21st International Conference on Computing in High Energy and Nuclear Physics **CHEP2015** Okinawa Japan: April 13 - 17, 2015

Track 4

Middleware,
software development and tools,
experiment frameworks,
tools for distributed computing

Marco Clemencic
on behalf of the conveners of Track 4

Disclaimer

I tried to summarize all oral contributions
and give a view of all the great work done.
I apologize if I have missed something.

I'm very sorry that I did not
manage to cover the posters.

Overview of Track 4

- 42 orals + 48 posters
- Heterogeneous contributions
- Roughly grouped in categories:
  - Middleware
  - Framework
  - Application
  - Software
- Boundaries often fuzzy, so I reorganized them

Overview of Track 4

- LHC experiments dominating the scene
- Very valuable contributions from
  - Non-LHC/HEP experiments
  - Service providers
  - Computing centers
  - Etc.

Overview of Track 4

- We discussed mainly about middleware
- But also about frameworks and tools
- A lot of work has been done on improvements
  - “rewrite” is not a bad word, don't be afraid
- Sharing efforts seems the key to success

Middleware

- Contributions on
  - Job Management/Pilots
  - Data Management
  - Network Awareness
  - Multicore

Job Management / Pilots

- CMS and LHCb showed how pilots can bring uniformity to the Grid

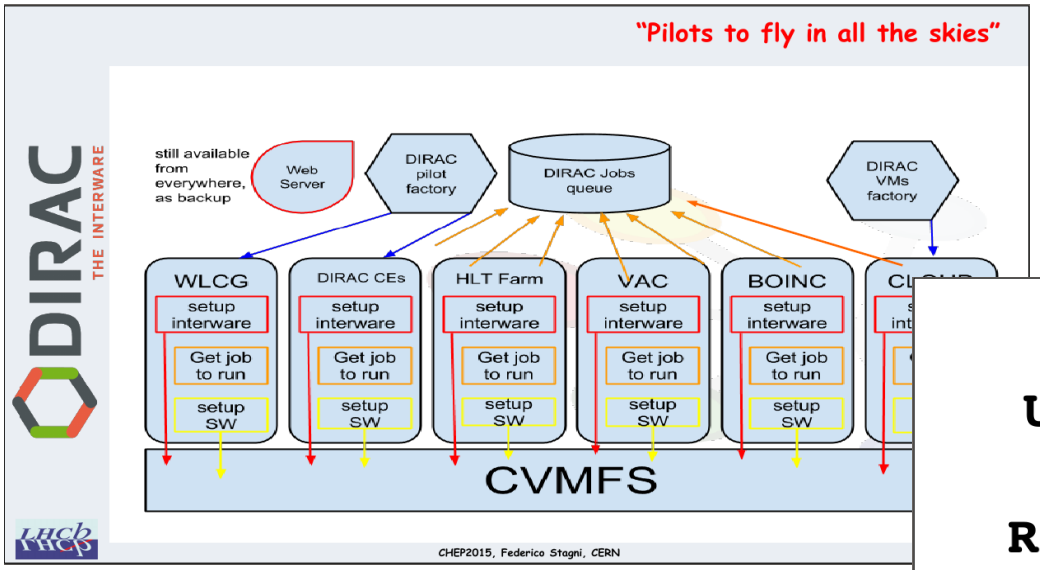

#113

#289

Using the glideinWMS System as a Common Resource Provisioning Layer in CMS

J. Balcas[1], S. Belforte[2], B. Bockelman[3], D. Colling[4], O. Gutsche[5], D. Hufnagel[5], F. Khan[6], K. Larson[5], J. Letts[7], M. Mascheroni[8], D. Mason[5], A. McCrea[7], S. Piperov[9], M. Saiz-Santos[7], I. Sfiligoi[7], C. Wissing[10]

1. Vilnius Univ. (LT) 2. INFN-Trieste (IT) 3. Univ. Nebraska Lincoln (US) 4. Imperial College London (UK) 5. FNAL (US) 6. NCP (PK) 7. UCSD (US) 8. INFN-Milano (IT) 9. Brown (US) 10. DESY (DE)

Middleware

Job Management / Pilots

- CMS commissioned CRAB3

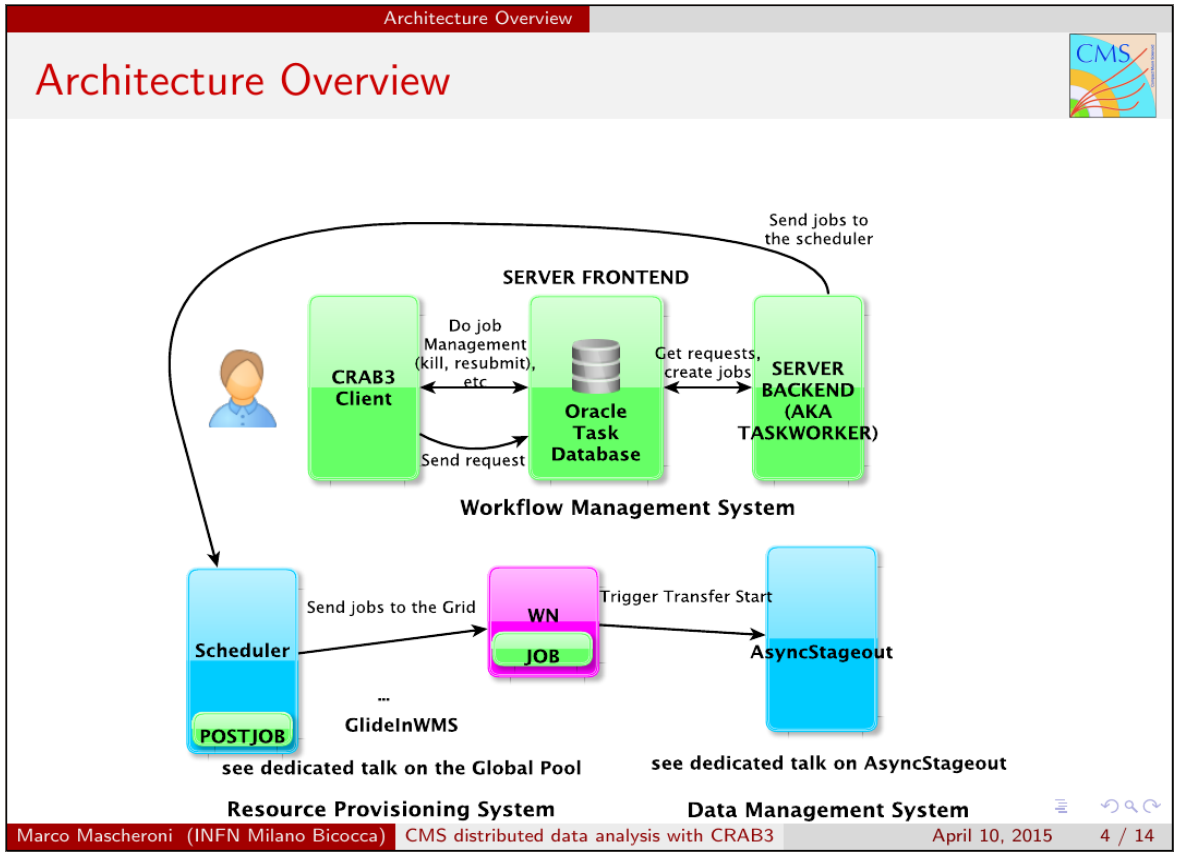

#345

- Complete re-implementation
- Integrates with CMS new developments
  - GlideInWMS Global Pool
  - Asynchronous Stage-Out

Middleware

Job Management / Pilots

- ATLAS presented an overview of the evolution of PanDA in preparation for RUN2

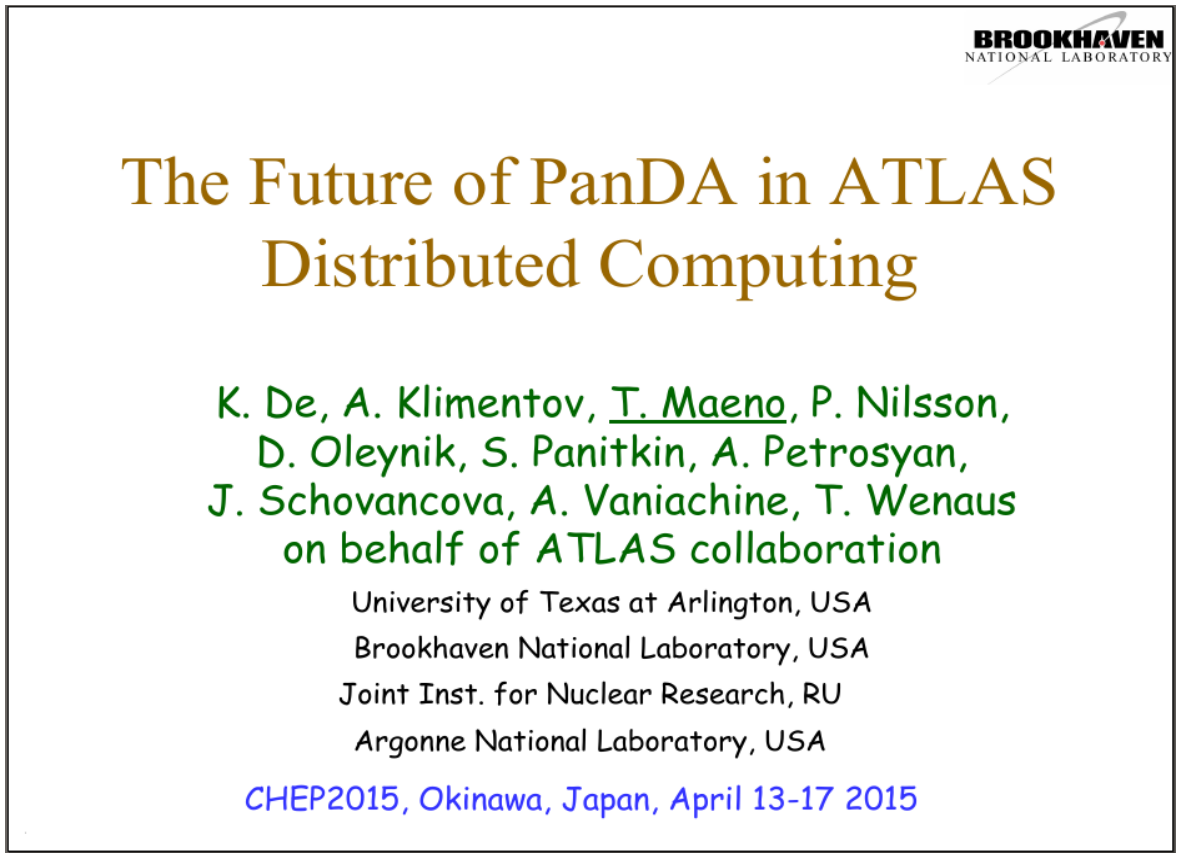

BROOKHAVEN NATIONAL LABORATORY

The Future of PanDA in ATLAS Distributed Computing

K. De, A. Klimentov, T. Maeno, P. Nilsson, D. Oleynik, S. Panitkin, A. Petrosyan, J. Schovancova, A. Vaniachine, T. Wenaus on behalf of ATLAS collaboration

University of Texas at Arlington, USA
Brookhaven National Laboratory, USA
Joint Inst. for Nuclear Research, RU
Argonne National Laboratory, USA

CHEP2015, Okinawa, Japan, April 13-17 2015

#144

- Dynamic Jobs
- Network Awareness
- Event Service
- New Pilot
- Support for HPC
- New Monitoring

Middleware

Job Management / Pilots

- We have seen how MIRA became the primary Alpgen event generation site for ATLAS

#537

via integration with PanDA

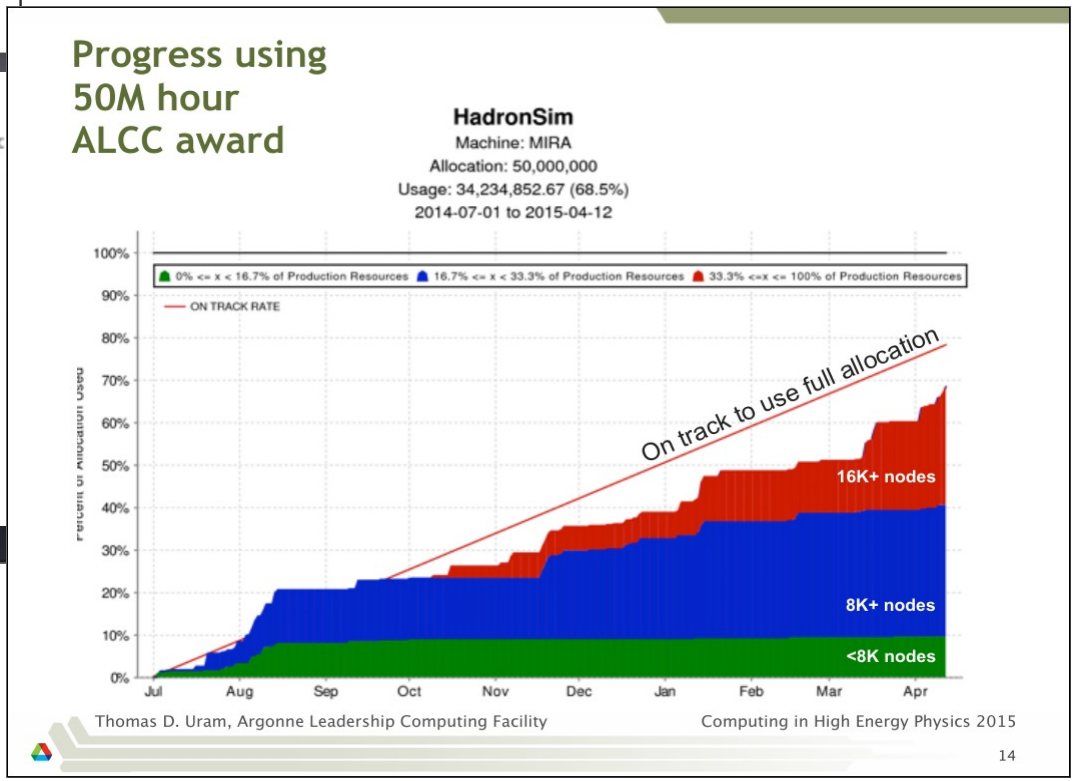

Job Management / Pilots

- Many contributions on ARC Control Tower
  - Job Management Layer in front of ARC-CE

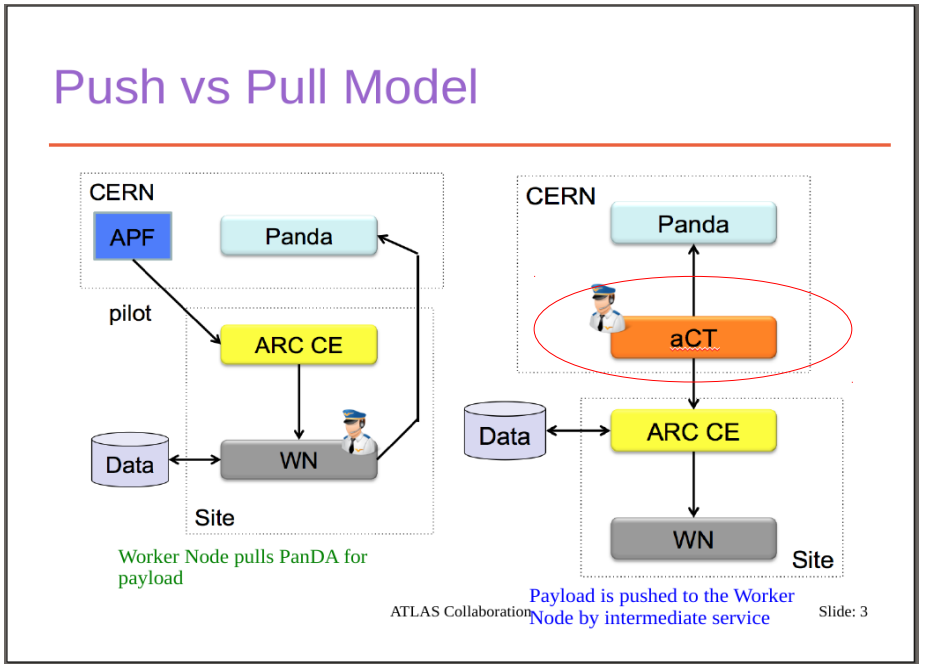

#145

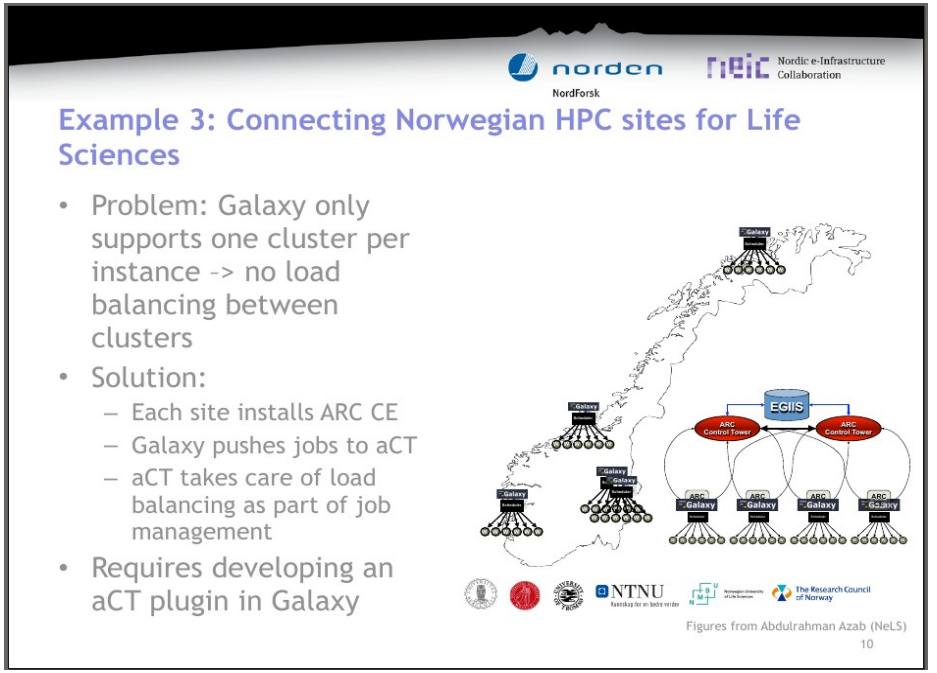

#263

Job Management / Pilots

- HTCondor-CE
  - use HTCondor to provide a CE interface

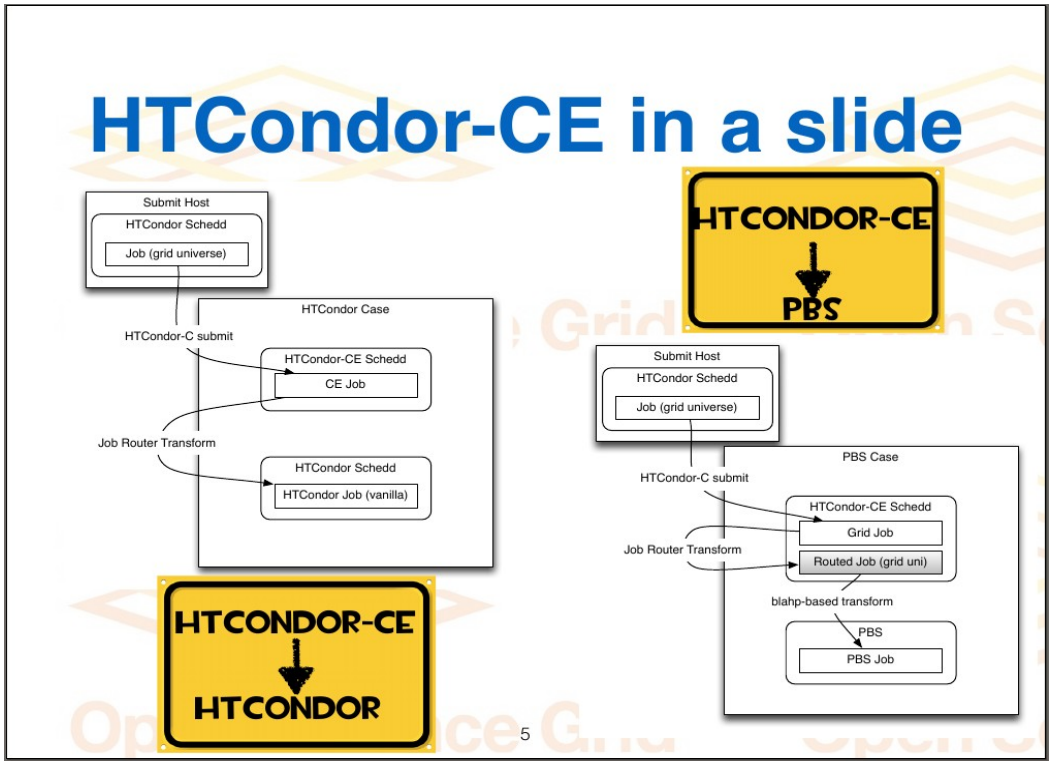

- It's a special configuration of HTCondor
- Choice strategic and technical

#519

Job Management / Pilots

- BelleII adopted DIRAC for their Production System

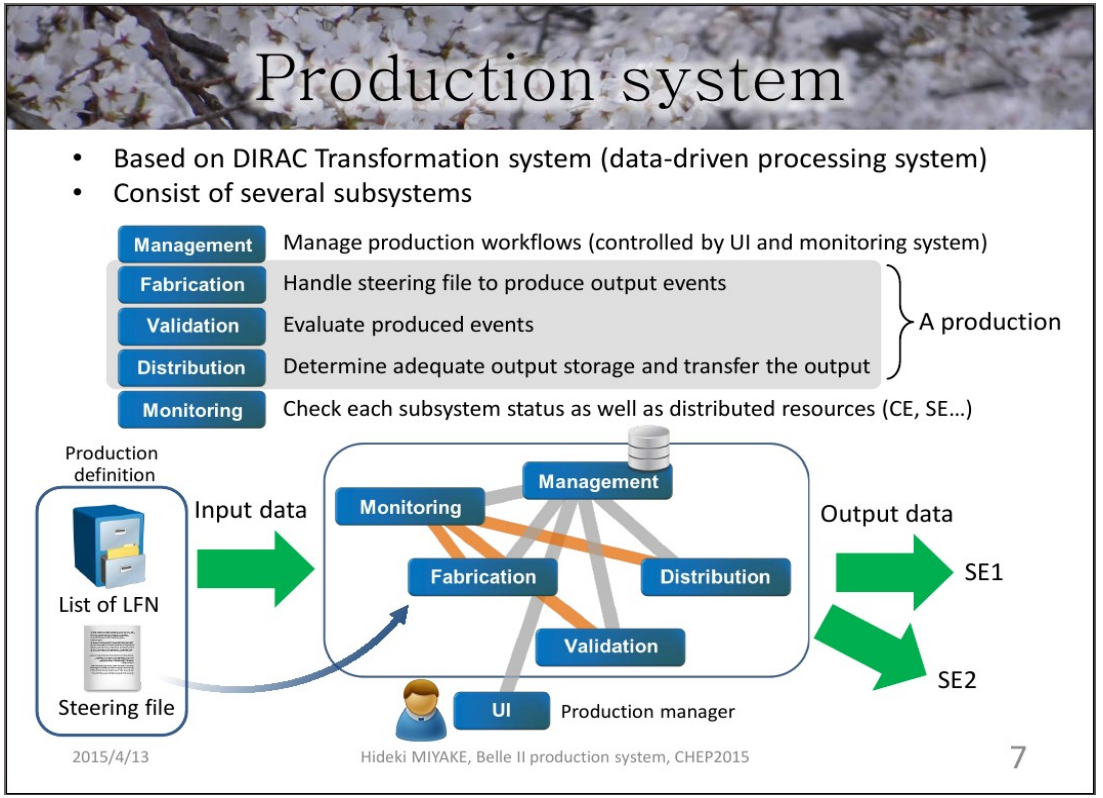

#329

Job Management / Pilots

- IHEP and GridPP extended DIRAC to support their many (small) VOs

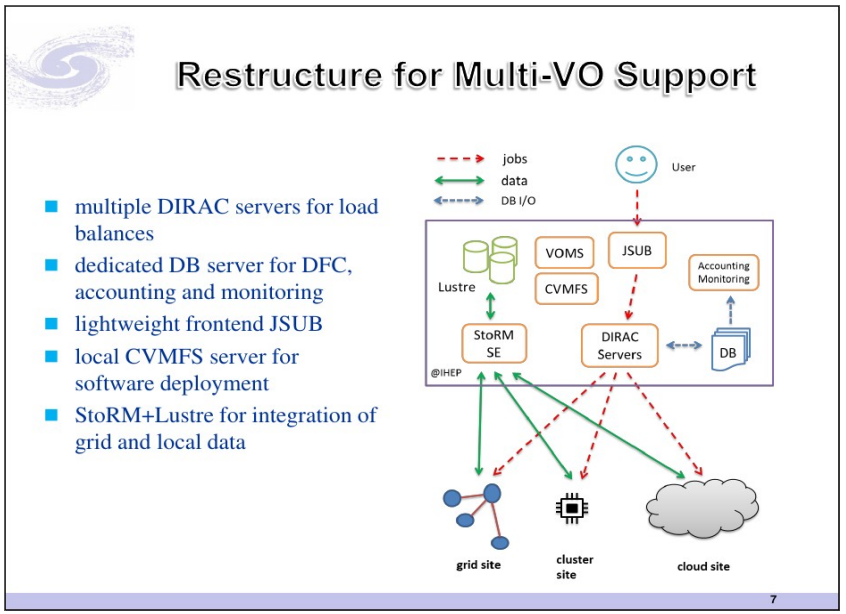

#479

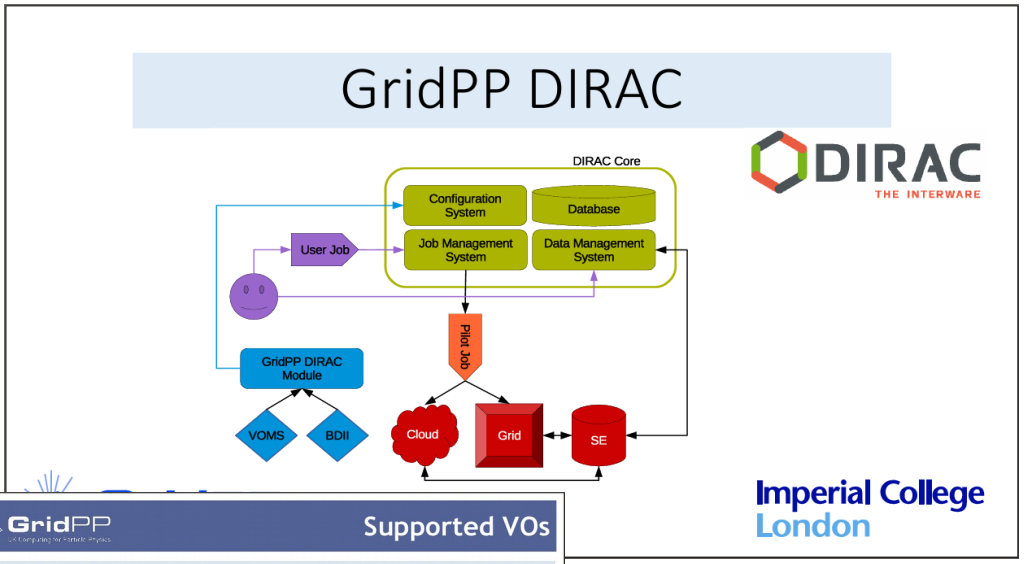

#346

#334

Middleware

Job Management / Pilots

- Fermilab combined existing tools to provide a new Distributed Computing system: FIFE

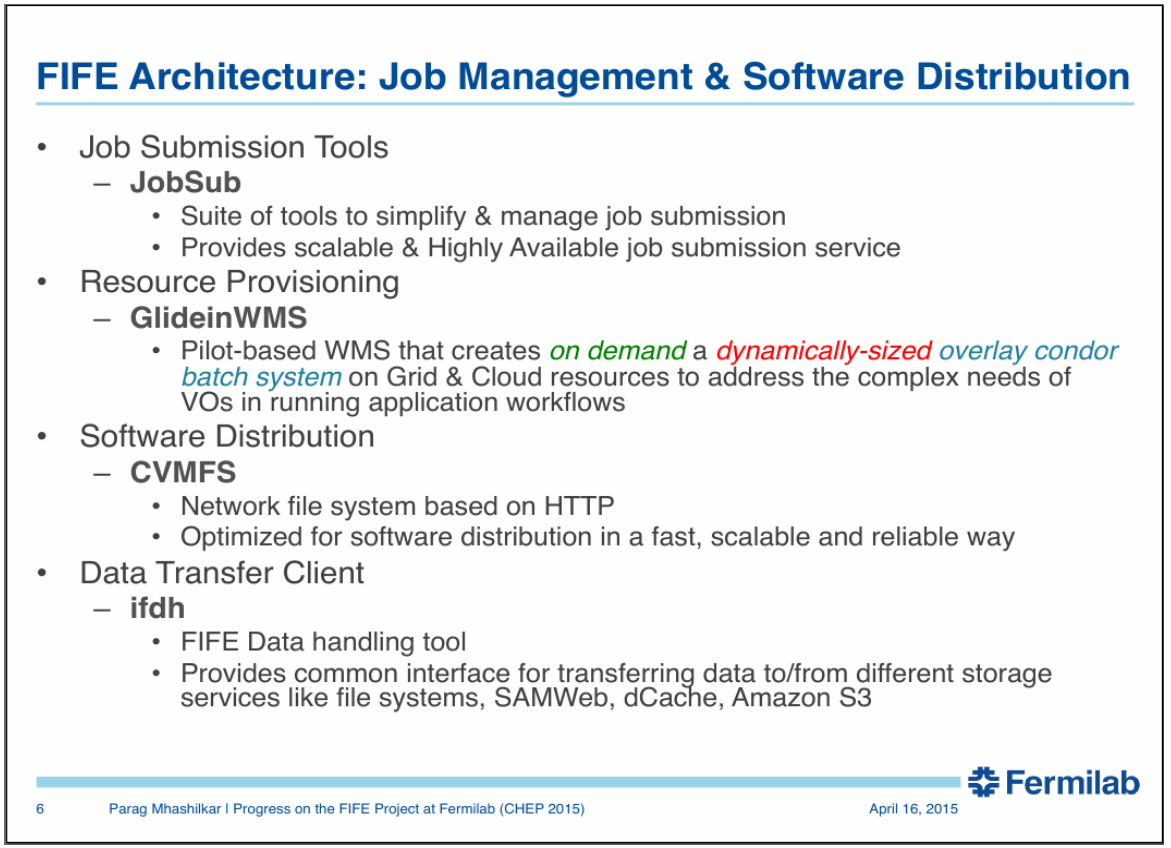

#444

Job Management / Pilots

- LHCb and ATLAS addressed the “masonry problem”

#112

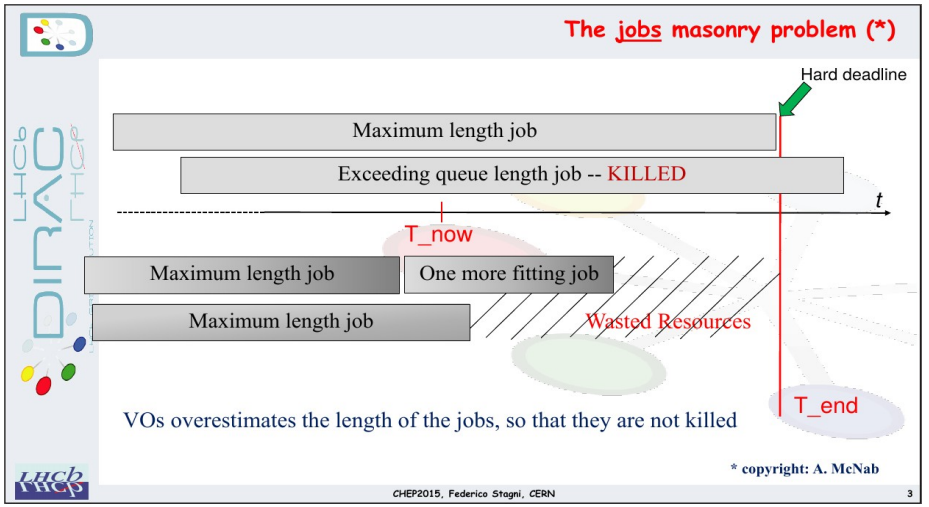

#112

LHCb can gracefully stop simulation jobs just before the allocated time is over.

ATLAS distributes single events to workers.

#183

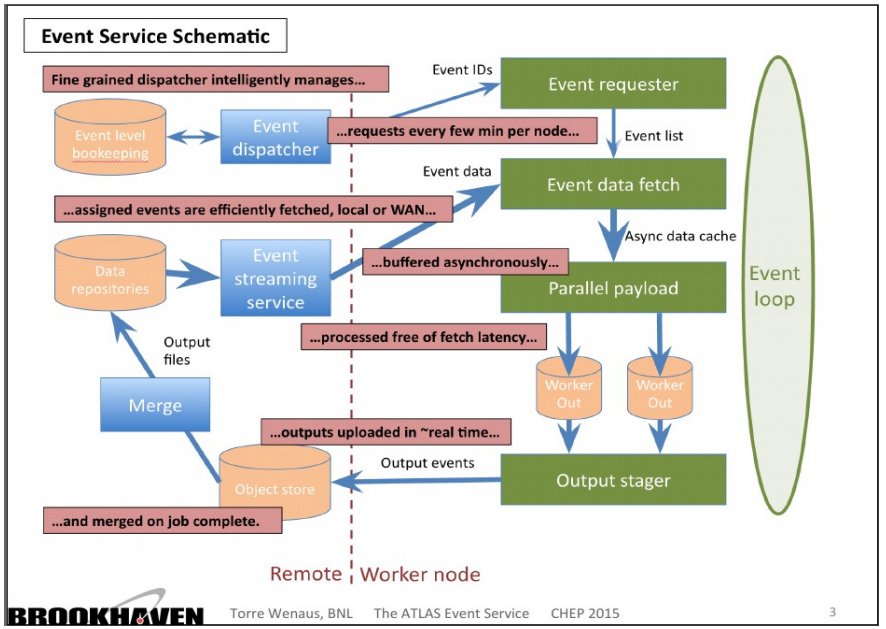

Job Management / Pilots

- LHCb showed how to predict required resources
- ALICE studied how to increase security on the Grid

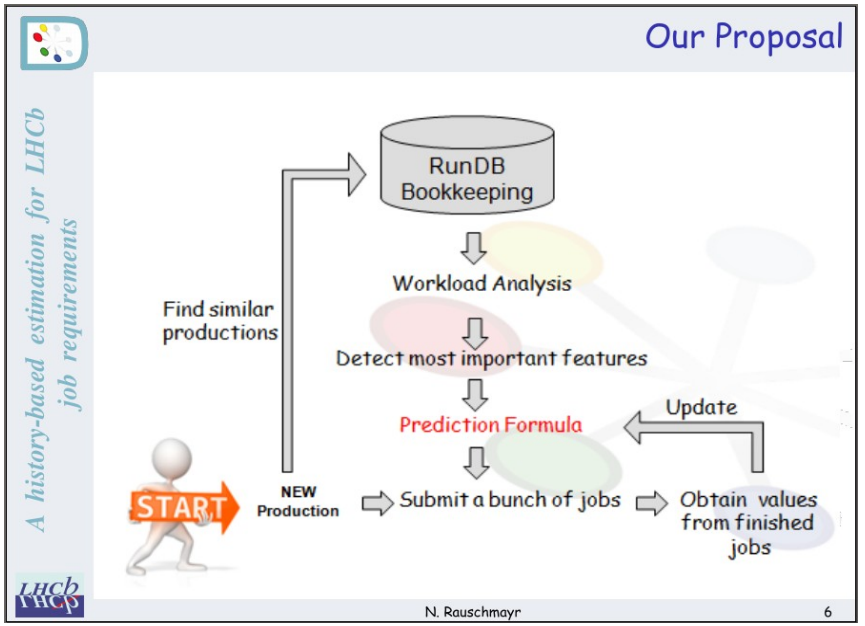

#96

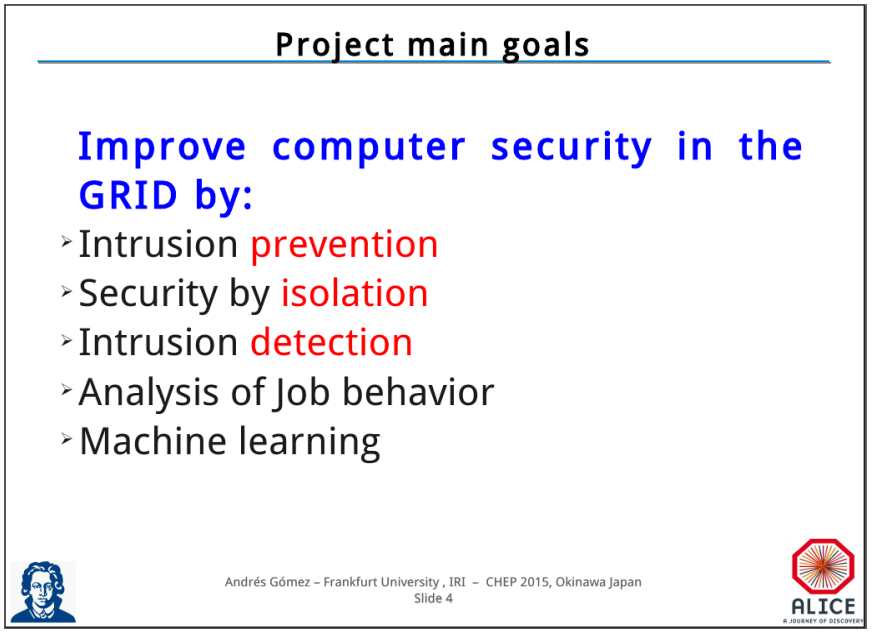

#14

Middleware

Data Management

- ATLAS implemented Rucio
  a new Distributed Data Management tool

Replacement for DQ2

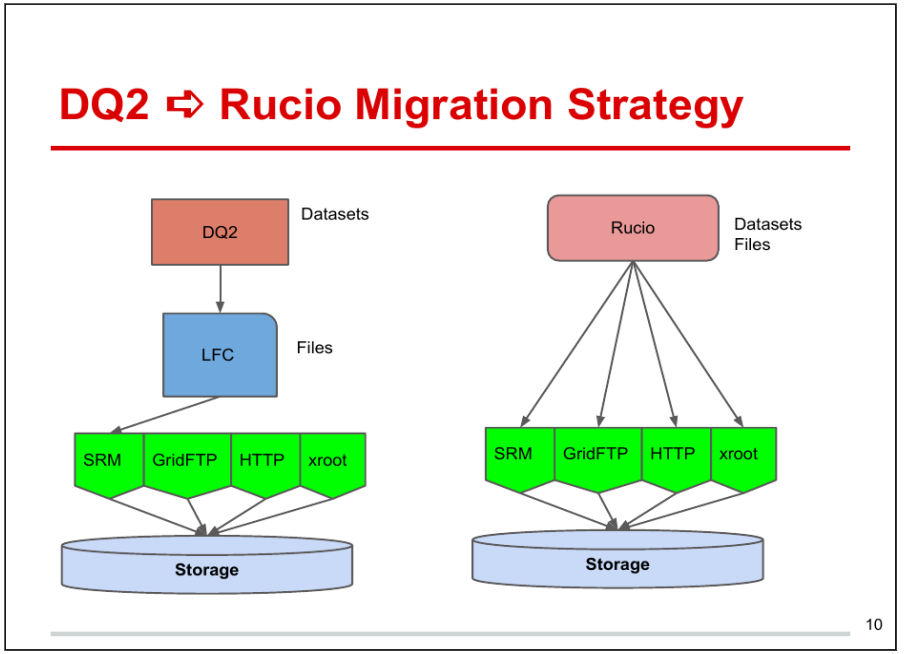

#205

Flexible Quotas and Accounting

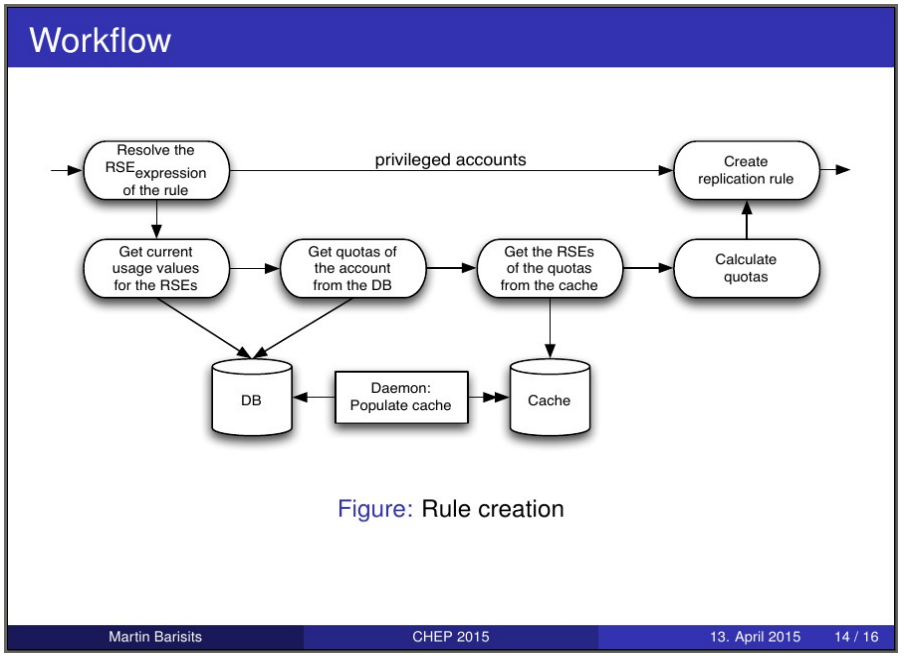

#207

Middleware

Data Management

- CMS implemented Asynchronous stage-out to avoid that jobs fail during data transfer

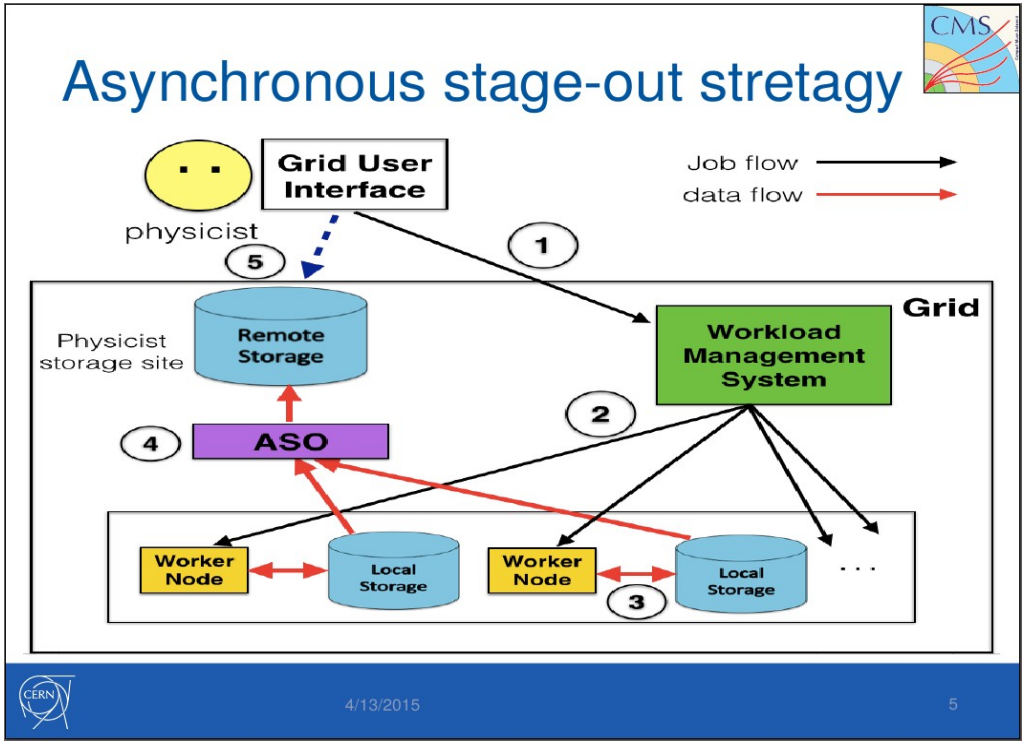

#225

Middleware

Data Management

- Fermilab re-engineered SAM with new interface while maintaining operations

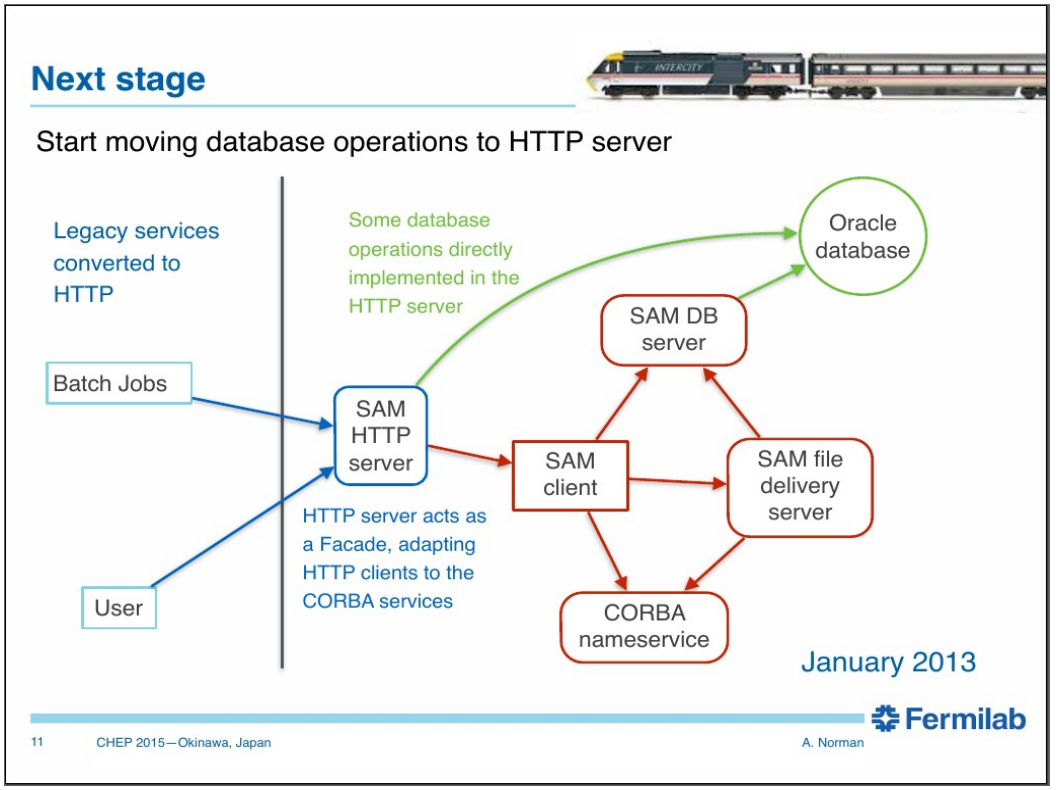

#463

Middleware

Network Awareness

- CMS and ATLAS showed uses of Network Awareness

CMS improves data transfers

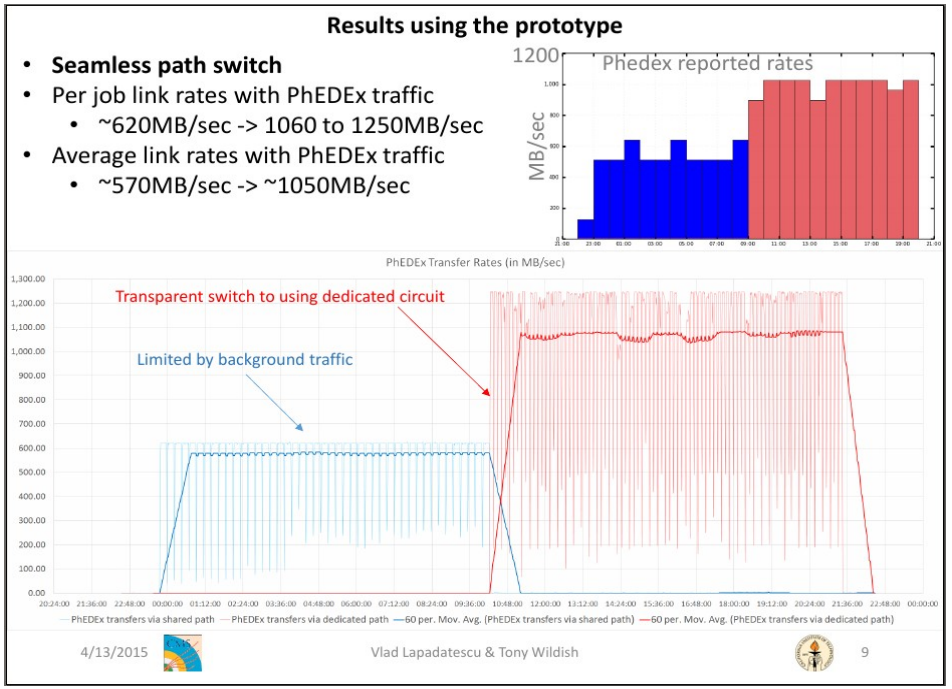

#191

ATLAS improves job submission

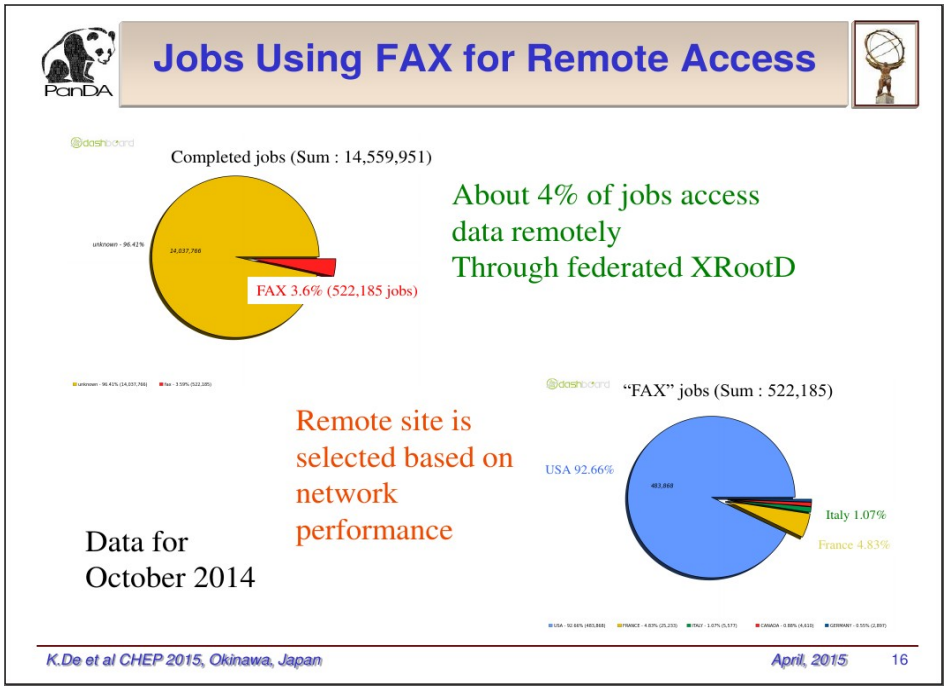

#237

Middleware

Multicore

- Report from WLCG Multicore Task Force
  - ATLAS and CMS cases

MANCHESTER 1824

Experiments submission

- CMS move the scheduling within the pilot
  - Predictability
  - Shared sites still have single core to handle
- ATLAS: mcore and score in parallel with 1 payload per pilot and let the scheduler do the job.
  - Entropy
  - Predictability still helps
    - Backfilling not an option yet

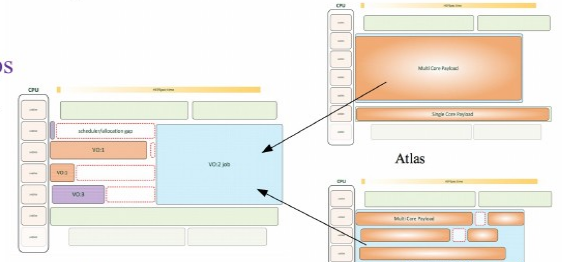

5

#225

- Good progress
- It works already
  - Fine tuning needed

Middleware

Multicore

- CMS reported on their successful use of multicore jobs on the Grid
- Ready for RUN2

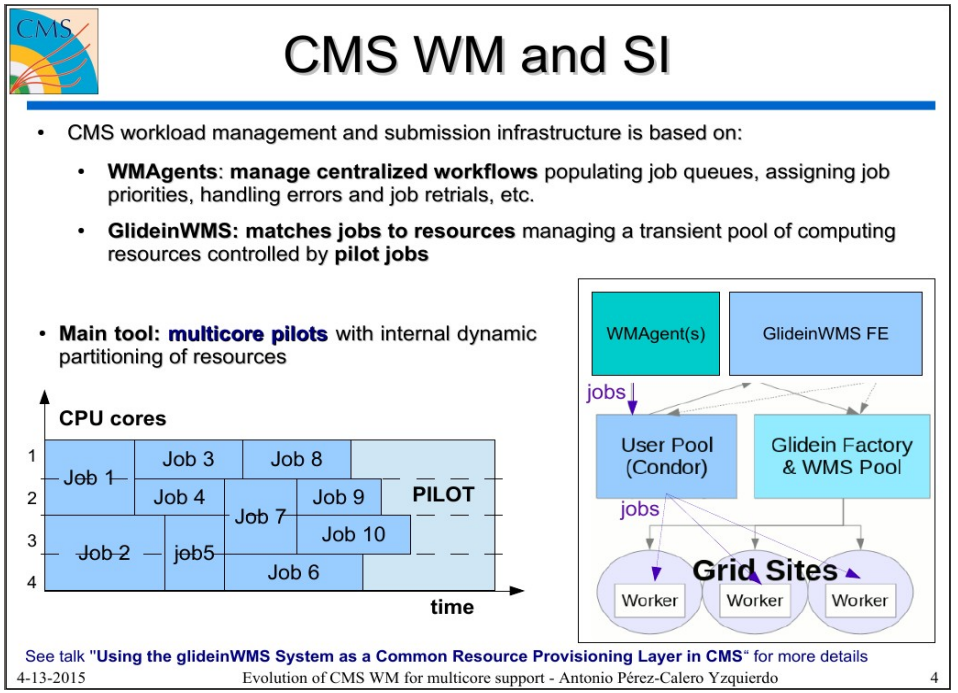

#409

Middleware

Multicore

- Interesting report of the importance of multicore awareness in data transfers

MDTM Early Test Evaluation

- Wide-area network links, end-to-end tests
- Performance comparison with GridFTP, bandwidth captured at ESnet's edge router

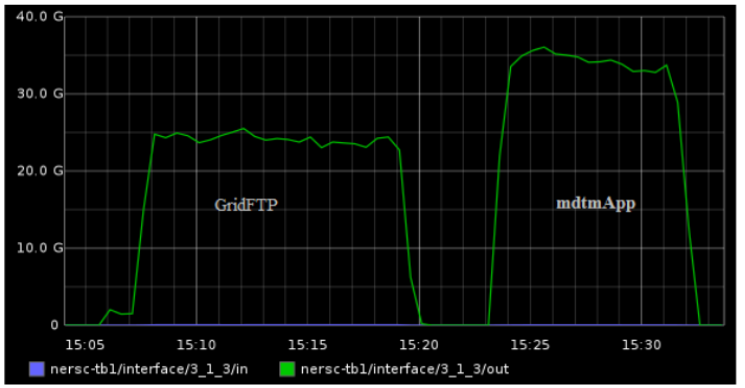

Parallel large file transfers(16 streams, 2TB, 8MB blocks), from SSDs to /dev/null, with 40Gbps links and 50Gbps aggregate disk bandwidth

Fermilab

#457

Frameworks

- Contributions on
  - Experiment frameworks
  - Monitoring frameworks
  - Validation frameworks
  - Analysis frameworks

Experiment Frameworks

- IceCube presented their new IceProd2

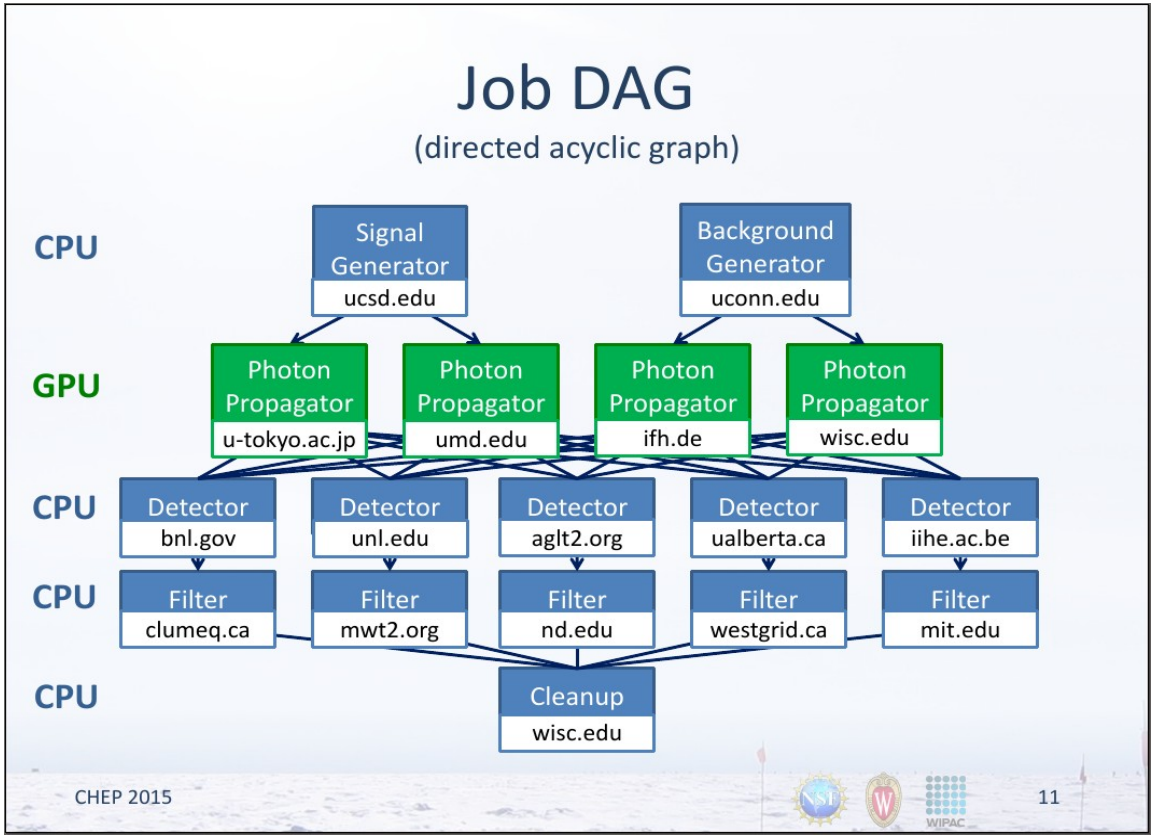

- Complete rewrite on
  - Python
  - SQLite
  - CVMFS
  - Web API
- Pilot jobs
- User permissions

#496

Experiment Frameworks

- Reports on ROOT 6 and beyond
  - Impossible to summarize all the changes, see #441
  - Impressive work on optimization and validation

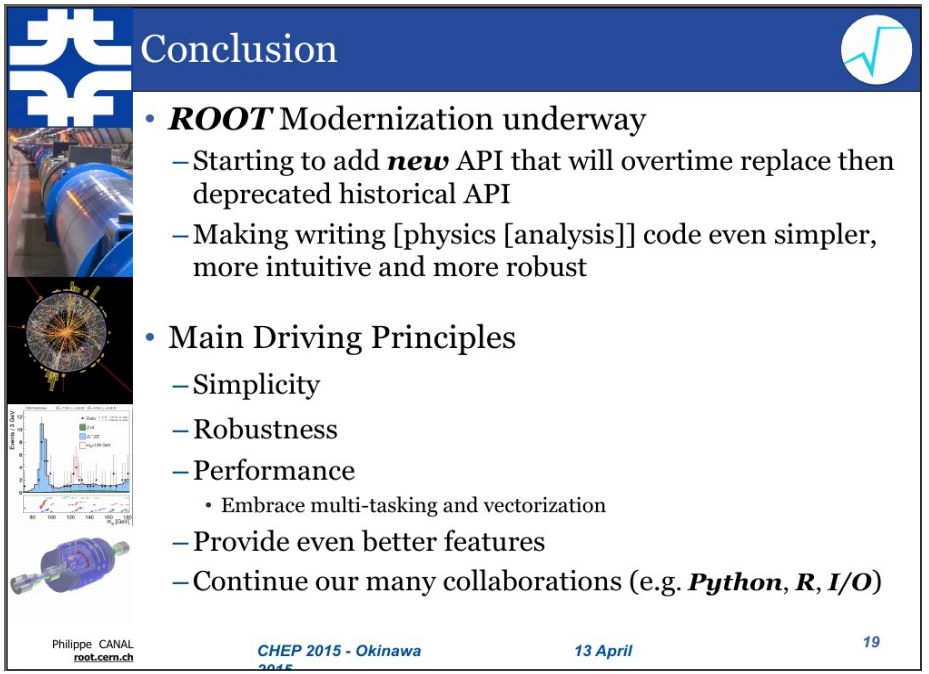

#441

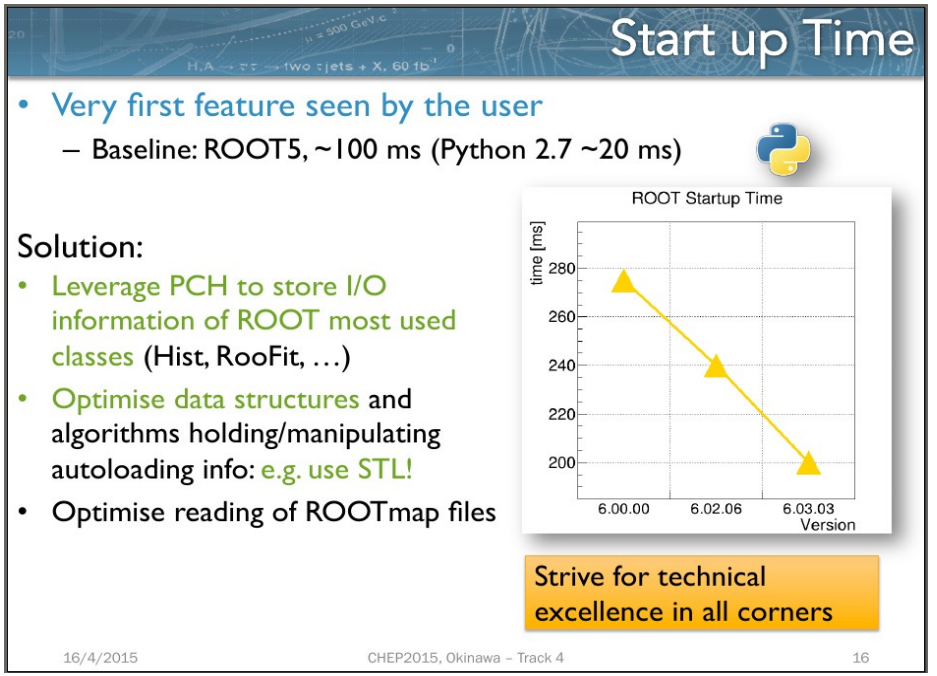

#381

Monitoring Frameworks

- From DESY we saw how to display dCache accounting informations

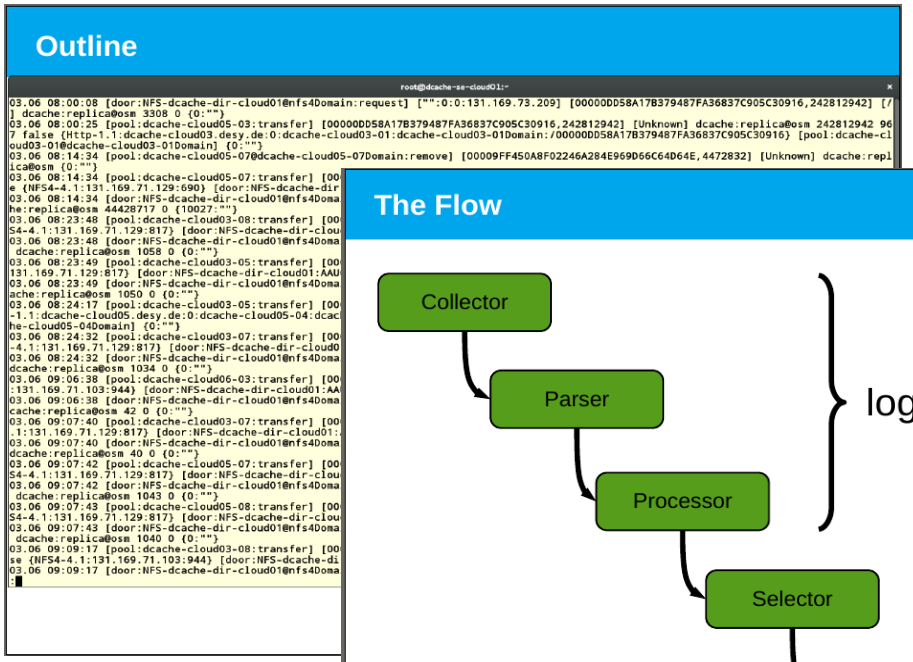

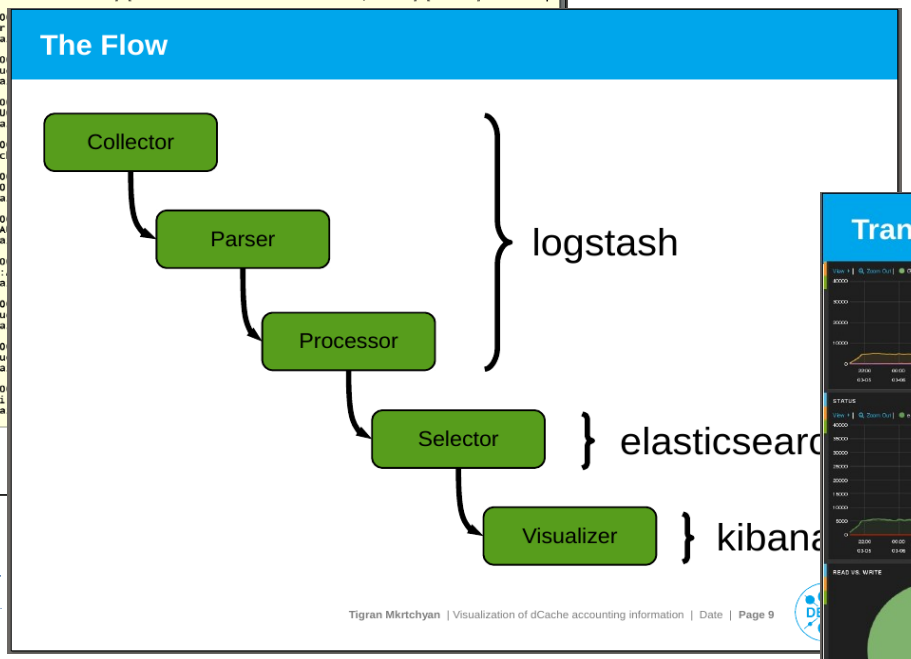

#45

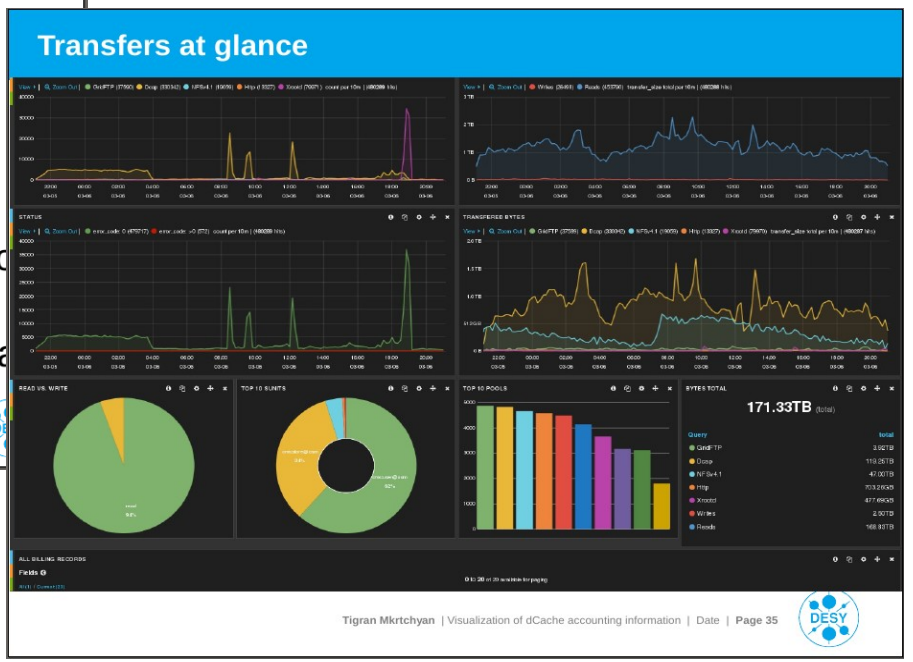

Monitoring Frameworks

- The quantitative monitoring of FTS3 has been crucial for commissioning and production

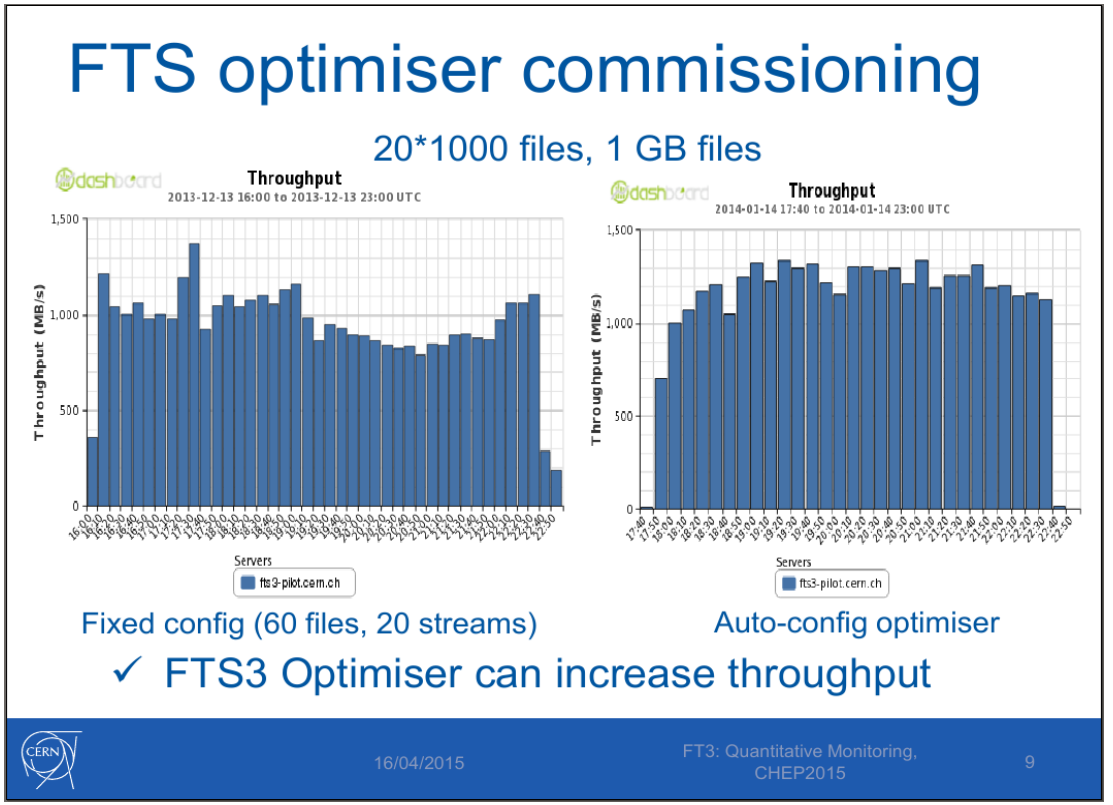

#232

Validation Frameworks

- Geant4 has been used as a test bench to study
  - Testability
  - Statistical analysis of software quality

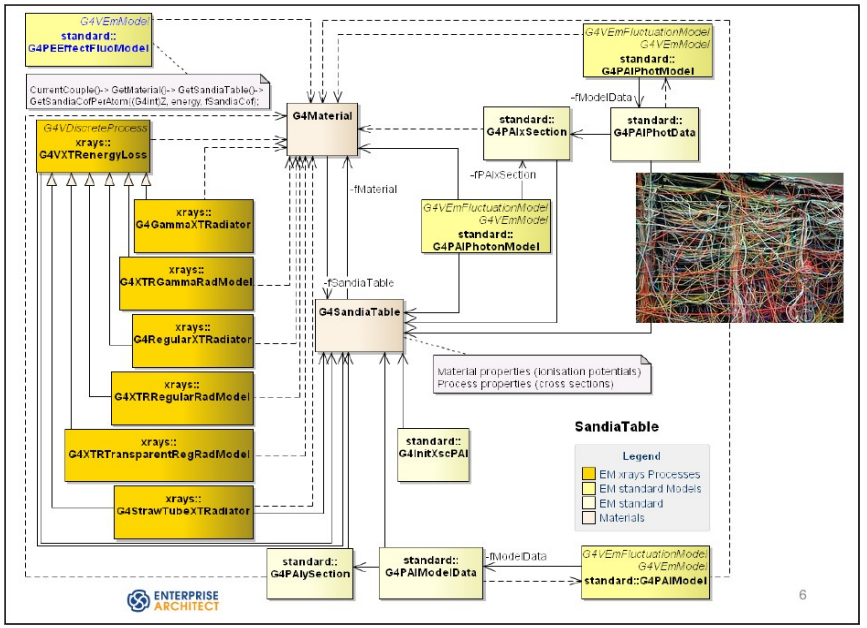

#348

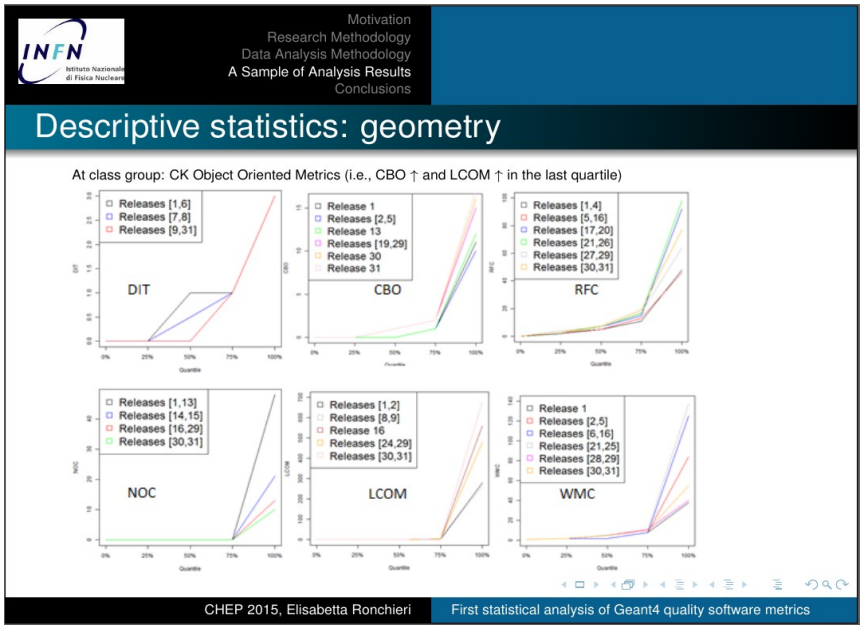

#485

Validation Frameworks

- Validate Puppet configuration in Jenkins-CI
- Validate ATLAS Shifter Assistant directives

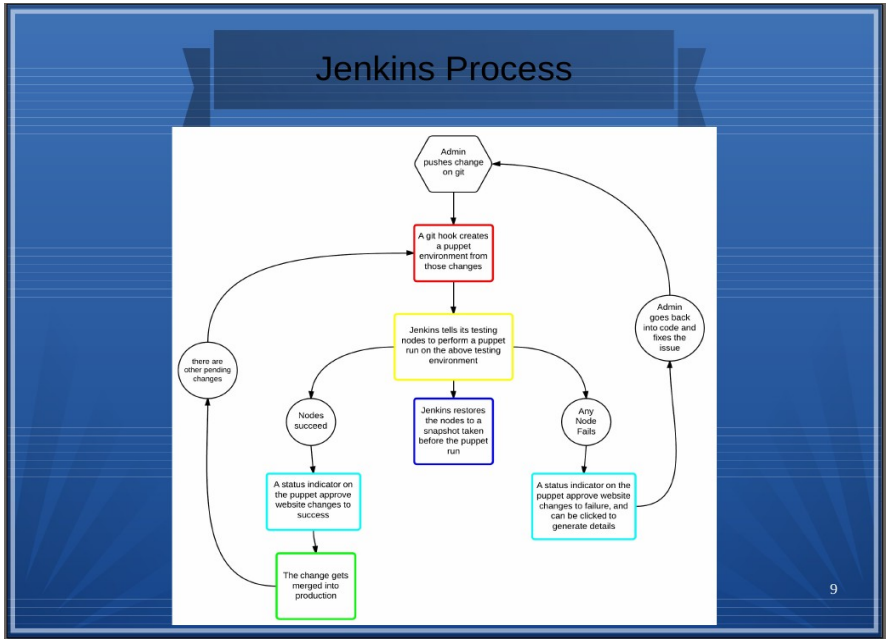

#28

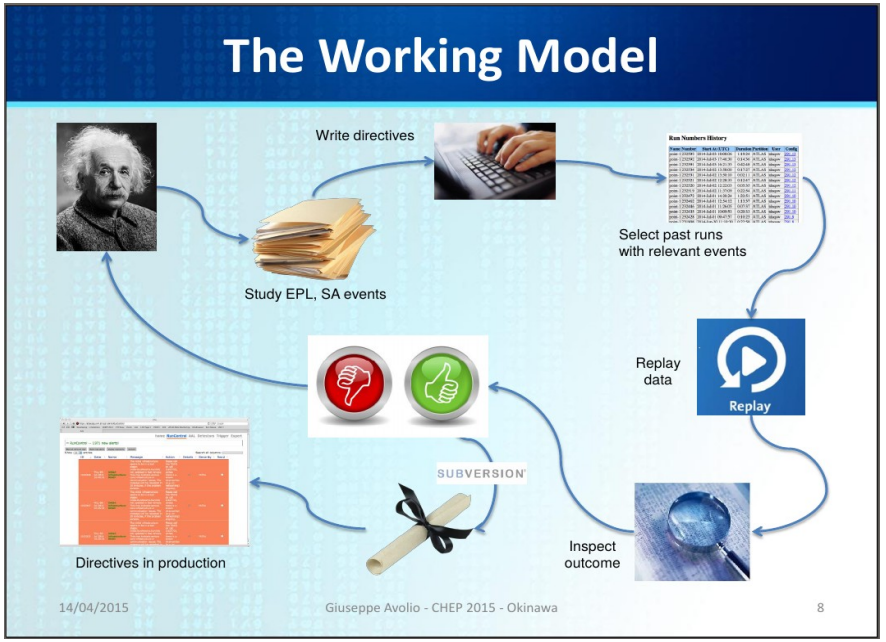

#36

Analysis Frameworks

- LHCb showed that it is possible to apply the SCRUM agile methodology to physics analysis

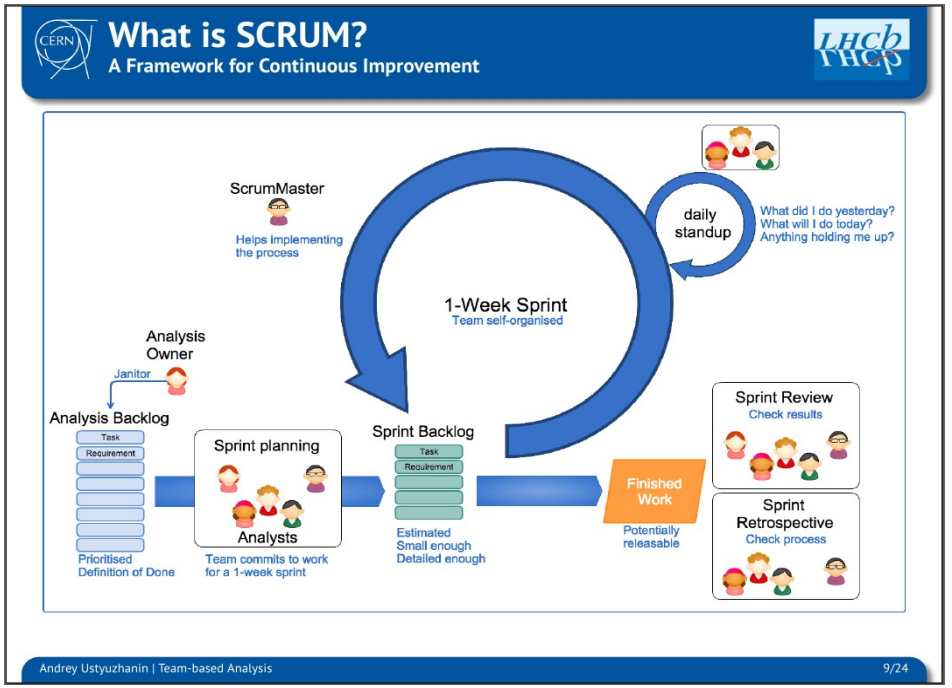

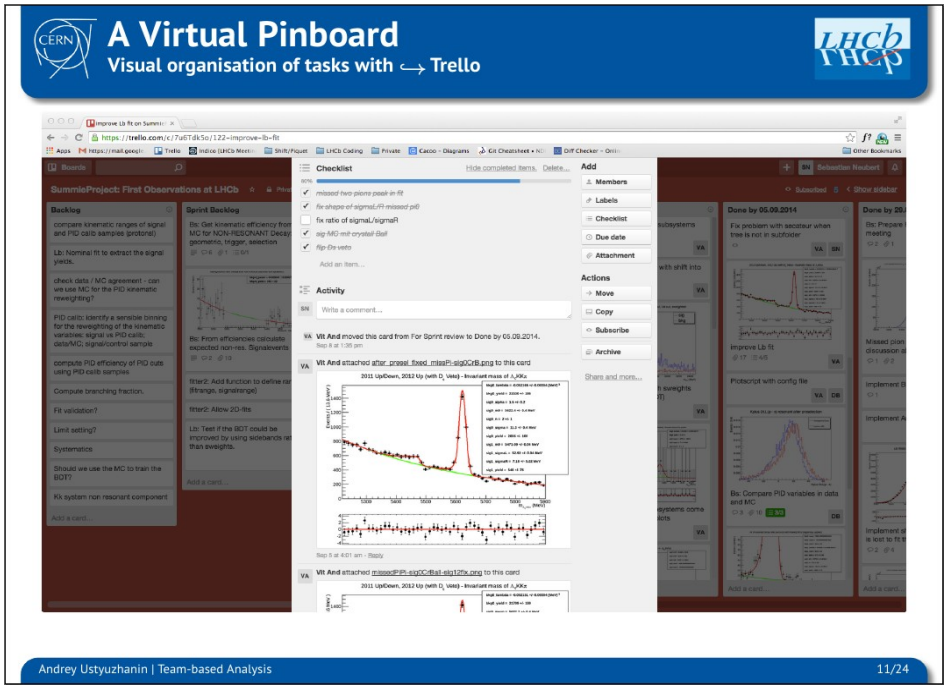

#527

Tools

Tools

- ATLAS presented the new incarnation of their Software Installation System

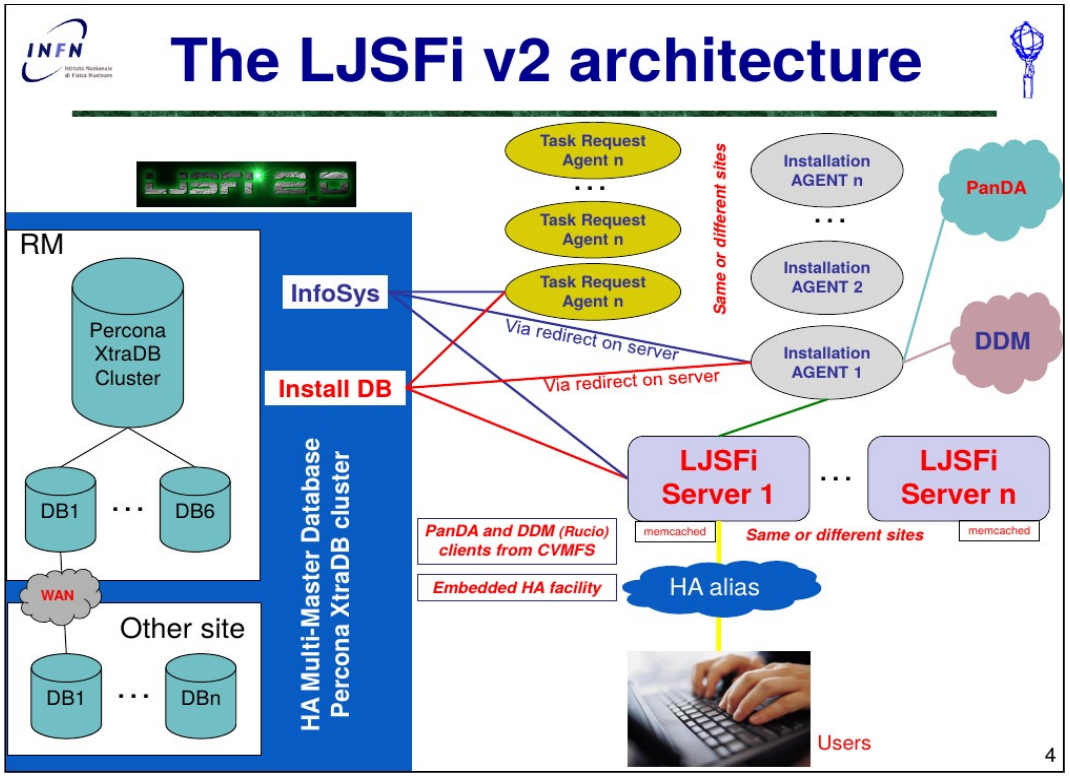

#204

Tools

- ATLAS decided to leverage on standard Message Queue technologies to synchronize Data Quality Monitoring tasks

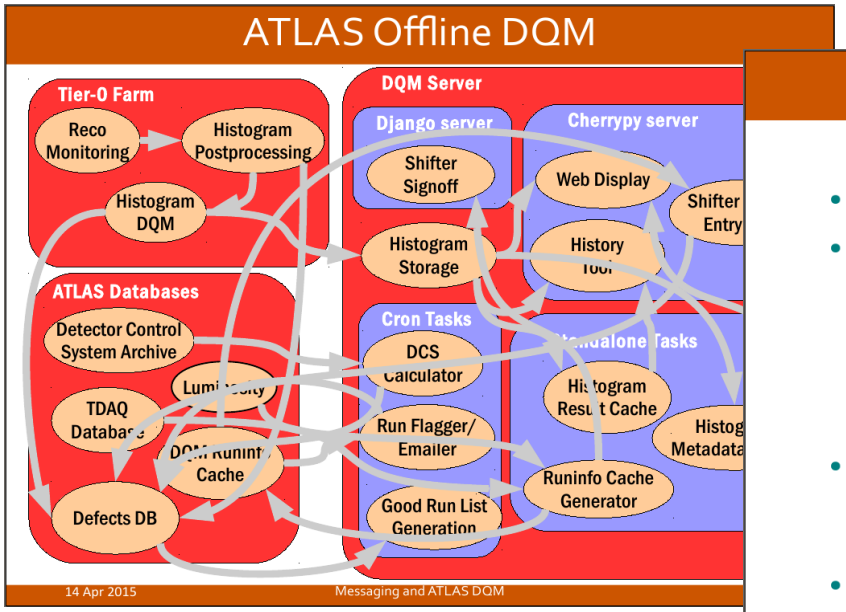

#176

Messaging Queues at CERN

ØMQ in bad location on complexity/benefit curve

Would need to support RabbitMQ ourselves

- Many options for messaging brokers
- CERN has standardized on **ActiveMQ**
  - Wrinkle: doesn't integrate at all with standard CERN auth mechanisms (must use app-specific passwords or certificates)
  - for security reasons, creating queues requires coordination with CERN IT
- Use **STOMP** protocol
  - near-universal availability of client libraries for different languages
- Piggybacking on servers set up for ATLAS Event Server project
  - message rate of few/minute is negligible perturbation

Tools

- From ATLAS we saw an interesting new way of developing web interfaces

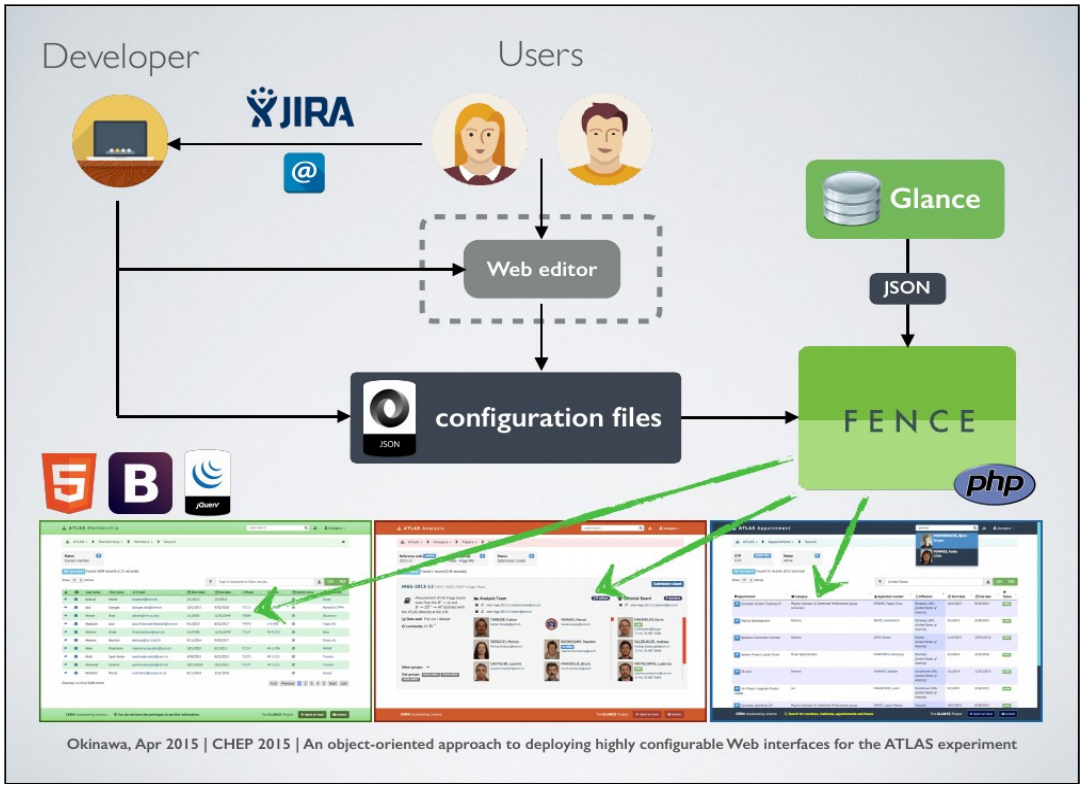

#167

Tools

- From SLAC a great contribution for debugging Grid jobs

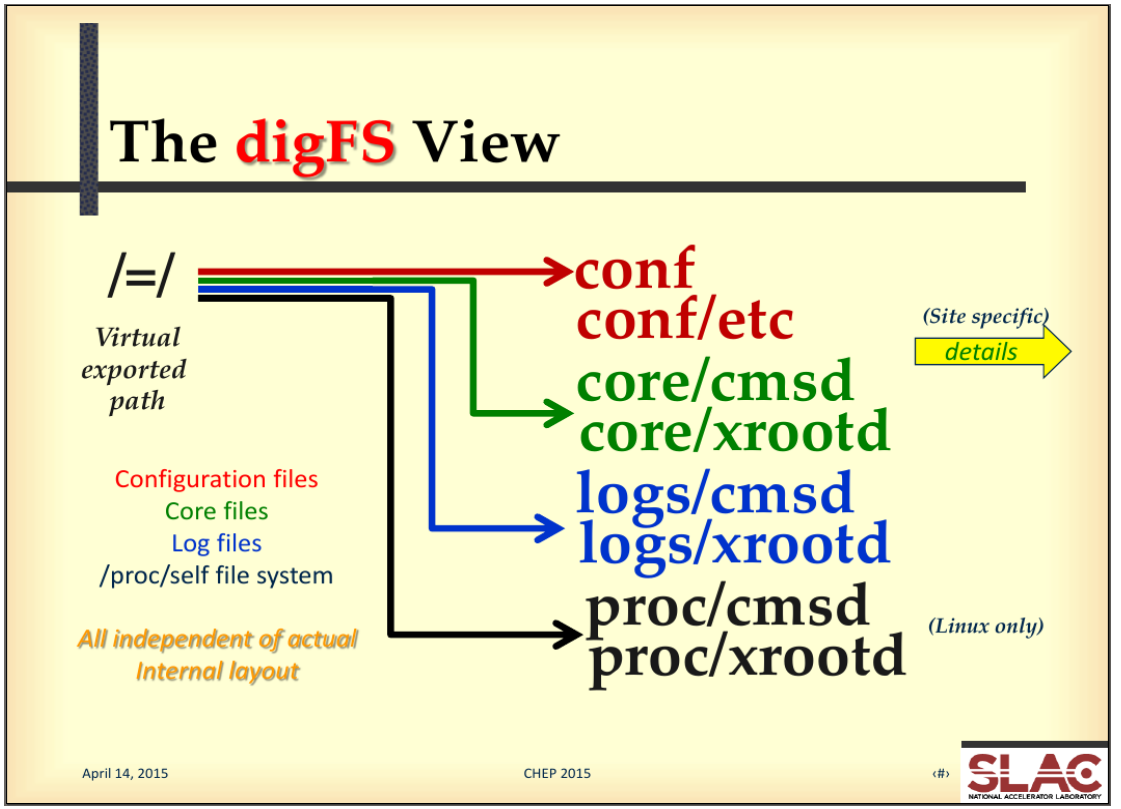

#310

Tools

- Deep insight on the features of IgProf
  - a feature rich profiler for HEP
  - including power monitoring for energy efficient code

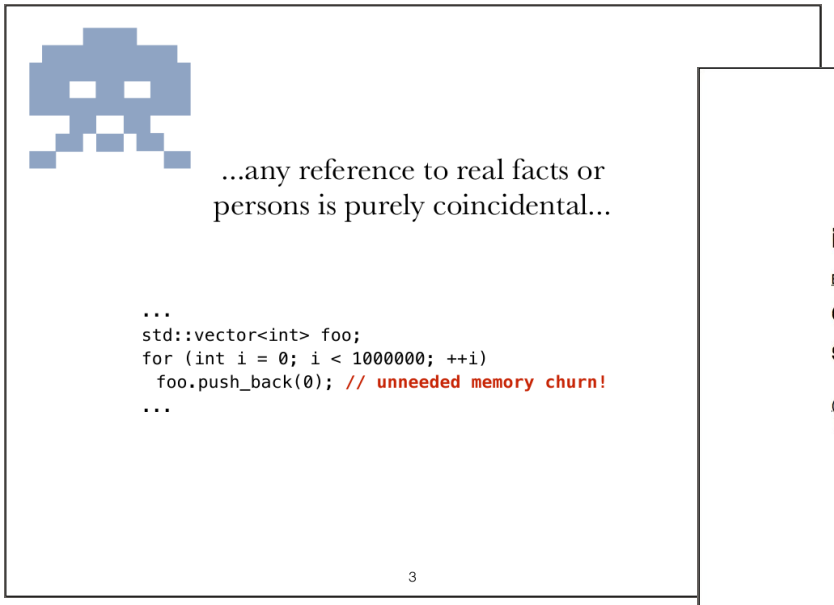

...any reference to real facts or persons is purely coincidental...

```
...
std::vector<int> foo;
for (int i = 0; i < 1000000; ++i)
 foo.push_back(0); // unneeded memory churn!
...
```

3

#478

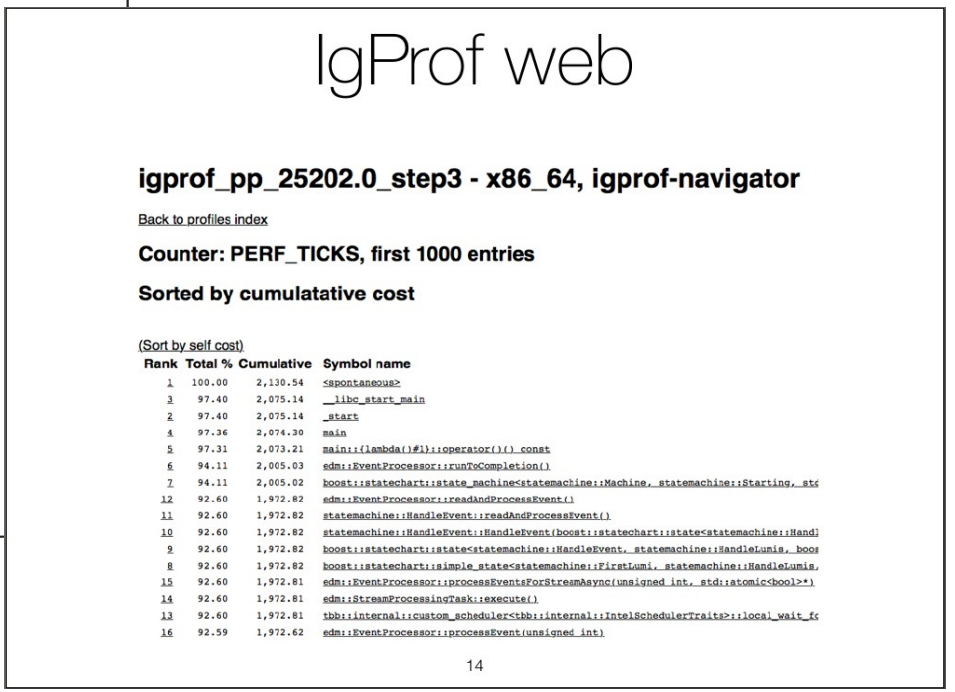

IgProf web

igprof_pp_25202.0_step3 - x86_64, igprof-navigator

Back to profiles index

Counter: PERF_TICKS, first 1000 entries

Sorted by cumulatative cost

(Sort by self cost)

Rank	Total %	Cumulative	Symbol name
1	100.00	2,130.54	<spontaneous>
3	97.40	2,075.14	__libc_start_main
2	97.40	2,075.14	_start
4	97.36	2,074.30	main
5	97.31	2,073.21	main::{lambda()#1}::operator()() const
6	94.11	2,005.03	edm::EventProcessor::runToCompletion()
7	94.11	2,005.02	boost::statechart::state_machine<statemachine::Machine, statemachine::Starting, std
12	92.60	1,972.82	edm::EventProcessor::readAndProcessEvent()
11	92.60	1,972.82	statemachine::HandleEvent::readAndProcessEvent()
10	92.60	1,972.82	statemachine::HandleEvent::HandleEvent(boost::statechart::state<statemachine::Handl
9	92.60	1,972.82	boost::statechart::state<statemachine::HandleEvent, statemachine::HandleLumis, boos
8	92.60	1,972.82	boost::statechart::simple_state<statemachine::FirstLumi, statemachine::HandleLumis,
15	92.60	1,972.81	edm::EventProcessor::processEventsForStreamAsync(unsigned int, std::atomic<bool>*)
14	92.60	1,972.81	edm::StreamProcessingTask::execute()
13	92.60	1,972.81	tbb::internal::custom_scheduler<tbb::internal::IntelSchedulerTraits>::local_wait_fo
16	92.59	1,972.62	edm::EventProcessor::processEvent(unsigned int)

14

Tools

- CERN developed an EOS-based Dropbox alternative: CERNBox

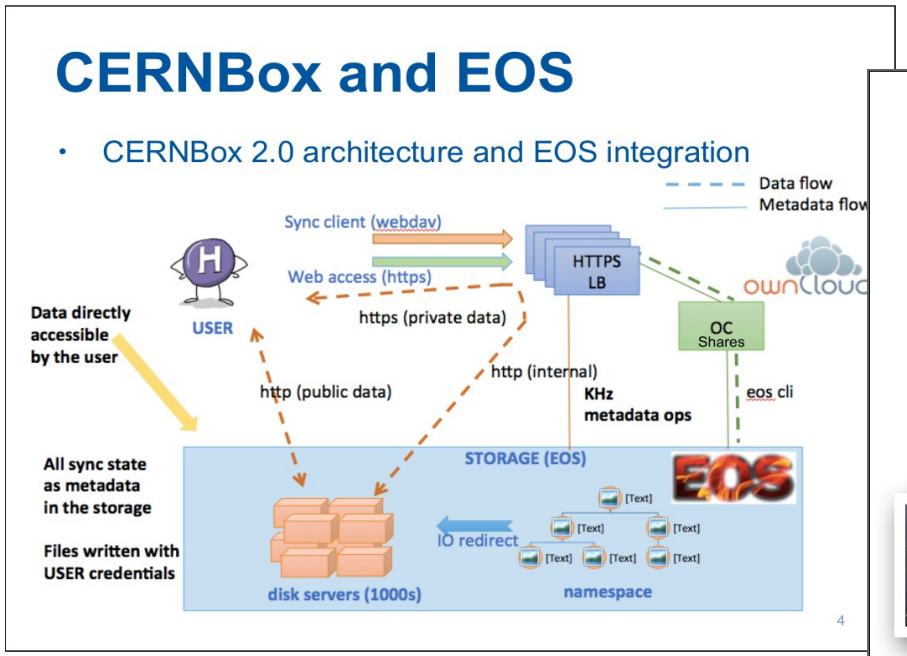

#327

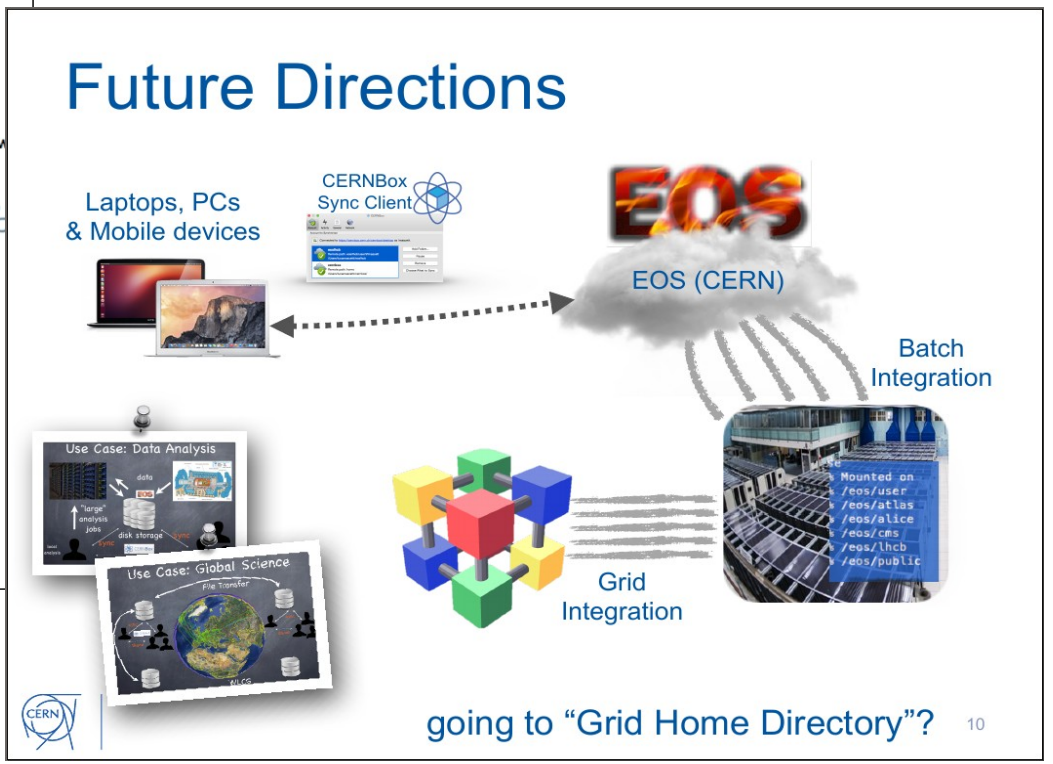

Conclusions

- We are working for a better (computing) world
  - Continuous efforts towards improvements
  - Sometimes “rewrite” is good
- Common solutions are beneficial to many
- Many are beneficial to common solutions

I want to thank the organizers
for the great work they have done
to make CHEP 2015 a success

Thank you!