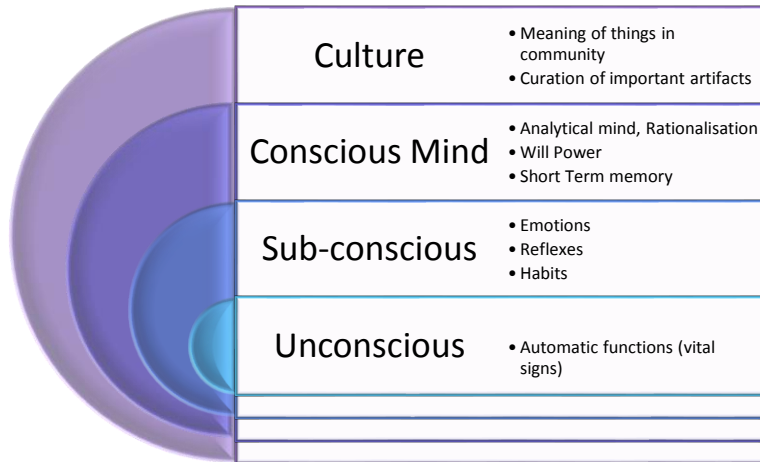


The story so far ...

Jeremy Coles

3rd November 2016

SKA background

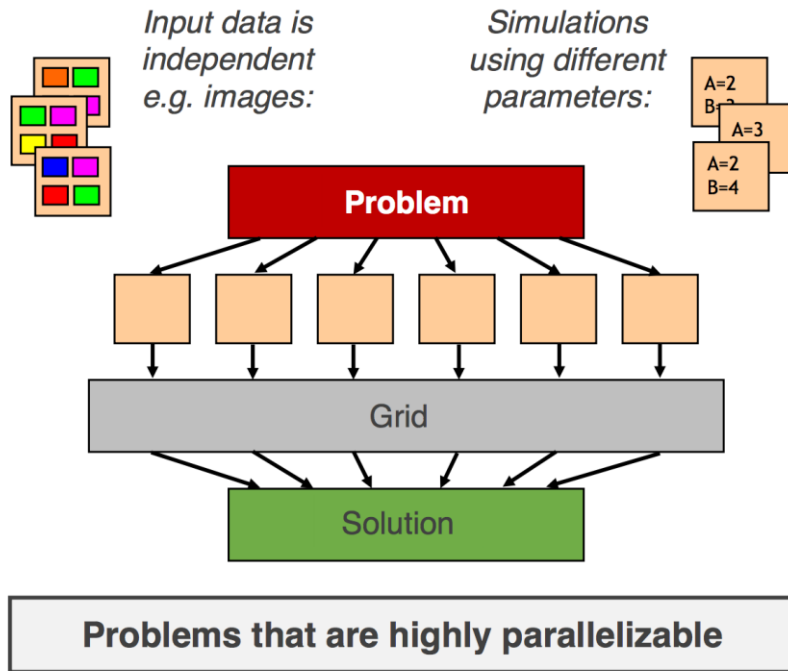


- Develop a concept and design for a distributed, federated European Science Data Centre (ESDC) to support the astronomical community in achieving the scientific goals of the SKA


To Note:

- Want to hide the complexities.
 - AIMS: Better understand expertise in UK; routes to collaboration (who?); understand mutual benefits (i.e. funding!); “Where’s the dictionary?”. What is the national view.
 - AENEAS:
 - ***Inventory of SKA science cases and post-SDP computing and data storage requirements***
 - ***Evaluation of existing HPC, cloud and distributed computing technologies***
 - ***Design and costing for distributed ESDC computing architecture***
 - ***Requirements for interfaces to SKA Science Archives & Other Repositories***
- Next -> Identify contributors; check coverage; what can be in-kind... early start options.

GridPP infrastructure & approaches



- What is GridPP
- Key concepts of a Grid
- The 10% offer!
- Needed: Software deployment + Job & data management + User Interface.

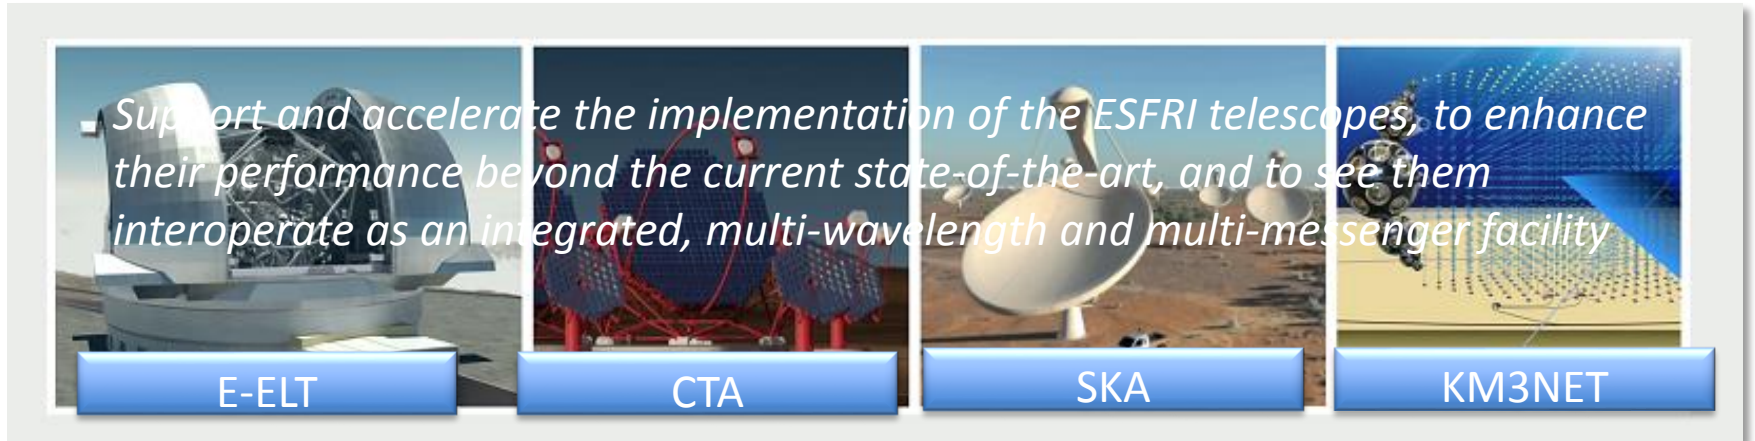


- Links to use-cases and documentation

Q&A:

- 1PB split over two sites 50:50. Note CERN-Wigner link on failover several times recently.
- GridPP will be a peer in UK-T0

ASTERICS & typical analysis



Analysis	Requires	Note
Image based	HTC or HPDA	Like LSST
Source finding	HTC	
... if full Bayesian approach	HPC	N -> N+1 comparisons
Stacking	HPDA	Uses different sources
Extraction dynamics	HTC then HPDA	Machine learning
Visibility based	Extreme HPDA/HPC	Pipeline implemented

Q&A: Signed up to OpenData but need to build telescope first!

GridPP funded for HEP. Infrastructure Influences.

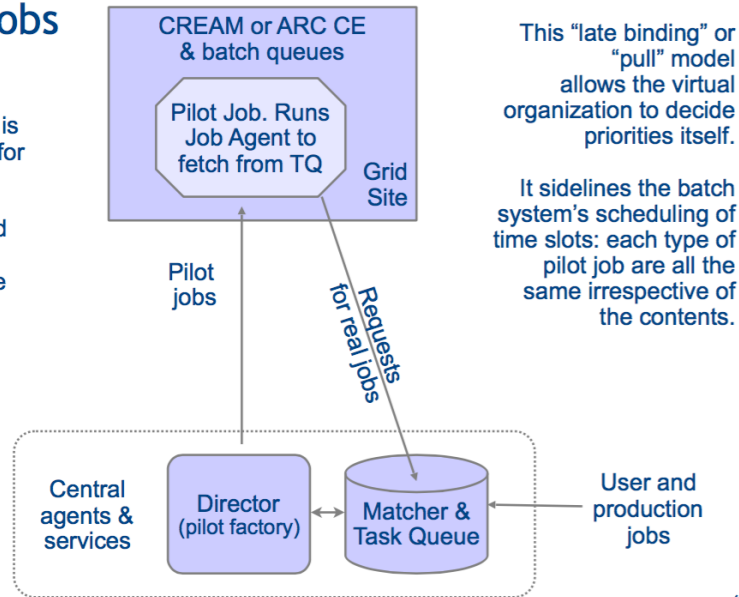
Challenge for commercial cloud:
Cost & capability⁴

GridPP DIRAC, VMs, OpenStack, HTC, HPC

The Grid with Pilot Jobs

The Grid + pilot jobs is the dominant model for running HEP jobs.

Well established, and gives access to resources around the world.



This "late binding" or "pull" model allows the virtual organization to decide priorities itself.


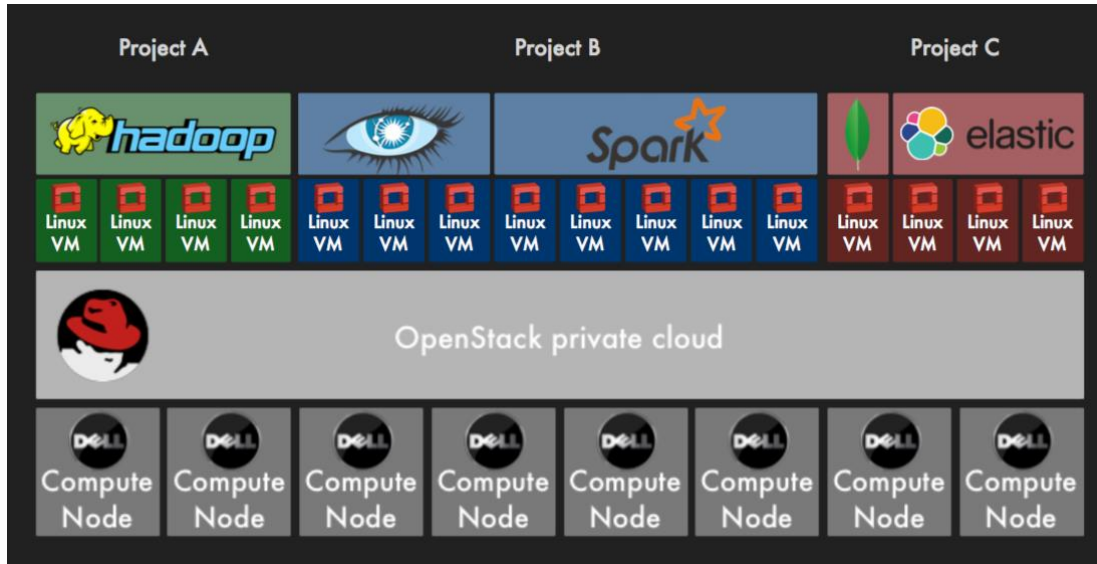
It sidelines the batch system's scheduling of time slots: each type of pilot job are all the same irrespective of the contents.

6

- HEP HTC model. How it evolved over 15 years.
- Late binding and pilot jobs
- HEP HTC no need for coherence
- HPC growing in HEP. Mixed workflows.
- Could add HPC to DIRAC production requests.
- Q: How to load balance for aggregation. A: Chained.

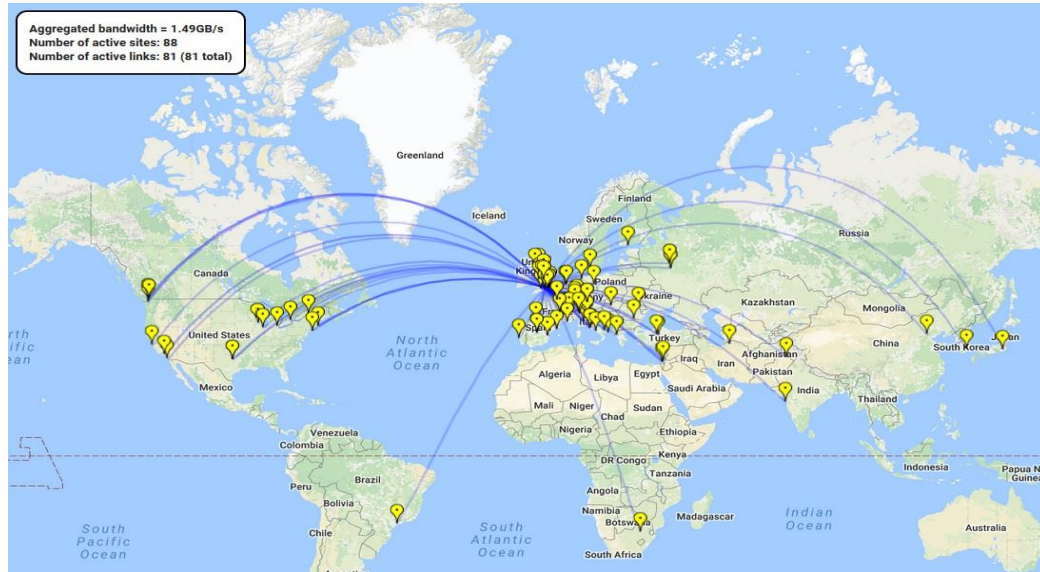
- Could more complex data dependencies be expressed at job level (dependency graph)
- Bridging – populate file catalogue directly?
- What about working with limited local resources? Pilots and matching.

OpenStack - HPC



- The overhead of virtualisation – remediation strategies
- Emerging capabilities for HPC
- Remaining gaps and options

Data transfer approaches

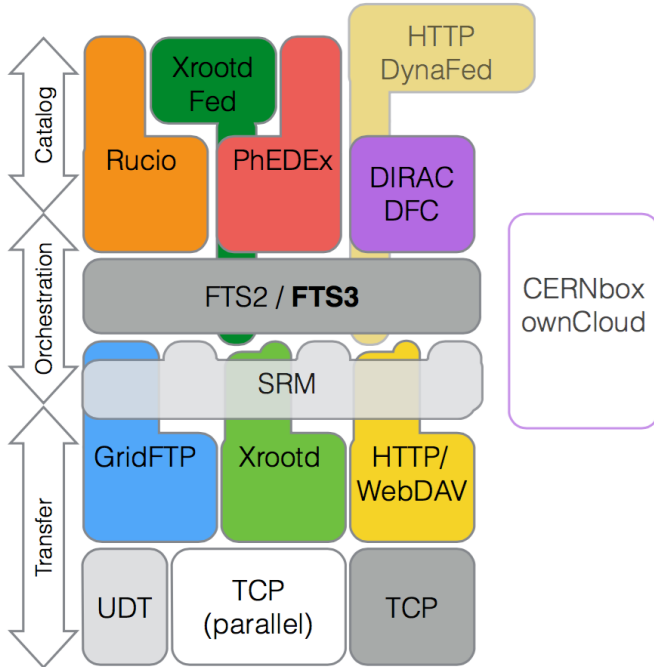


- Multi-TB/day routine. Reasonable intercontinental rates
- Use File Transfer Service (FTS)
- Helped DiRAC. Watch those file sizes.

To Note:

- Have a minimum & maximum file size specified
- Be able to handle transfer interrupts
- Consider carefully the storage middleware & transfer protocols
- Checksum validation important
- Now looking at federated redirection services

Data management



To Note:

- Abstraction layers can be enemies of performance
- Think about deletion times!

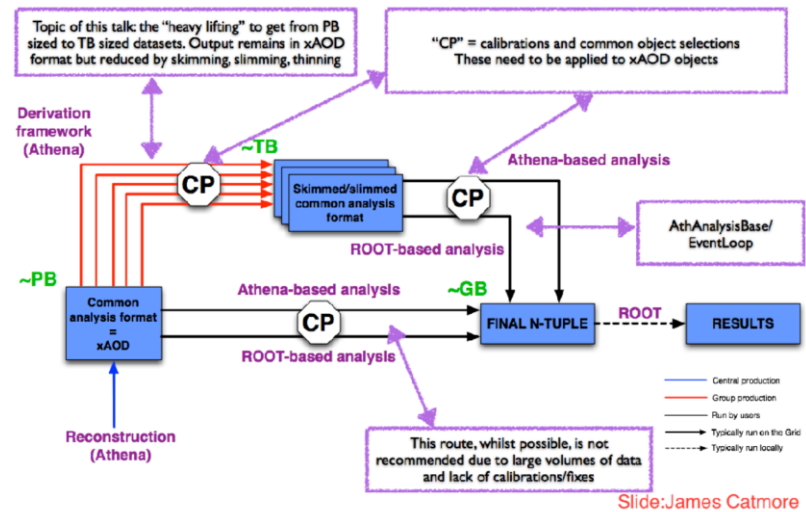
Considerations:

- What are the objectives? Testing for the purposes of planning, R&D. To build-up experience with tools/techniques. Help with pre-construction/commissioning tasks?
- AENEAS is not doing prototyping – it is forbidden!
- WLCG got started in some form around 2003 with challenges. An infrastructure wide on in 2006... but years ahead of actual target production dates.
- Many things developed were wrong/changed/evolved, but that step had to be done.
- How much pain can GridPP absorb?
- There are things that it is difficult to write down. To learn one must try things out.
- Early on WLCG required heroic efforts.
- Can we run workflows on the same infrastructure?
- Are there network QoS criteria? Not really for HEP.
- Evolution of resources and approach partly driven by resource limitations... also requirements, finance...

Why did the LFC get replaced!?

Data analytics - HEP




The Run-II analysis model for ATLAS 4



- Overview of data reduction operations; augmentation; derivation datasets.
- RUCIO for data management

Analytics:


- Logs collection, storage & analysis
- Ops analysis, insight & steering
- Physics analysis



NOTES:

- Don't forget about deletion!
- Data lifetime considerations
- Pledged vs used resources – almost x2
- Sharing resource – issue is the effort not the resources!

Data analytics - Spark



- Why Spark?
- Spark use cases
- Performance depends on problem scale and level of parallelism

- Linear algebra case study – classical MPI-based more efficient but development, interface, ecosystem considerations balance.
- Look at commodity cluster vs HPC platform
- Science drivers: Climate science; Nuclear physics & Mass Spectrometry
- Study Spark vs C+MPI scaling.
- Mitigate drawbacks of GFS.

Coming up....

- Sketelescope.eu VO
- Manchester test case & experiences
- AENEAS objectives & what is needed
- SRCCG?
- Network considerations

- Planning and next steps
- SRC (not SKA!)