


Track 3 summary

Distributed Computing



22nd International Conference on Computing in High Energy and Nuclear Physics, Hosted by SLAC and LBNL, Fall 2016

Latchezar Betev for T3
14 October 2016

Track 3

- Conveners
 - Tanya Levshina (FNAL), Weidong Li (IHEP), Latchezar Betev (CERN)
- Session chairs
 - Andreas Peters, Barthelemy von Haller, Maarten Litmaath, Miguel Pedreira, Xavier Espinal Curull
- Many thanks to the student helpers for the efficient support!



T3 stats

- 55 abstracts
- 38 oral, 17 posters
 - As usual, very difficult to make the oral/poster decision
- Room occupancy
 - **almost full** –
 - 45-55 participants




General notes on T3

- ***What it was not***
 - Presentation of high-level all-encompassing frameworks design
- ***Common themes:***
 - Adaptation and tuning of said ‘high-level...’ to targeted local / opportunistic / someone else’s (LHC@HOME) / Commercial Cloud resources + making the Grid work better
 - Introduction of more efficient and compact workflows
 - Inclusion of external information sources to boost performance

General notes (2)

- **Goals**

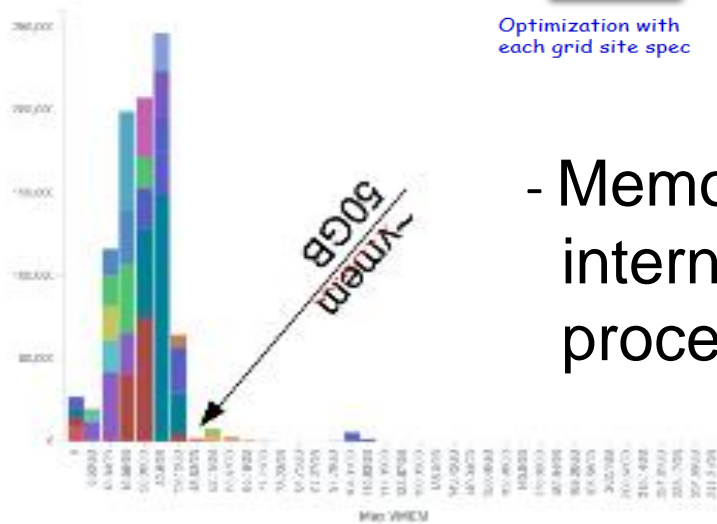
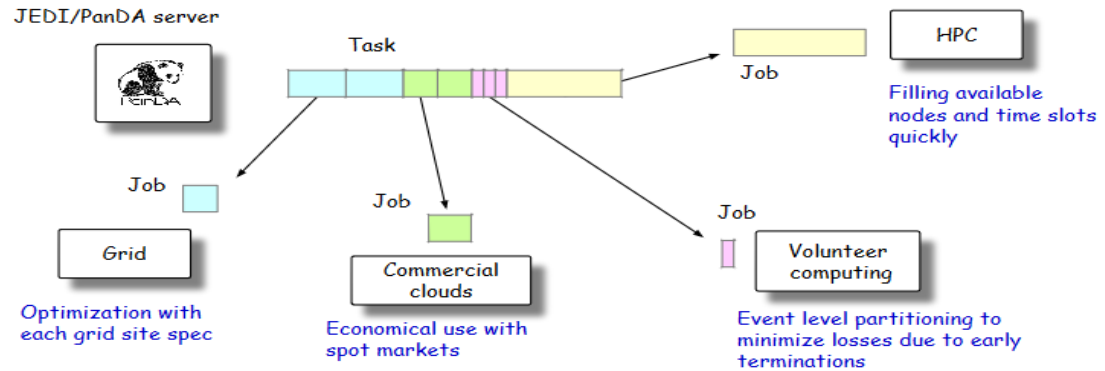
- Getting the remaining 10-15% efficiency out of the existing Grid
- Aggressive inclusion of new resource types
- Looking into all corners to get what is needed for LHC and non-LHC computing...
- ***... in a progressively tighter resource market***
- as more computing is needed to get the job done
 - ***Thanks a lot LHC*!***
- ***...and the upcoming upgrades***
- ***...and other projects***




* - a very welcome problem

Central and internal job brokering

- Matching to fit resources types, availability and limits
- HLT farms, HPC+Supercomputers, Cloud, @HOME (and the Grid)

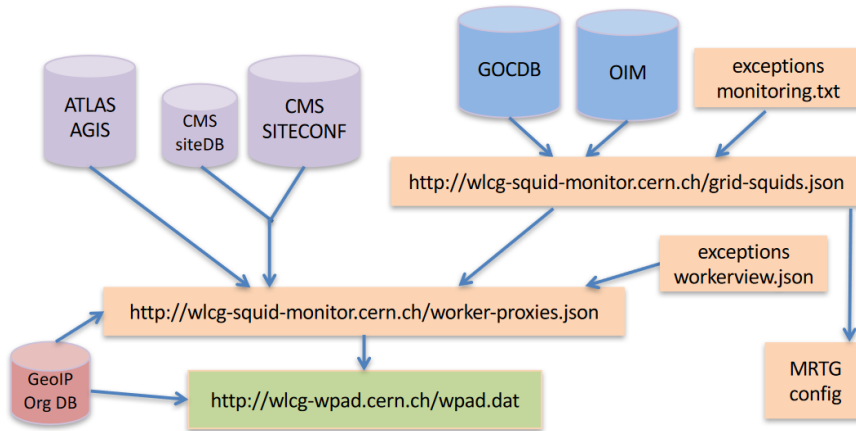


- Memory limits management and internal brokering for multicore processing

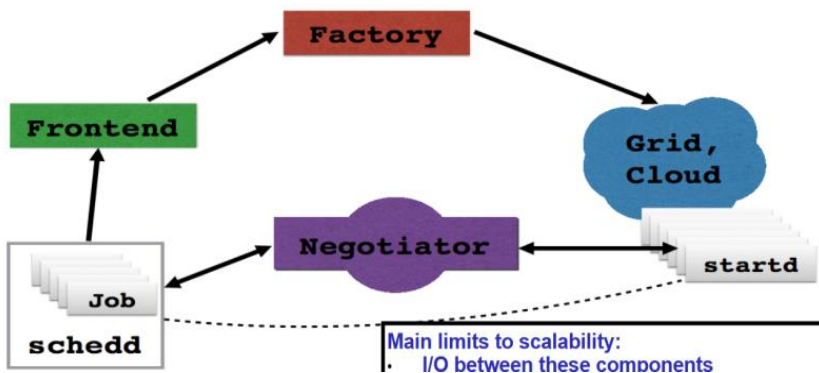


Fabric management and global pools

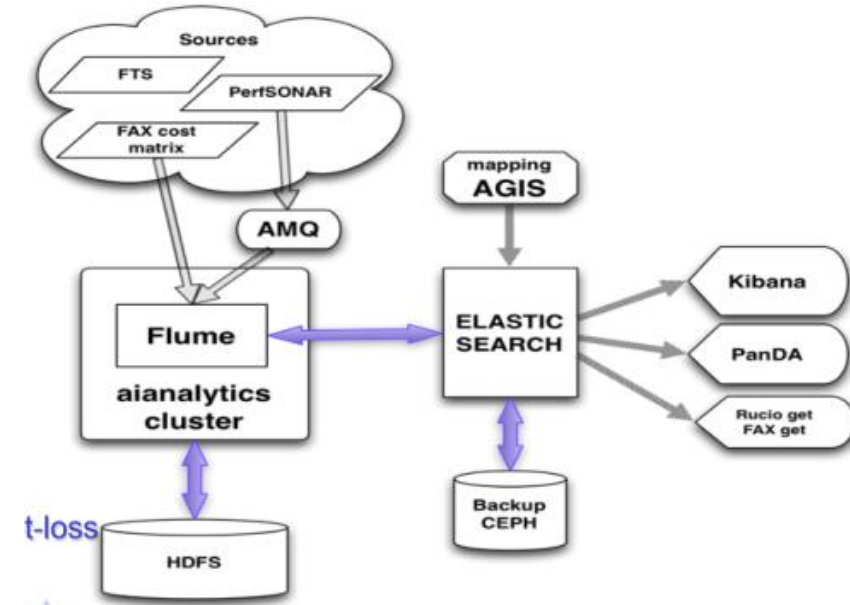
- WEB proxies discovery standardization



- ‘Global pools’ for all resources and tasks




- Network-aware brokering: treat as resource and adjust the workload





Non-LHC VO support models

- Medium- and small collaborations, no knowledge of distributed computing, no manpower to develop own





- Robust, common, modular set of tools – reuse existing set as much as possible



Cherenkov Telescope Array

User support


- Advanced tools for end-user analysis and small site operation as solution to many of the large-frame systems shortcomings
- Adapting the existing tools to support better central analysis activities – output formats and job priorities



Fabric evolution in response to evolving workflow

- Analysis of fabric parameters, feedback to the experiments to improve the upstream submission tools

$$U_{prio} = \frac{U_{share}}{\varepsilon CPT + \alpha WCT + \beta(1 + SLOTS) + \gamma ADJUST}$$




- Coordinated data caches to improve analysis efficiency (between batch and data servers)



Bye Tiers!

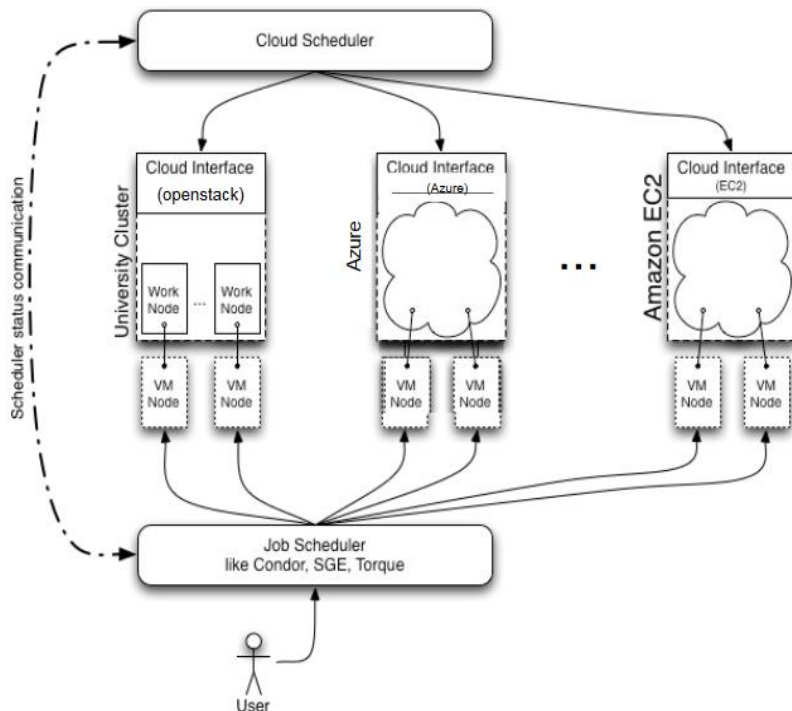
- Tiers not mentioned a lot, but the dissolution is almost complete
- Excellent site and network performance are to blame
- ...and that's all I have to say about that





To Clouds!

- Clouds from single site, several geographically close centres, Cloud of clouds and heterogeneous resources (HPCs)
- Common theme – enabled for multiple user communities

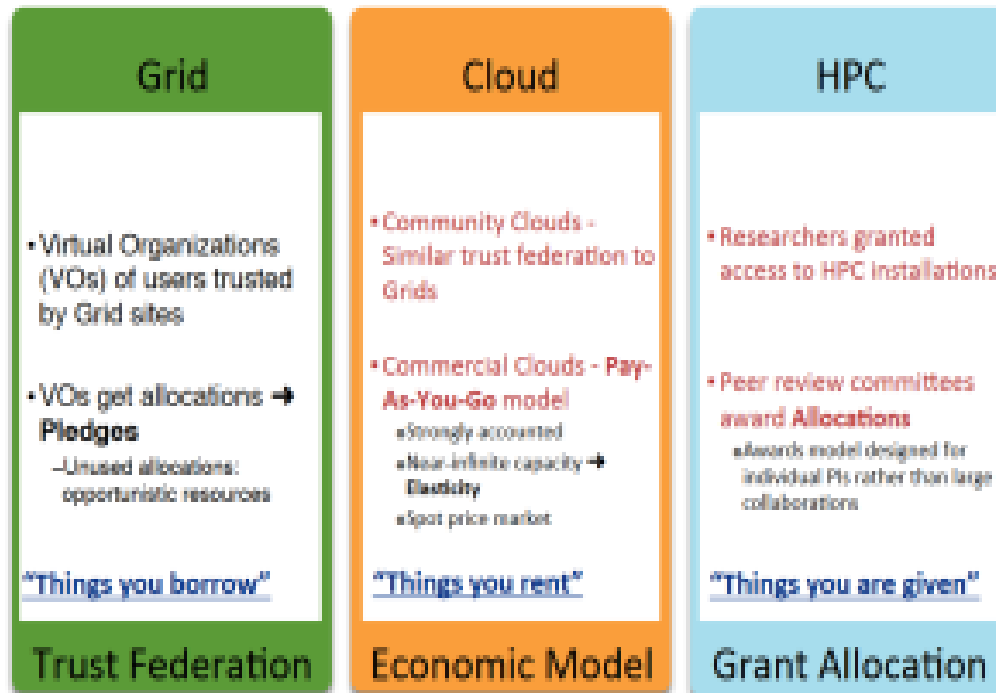


- CERNVM, Containers, Mesos, Marathon, Calico, Docker, HTCondor - plenty of enabling technology
- New context-aware cloud scheduling
- Cloudy (bright) future



Cost of (commercial) Clouds

- From theory to practice - find the right bidding strategy for commercial clouds (AWS)
- Apply in practice – get many cores for full simulation workflow

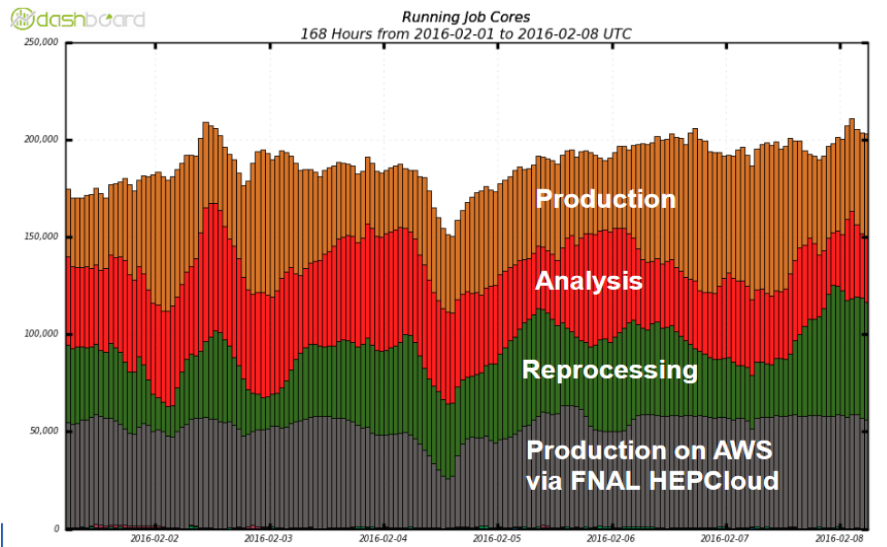


- On premise resources still cheaper, but Cloud is a viable strategy if properly approached – need to do smart bidding
- Commercial Cloud prices are decreasing



More clouds

- Significant bursting capabilities, on par with large T1s (not really news)



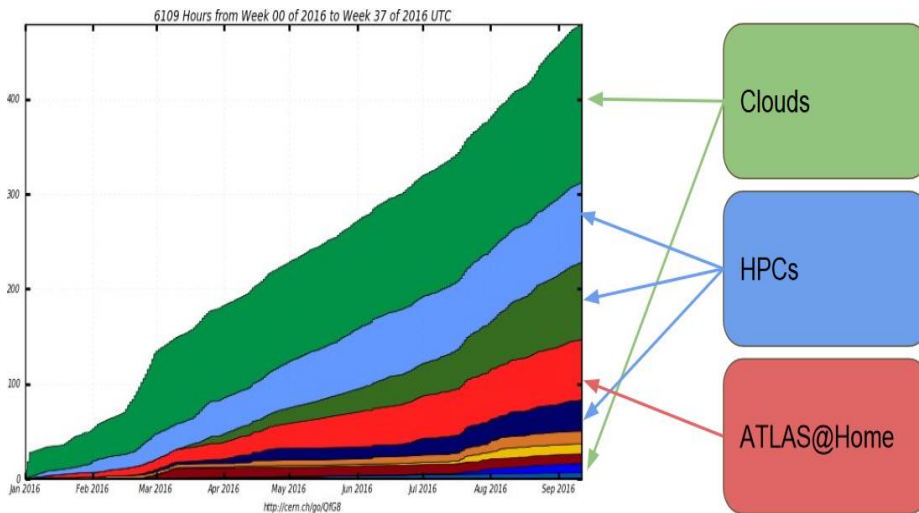
- Acquiring the resources is becoming simpler and faster

- Several large-scale projects for commercial cloud procurement – to be used directly in production



Supercomputers and volunteering

- Sparse network (if any), may require specific SW porting, missing software delivery, non-x86 resources
- Access for non-LHC VOs through same tools
- Challenges abound, but developers are optimistic



- Getting resources 'for free'
- Very willing participants, great PR
- Up to 1-2% of total CPU resources
- Steady growth

From past



'The End of the Day' by Harold Septimus Powers

...to today





Thanks to all presenters for the interesting and engaging talks and posters (and largely sticking to the time limit)

Many thanks to our LBNL and SLAC hosts for the great organization and support throughout the T3 sessions