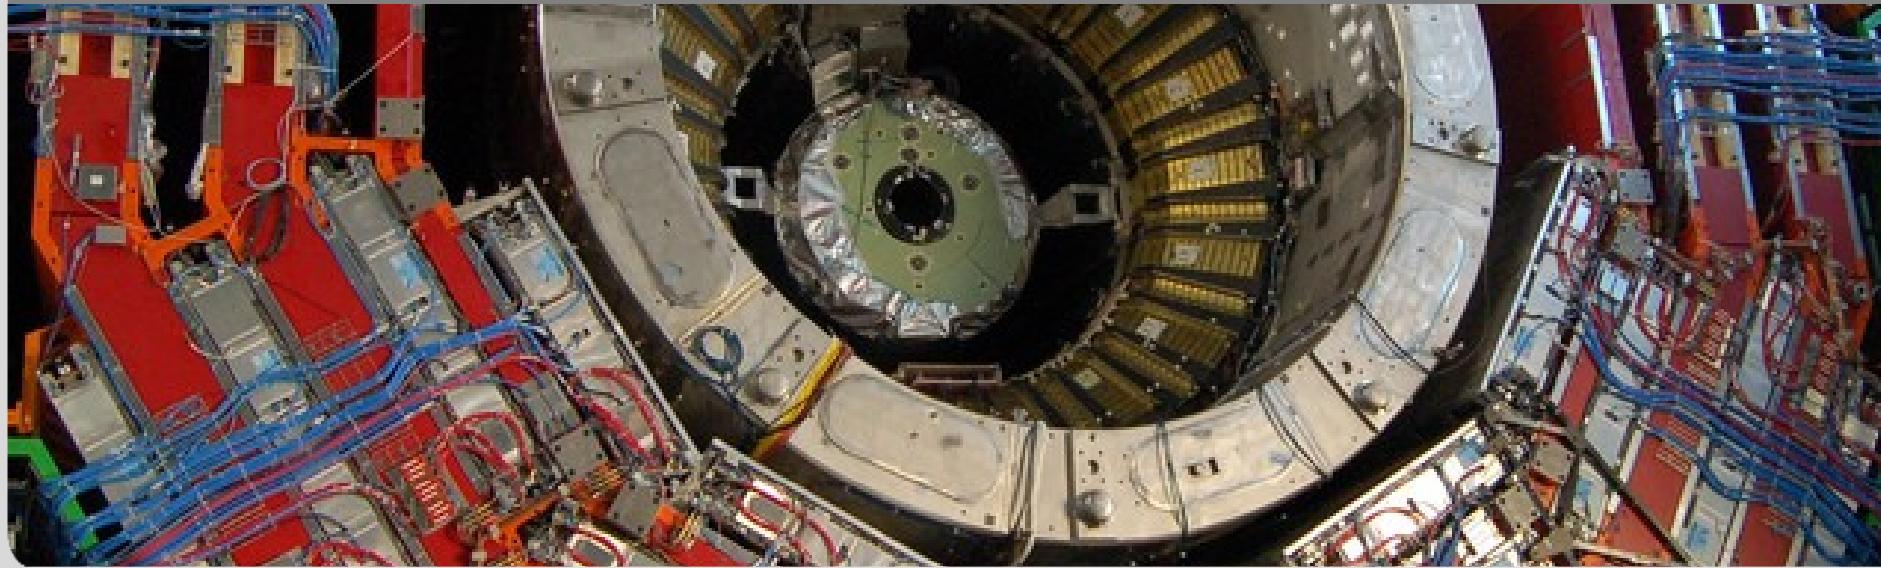


# The HappyFace Project – A Meta Monitoring Tool

## Monitoring Sites, Services, Jobs

Armin Burgmeier, for the HappyFace Developers

INSTITUT FÜR EXPERIMENTELLE KERNPHYSIK, KIT



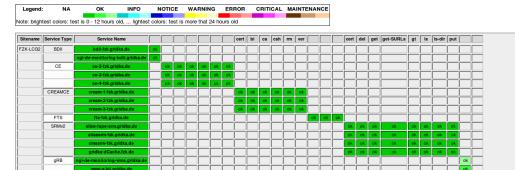
# Outline

- Introduction
  - Current Monitoring Situation
  - How Meta Monitoring improves the Situation
- The HappyFace Project
  - Overview and Features
  - Architecture
  - Installation
  - Recent and ongoing development
- Live Demonstration
- Summary and Outlook

# What Needs to be Monitored?

A Grid site consists of **many critical components** each of which needs to be **monitored** to make sure the center performs well.

- Grid Infrastructure (Software failures)
- Batch System (Inefficient jobs, failing jobs)
- Storage System (Staging failures, Orphaned files, Replicas, Pools utilization)
- Data Transfer (Slow transfers, failed transfers)



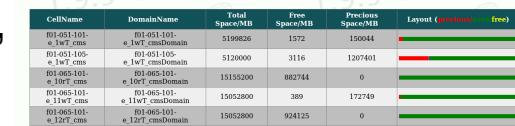
Date: 2010-09-03 15:34:18

Site	DDM / dh		Jobs	Inst. Prod. SW	SAM	GangaRobot
	Eff. (%) / Rate(MB/s)	Eff. (%) / Done / Run				
CSCS(Contact)	100.0 / 28	0.0 / 0 / 119	[15.9.0.3 / 105171]	Warn	OK	ok
CYFRONET(Contact)	100.0 / 37	0.0 / 0 / 0	[15.9.0.3 / 11713171]	OK	OK	ok
DESY-HH(Contact)	99.6 / 99	100.0 / 253 / 1167	[15.9.0.3 / 107171]	OK	OK	ok
DESY-ZN(Contact)	99.9 / 80	0.0 / 0 / 242	[15.9.0.3 / 1061971]	OK	OK	ok
FZK(Contact)	95.1 / 24	99.6 / 2702 / 3471	[15.9.0.3 / 105171]	OK	OK	ok
GOREGRID(Contact)	100.0 / 24	89.3 / 192 / 722	[15.9.0.3 / 90171]	OK	OK	ok
HEPHY-UIBK(Contact)	100.0 / 0	0.0 / 0 / 220	[15.9.0.3 / 194171]	Warn	OK	warn

Birds Home



## Disk Space Usage

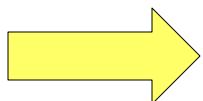


## PhEDEx – CMS Data Transfers



Time span: Left Now □ Include time sets nothing but errors □ Update

Last Hour	File	Total File Name	Errors	Elapsed	Avg. Elap.	Rate	Avg. Elap.	Rate	Elapsed
T1_FZK_CERNF1_0008	T1_FZK_CERNF1_Expert	26 1.77B 400.2 MB/s	-	2	506.74MB/s	3940			
T1_FZK_CERNF1_0008	T1_FZK_CERNF1_Buffer	141 1.77B 400.2 MB/s	-	1	114.34MB/s	3940			
T1_FZK_CERNF1_0008	T1_FZK_CERNF1_Buffer	150 1.69B 399.9 MB/s	-	1	68.45MB/s	3940			
T1_FZK_CERNF1_0008	T1_FZK_CERNF1_Expert	111 4.69B 130.0 MB/s	-	8	73.24MB/s	1040			
T1_FZK_CERNF1_0008	T1_FZK_CERNF1_Buffer	141 1.77B 400.2 MB/s	-	1	114.34MB/s	3940			
T1_FZK_CERNF1_0008	T1_FZK_CERNF1_Buffer	45 20.0B 304.0 MB/s	2	-	23.44MB/s	3940			
T2_DE_DESY	T2_DE_CERNF1_Buffer	89 17.6B 304.0 MB/s	1	-	52.04MB/s	3940			
T1_FZK_CERNF1_MSS	T1_FZK_CERNF1_Buffer	36 40.5B 47.1 MB/s	-	1	114.45MB/s	3940			
T1_FZK_CERNF1_MSS	T1_FZK_CERNF1_Buffer	141 1.77B 400.2 MB/s	-	1	114.34MB/s	3940			
T2_PR_DAF_1LB	T2_TW_Tunnel	40 44.6B 41.5 MB/s	-	-	22.34MB/s	2160			
T1_FZK_CERNF1_Buffer	T1_CERN_CERN_Expert	27 19.0B 304.0 MB/s	2	-	23.45MB/s	3940			
T1_DKE_PSI_MSS	T1_DKE_PSI_Buffer	35 17.1B 304.0 MB/s	-	-	26.04MB/s	3940			
T1_DKE_PSI_MSS	T1_CERN_CERN_Expert	27 19.0B 304.0 MB/s	-	-	24.44MB/s	3940			



All Information is **available** but **distributed**

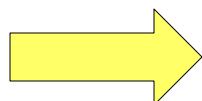
# So what's the Problem?

Overall Monitoring is **inconvenient** at best:

- Manage **many browser tabs** or windows
- Each website has its **own settings** (time ranges, Grid site, ...) that needs to be set up



- **Long loading times**, can be more than 30 seconds
- Difficult to identify **correlations** between failures at different services or sites
- Hard to get a **quick overview** of a site's status, especially for non-experts



**Solution: Important information available **at a single place****

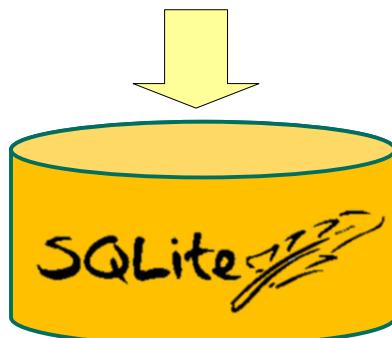
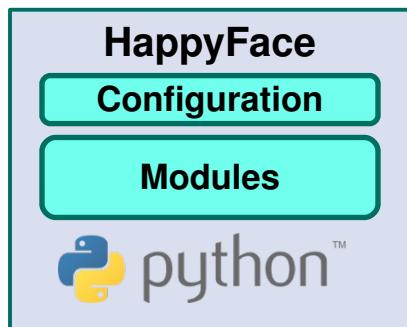
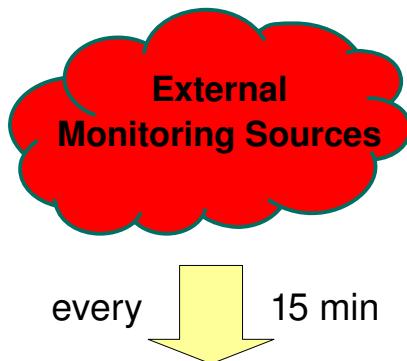
# The solution: The HappyFace Project Ver. 2

## HappyFace **features**:

- Collect important Monitoring information from multiple services
  - No generation of new data, only process **available data**
  - Runs in a **short time interval**, O(15 min)
- Present the current status of the services in a **compact way**
- Easy **rating system** to quickly see whether a service is good or critical
  - Detailed information can then be obtained via the source website
- **Modular Layout**: Individual tests or plots can easily be enabled or disabled; new ones can be developed without interfering with the rest of the system
- **History navigation**: It is possible to go back in time and check each module's status and other output, including plots
- Generate **time-dependent plots** of recorded quantities



# The HappyFace Project Architecture

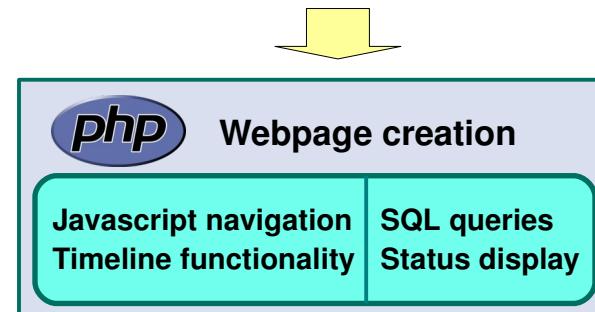


## HappyFace Core:

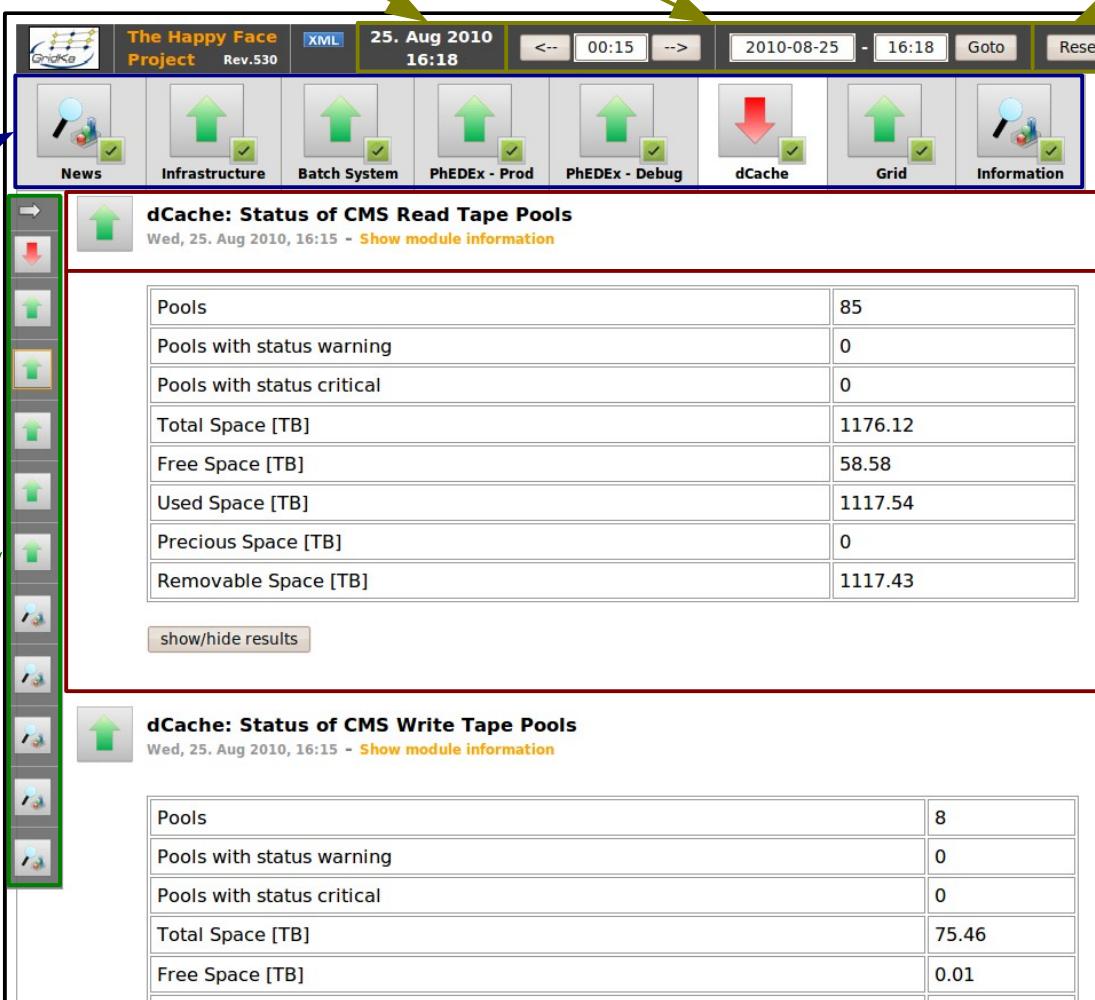
- Read **global configuration** and execute all enabled **modules**

## HappyFace Modules:

- Read **module configuration**
- Download **monitoring data** (XML, Plots, ...)
- Run the test and compute module **rating**
- Write results into **database** and plots into **filesystem**
- Write PHP code for **web output**



# Overview of the HappyFace website



The screenshot shows the HappyFace website interface with various annotations:

- Date and Time of current view**: Points to the top center showing "25. Aug 2010 16:18".
- History navigation**: Points to the navigation buttons above the main content area.
- Go to most recent version**: Points to the "Reset" button in the top header.
- Category Selection and Status**: Points to the left sidebar with category icons and status indicators.
- Module Title, Status and Execution Time**: Points to the title and data of the first module, "dCache: Status of CMS Read Tape Pools".
- Module Output**: Points to the table data of the second module, "dCache: Status of CMS Write Tape Pools".
- Category Navigation and Module Status**: Points to the left sidebar with category icons and status indicators.

**dCache: Status of CMS Read Tape Pools**  
Wed, 25. Aug 2010, 16:15 - [Show module information](#)

Pools	85
Pools with status warning	0
Pools with status critical	0
Total Space [TB]	1176.12
Free Space [TB]	58.58
Used Space [TB]	1117.54
Precious Space [TB]	0
Removable Space [TB]	1117.43

[show/hide results](#)

**dCache: Status of CMS Write Tape Pools**  
Wed, 25. Aug 2010, 16:15 - [Show module information](#)

Pools	8
Pools with status warning	0
Pools with status critical	0
Total Space [TB]	75.46
Free Space [TB]	0.01

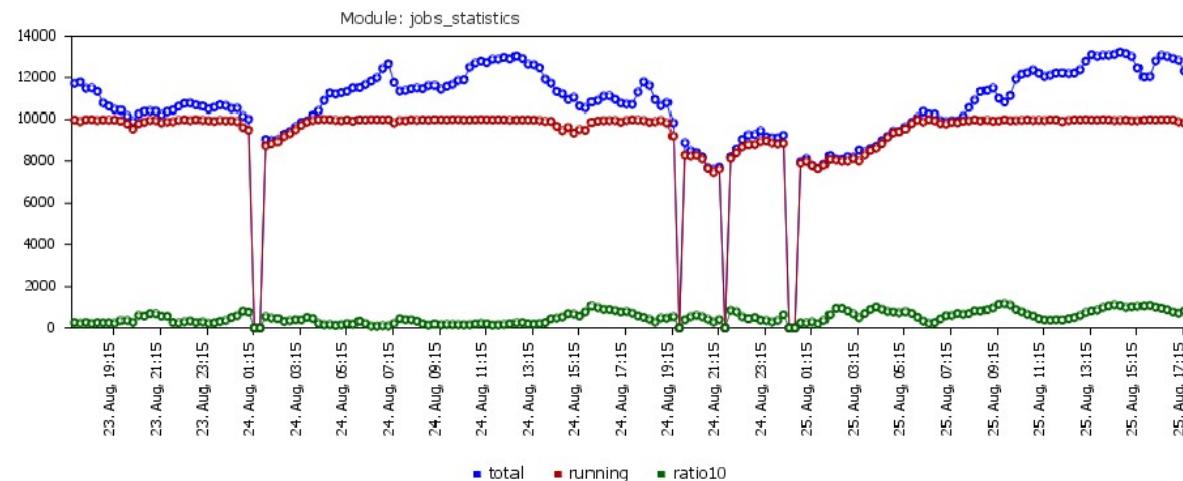
# History Plots

The **history** of each module's status and, depending on the module, other numerical variables can **be plotted**.

## Module: jobs\_statistics

Start: 2010-08-23 17:23 Stop: 2010-08-25 17:23 Variable: **total** + 2 more  
Interval: 48 hours Stop: now Legend: ● bottom ○ inside ○ right  
Interval: 48 hours Stop: 2010-08-25 17:23 Show Plot

### Variable(s): total, running, ratio10



Here: Number of Jobs at GridKa

- All jobs
- Running jobs
- Jobs with walltime ratio below 10%

# HappyFace Installation

Setting up an own HappyFace instance is **easy**:

- Make sure a **webserver**, **Python**, **PHP** and **SQLite** are installed
- Check out the **HappyFace code**:

```
svn co https://ekptrac.physik.uni-karlsruhe.de/public/HappyFace/trunk myHFinstance
```

- Add it to the **crontab** so that it runs every 15 minutes:

```
*/15 * * * * cd /path/to/myHFinstance/HappyFace && ./run.py >/dev/null 2>&1
```

- Configure **categories** in local/cfg/run.local
- Configure **modules** in local/cfg/myModule.local
- **Example modules** and **configuration** files available
- Full **documentation** available at  
[https://ekptrac.physik.uni-karlsruhe.de/trac/HappyFace/wiki/Version\\_2](https://ekptrac.physik.uni-karlsruhe.de/trac/HappyFace/wiki/Version_2)



# Example modules: dCache Data Transfers module

Watch **file transfers** from dCache pools

 **dCache Transfers**

Thu, 26. Aug 2010, 16:00 - [Show module information](#)

Total number of transfers	102
Speed average [KB/s]	945
Standard deviation of speed distribution [KB/s]	3703
Number of transfers with warnings	83
... due to time limit	0
... due to speed limit	83
Number of critical transfers	66
... due to time limit	0
... due to speed limit	66

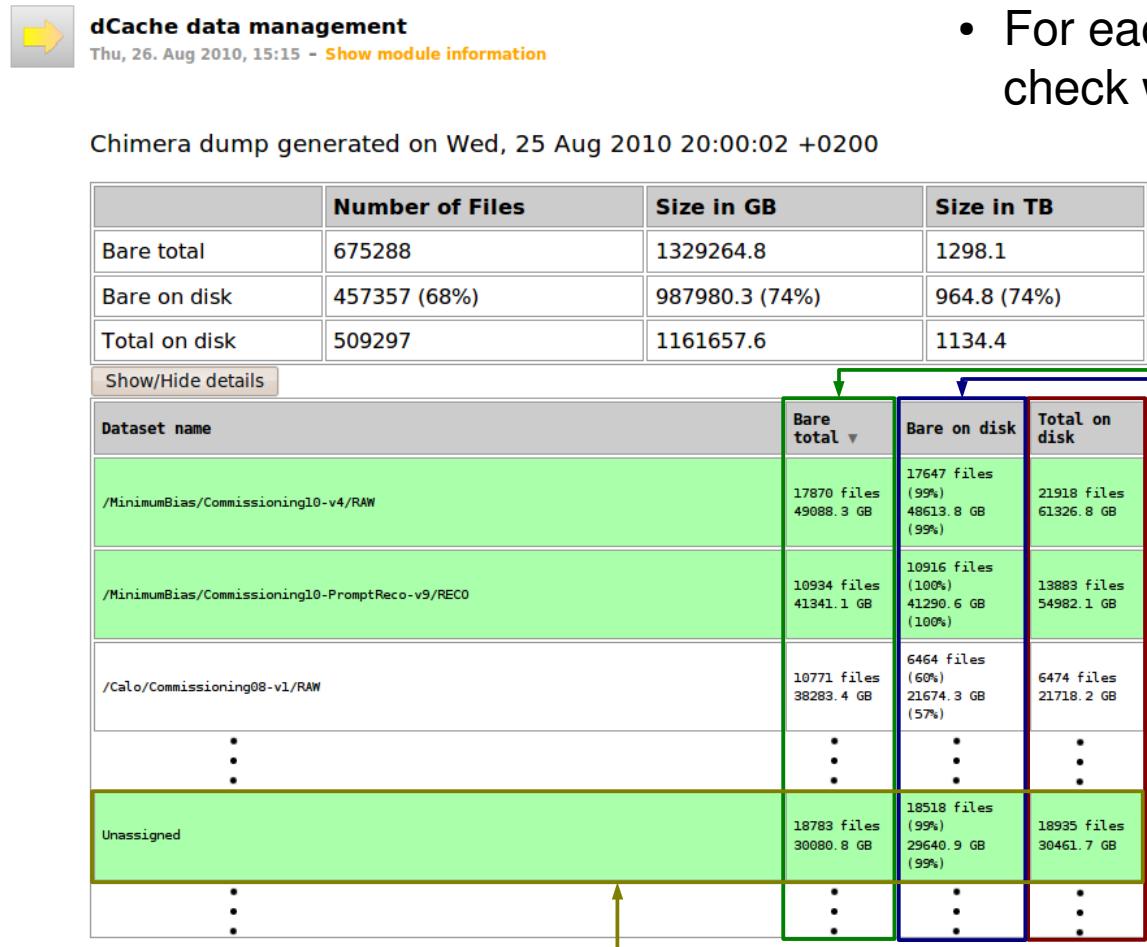
[show/hide results](#)

pnfsID	Pool	Host	Status	Time	Trans. [GB]	Speed [KB/s]
00002C9BE1087FF34E03B82E51AFCA35FDA7	f01-070-101-e_7rT_cms	c01-001-120.gridka.de	WaitingForDoorTransferOk	0 d 09: 33: 54	1. 6	50
0000301FC8B7FFE84CD1BF3BF1088D6C7BF0	f01-140-130-e_lrT_cms	c01-016-116.gridka.de	WaitingForDoorTransferOk	0 d 07: 45: 32	1. 4	51
00001500FCBA02764D5FA350311CEE92D9B6	f01-120-109-e_lrT_cms	c01-010-108.gridka.de	WaitingForDoorTransferOk	0 d 08: 06: 38	1. 5	55
0000B22CD76A7EB5403682ED7774C097365C	f01-140-105-e_10rT_cms	c01-013-168.gridka.de	WaitingForDoorTransferOk	0 d 10: 59: 35	2. 1	55

- Warns if transfers take **much time**
- Warns if transfers are **slow**
- Indicator for inefficient jobs
- Correlation with **problematic pools or host systems** is revealed

# Example modules: dCache Data Management module

Compare files known by **dCache** with the CMS DBS database



The screenshot shows a comparison between dCache and CMS DBS datasets. It includes a summary table and a detailed dataset comparison table.

**Summary Table:**

	Number of Files	Size in GB	Size in TB
Bare total	675288	1329264.8	1298.1
Bare on disk	457357 (68%)	987980.3 (74%)	964.8 (74%)
Total on disk	509297	1161657.6	1134.4

**Detailed Dataset Comparison:**

Dataset name	Bare total	Bare on disk	Total on disk
/MinimumBias/Commissioning10-v4/Raw	17870 files 49088.3 GB	17647 files (99%) 48613.8 GB (99%)	21918 files 61326.8 GB
/MinimumBias/Commissioning10-PromptReco-v9/RECO	10934 files 41341.1 GB	10916 files (100%) 41290.6 GB (100%)	13883 files 54982.1 GB
/Calo/Commissioning08-v1/Raw	10771 files 38283.4 GB	6464 files (60%) 21674.3 GB (57%)	6474 files 21718.2 GB
⋮	⋮	⋮	⋮
Unassigned	18783 files 30080.8 GB	18518 files (99%) 29640.9 GB (99%)	18935 files 30461.7 GB
⋮	⋮	⋮	⋮

- For each file known by dCache check which **dataset** it belongs to
  - Serves as a **consistency cross check**
  - Datasets **fully on disk** are interesting to end-user analyses

Size of dataset in #files and GB

Part of dataset present on disk

Total space occupied on disk, including replicas

Orphaned files (available in dCache but not in any dataset)

# Example modules: Jobs Statistics module

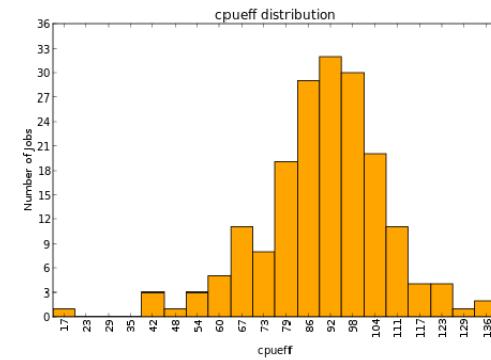
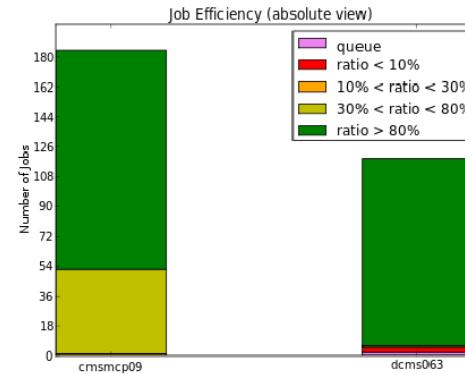
Monitors **Jobs** running at a Grid site

- Standalone producer generates **common XML** file by querying the batch system; producers are already available for PBS, Condor.
- HappyFace module reads the XML, then **visualizes** the information and stores it into its database
- **Decoupling** of data generation and data visualization
- Warn if there is a significant number of **inefficient jobs** in a group

GridKa Jobs Statistics  
Fri, 27. Aug 2010, 10:00 ~ Show module information

Start:	2010-08-25	10:06	End:	2010-08-27	10:06	Show Trend plot
<b>Group</b>	<input checked="" type="checkbox"/> Total jobs	<input type="checkbox"/> Running jobs	<input checked="" type="checkbox"/> Jobs with wallratio < 10%	<input type="checkbox"/> Plot jobs		
	Toggle Selection	Plot Col	Plot Col	Plot Col	Plot Selected	
<input checked="" type="checkbox"/> all	9746	9719	1951		Plot Row	
<input checked="" type="checkbox"/> cms	303	301	3		Plot Row	
<input checked="" type="checkbox"/> cmsother	0	0	0		Plot Row	
<input checked="" type="checkbox"/> cmsproduction	184	184	0		Plot Row	
<input checked="" type="checkbox"/> cmst1p	0	0	0		Plot Row	
<input checked="" type="checkbox"/> cmsmcp	184	184	0		Plot Row	
<input checked="" type="checkbox"/> dcms	119	117	3		Plot Row	

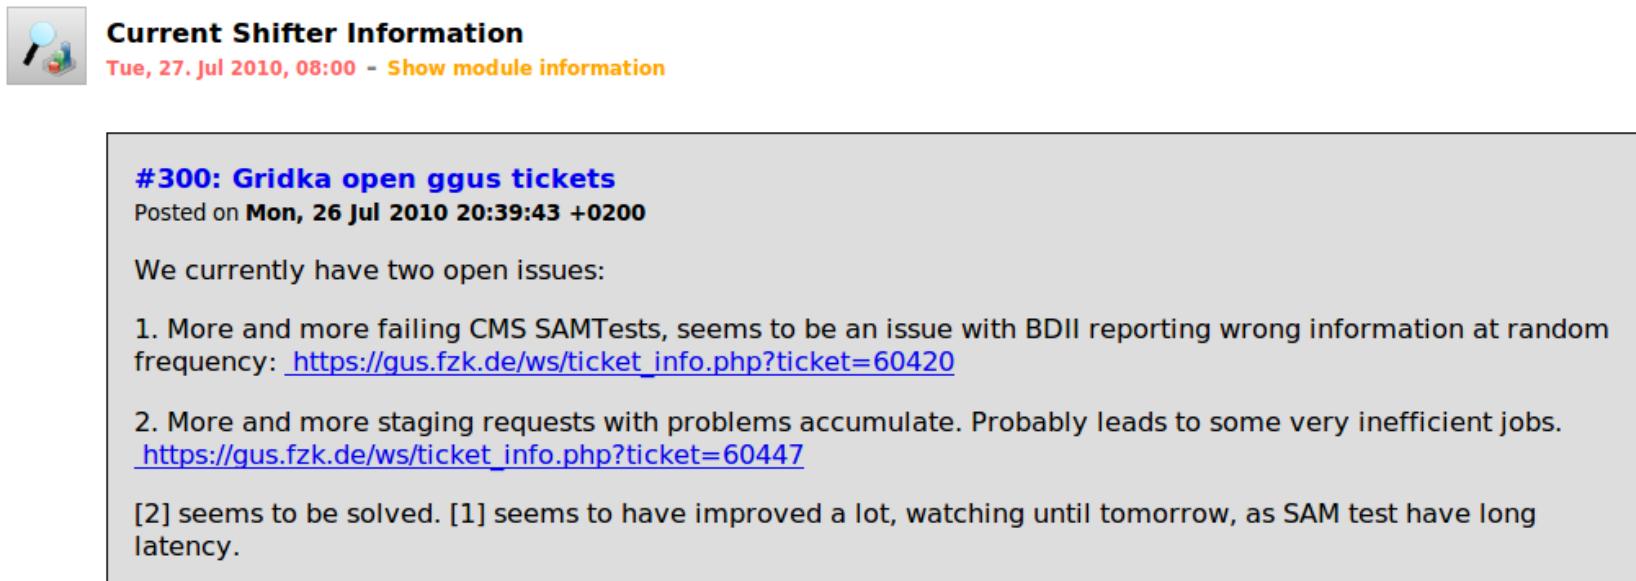
show/hide details



# Example modules: RSS module

Allows to display **RSS feeds** as a non-rated HappyFace module

- Inform shifters about general news or **known problems**
- Useful for **external monitoring sources** and service portals which use RSS
  - For example, show **open tickets** of relevant services



**Current Shifter Information**  
Tue, 27. Jul 2010, 08:00 - Show module information

**#300: Gridka open gus tickets**  
Posted on Mon, 26 Jul 2010 20:39:43 +0200

We currently have two open issues:

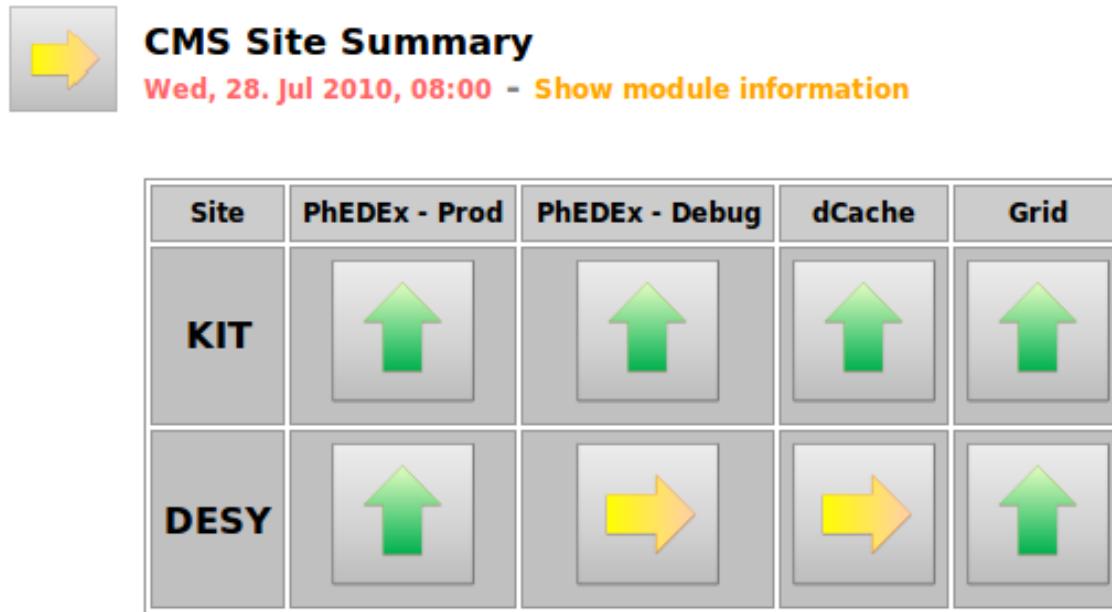
1. More and more failing CMS SAMTests, seems to be an issue with BDII reporting wrong information at random frequency: [https://gus.fzk.de/ws/ticket\\_info.php?ticket=60420](https://gus.fzk.de/ws/ticket_info.php?ticket=60420)
2. More and more staging requests with problems accumulate. Probably leads to some very inefficient jobs. [https://gus.fzk.de/ws/ticket\\_info.php?ticket=60447](https://gus.fzk.de/ws/ticket_info.php?ticket=60447)

[2] seems to be solved. [1] seems to have improved a lot, watching until tomorrow, as SAM test have long latency.

# Example modules: Summary module

HappyFace can export aggregated information as **XML**

- The **Summary** module makes use of it by showing a quick overview of **many HappyFace instances**: *Meta<sup>2</sup> Monitoring*

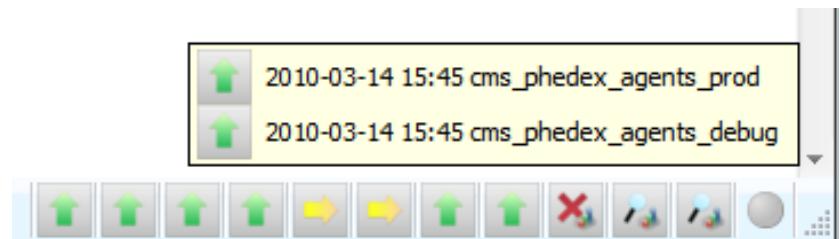


- Ideally this allows a **small shift crew** to supervise a **large number of sites**

# More modules

There are many more modules available in HappyFace:

- dCache pools information:
  - Rating on various variables, for example **free space/total space**
  - Verify **load balancing**
- dCache dataset restore monitor (**staging**)
- SAM Tests:
  - Test that software is **available** and **well-behaving** on worker nodes
  - OPS and Experiment specific tests
- Site Readiness statistics
- PhEDEx **transfer errors** and statistics
- Firefox Plugin



# In development: T2 User Space Monitoring

Information about used **disk space per user**

- Security: **Authentication** is realized via CA
- User mode: Only user information available
- Admin mode: Detailed information about each user
- Rating based on number of users **over quota**

## User Space Monitoring

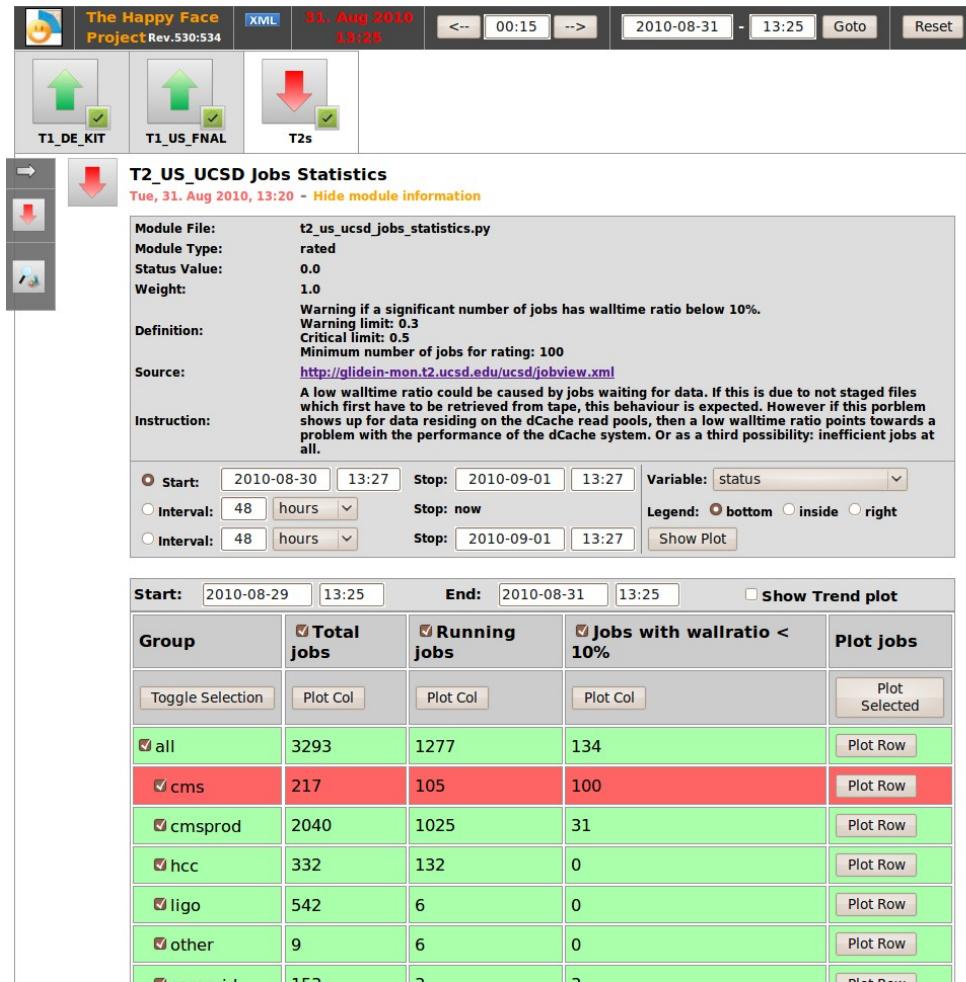
Thu, 11. Mar 2010, 11:50

used disk space	22.9 TB
sites: T2_DE_RWTH, T2_DE_DESY	
users exceeding quota	1
quota of 2.0 TB per user, individual limits for power users	
unmatched directories	
<a href="#">show/hide details</a>	

User	T2_DE_RWTH	T2_DE_DESY	Total Usage
Power User1	10318 GB	318 GB	10636 GB
Power User2	104 GB	9818 GB	9922 GB
DCMS User	—	681 GB	681 GB
CMS User	307 GB	—	307 GB

# In development: T1 CMS Production Jobs Monitoring

CMS to use HappyFace for production job monitoring of CMS jobs at **all T1s**



The screenshot shows the HappyFace interface with the following details:

- Top Bar:** The Happy Face Project Rev.530:534, XML, 31. Aug 2010 13:25, navigation buttons (back, forward, search), date range (2010-08-31 - 13:25), Goto, Reset.
- Site Status:** Icons for T1\_DE\_KIT (green up arrow), T1\_US\_FNAL (green up arrow), and T2s (red down arrow).
- T2\_US\_UCSD Jobs Statistics:** Tue, 31. Aug 2010, 13:20. Module File: t2\_us\_ucsd\_jobs\_statistics.py, Module Type: rated, Status Value: 0.0, Weight: 1.0. Definition: Warning if a significant number of jobs has walltime ratio below 10%. Warning limit: 0.3, Critical limit: 0.5, Minimum number of jobs for rating: 100. Source: <http://glidein-mon.t2.ucsd.edu/ucsd/jobview.xml>. Instruction: A low walltime ratio could be caused by jobs waiting for data. If this is due to not staged files which first have to be retrieved from tape, this behaviour is expected. However if this problem shows up for data residing on the dCache read pools, then a low walltime ratio points towards a problem with the performance of the dCache system. Or as a third possibility: inefficient jobs at all. Time controls: Start: 2010-08-30 13:27, Stop: 2010-09-01 13:27, Variable: status; Interval: 48 hours, Stop: now; Legend: bottom, inside, right; Interval: 48 hours, Stop: 2010-09-01 13:27, Show Plot.
- Summary Table:** Shows counts of total and running jobs, and jobs with wallratio < 10% for various groups: all, cms, cmsprod, hcc, ligo, other, and emcprod. The table includes checkboxes for selecting rows and buttons for plotting.

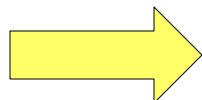
Group	Total jobs	Running jobs	Jobs with wallratio < 10%	Plot jobs
all	3293	1277	134	Plot Row
cms	217	105	100	Plot Row
cmsprod	2040	1025	31	Plot Row
hcc	332	132	0	Plot Row
ligo	542	6	0	Plot Row
other	9	6	0	Plot Row
emcprod	152	2	2	Plot Row

- Each T1 sets up an XML producer for their respective batch system
- HappyFace aggregates the information and visualizes it all **at one place**
- One category for each site
- T2 jobs can also be shown in another category

# Summary

Monitoring	Meta Monitoring
Many browser windows/tabs	All information available on a <u>single website</u>
Each monitoring site requires individual settings	Individual site settings are configured once in <u>HF config</u>
Long loading times	Single website, <u>no complex database backend</u>
Difficulty to find correlations	Easy due to having <u>all information available immediately</u> , plus <u>history</u>
Hard to get quick overview	Just see whether there are <u>red arrows</u> or not; each module has <u>instructions</u>

Modular design allows for easy **Development** and **Deployment** of new modules



- Contributions by Aachen, Göttingen, Hamburg, Karlsruhe
- HappyFace deployed at **many sites**
- **Active development** is ongoing (USM, Mail notification)

# Thank you!

Thanks to  
all HappyFace developers:

- Aachen: M. Edelhoff, P. Sauerland, O. Tsigenov
- Göttingen: C. Ay, S. Birkholz, J. Meyer, A. Quadt
- Hamburg: F. Nowak, P. Schleper, H. Stadie
- Karlsruhe: V. Büge, A. Burgmeier, V. Mauch, G. Quast, N. Ratnikova, A. Scheurer, M. Zvada

and the Helmholtz Alliance “Physics at the Terascale”

[https://ekptrac.physik.uni-karlsruhe.de/trac/HappyFace/wiki/Version\\_2](https://ekptrac.physik.uni-karlsruhe.de/trac/HappyFace/wiki/Version_2)

# Questions?