


Site Meta-Monitoring – The HappyFace Project

A. Scheurer on behalf of the HappyFace Developers

4th Annual Workshop of the Helmholtz Alliance "Physics at the Terascale"

Dresden, 01. – 03. December 2010

Institut für Experimentelle Kernphysik



HappyFace (HF) Basics

■ HF - **What it is:**

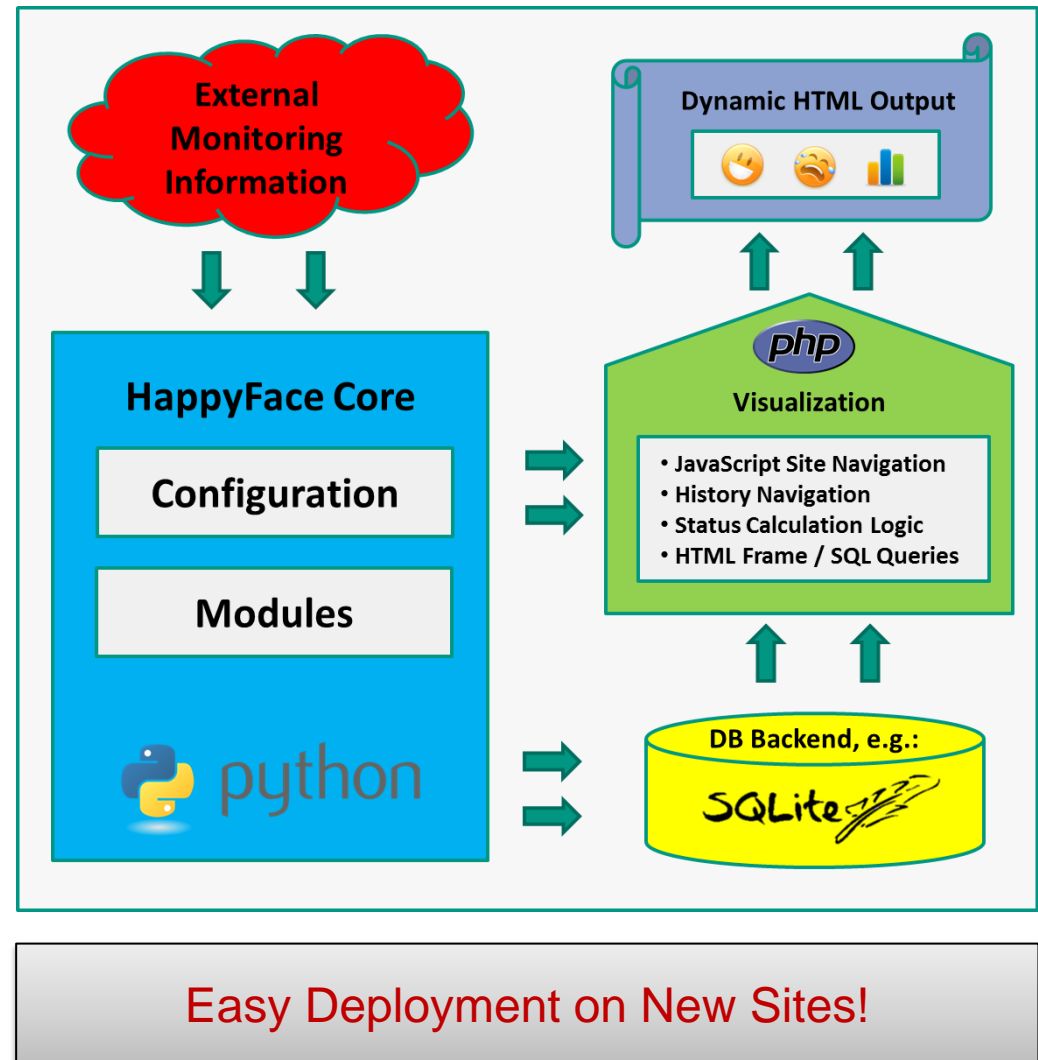
- Allows real-time site monitoring
- Acquires information automatically, not on demand!
- Can be used as sophisticated and modular shift tool
- Provides as well detailed information for admins (if required)
- Auto-refresh system, no user intervention necessary
- Provides rating system (Keyword: non-expert shift crews)
- Allows to correlate information
- Can trigger automatic alarms/notifications
- Highly configurable and adjustable

→ **Meta-monitoring Suite**

Technical Background

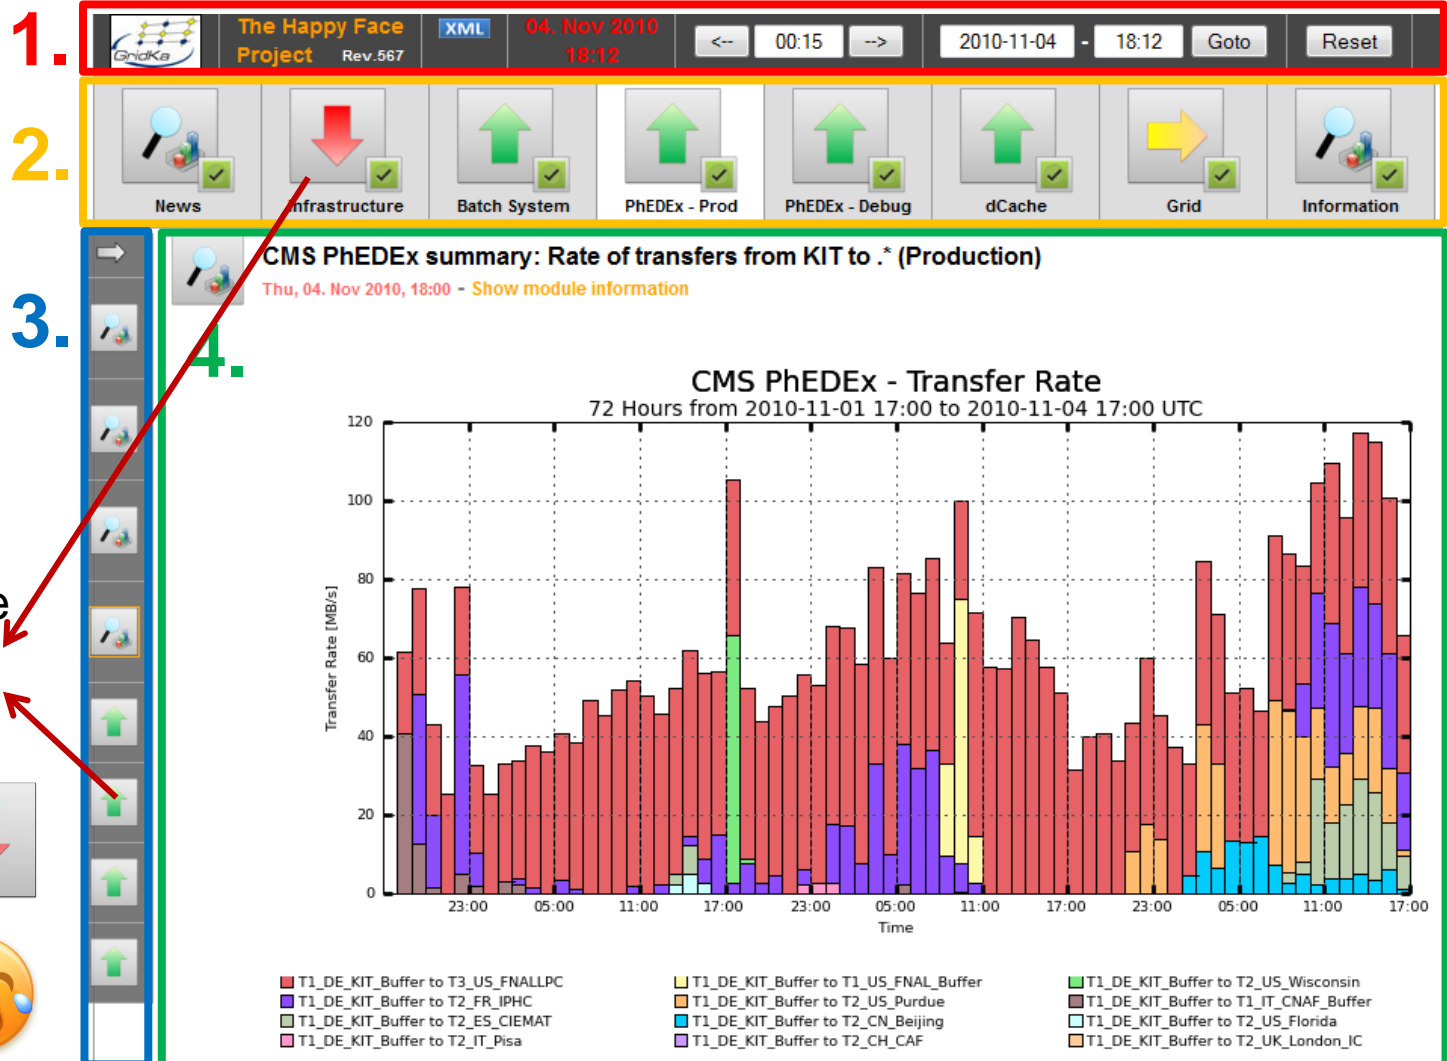

■ HF is:

- Written in Python
- Highly modular
 - Similar modules inherit functionality
 - Configuration possible on each inheritance level, e.g. 5 modules which provide CMS Dashboard information for T1_DE_KIT can be changed to monitor the site T2_DE_DESY by just altering one line in the parent config file.
- DB assisted
 - Intrinsic history functionality
- Lightweight, fast and reliable
 - In production use since more than 2 years now



The Interface & Rating System

- 1. History Navigation
- 2. Category Navigation
- 3. Module Navigation
- 4. Module Content
- Simple module and category rating system



HF – CMS and ATLAS partners

■ HF core and module development:

■ KIT Karlsruhe

■ Module development/usage

■ University of Hamburg

■ DESY Hamburg

■ RWTH Aachen

■ University of Göttingen



Some Selected Modules – Batch System Karlsruhe

- **Real-time** batch system monitoring
 - Use batch system xml provider (CMS has providers for e.g. PBS, LSF, Condor)
 - See currently running batch jobs
 - Calculate current job efficiency
 - Define warning/error thresholds

Start: 2010-11-05 09:18 End: 2010-11-07 09:18 ☐ Show Trend plot

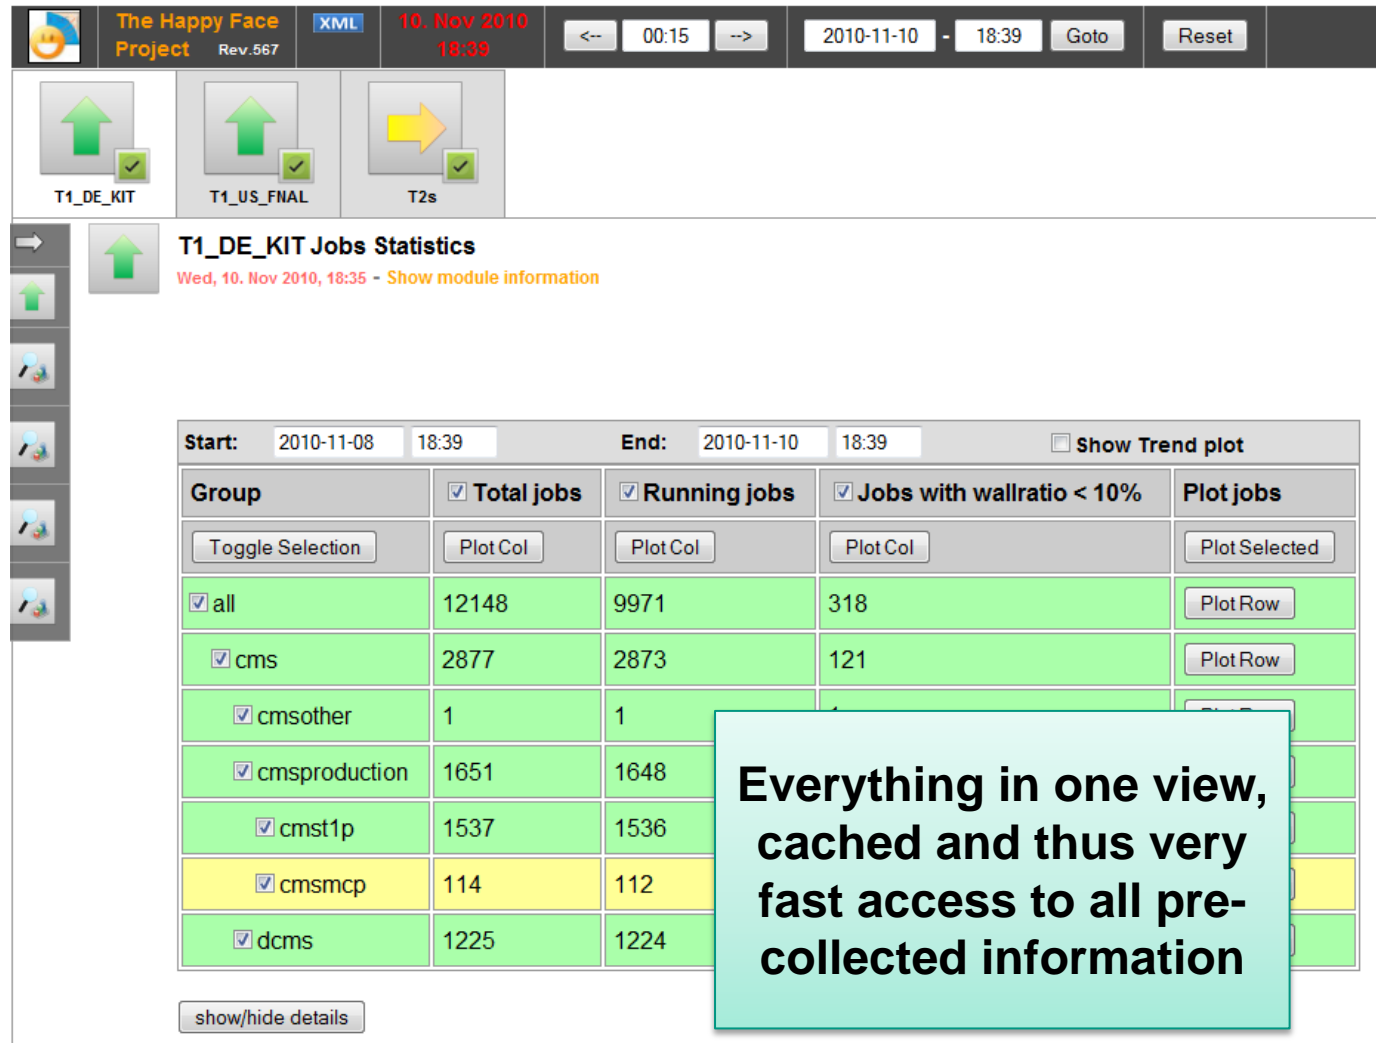
Group	<input checked="" type="checkbox"/> Total jobs	<input checked="" type="checkbox"/> Running jobs	<input checked="" type="checkbox"/> Jobs with wallratio < 10%	Plot jobs
<input type="button" value="Toggle Selection"/>	<input type="button" value="Plot Col"/>	<input type="button" value="Plot Col"/>	<input type="button" value="Plot Col"/>	<input type="button" value="Plot Selected"/>
<input checked="" type="checkbox"/> all	12287	6606	1424	<input type="button" value="Plot Row"/>
<input checked="" type="checkbox"/> cms	990	973	403	<input type="button" value="Plot Row"/>
<input checked="" type="checkbox"/> cmsother	16	0	0	<input type="button" value="Plot Row"/>
<input checked="" type="checkbox"/> cmsproduction	965	965	398	<input type="button" value="Plot Row"/>
<input checked="" type="checkbox"/> cmst1p	751	751	351	<input type="button" value="Plot Row"/>
<input checked="" type="checkbox"/> cmsmcp	214	214	47	<input type="button" value="Plot Row"/>
<input checked="" type="checkbox"/> dcms	9	8	5	<input type="button" value="Plot Row"/>

Note: This table shows users of the group cms only

User	Total	Running	Waiting	Queued	Eff. > 80%	80% > Eff. > 30%	30% > Eff. > 10%	10% > Eff.
cmsmcp33	192	192	0	0	156	8	3	25
dcms118	2	2	0	0	1	1	0	0
dcms013	1	1	0	0	0	0	0	1
cmssgm	16	0	0	16	0	0	0	0
cmst1p31	5	5	0	0	1	0	1	3
cmst1p23	746	746	0	0	397	0	1	348
cmsmcp41	22	22	0	0	0	0	0	22
dcms131	3	3	0	0	0	1	0	2

Prototype: CMS Tier1 Batch System Monitoring

- Individual categories for each Tier1, e.g.
 - KIT
 - FNAL
 - etc.
- Combined or individual categories for Tier2s
- Easily extendable to integrate e.g. storage system monitoring, etc.



Some Selected Modules – Storage System Karlsruhe



dCache: Status of CMS Disk Only Pools

Mon, 15. Nov 2010, 16:00 - [Show module information](#)

- HF provides modules to monitor **dCache systems**
 - Use standard dCache xml provider
 - Can be used by any dCache site
 - Monitor overall dCache status, down to the point of individual pools
 - Monitor current I/O throughput and status of active/queued transfers
 - “Free space”
 - And many more features ...

Pools	6
Pools with status warning [%]	0
Pools with status critical [%]	0
Total Space [TiB]	174.02
Free Space [TiB]	137.07
Used Space [TiB]	36.95
Precious Space [TiB]	0
Removable Space [TiB]	0
Free Space/Total Space [%]	78.8
Pools with status critical [%]	0
Pools with status warning [%]	0

show/hide results

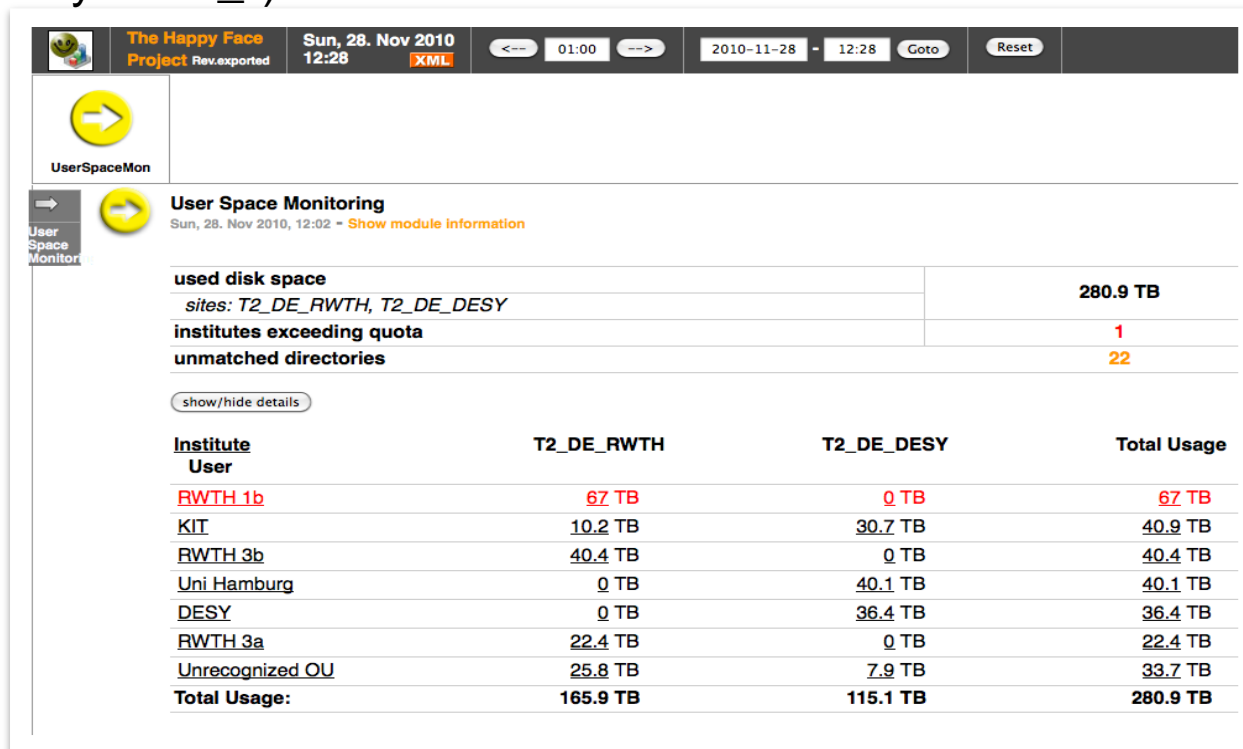
Start: 2010-11-13 16:05 End: 2010-11-15 16:05 ☐ Show different variables in the same plot

Poolname	<input checked="" type="checkbox"/> Total Space [TiB]	<input checked="" type="checkbox"/> Free Space [TiB]	<input checked="" type="checkbox"/> Used Space [TiB]	<input checked="" type="checkbox"/> Precious Space [TiB]	<input checked="" type="checkbox"/> Removable Space [TiB]	<input checked="" type="checkbox"/> Free Space/Total Space [%]	Pool Plot
<input type="button" value="Toggle Selection"/>	<input type="button" value="Plot Col"/>	<input type="button" value="Plot Col"/>	<input type="button" value="Plot Col"/>	<input type="button" value="Plot Col"/>	<input type="button" value="Plot Col"/>	<input type="button" value="Plot Col"/>	<input type="button" value="Plot Selected"/>
<input checked="" type="checkbox"/> f01-070-134-e_1D_cms	29	24.11	4.89	0	0	83.1	<input type="button" value="Plot Row"/>
<input checked="" type="checkbox"/> f01-070-120-e_1D_cms	29	21.7	7.3	0	0	74.8	<input type="button" value="Plot Row"/>

User Storage Consumption - I

Slide by Aachen HappyFace Team

- Contributions (from all three Aachen institutes) financed by BMBF FSP-CMS
- In CMS model a user gets a few TB „Grid home storage space“ at national T-2s
- dCache does not (yet) offer a usable quota system, so we use HappyFace to monitor the consumption at T2_DE_DESY and T2_DE_RWTH (can be extended easily to T3_*)
- Summary



User Storage Consumption - II

Slide by Aachen HappyFace Team

■ User's view (authentication by cert.p12 in browser)

Thomas Kress	940.3 GB	—	940.3 GB
Directory Name: kress			
DN: /O=GermanGrid/OU=RWTH/CN=Thomas Kress			
email: Thomas.Kress@cern.ch			
VOMS Group: /cms/dcms			
Storage Path: 940.3GB @ [T2_DE_RWTH] grid-srm.physik.rwth-aachen.de:/pnfs/physik.rwth-aachen.de/cms/store/user/kress			
Storage Path: 0.0GB @ [T2_DE_DESY] dcache-se-cms.desy.de:/pnfs/desy.de/cms/tier2/store/user/kress			
Status: Use 2.3%			

- User can use both German CMS Tier-2s to place data
- Includes user's job „output“ data and official CMS „input“ data sets transferred (by PhEDEx) to our Tier-2s on behalf of the user
- Admins can see the details of all users, grouped by institutes/labs
- We can find (most) users not associated to German CMS groups and will start soon to automatically inform and ask them to delete data (if above a threshold)
- History view available

Physics Group Space Usage (CMS)

Slide by Hamburg HappyFace Team

- In CMS, each dataset skim is allocated to a specific “physics group”
- Each of these groups has dedicated T2 sites for storage purposes
- The DESY T2 hosts 4 groups: Top Physics, Jets + Missing E_T , Forward Physics, and QCD
- Module shows **warnings** for missing files in the group space
- Module shows **errors** for exceeding space budgeted

details

Group	resident data [TiB]	subscribed data [TiB]	resident files	subscribed files
top	32.28	32.29	19259	19260
jets-met_hcal	44.42	44.42	22823	22823
forward	62.95	62.95	25862	25862
AnalysisOps	45.33	45.33	12305	12305
local	174.4	174.4	66159	66159
DataOps	6.1	6.1	3289	3289
FacOps	0.51	0.51	510	510
qcd	38.83	38.83	28308	28308
undefined	9.24	9.24	4878	4878



Storage Usage of the Physics Groups

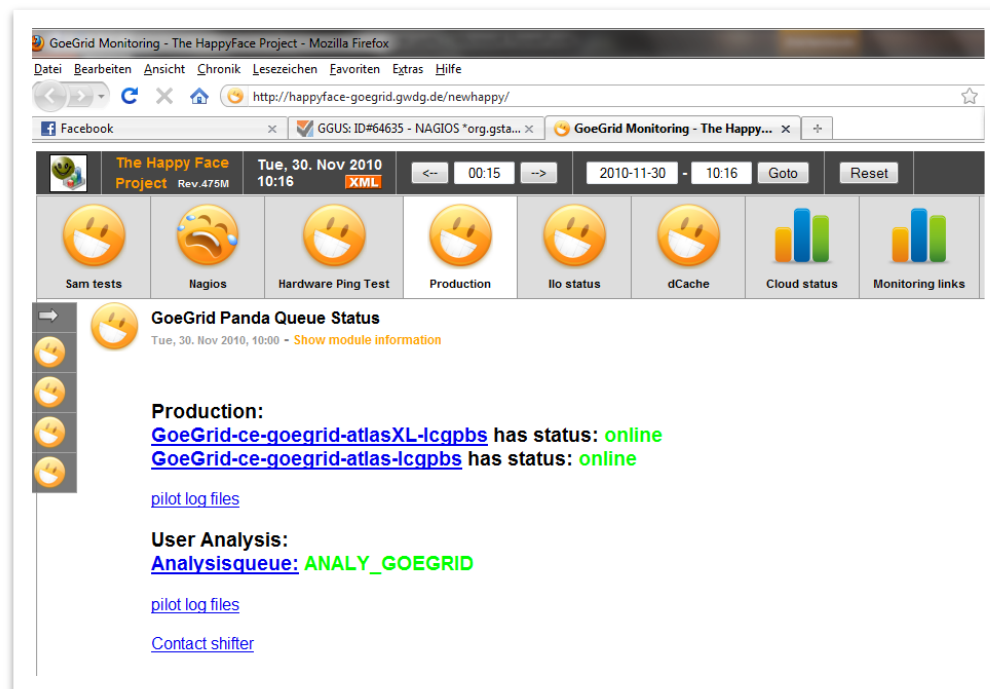
Mon, 29. Nov 2010, 11:00 - [Show module information](#)

Group	resident data [TiB]	percent of subscribed data
top	32.28	100
jets-met_hcal	44.42	100
forward	62.95	100
AnalysisOps	45.33	100
local	174.4	100
DataOps	6.1	100
FacOps	0.51	100
qcd	38.83	100
undefined	9.24	100

- Details are listing the number of subscribed files and datasize
- Also used frequently: dCache storage monitoring

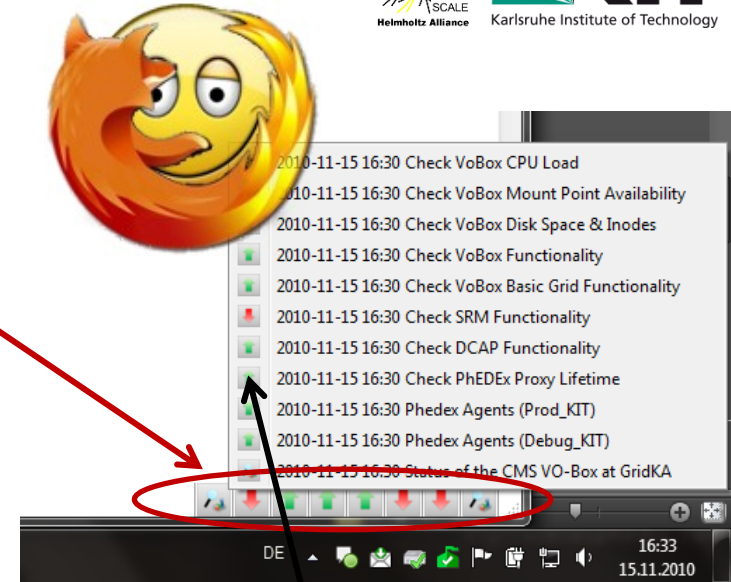
Some Modules Developed in Göttingen

- Developed modules in Göttingen
 - Interface to Nagios
 - Panda monitoring (e.g. queue status)
 - Hardware monitoring (e.g. ILO, network)
 - ATLAS GridKa-Cloud monitoring module



Extended Functionality

- HF publishes its current status via xml
 - Used as input for a plugin available for the Firefox web browser statusbar
 - Used as input for smartphone apps (e.g. iPhone, Android)
 - Used for Meta² - Monitoring, ideal for central shifts – the HF matrix



Site	PhEDEx - Prod	PhEDEx - Debug	dCache	Grid
KIT	↑	↑	↓	↓
DESY	→	→	↓	↑

By clicking on the individual arrows, directly jump to the proper HF instance and module

Summary



■ HF – What it is:

- Modular and easily configurable tool for shifters and admins
- All information is pre-collected (time interval: ~10 min)
 - No waiting time, “live feeling”
- Stores information from external sources (plots, xml/html, text files...)
- Stores configuration parameters with the data to allow a consistent history view even after threshold changes, etc.
- Identify problems: possibility to “get exact state of my site on Sunday night”
- Provides rating system, allows to define complex correlations
- Hides unnecessary information but provides access to all details if required
- Possibility to automatically trigger alarms/notifications
- Exports it's status via xml for further usage/harvesting (e.g. iPhone app, Firefox plugin, Meta²-Monitoring)
- Used by German CMS and Atlas Sites for more than 2 years now
 - Stable, reliable, tested
- Many modules available – designed for collaboration via central repository

The HappyFace Project

More Information & Documentation:

<https://ekptrac.physik.uni-karlsruhe.de/trac/HappyFace>