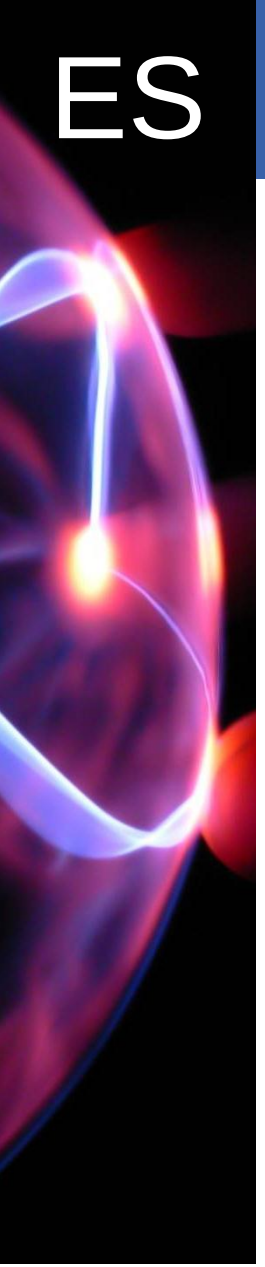

The Common Solutions Strategy of the Experiment Support group at CERN for the LHC Experiments

Maria Girone, CERN
On behalf of the CERN IT-ES Group

CHEP, New York City, May 2012

CERN IT Department
CH-1211 Geneva 23
Switzerland
www.cern.ch/it

Motivation

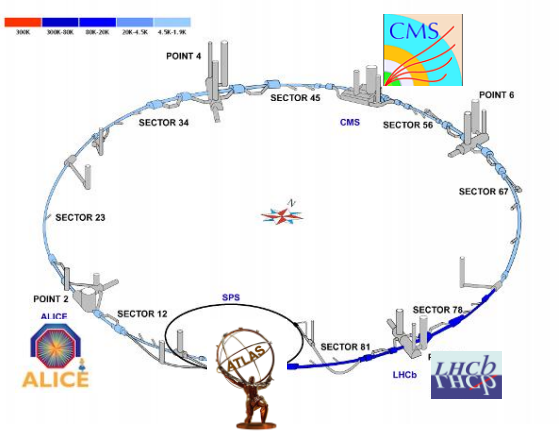

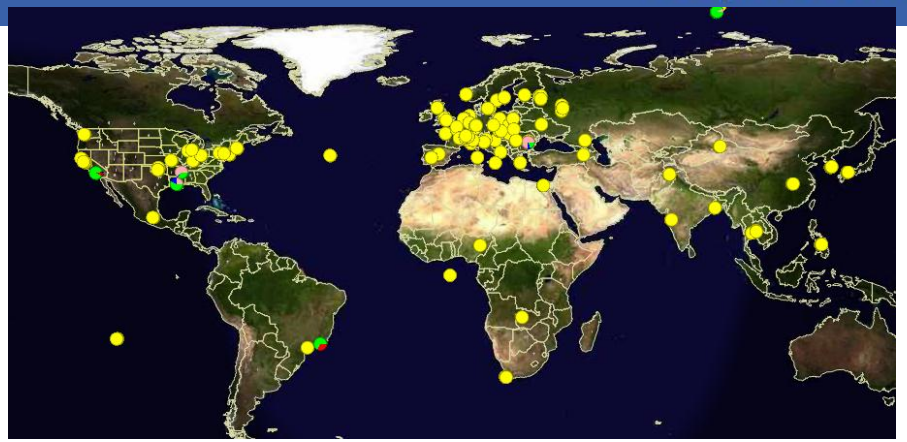

- Despite their differences as experiments at the LHC, from a **computing perspective** a lot of the workflows are similar and can be done with **common services**
- While the collaborations are huge and highly distributed, effort available in ICT development is **limited** and **decreasing**
  - Effort is focused on analysis and physics
- Common solutions are a more **efficient** use of effort and more **sustainable** in the long run

CERN IT Department
CH-1211 Geneva 23
Switzerland
www.cern.ch/it

Anatomy of Common Solution

Most common solutions can be diagrammed as the interface layer between common infrastructure elements and the truly experiment specific components

- One of the successes of the grid deployment has been the use of common grid interfaces and local site service interfaces
- The experiments have a environments and techniques that are unique
- In common solutions we target the box in between. A lot of effort is spent in these layers and there are big savings of effort in commonality
  - not necessarily implementation, but approach & architecture
- LHC schedule presents a good opportunity for technology changes

CERN IT Department
CH-1211 Geneva 23
Switzerland
www.cern.ch/it

The Group

- IT-ES is a unique resource in WLCG
  - The group is currently supported with substantial EGI-InSPIRE project effort
  - Careful balance of effort embedded in the experiments & on common solutions
  - Development of expertise in experiment systems & across experiment boundaries
  - People uniquely qualified to identify and implement common solutions
    - Matches well with the EGI-InSPIRE mandate of developing sustainable solutions
    - A strong and enthusiastic team

Activities

- Monitoring and Experiment Dashboards
  - Allows experiments and sites to monitor and track their production and analysis activities across the grid
    - Including services for data popularity, data cleaning and data integrity and site test stressing
- Distributed Production and Analysis
  - Design and development for experiment workload management and analysis components
- Data Management support
  - Covers development and integration of the experiment specific and shared grid middleware
- The LCG Persistency Framework
  - Handles the event and detector conditions data from the experiments

CERN IT Department
CH-1211 Geneva 23
Switzerland
www.cern.ch/it

ES Examples: Data Popularity

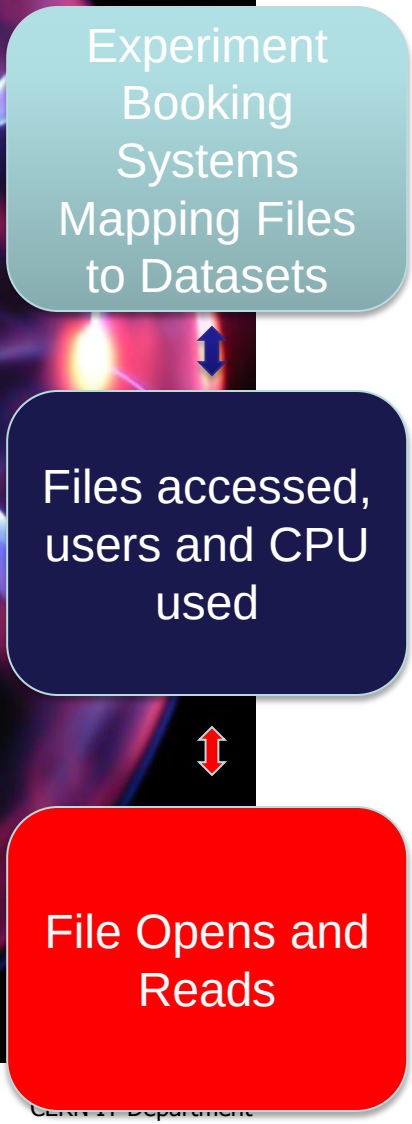

- Experiments want to know which datasets are used, how much, and by whom
  - Good chance of a common solution
- Data popularity uses the fact that all experiments open files and access storage
- The monitoring information can be accessed in a common way using generic and common plug-ins
- The experiments have systems that identify how those files are mapped onto logical objects like datasets, reprocessing and simulation campaigns

CERN IT Department
CH-1211 Geneva 23
Switzerland
www.cern.ch/it

Popularity Service

- Used by the experiments to assess the importance of computing processing work, and to decide when the number of replicas of a sample needs to be adjusted either up or down

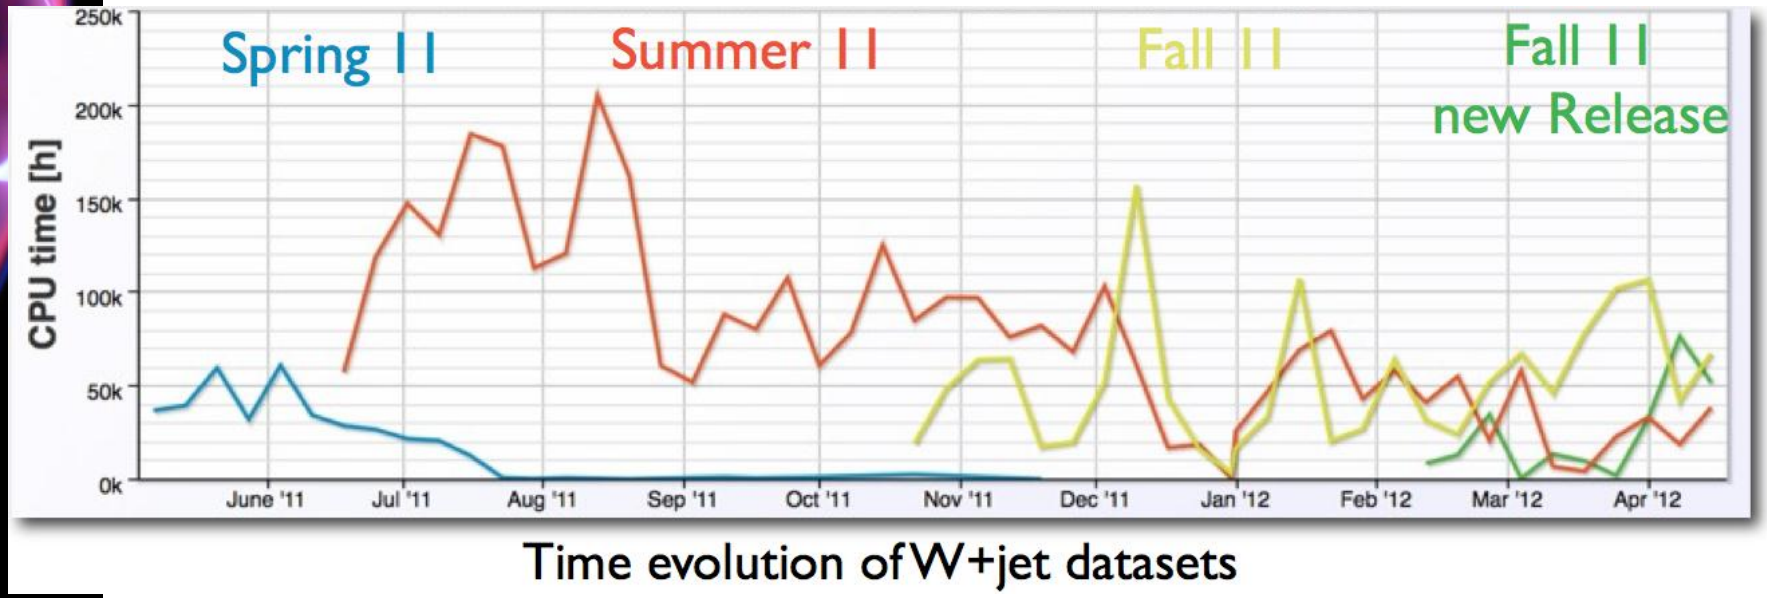

Time evolution of W+jet datasets

See D. Giordano et al., [176] *Implementing data placement strategies for the CMS experiment based on a popularity model*

CERN IT Department
CH-1211 Geneva 23
Switzerland
www.cern.ch/it

ES Cleaning Service

CERN IT Department

- The Site Cleaning Agent is used to suggest obsolete or unused data that can be safely deleted without affecting analysis.
- The information about space usage is taken from the experiment dedicated data management and transfer system

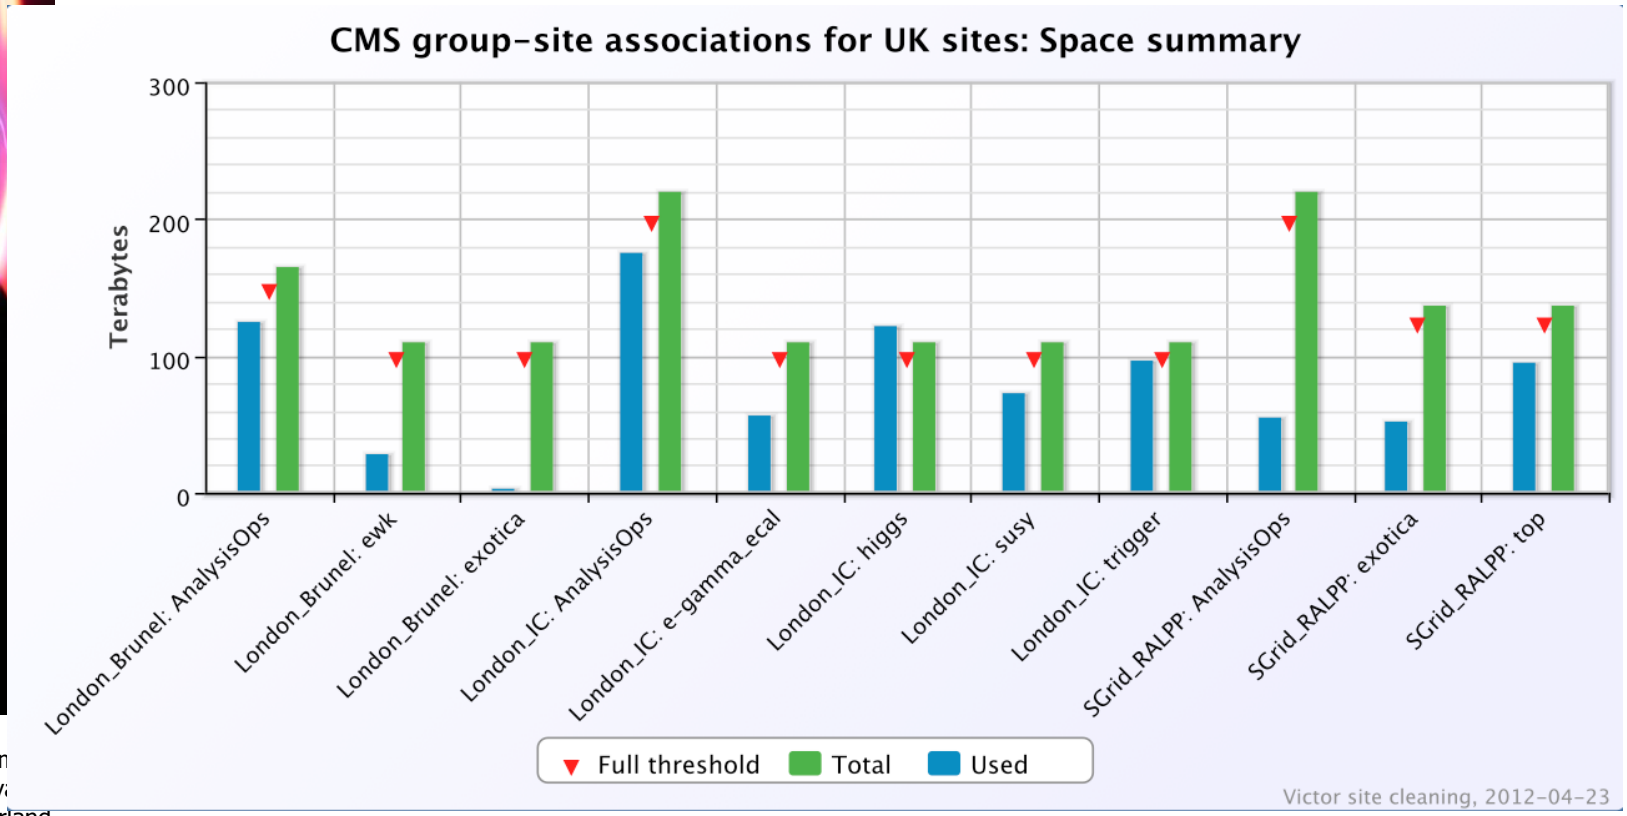

ES Dashboard Framework and Applications

CERN IT Department

Sites and activities

Framework & visualization

Job submission & data transfers

Dashboard is one of the original common services

- All experiments execute jobs and transfer data
- Dashboard services rely on experiment specific information for site names, activity mapping, error codes
- The job monitoring system collects centrally information from workflows about the job status and success
  - Database, framework and visualization are common

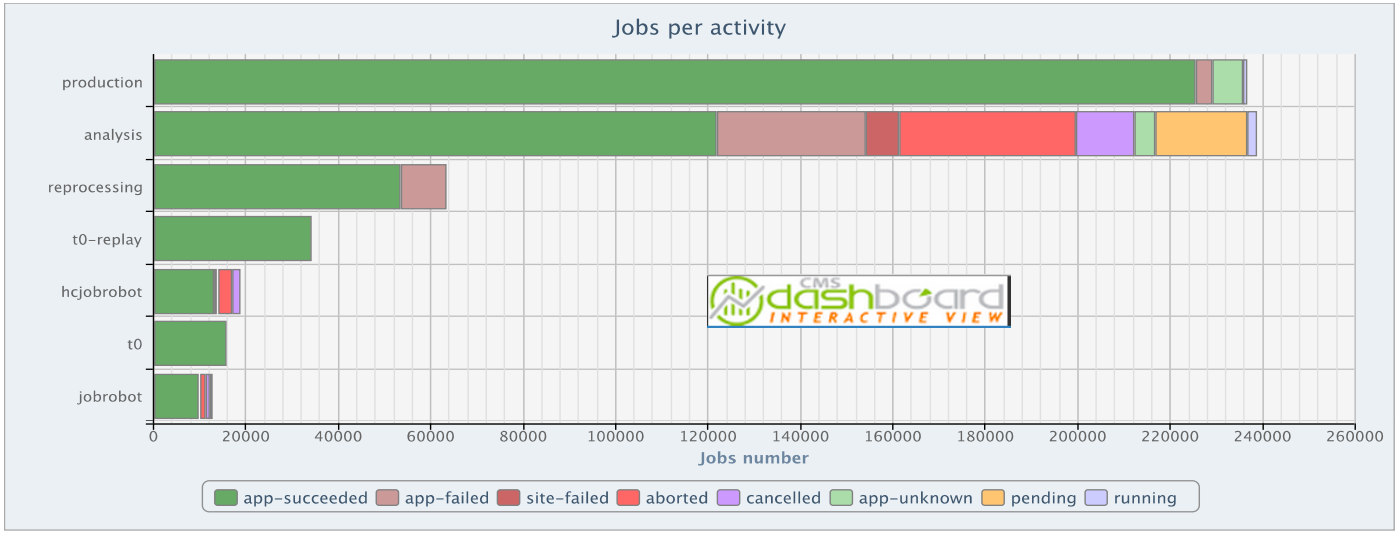

00], *Designing and developing portable large-scale JavaScript web applications within the Experiment Dashboard framework*

Site Status Board

CERN IT Department

- Another example of a good common service
  - Takes specific lower level checks on the health of common services
  - Combines with some experiment specific workflow probes
  - Includes links into the ticketing system
  - Combines to a common view

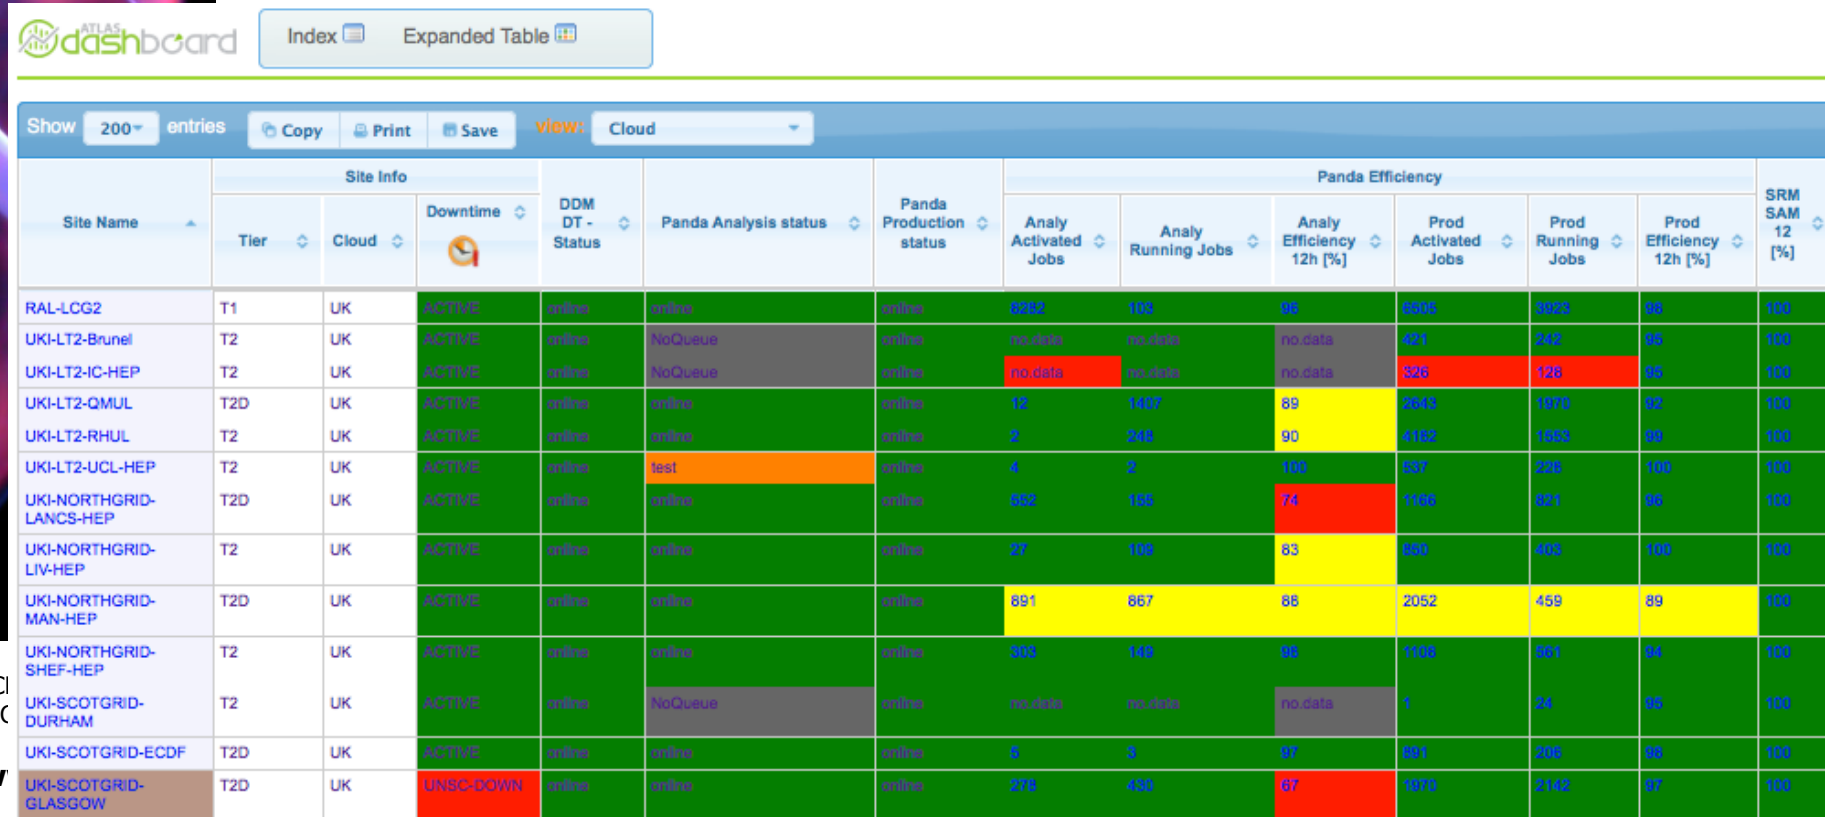

Site Name	Site Info			DDM DT - Status	Panda Analysis status	Panda Production status	Panda Efficiency						SRM SAM 12 [%]
	Tier	Cloud	Downtime				Analy Activated Jobs	Analy Running Jobs	Analy Efficiency 12h [%]	Prod Activated Jobs	Prod Running Jobs	Prod Efficiency 12h [%]	
RAL-LCG2	T1	UK	ACTIVE	online	online	online	8282	103	96	6605	3923	98	100
UKI-LT2-Brunel	T2	UK	ACTIVE	online	NoQueue	online	no data	no data	no data	421	242	95	100
UKI-LT2-IC-HEP	T2	UK	ACTIVE	online	NoQueue	online	no data	no data	no data	326	128	95	100
UKI-LT2-QMUL	T2D	UK	ACTIVE	online	online	online	12	1407	89	2643	1970	92	100
UKI-LT2-RHUL	T2	UK	ACTIVE	online	online	online	2	248	90	4162	1663	99	100
UKI-LT2-UCL-HEP	T2	UK	ACTIVE	online	test	online	4	2	100	537	226	100	100
UKI-NORTHGRID-LANCS-HEP	T2D	UK	ACTIVE	online	online	online	562	155	74	1166	821	96	100
UKI-NORTHGRID-LIV-HEP	T2	UK	ACTIVE	online	online	online	27	109	83	860	403	100	100
UKI-NORTHGRID-MAN-HEP	T2D	UK	ACTIVE	online	online	online	891	867	88	2052	459	89	100
UKI-NORTHGRID-SHEF-HEP	T2	UK	ACTIVE	online	online	online	303	149	98	1108	561	94	100
UKI-SCOTGRID-DURHAM	T2	UK	ACTIVE	online	NoQueue	online	no data	no data	no data	1	24	95	100
UKI-SCOTGRID-ECDF	T2D	UK	ACTIVE	online	online	online	5	3	97	891	206	98	100
UKI-SCOTGRID-GLASGOW	T2D	UK	UNSC-DOWN	online	online	online	278	430	67	1970	2142	97	100

ES HammerCloud

CERN IT Department

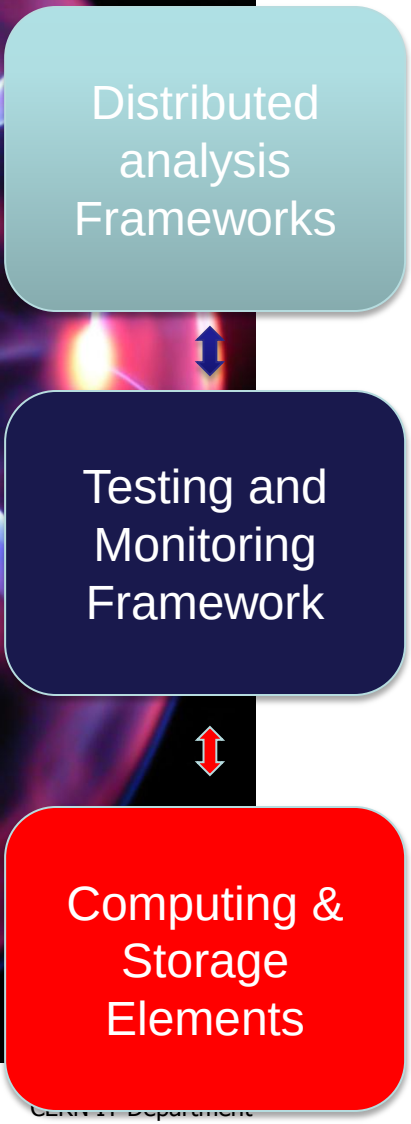

- HammerCloud is a common testing framework for ATLAS (PanDA), CMS (CRAB) and LHC*b* (Dirac)
- Common layer for functional testing of CEs and SEs from a user perspective
- Continuous testing and monitoring of site status and readiness. Automatic Site exclusion based on defined policies
- Same development, same interface, same infrastructure → less workforce

ES HammerCloud

CERN IT Department

HammerCloud | LHCB

Home | Tests | Statistics

HammerCloud | ATLAS

You are connected as gangarbt, click here to Logout!

Home | Tests | Statistics | Robot | PanDA Dashb. | More HC... | Help | Administration

Welcome to HammerCloud-ATLAS.

Running and Scheduled AFT Tests

State	Id	Host	Template	Start (CET)	End (CET)	Cloud	Sites	subm jobs	run jobs	comp jobs	fail jobs	tot jobs
						CA_PANDA, DE_PANDA, ES_PANDA, 8 more...	ANALY_TRIUMF, ANALY_AUSTRALIA, ANALY_SFU_bugaboo, 94 more...	55	55	70	2	243
						CA_PROD, DE_PROD, ES_PROD, 8 more...	CA-ALBERTA-WESTGRID-T2, CA-VICTORIA-WESTGRID-T2, SFU-LCG2, 111 more...	34	45	79	2	169
						CA_PROD, DE_PROD, ES_PROD, 8 more...	CA-ALBERTA-WESTGRID-T2, CA-VICTORIA-WESTGRID-T2, SFU-LCG2, 111 more...	56	121	526	5	751
						CA_PANDA, DE_PANDA, ES_PANDA, 8 more...	ANALY_TRIUMF, ANALY_AUSTRALIA, ANALY_SFU_bugaboo, 93 more...	71	78	717	19	950
						CA_PANDA, DE_PANDA, ES_PANDA, 8 more...	ANALY_TRIUMF, ANALY_AUSTRALIA, ANALY_SFU_bugaboo, 94 more...	29	11	1251	18	1325
						CA_PROD, DE_PROD, ES_PROD, 8 more...	CA-ALBERTA-WESTGRID-T2, CA-VICTORIA-WESTGRID-T2, SFU-LCG2, 112 more...	64	135	1530	34	1801
						CA_PROD, DE_PROD, ES_PROD, 8 more...	CA-ALBERTA-WESTGRID-T2, CA-VICTORIA-WESTGRID-T2, SFU-LCG2, 111 more...	35	38	2498	88	2667
						CA_PANDA, DE_PANDA, ES_PANDA, 8 more...	ANALY_TRIUMF, ANALY_AUSTRALIA, ANALY_SFU_bugaboo, 95 more...	68	87	6633	210	7049
						CA_PROD, DE_PROD, ES_PROD, 8 more...	CA-ALBERTA-WESTGRID-T2, CA-VICTORIA-WESTGRID-T2, SFU-LCG2, 112 more...	79	144	7061	63	7385

HammerCloud | CMS

You are connected as gangarbt, click here to Logout!

Home | Tests | Statistics | Robot | More HC... | Help | Administration

Welcome to HammerCloud-CMS.

Running and Scheduled Stress Tests

No entries

Running and Scheduled Functional Tests

State	Id	Host	Template	Start (CET)	End (CET)	Region	Sites	subm jobs	run jobs	comp jobs	fail jobs	tot jobs
running	6025	vocms207	17: default functional T2 IT	14/May, 10:46 a.m.	16/May, 7:21 a.m.	CNAF, OTHER	T2_IT_Bari, T2_IT_Legnaro, T2_IT_Pisa, 2 more...	120	117	1	0	238
running	6024	vocms207	20: default functional T2 UK	14/May, 10:44 a.m.	15/May, 8:48 a.m.	RAL	T2_UK_London_Brunel, T2_UK_London_IC, T2_UK_SGrid_Bristol, 3 more...	225	98	6	11	340
running	6022	vocms06	19: default functional T2 US	14/May, 9:56 a.m.	15/May, 9:31 a.m.	FNAL	T2_US_Caltech, T2_US_Florida, T2_US_MIT, 8 more...	559	197	31	29	816
running	6021	vocms207	18: default functional T2 ES	14/May, 7:34 a.m.	15/May, 9:54 a.m.	PIC	T2_ES_CIEMAT, T2_ES_IFCA, T2_PT_LIP_Lisbon, 1 more...	207	65	165	107	544
running	6020	vocms207	3: default functional T2 DE	14/May, 5:56 a.m.	15/May, 8:06 a.m.	KIT	T2_DE_DESY, T2_DE_RWTH, T2_AT_Vienna, 2 more...	280	34	127	511	952
running	6019	vocms207	4: default functional T2 FR	14/May, 3:20 a.m.	15/May, 4:20 a.m.	CC-IN2P3	T2_FR_CCIN2P3, T2_FR_GRIF_IRFU, T2_FR_GRIF_LLR, 4 more...	419	25	804	248	1496
running	6018	vocms06	21: default functional T2 TW	14/May, 12:04 a.m.	15/May, 2:18 a.m.	ASGC	T2_TW_Taiwan, T2_KR_KNU, T2_IN_TIFR, 1 more...	136	39	254	557	986
running	6017	vocms06	30: default functional T2 CH	13/May, 3:12 p.m.	14/May, 3:43 p.m.	CERN Tier-0	T2_RU_IHEP, T2_RU_INR, T2_RU_ITEP, 6 more...	473	73	1299	841	2686
running	6011	vocms06	22: default functional T1	13/May, 12:58 a.m.	15/May, 2:54 a.m.	KIT, PIC, CC-IN2P3, 5 more...	T1_DE_KIT, T1_ES_PIC, T1_FR_CCIN2P3, 5 more...	455	17	2882	216	3570

D. van der Ster et al. [283], *Experience in Grid Site Testing for ATLAS, CMS and LHCb with HammerCloud*

New Activities – Analysis Workflow

Data discovery, environment configuration, and job splitting

Job Tracking, Resubmission, and scheduling

Job submission and Pilots

- Up to now services have generally focused on monitoring activities
  - All of these are important and commonality saves effort
  - Not normally in the core workflows of the experiment
- **Success with the self contained services has provided confidence moving into a core functionality**
  - Looking at the Analysis Workflow
    - Feasibility Study for a Common Analysis Framework between ATLAS and CMS

CERN IT Department
CH-1211 Geneva 23
Switzerland
www.cern.ch/it

ES Analysis Workflow Progress

Data discovery, job splitting and packaging of user environment

Job Tracking, Resubmission, and scheduling

Job submission and Pilots

- Looking at ways to make the workflow engine common between the two experiments
  - Improving the sustainability of the central components that interface to low-level services
    - A thick layer that handles prioritization, job tracking and resubmission
  - Maintaining experiment specific interfaces
    - Job splitting, environment, and data discovery would continue to be experiment specific

CERN IT Department
CH-1211 Geneva 23
Switzerland
www.cern.ch/it

Proof of Concept Diagram

- Feasibility Study proved that there are no show-stoppers to design a common analysis framework
- Next step is a proof of concept

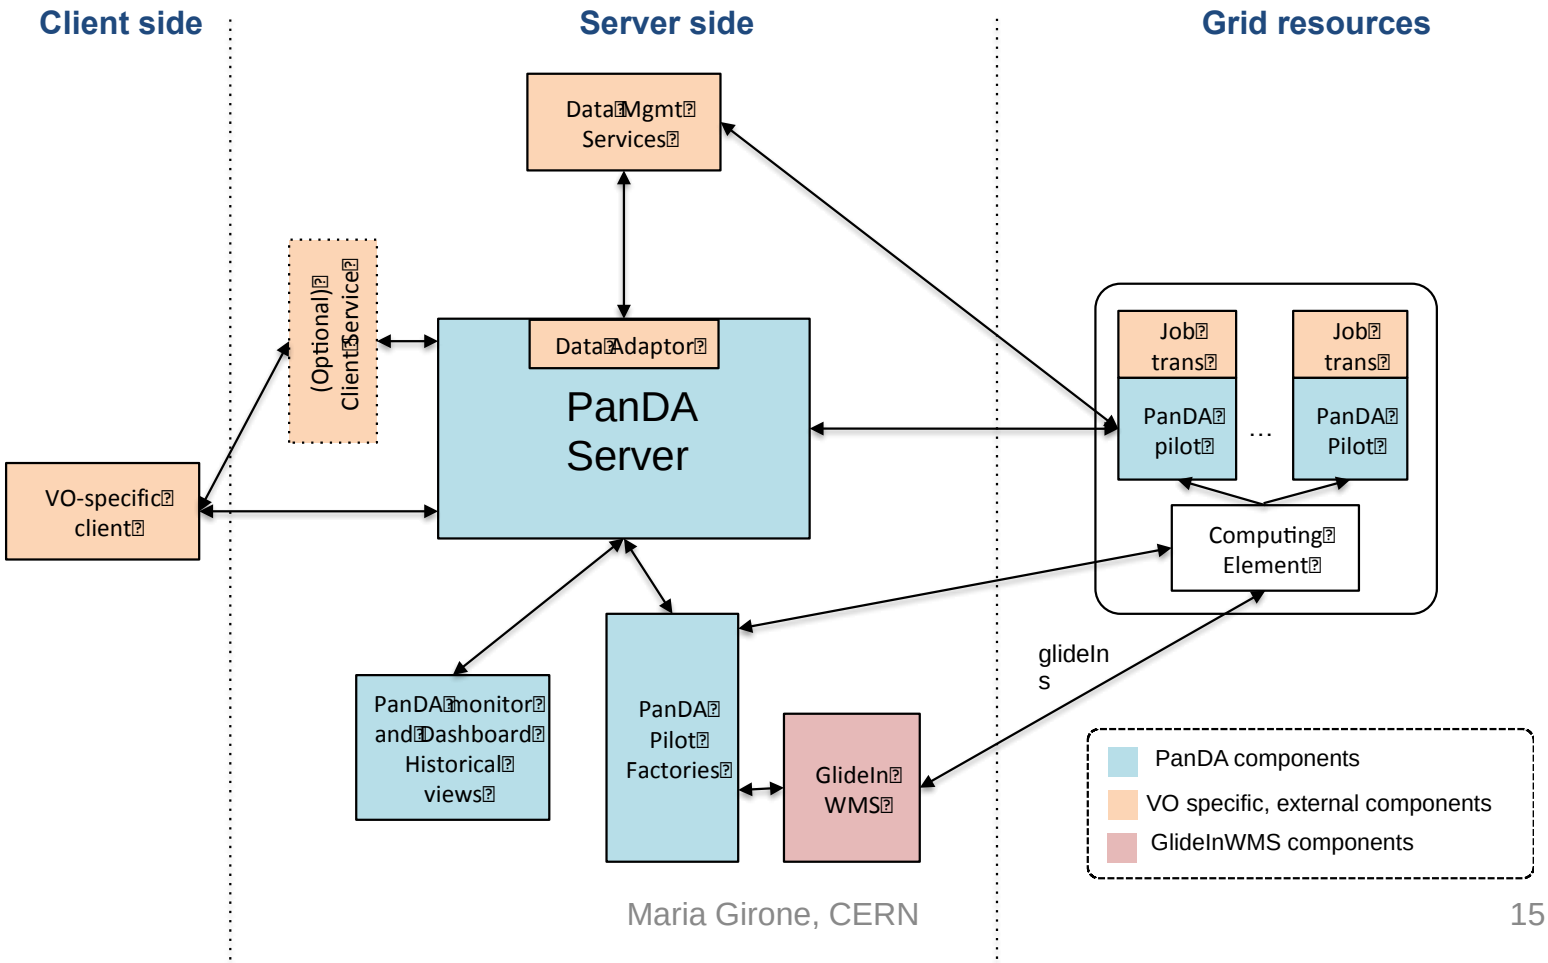

ES Even Further Ahead

CERN IT Department

- As we move forward, we would also like to assess and document the process
  - This should not be the only common project
- The diagram for data management would look similar
  - A thick layer between the experiment logical definitions of datasets and the service that moves files
    - Deals with persistent location information and tracks files in progress and validates file consistency
- Currently no plans for common services, but has the right properties

Outlook

- ✓ IT-ES has a good record of identifying and developing common solutions between the LHC experiments
  - Setup and expertise of the group have helped
- ✓ Several services focused primarily on monitoring have been developed and are in production use
- ✓ **As a result, more ambitious services that would be closer to the experiment core workflows are under investigation**
  - ✓ **The first is a feasibility study and proof of concept of a common analysis framework between ATLAS and CMS**
- ✓ **Both better and more sustainable solutions could result – with lower operational and maintenance costs**

CERN IT Department
CH-1211 Geneva 23
Switzerland
www.cern.ch/it