


Novosibirsk
Новосибирск
Novosibirsk Oblast
Russia
Mostly cloudy -19°C
7:57 AM



Photos



Quick facts

Novosibirsk is a city in Siberia, southern Russia, bisected by the Ob River. The Trans-Siberian Railway fueled much of the city's 19th-century growth, symbolized by the Novosibirsk Rail Bridge, which still stands today. In the city center is the 19th-

Accelerator focused Computing in Australia

Martin Sevier (University of Melbourne)

(Thanks to Prof. Ulrik Egede (Monash University))

Martin Sevier, AFAD, Novosibirsk, Russia, March 2021



Australian Accelerator based Computing

Australia contributes computing resources as follows:


- ATLAS (CERN) Tier-2 site
 - Contributes to ATLAS grid: ~1200 CPU cores ~ 1.2 PB SE
- Belle II (KEK) Tier-2 site
 - Contributes to Belle II grid: ~300 CPU cores, 20 TB SE
 - Also 250 – 900 cpu's via Melbourne Research Cloud
- ATLAS and Belle II Tier-3
 - Provided via University of Melbourne HPC service.
 - CVMFS access to Research Software
 - 60TB/100TB online/offline storage (Spectrum Scale for online)
 - 5500 CPU/300 GPU worker nodes
- LHCb (CERN)
 - DIRAC Cloud site: 100 – 1000 CPU cores
- COMET (KEK) + LHCb
 - MC Simulations on Monash University M3 HPC service.
 - CPU/GPU cluster provides throughput $\sim 10^3$ jobs per day

Migration of ATLAS and Belle II

- Australian ATLAS and Belle II sites had been maintained by the ARC Center of Excellence for Particle Physics at the Terascale (CoEPP)
- Since CoEPP has concluded, these are no longer viable
- Migrate to University of Melbourne Research Cloud (MRC) for contributions to ATLAS and Belle II grid's
- MRC operates an OpenStack private cloud:
 - Consists of 20000 CPU cores and 3 PB Object store
- Australia ATLAS SE of 1.5 PB is now below threshold for “useful” contribution to ATLAS-grid
- Open negotiations with ATLAS-grid for migration to sustainable contributions

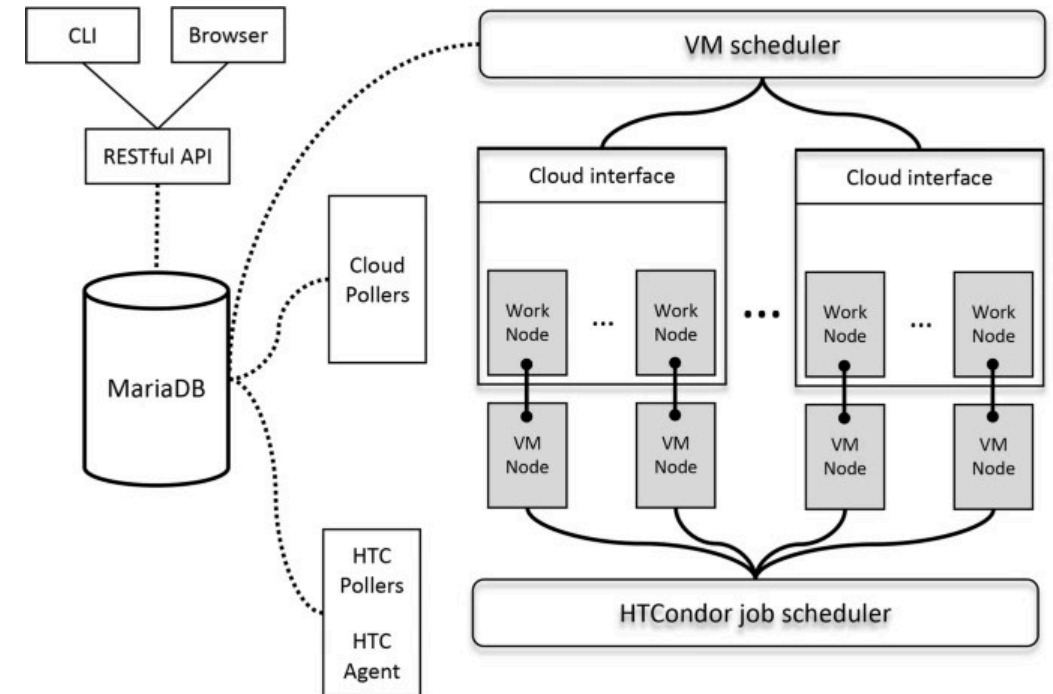
Australian contributions to Belle II grid

- Australian contribution ~900 CPU core and SE of 200 TB in 2020
- Look to provide this via MRC and UVic Cloud-scheduler + Dynafed
- Currently running 250 CPU test site, will upgrade to 900 slots soon.
- Implemented http-based Dynafed site with 50 TB test.
- Needs some fixes in Belle-Dirac before going live



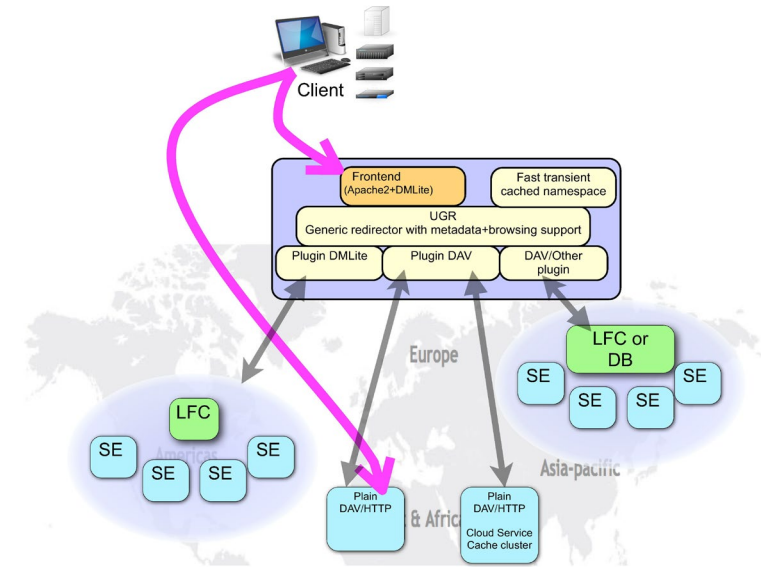
University of Victoria Cloud- Scheduler

- Belle II employs the Uvic instance of CloudScheduler to launch Cloud jobs
- Pilots are sent from the DIRAC instance in KEK to Victoria, Canada, which in turn launches jobs in Clouds around the world (including Melbourne)



http-Dynafed


- Provides a unified names-pace to geographically distributed data.
- Data access and transfer via http (webdav)
- Plugin for Cloud-based Object-Store
- Belle II has a working dynafed
- Melbourne has a test site implemented via MRC



```
mebert@heplw65:~$ gfal-ls -l davs://dynafed02.heprc.uvic.ca:8443/me
-rwxr-xr-x 00 0 93593105 Nov 5 23:27 B_2K+K-_raw_MVAPhoto
-rwxr-xr-x 00 0 93693725 Nov 5 23:25 B_2K+K-_raw_MVAPhoto
-rwxr-xr-x 00 0 164692397 Nov 5 23:27 B_2KK_1M_test.root
-rwxr-xr-x 00 0 164605095 Nov 5 23:26 B_2KK_1M_train.root
-rwxr-xr-x 00 0 6100556 Nov 5 23:25 B_2KK_test.root
-rwxr-xr-x 00 0 6099687 Nov 5 23:25 B_2KK_train.root
-rwxr-xr-x 00 0 3087903 Nov 5 23:25 B_2KK_qq_test.root
-rwxr-xr-x 00 0 3088779 Nov 5 23:25 B_2KK_qq_train.root
-rwxr-xr-x 00 0 16248334 Nov 5 23:25 B_2KK_raw_test.root
-rwxr-xr-x 00 0 16266934 Nov 5 23:26 B_2KK_raw_train.root
-rwxr-xr-x 00 0 3027954 Nov 5 23:27 B_2KK_sig_test.root
-rwxr-xr-x 00 0 3028031 Nov 5 23:26 B_2KK_sig_train.root
```

Australian Network connectivity

- Provided by AARNet who continue to provide outstanding service.
- Substantial upgrade for Asian connectivity with new 1 Tb/sec links to Singapore (via indigo) and Japan (via JGA) in 2020
- Substantial upgrades in domestic service including delivery of zoom as Universities preferred video-conferencing service
- Multiple 100 Gbs links within Australia



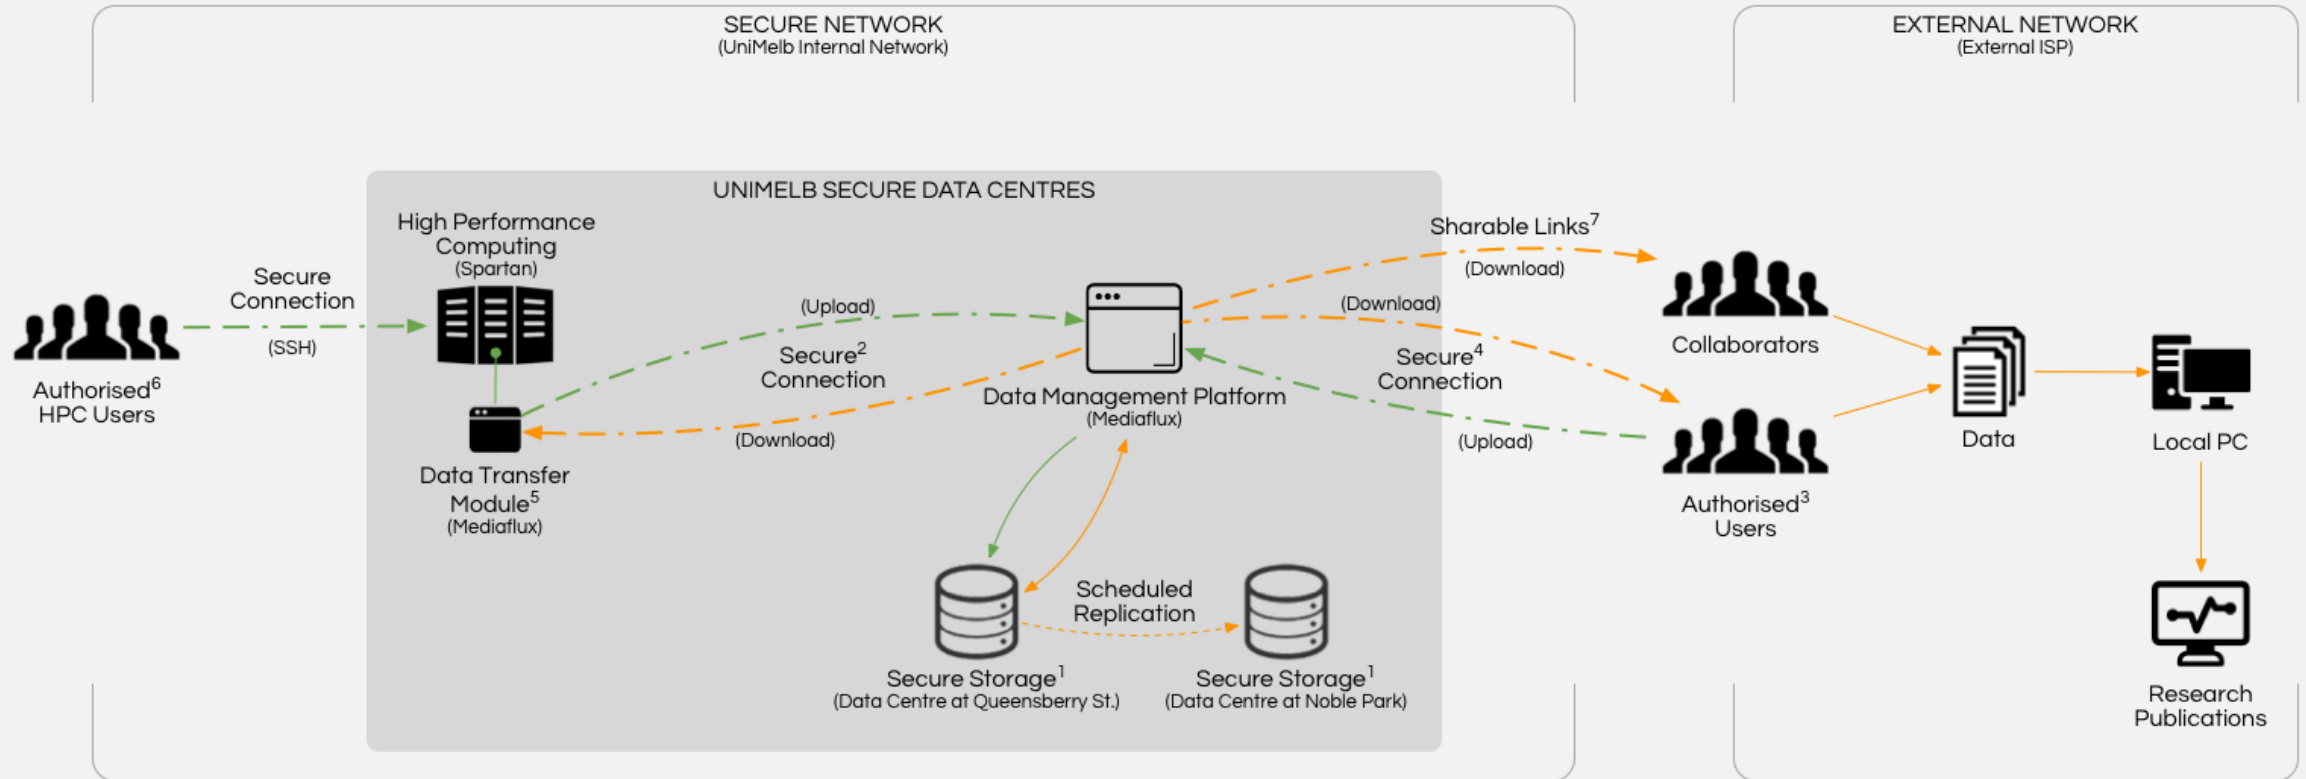
New zoom users in 2020. Note COVID impact in March

University of Melbourne Spartan HPC facility

- 5,500 CPU + 324 NVIDIA P-100 GPU cores
- 2.1 PB of Spectrum-Scale disk
- SLURM batch system
- CENTOS 7 base OS
- CVMFS for access to ATLAS, Belle II and other WLCG software (GEANT4)
- 20 PB of “Mediaflux” (CEPH) for offline, long term storage & data sharing
- General purpose data analysis and simulation facility

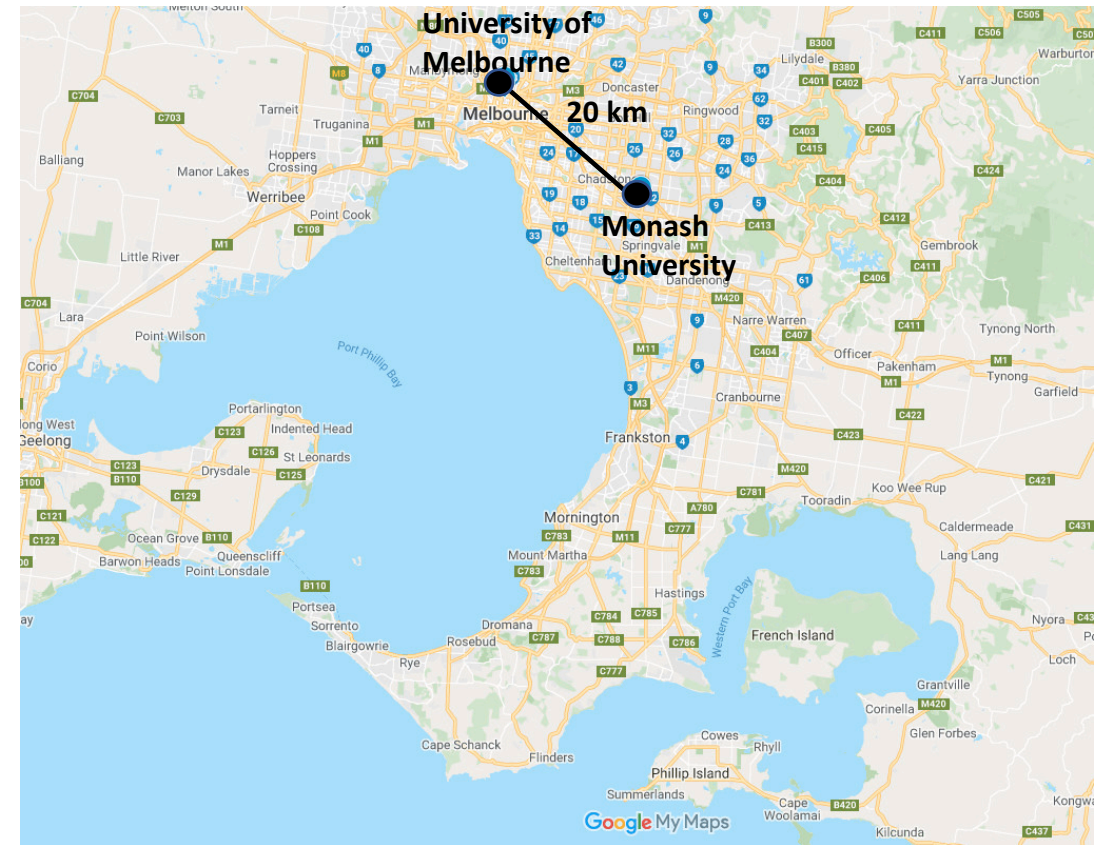
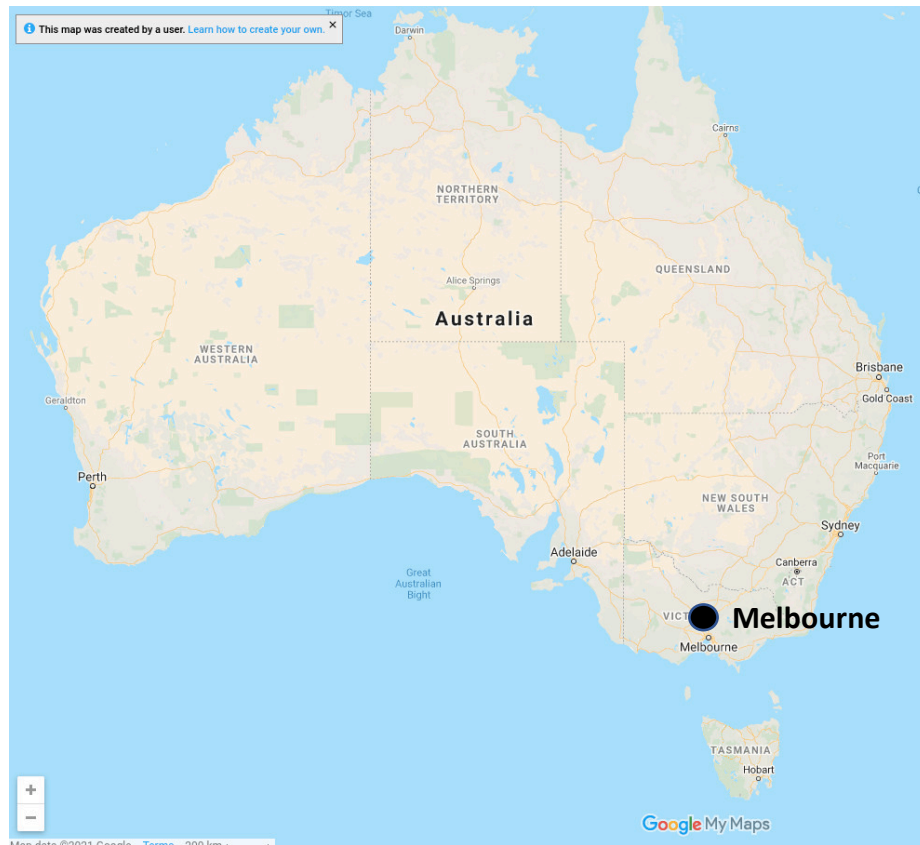
Mediaflux for Offline storage

Use Case: Long-term Data Storage for HPC/Spartan Users



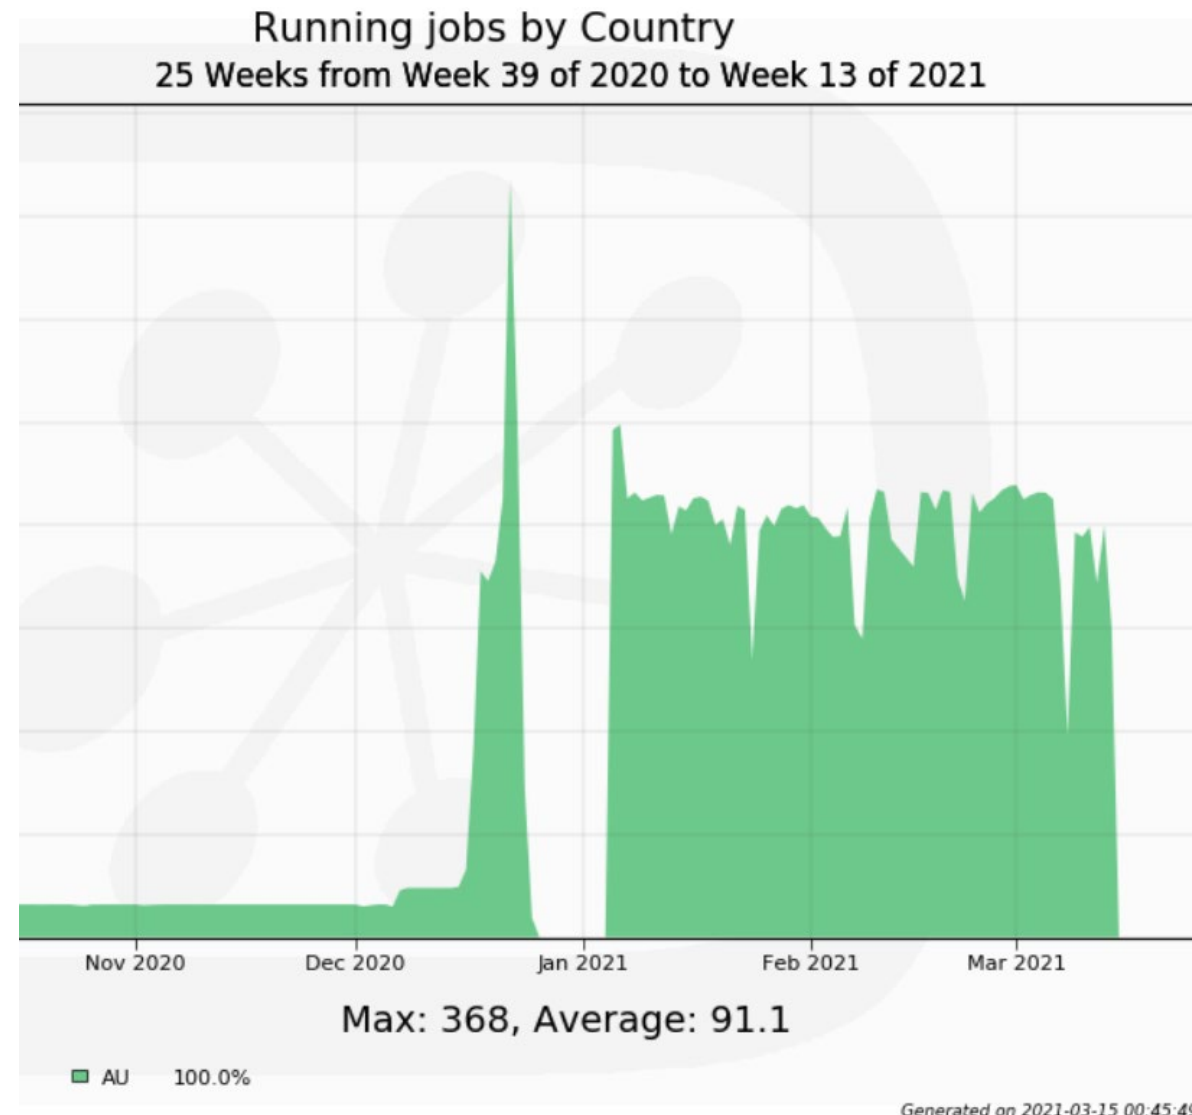
Melbourne & Monash Universities

Similar in size ~ 50,000 students



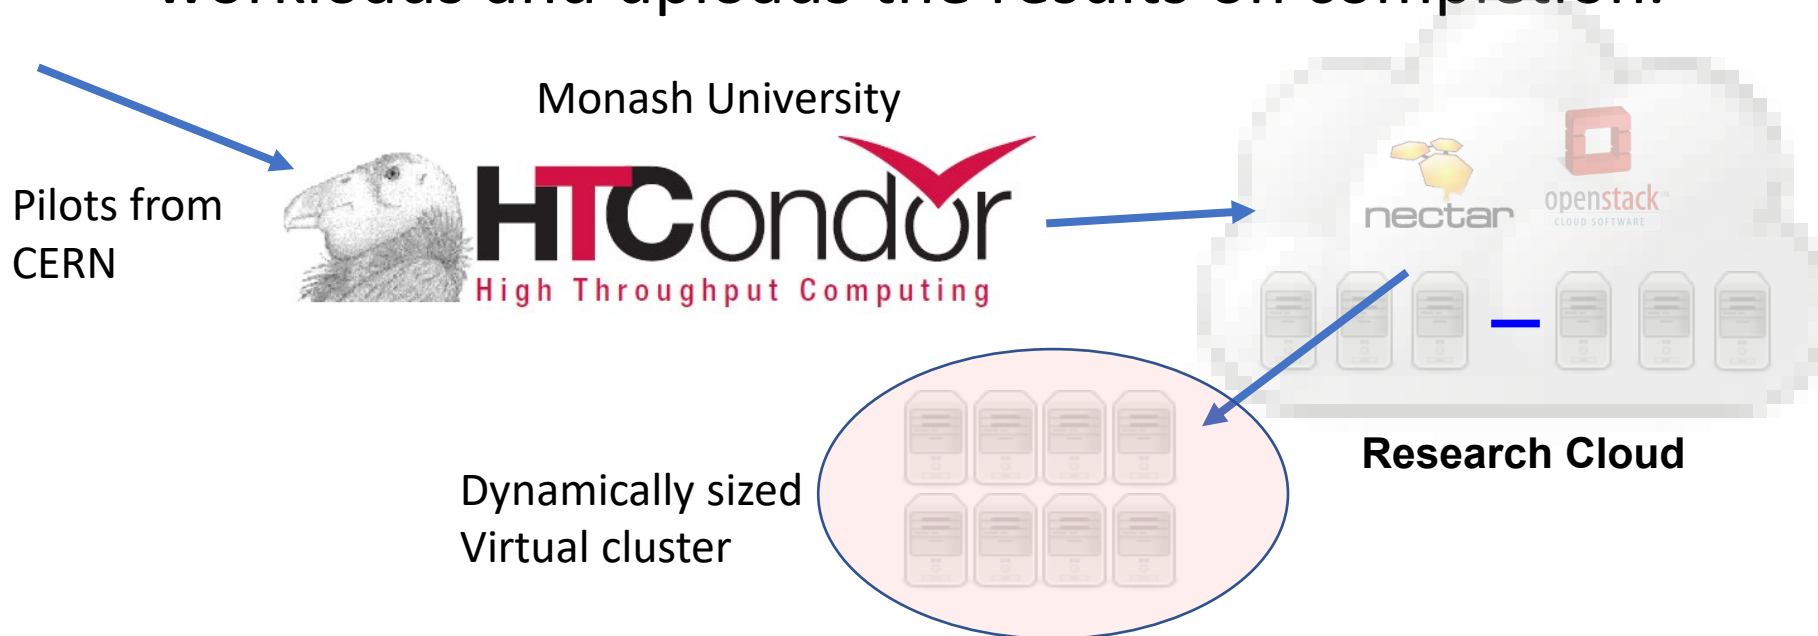
Monash University contributions to LHCb & COMET

- Federal Nectar (OpenStack) cloud resources
- Since October 2020, Australia is contributing to the core computing of LHCb
- Small permanent allocation with spot allocation of nodes on a 24h basis since late December
- HTCondor nodes booted which then pulls workloads through a pilot system
- Contributes the fair share of simulation jobs for collaboration



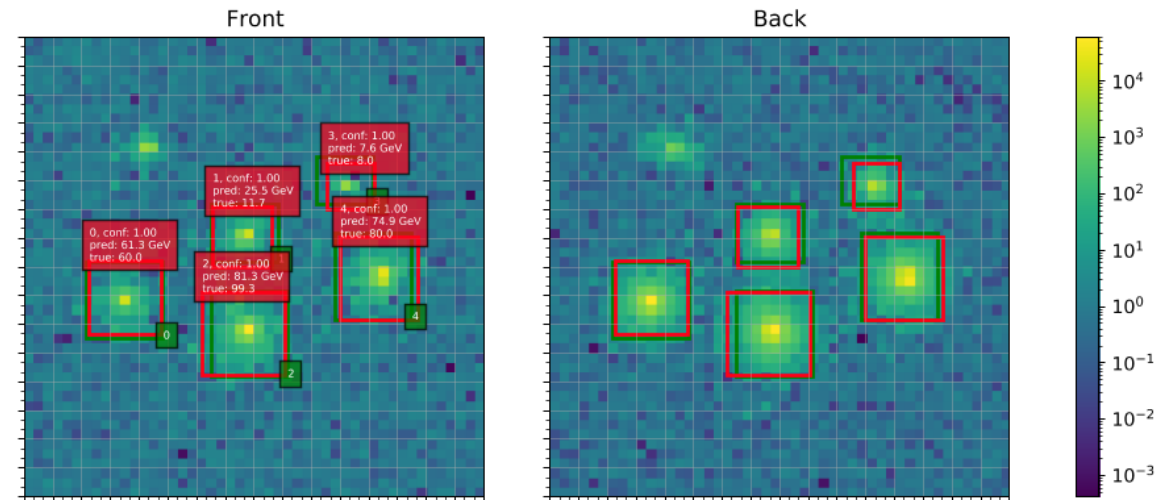
HTC-Condor + DIRAC

- Local HTC-Condor is started in the NeCTAR OpenStack Cloud
- This is configured to launch new VM's the NeCTAR Cloud on demand
- Pilot jobs from the LCHb DIRAC instance at CERN are sent to Monash
- HTC-Condor creates the VM's on demand and the Pilot jobs pull in the workloads and uploads the results on completion.



Monash University M3 farm

- 5,200 CPU cores, 1,700 GPU Tesla cores
- 2.9 PB Luster Filesystem
- Used for COMET Phase-II simulations. Can process $O(10^3)$ jobs per day.
- Data stored locally
- Use the GPU capabilities of the farm to develop machine learning algorithms for cluster identification in an upgraded calorimeter for LHCb



Summary

- Migrate to University-based general purpose IT solutions
- Extensive use of Cloud resources
- Move of tailor-made HEP software to mainstream solutions
- Payoff is substantially reduced manpower costs
- CVMFS is a core enabling technology