HammerCloud Use Cases & Plans

D-HEP Computing Strategieworkshop

28./29. Apr 2014

G.Duckeck on behalf of HC team:
M. Boehler
J.Elmsheuser
F.Hoenig
F.Legger

- Introduction
- Use cases
  - Site functional & performance test
  - WAN data access
- Plans

Hammercloud Overview

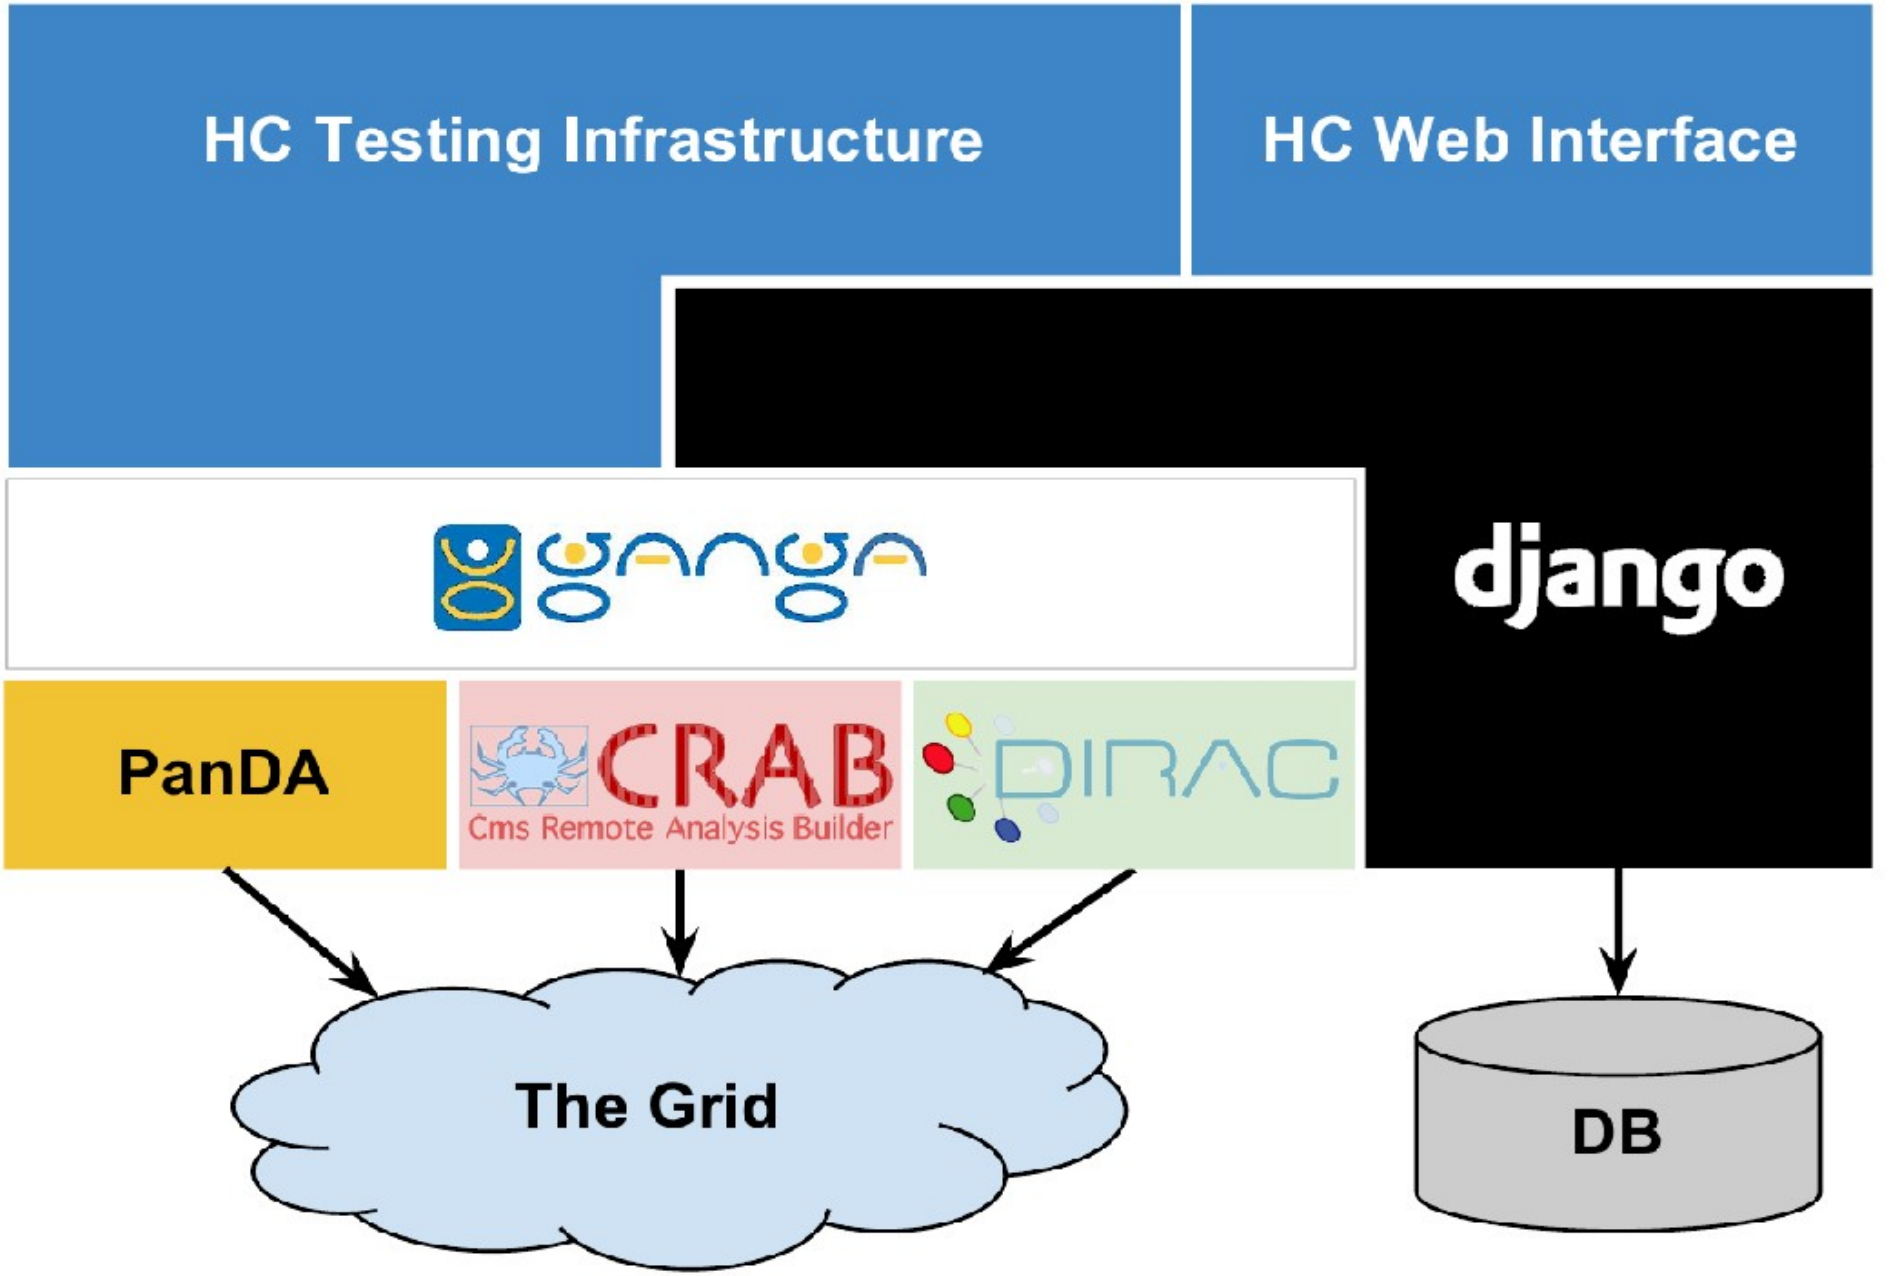

HammerCloud ATLAS use cases I

- Panda queue black-listing:
  - Continuous flow of 6 different short tests to verify Panda analy & prod queues
    - 3 AFT (analysis functional tests)
    - 3 PFT (production functional tests)
  - Real analysis and Geant4 workload
  - Automatically black-list and white-list Panda queues with email and SAM/Nagios notifications
- Functional test: (in use):
  - Continuous flow of very low frequency short jobs at multiple or handful of sites
    - 4 different Panda pilot development tests & 8 Athena nightly build tests
    - 2 tests for FAX WAN cost matrix and ROOT I/O measurements
    - 3 FAX functional tests
    - 1 multi-core production queue test & 1 Production queue G4 bench-marking
    - 1 Dedicated DPM SE functional test

HC Panda queue setting example for DE

Historic view for "panda analy status NEW"
from 00:00 01.03.2014 to 00:00 31.03.2014

Show 100 entries

SITE Name	TIER	CLOUD	History plot time bin = 60 hours	offline		brokeroff		online	
				%	count	%	count	%	count
CSCS-LCG2	T2D	DE		0	0	0	0	93.97	16
CYFRONET-LCG2	T2	DE		0	0	0	0	96.72	9
DESY-HH	T2D	DE	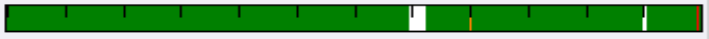	0.28	1	0	0	96.19	9
DESY-ZN	T2D	DE		0	0	2.27	2	94.73	8
FMPhI-UNIBA	T2	DE	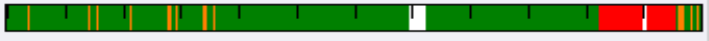	10.51	2	0	0	81.7	17
FZK-LCG2	T1	DE		0	0	0	0	97	6
GoeGrid	T2D	DE	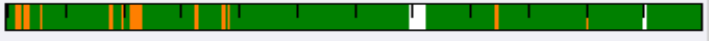	0	0	0	0	90.16	18
HEPHY-UIBK	T2	DE		1.71	1	0.51	1	69.57	11
IEPSAS-Kosice	T2	DE	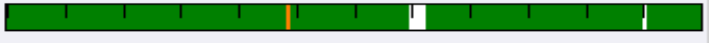	0	0	0	0	96.22	9
LRZ-LMU	T2D	DE	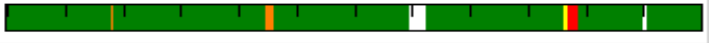	1.39	1	0.53	1	92.95	12
MPPMU	T2D	DE	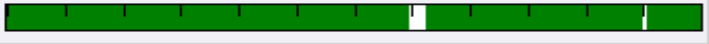	0	0	0	0	96.86	7
praguelcg2	T2D	DE		0	0	4.91	10	92.66	16
PSNC	T2	DE		97	6	0	0	0	0
UNI-FREIBURG	T2D	DE		3.47	2	0	0	91.71	7
wuppertalprod	T2D	DE	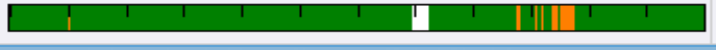	0	0	0	0	92.57	14

- Site performance used to determine data distribution shares
  - → Site categories A, B, C, D

Site Overview

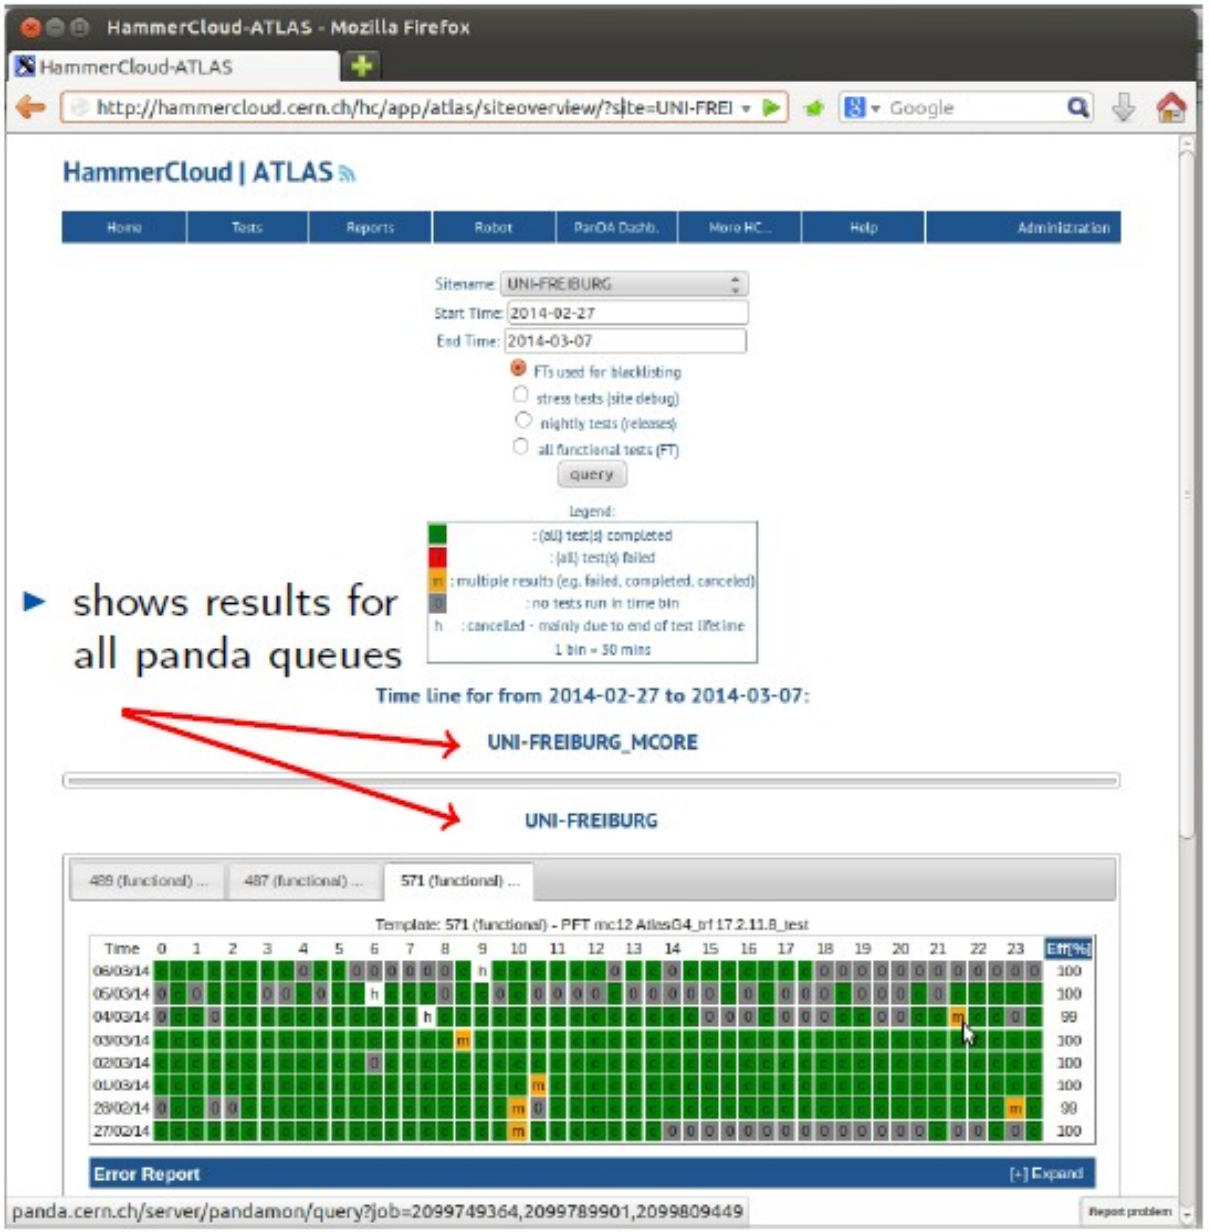

- shows results for all panda queues

- *here* no *isGolden* functional tests for MCORE queue

Site Overview - More Details

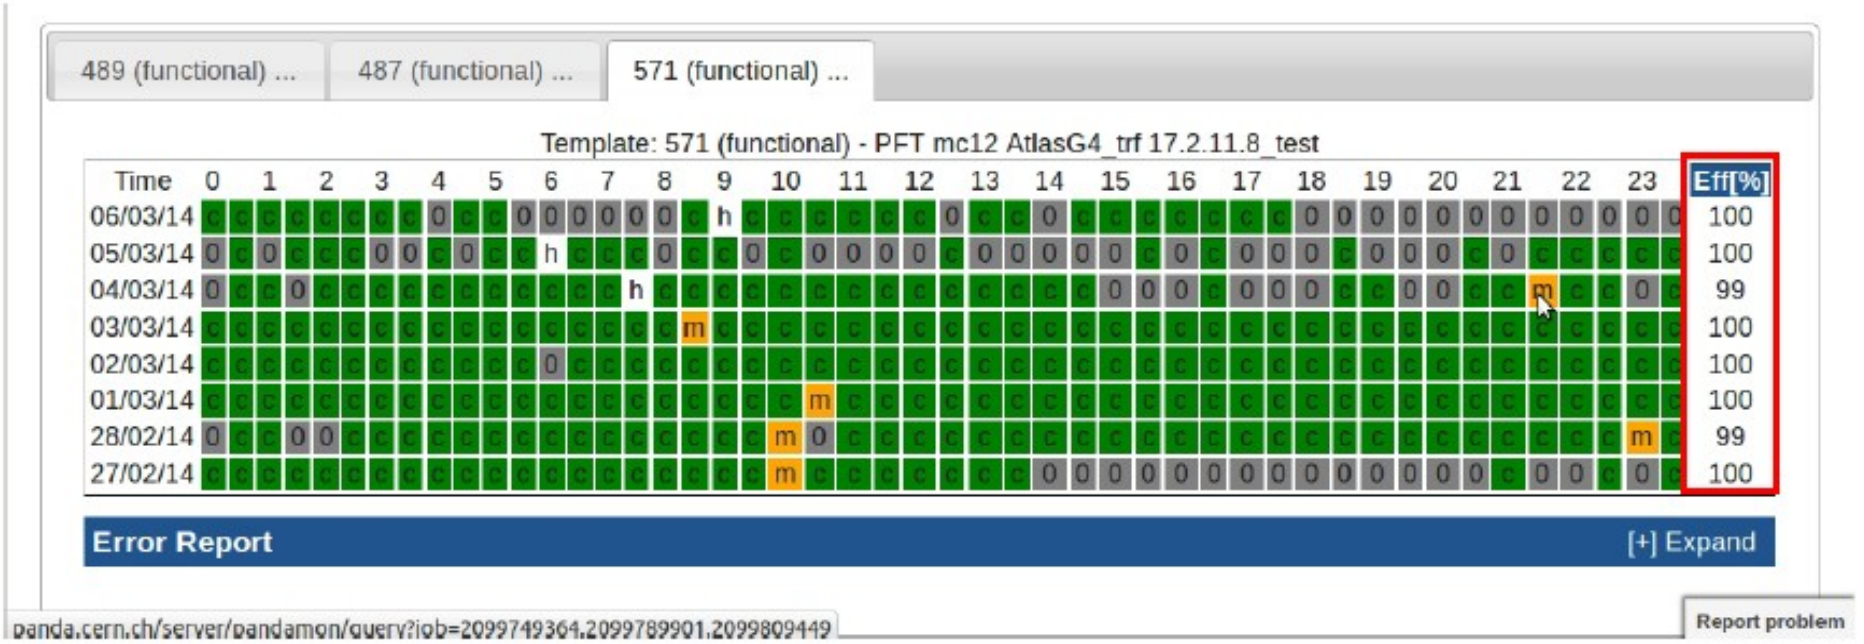

- one table per template - test type (usually 3 templates per AFT and 3 per PFT)
  - *here*: PFT mc12_AtlasG4_trf_17.2.11.8 → Geant4 Simulation test
- each line shows job status in 30 min bin
  - binning allows to spot if no test arrives in a certain period
  - if more than one job finished within a bin, all jobs are considered (see status line)
  - clicking on bin → link to panda page for all jobs within this time interval
- test efficiency = (completed jobs) / (completed + failed jobs) per day

Site Overview - More Details

489 (functional) ... | 487 (functional) ... | 571 (functional) ...

Template: 571 (functional) - PFT mc12 AtlasG4_trf 17.2.11.8_test

Time	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	Eff[%]
06/03/14	c	c	c	c	c	c	c	0	c	c	0	0	0	0	0	0	c	h	c	c	c	c	c	c	0	c	c	0	c	c	c	c	c	c	0	0	0	0	0	0	0	0	0	0	0	100
05/03/14	0	c	0	c	c	c	0	0	c	0	c	c	h	c	c	c	0	c	c	0	c	0	0	0	0	c	0	0	0	0	0	c	0	c	0	0	0	c	0	0	0	c	0	c	c	c	c	c	100
04/03/14	0	c	c	0	c	c	c	c	c	c	c	c	c	c	h	c	c	c	c	c	c	c	c	c	c	c	c	c	c	0	0	0	c	0	0	0	c	c	0	0	c	c	m	c	c	0	c	99
03/03/14	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	m	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	100
02/03/14	c	c	c	c	c	c	c	c	c	c	c	c	0	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	100
01/03/14	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	m	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	100
28/02/14	0	c	c	0	0	c	c	c	c	c	c	c	c	c	c	c	c	c	c	m	0	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	m	c	99
27/02/14	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	m	c	c	c	c	c	c	0	0	0	0	0	0	0	0	0	0	0	0	0	0	c	0	0	c	0	c	100

Error Report [-] Collapse

Time	Node	Job ID	Error message
4/3/2014 21:48	None	2099789901	pilotErrorDiag: Failed to import class: Exception caught: 15, transExitCode: 0
28/2/2014 23:25	None	2096051480	pilotErrorDiag: Put error: file:/tmp/20017275.pbs.bfg.uni-freiburg.de/condorg_OeJgaysZ /pilot3/Panda_Pilot_11756_1393628787/PandaJob_2096051480_1393628788 /Hits.hc_20033136.HITS.pool.root.UNI-FREIBURG.1393628482.d8ec: zero number of replicas lcg_cp: Invalid argument, exeErrorDiag: OK, transExitCode: NULL

- expandable Error Report
- error overview sorted by time
- worker node ID is following soon
- by clicking the row in the error report
  → link to panda page is activated (for more information)

Cloud overview - Monitoring a complete cloud

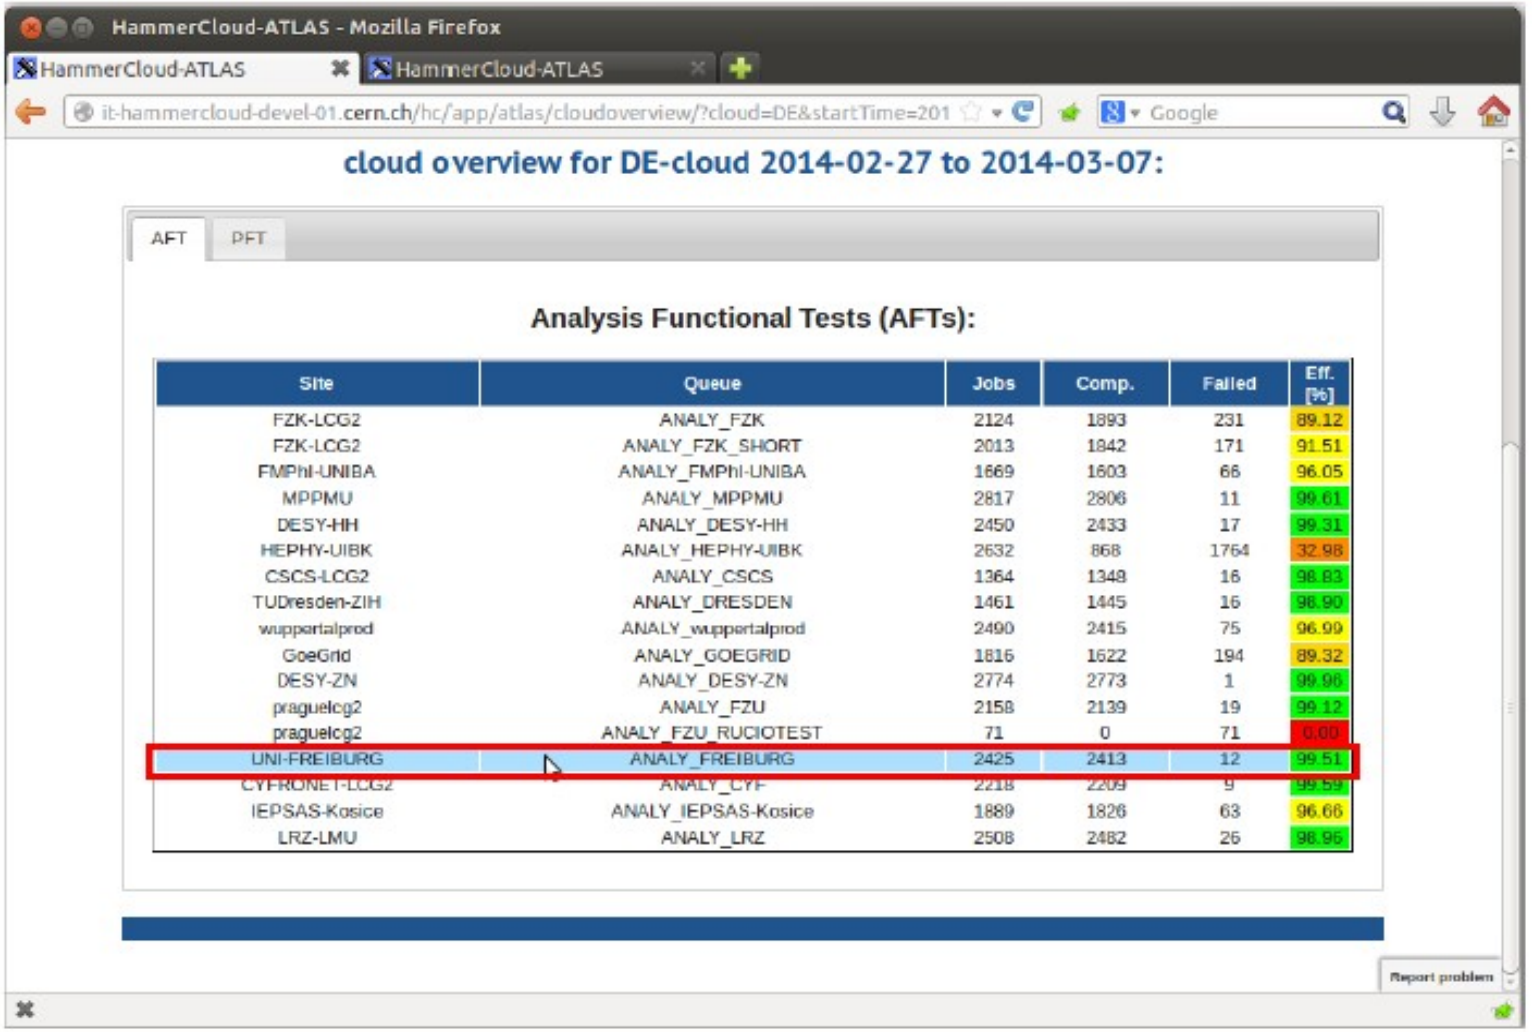

Site	Queue	Jobs	Comp.	Failed	Eff. [%]
FZK-LCG2	ANALY_FZK	2124	1893	231	89.12
FZK-LCG2	ANALY_FZK_SHORT	2013	1842	171	91.51
FMPhI-UNIBA	ANALY_FMPhI-UNIBA	1669	1603	66	96.05
MPPMU	ANALY_MPPMU	2817	2806	11	99.61
DESY-HH	ANALY_DESY-HH	2450	2433	17	99.31
HEPHY-UIBK	ANALY_HEPHY-UIBK	2632	868	1764	32.98
CSCS-LCG2	ANALY_CSCS	1364	1348	16	98.83
TUDresden-ZIH	ANALY_DRESDEN	1461	1445	16	98.90
wuppertalprod	ANALY_wuppertalprod	2490	2415	75	96.99
GoeGrid	ANALY_GOEGRID	1816	1622	194	89.32
DESY-ZN	ANALY_DESY-ZN	2774	2773	1	99.96
praguelcg2	ANALY_FZU	2158	2139	19	99.12
praguelcg2	ANALY_FZU_RUCIOTEST	71	0	71	0.00
UNI-FREIBURG	ANALY_FREIBURG	2425	2413	12	99.51
CYFRONET-LCG2	ANALY_CYF	2218	2209	9	99.59
IEPSAS-Kosice	ANALY_IEPSAS-Kosice	1889	1826	63	96.66
LRZ-LMU	ANALY_LRZ	2508	2482	26	98.96

- Both AFT and PFT results are evaluated seperately
- color code [100%-98%, 98%-90%, 90%-75%, 75%-25%, 25%-0%]
- click on row → link to site oveview

HammerCloud ATLAS use cases II

- Planned Functional tests:
  - http(s)/davix access for ROOT input files - http(s) TURLs will be available with new data managment system Rucio - need to test/verify ROOT functionality
- On-demand stress tests:
  - Frequently used by several sites admins for site commissioning
  - FAX (federated xrootd) stress testing
  - CERN Wigner test
  - Cloud computing sites
  - HPC sites

The Federated ATLAS XrootD system (FAX)

- Storage federation aiming to treat Tier1 and Tier2 storage space as a *single distributed storage system;*

- FAX routes the client to the nearest site with the requested data
  - transparent for the user

- Use cases:
  - quickly scan through large samples of data without copying them or run jobs at remote location;
  - fetch an unexpected missing file instead of failing the job;
  - process data in the Cloud (use of non-ATLAS resources).

- Made possible through
  - improved network bandwidth,
  - reduced latency,
  - data structure aware caching mechanism like TTreeCache.

Testing FAX access with HammerCLoud

- To probe the infrastructure and to measure the various data access patterns
  - Two realistic, typical analyses (Higgs->WW, SUSY) testing different ROOT versions (5.30, 5.32, 5.34)
  - Pre-placed, site-unique data
- XrootD access modes:
  - Stage-in
  - Direct access
- Test Objectives:
  - Local performance
  - Nearby performance (e.g. within a cloud)
  - Far-away performance
  - Stress test: 1 job running, 100, 200, ...
  - Continuous functional tests

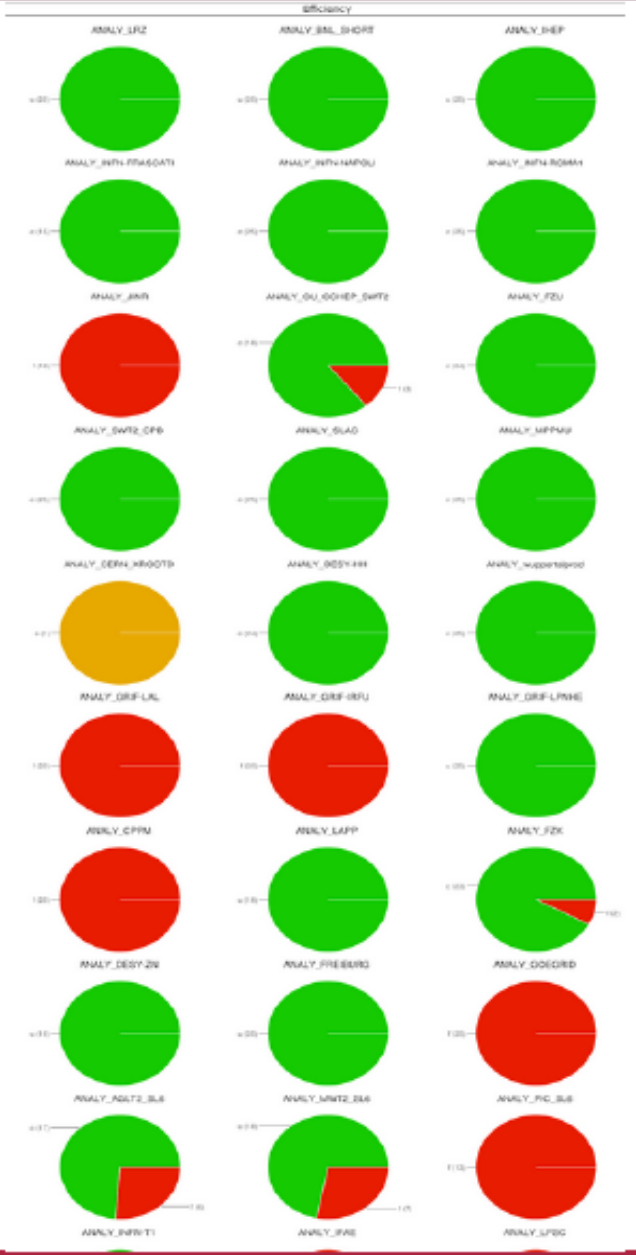

HC FAX test results: Event rate dCache sites

Triangles: stage-in
Dots: direct I/O

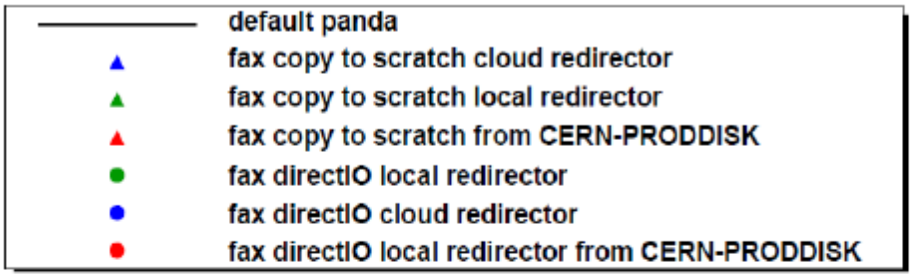

Green: local
Blue: cloud
Red: from CERN

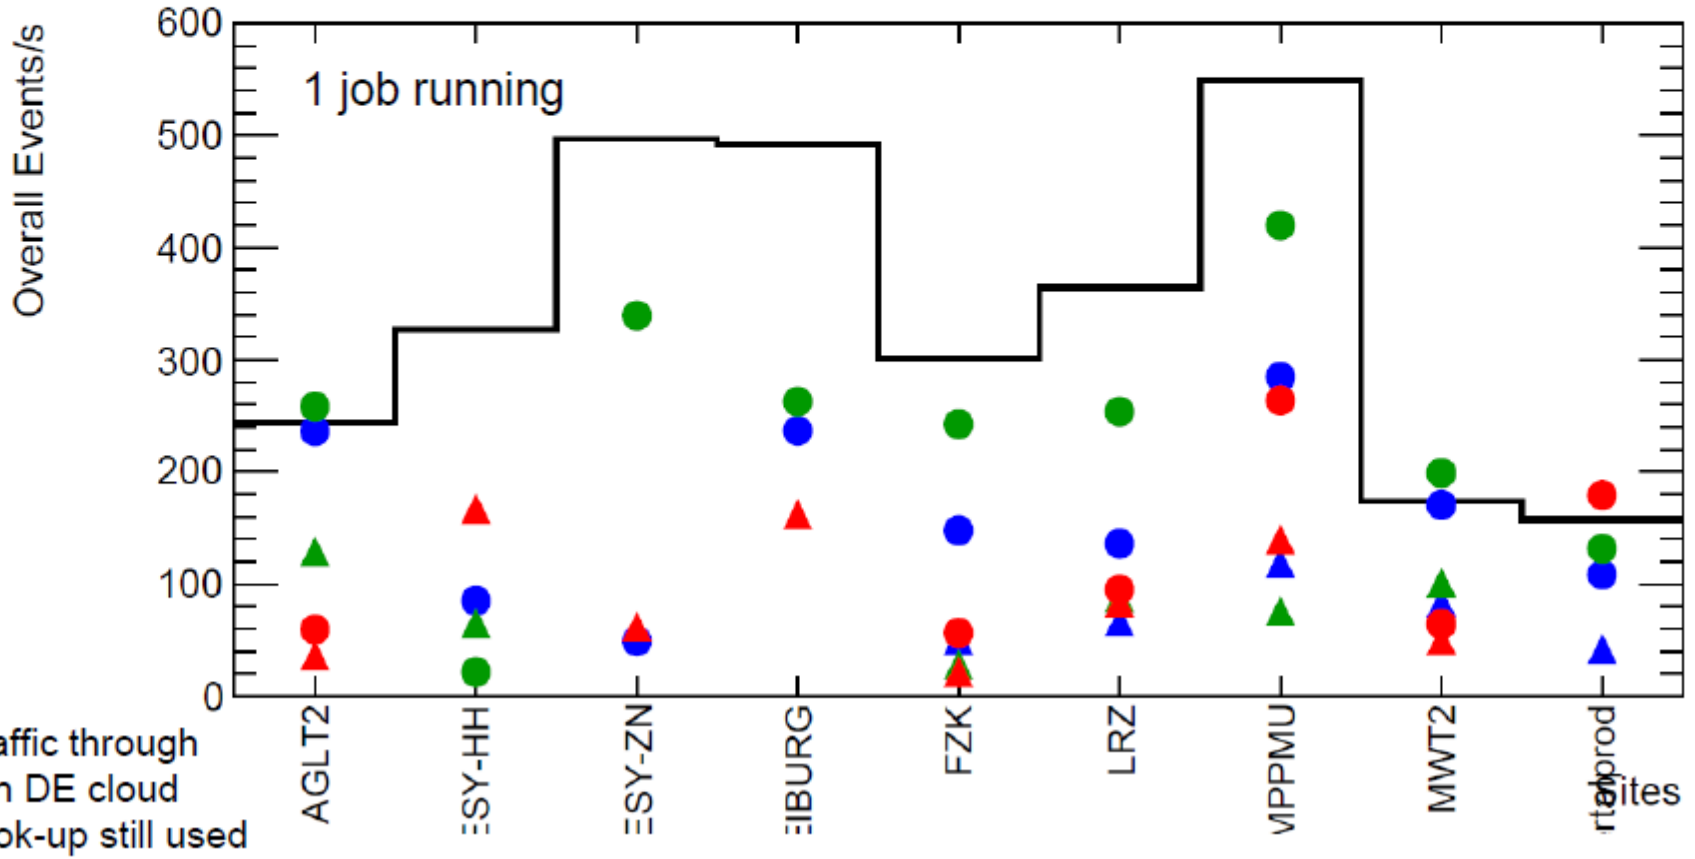

FAX traffic through
proxy in DE cloud
LFC look-up still used

HC FAX stress test results: far-away tests

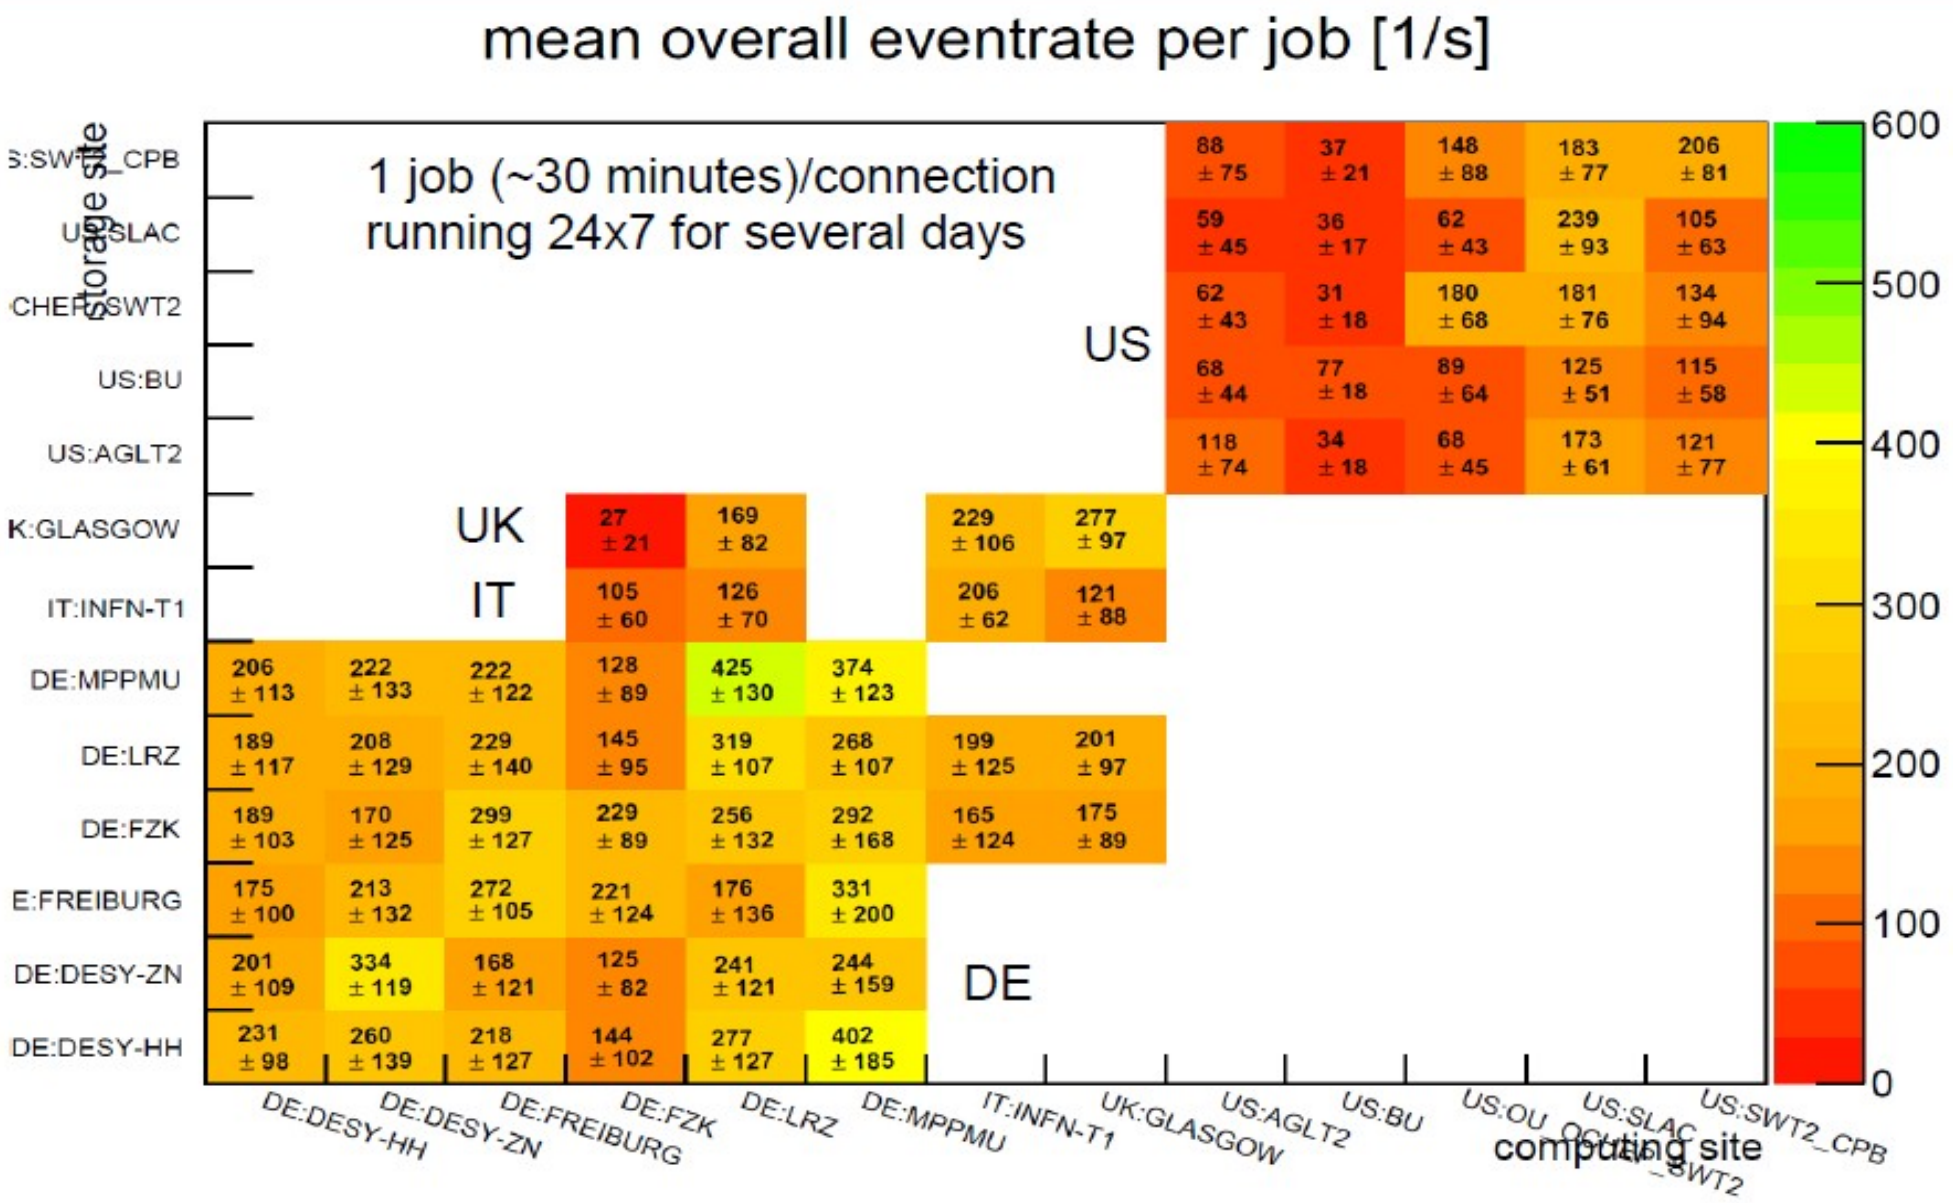

HC FAX stress test results: far-away tests

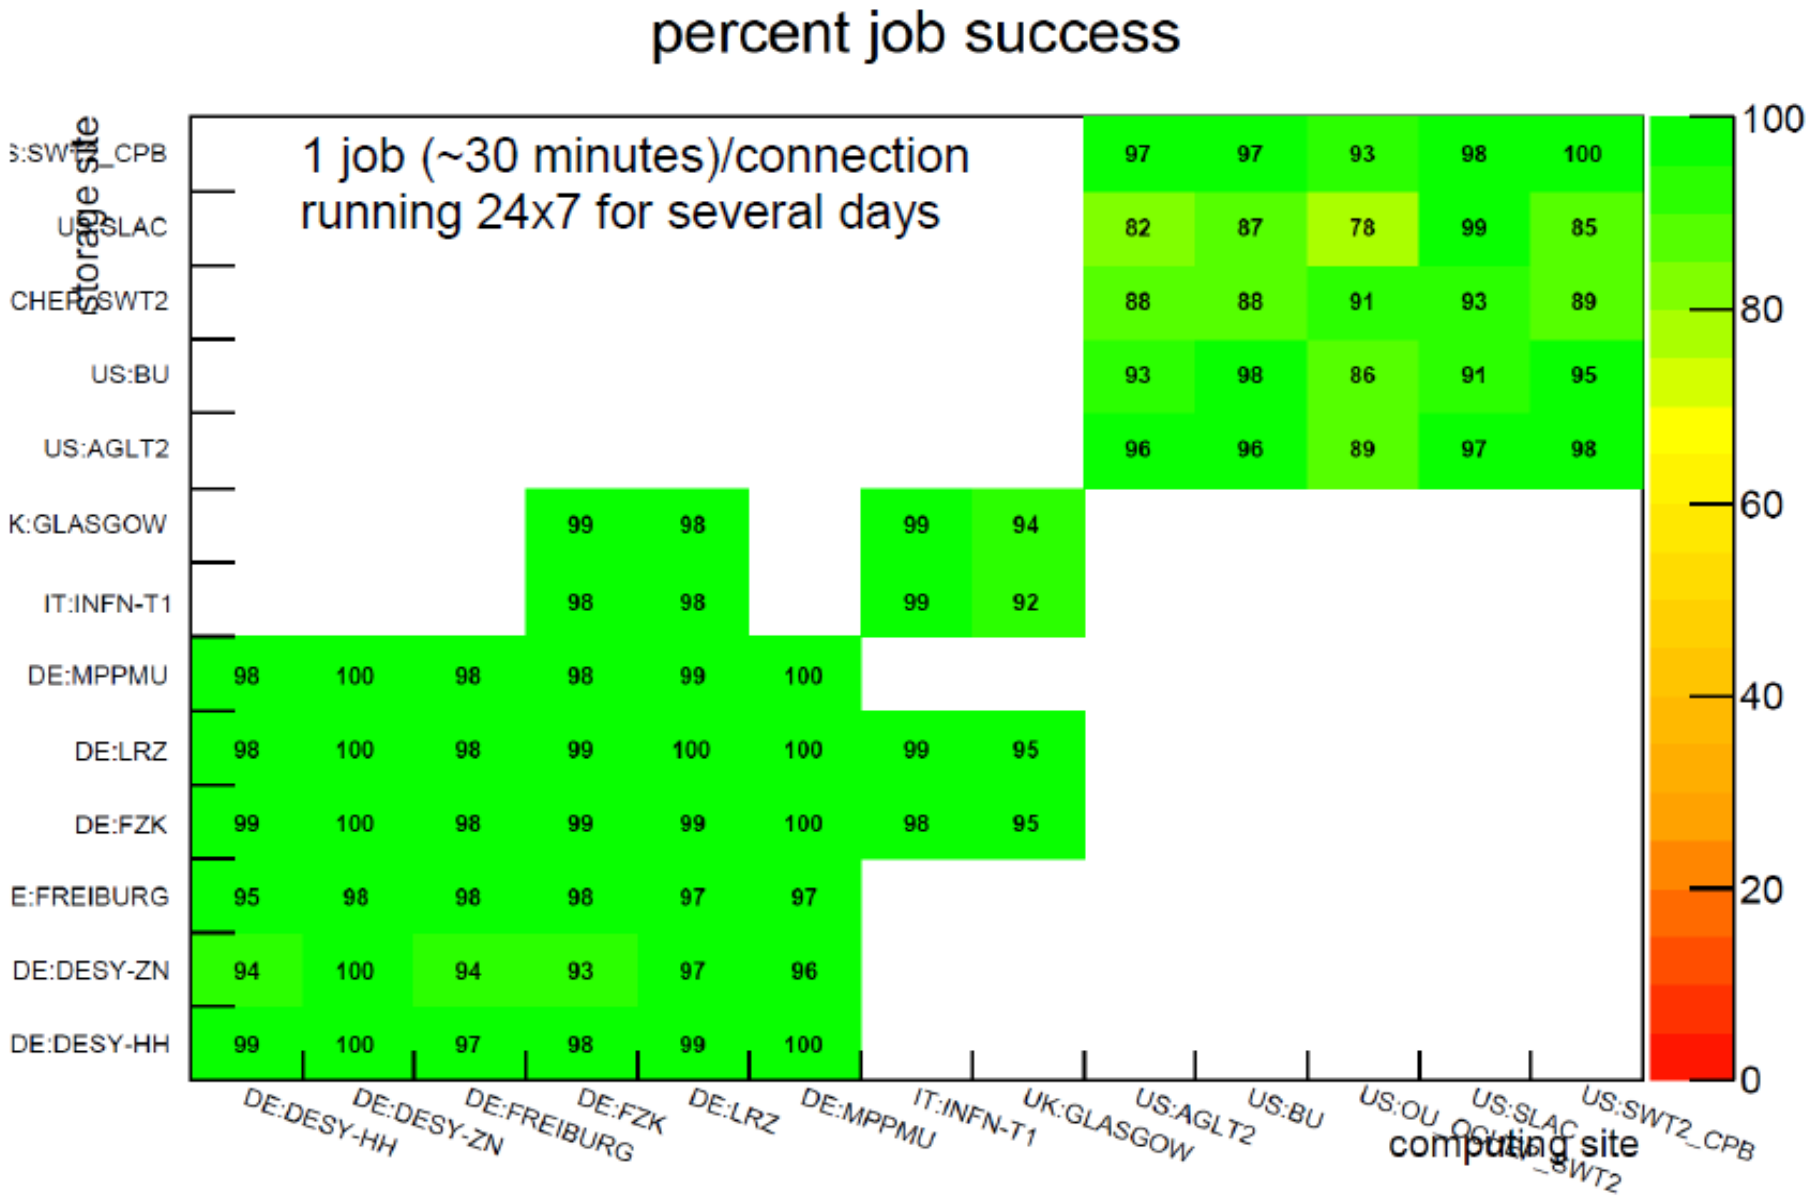

HammerCloud ATLAS further use cases

- CMS:
  - Functional tests:
    - Glidein functional tests, Site robot
  - On-demand stress tests:
    - CRAB3 scaling tests
- WLCG
  - Functional tests:
    - Evaluate HammerCloud as job robot for SAM/Nagios tests

HammerCloud Summary

- Broad range of applications
  - daily routing testing, site exclusion
  - evaluation of data acces modes & protocols
    - direct vs copy-to-scratch, remote IO
    - dcap, xroot, http, ...
  - deployment of new operation models & data layouts
    - Cern / Wigner T0 w/ mutual access
    - new data ATLAS format xAOD
- Indispensable tool for ATLAS/LHC distributed computing
- Very active development to adapt to new requirements and use-cases ongoing