A Computing Model for SuperBelle

This is an idea for discussion only!

- Large on-site CPU at KEK
- Analysis model based on GAUDI
- Employ Cloud Computing for MC Generation

Martin Sevior
University of Melbourne

Substantial CPU power outside KEK from New and Existing Collaborators

Why not use it?

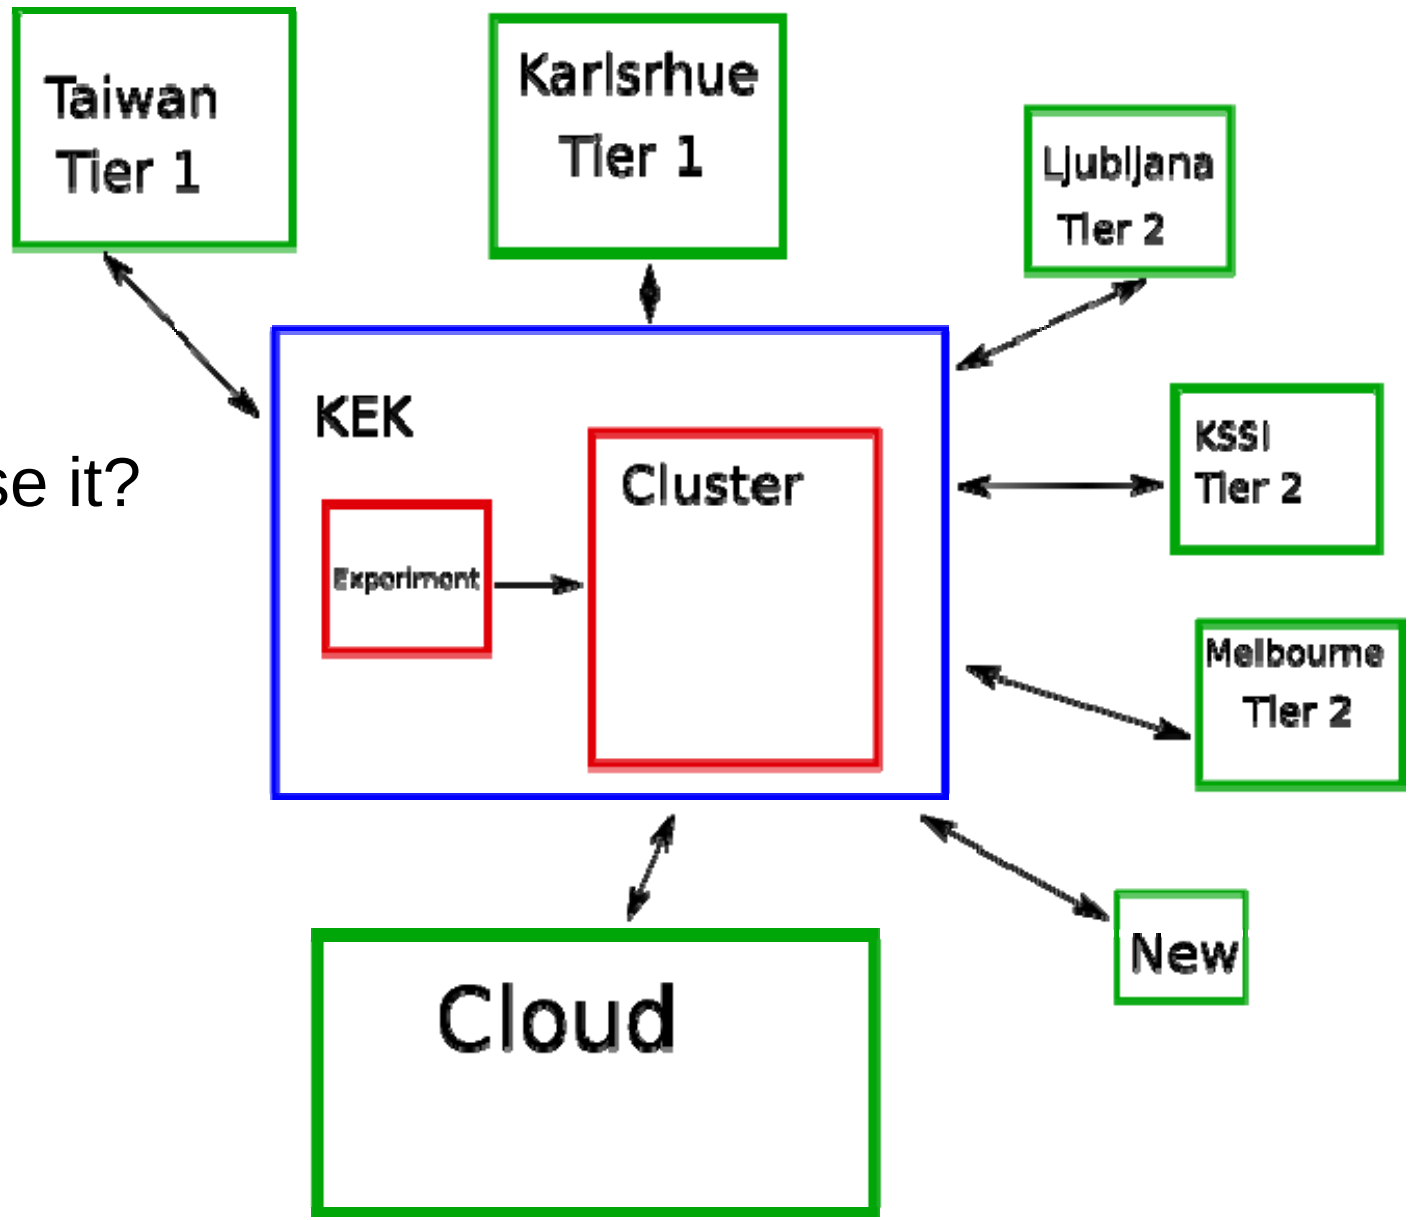

We will require a large on-site cluster

- Needed for data acquisition system
- Also use this for reconstruction and first pass analysis

Current KEKB Computer System

Data size ~ 1 ab^{-1}

New KEK Computer System has 4000 CPU cores

Storage ~ 2 PetaBytes

SuperBelle Requirements

Initial rate of $2x10^{35}$ cm^2sec^{-1}=> 4 ab^{-1} /year

Design rate of $8x10^{35}$ cm^2sec^{-1}=> 16 ab^{-1} /year

CPU Estimate 10 – 80 times current depending on reprocessing rate

So $4x10^4$ – $3.4x10^5$ CPU cores

Storage 10 PB in 2013, rising to 40 PB/year after 2016

Spreadsheet

CPU (8 -32)x10^4 cpus over 5 years ~ 500$ per core (2008)

Storage costs over 5 years (10 - 140) PB (Disk, no tape) $800/TB (2008)

Electricity ~ 100 W/CPU (2008), Price $0.2/KWhr (2008)

	A	B	C	D	E	F	G	H	I	J	K	L
1			PB Purchase	Price $								
2	2013	2	10	8000000								
3	2014	4	20	16000000								
4	2015	6	30	24000000								
5	2016	8	40	32000000								
6	2017	8	40	32000000								
7												
8	Total		140	112000000								
9												
10												
11			CPU purchase	Price $		Electricty KWHr	Costs ($0.2/KWHr)			Total 2008 $		Deflated 18 month/double
12	2013	2	80000	40000000		64000000	12800000			60800000		11111807.3637774
13	2014	4	80000	40000000		128000000	25600000			81600000		11500000
14	2015	6	80000	40000000		192000000	38400000			102400000		10079368.399159
15	2016	8	80000	40000000		256000000	51200000			123200000		8135430.39133702
16	2017	8	0	0		256000000	51200000			83200000		6849604.2078728
17												
18	Total		320000	320000000								
19										Total		47676210.3621462
20												

Price in 2008 of SuperBelle Cluster

(At best 100% uncertainty!)

CPU (8 -32)x10^4 cpus over 5 years ~ 500$ per core => $40 Million/Year

Storage costs over 5 years (10 - 140) PB (Disk, no tape) $800/TB => $(8 - 32) Million/Year

Electricity ~ 100 W/CPU (64 – 256) TWHr=> $(13 - 52) Million/year

Rough Estimate over 5 years $(61, 82,102,123,83) Million/Year

Moores Law – Double Performance every 18 months

Rough Estimate over 5 years $(11,12,10,8,7) Million/Year

Total Cost over 5 years ~ $50 Million

This is a defensible solution but needs more study...

Use of GRID Computing and Cloud Computing could substantially reduce the size of the KEK Cluster!

ATLAS analysis model has substantial strengths beyond current Panther Banks or Root-only based BASF

Rather than have a single in-memory framework like BASF, ATLAS employs the GAUDI-based collection of *services* which communicate via out-of-process mechanisms

This provides data persistency through an interface to a File-Catalog service.

Analysis Models – LHC (ATLAS)

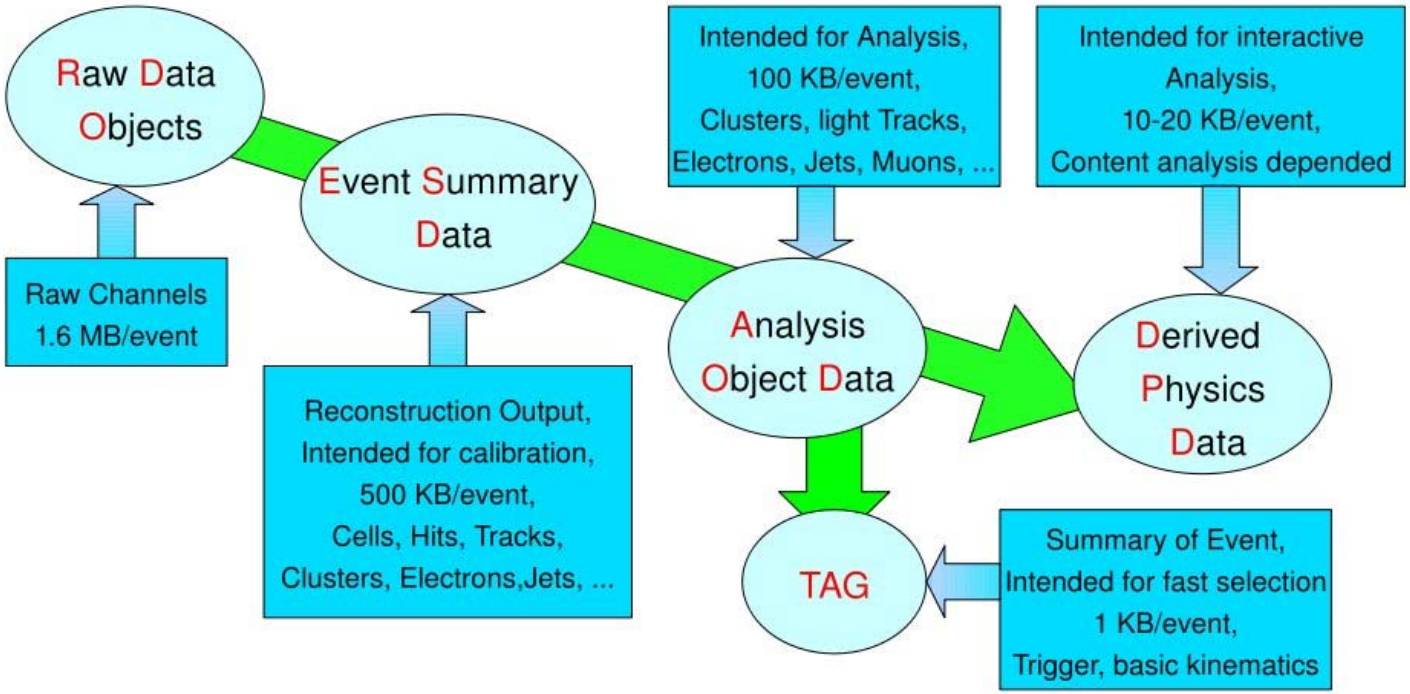

Johannes Elmsheuser (LMU M¨unchen) (ATLAS Users Analysis)

The ATLAS analysis model does not require data to be stored centrally.

Furthermore the ATLAS Athena analysis framework makes it possible to recover original data from derived data.

The GRID now basically works

Our Graduate students routinely use it to do Analysis.

School of Physics

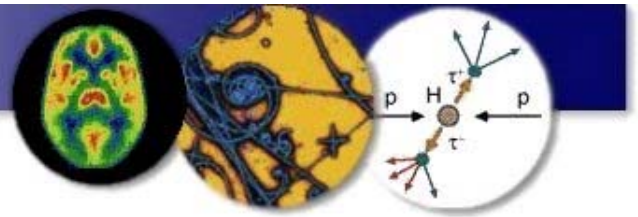

THE UNIVERSITY OF MELBOURNE

- Public Home
- Intranet Home
- SysAdmin Home

- Safety Info
- Travel Info
- HEPStay Database
- Group Meetings
- Mailing Lists
- Video Conferencing
- Computing HowTo's
- Computer Support
- Internal Documents
- Demonstrating
- Scholarships / Grants

ATLAS

- - Essentials
- - Notes & Tips
- - T2 Procurement Notes
- - T2 Cluster Notes
- - T2 Statistics

Belle

- - Presentations
- - Notes & Tips

Log Out

You are here: Physics > EPPprivate Web > WebMaintenance > WebMenu > ComputerHelp > AtlasNotes > PAthenaOnT2

r2 - 21 Oct 2008 - 01:47:07 - Main.TomFifield

Edit | WYSIWYG | Attach | Printable | Settings+Help

Using PAthena on the Tier 2

The purpose of this page is to provide a quick-start guide on using Pathena for Distributed Analysis, based on Australia-ATLAS in Melbourne. Much of the content has come from hard work by Nadia and DAOnPanda at CERN twiki. It's important to note that using Pathena will not allow you to run jobs on the Tier 2 with priority (like ganga/glite-wms-jobsubmit will) - but in general, jobs running at Taiwan have been reasonable in the past.

Getting started

1. Setup your favourite version of Athena as you normally would
2. cd to your test area
3. Fetch the panda tools

```
ui $ cmt co PhysicsAnalysis/DistributedAnalysis/PandaTools
ui $ cd PhysicsAnalysis/DistributedAnalysis/PandaTools/cmt
ui $ source setup.sh
ui $ make
```

Running PAthena

With pathena setup, change back to your test directory - you can check pathena's options out by using: pathena --help - let's run through a few of them:

```
--split=SPLIT         Number of sub-jobs to which a job is split
--nFilesPerJob=NFILESPERJOB
                      Number of files on which each sub-job runs
--nEventsPerJob=NEVENTSPERJOB
```

Adopting a GAUDI-based analysis system provides a natural means moving to a GRID environment.

Adopting GRID enables users to do data analysis outside of KEK.

Also keeps track of data collections with a natural means of providing metadata descriptors.

This metadata could be as detailed as constants file used in the reconstruction or the specific parameters and Physics parameters of MC generated data.

There is a significant overhead in moving to a GAUDI framework

Out-of-process communication creates a significant increase in complexity.

On the bright side most of the hard work has already been done and we can re-use other peoples hard work.

Cloud Computing

Commercial internet companies like Google and Amazon have computing facilities orders of magnitude larger than HEP.

They have established a Business based on CPU power on demand, one could imagine that they could provide the compute and storage we need at a lower cost than dedicated facilities.

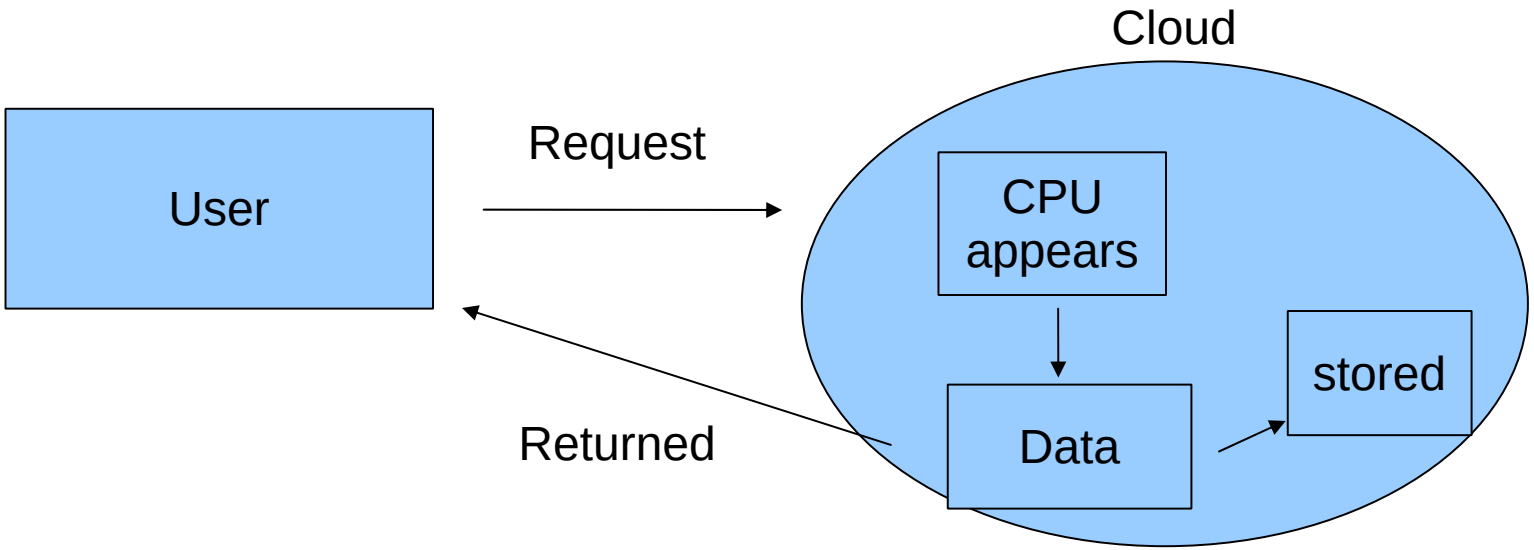

Resources are deployed as needed.

Pay as you go.

Essentially infinite CPU power

Initial overhead in creating a virtual machine instance configured as needed

After this however it is trivial to generate as many instances as needed

Cloud Computing http://aws.amazon.com/ec2/

Once the work to create a Amazon Machine Interface (AMI) is done, it is trivial to Deploy as many as needed.

Amazon EC2 charges 20 cents per 4 core CPU-hour

Assuming 10 events/minute for each core => 12,000 MC events per dollar

Then 10^9 events costs ~ $80,000

150×10^9 MC events (entire MC sample needed by SuperBelle) costs $12.5 Million today!

Let's do an experiment to benchmark it!

On the other hand 1 PB of storage costs $1.2 Million per Year, not competitive right now

Combine the benefits of Local large cluster, Data persistency, GRID and Cloud Computing

Martin Sevior
University of Melbourne