Clouds for research computing

Randall Sobie

Institute of Particle Physics

University of Victoria

Grand challenges

Why is the universe not made of equal amounts matter and antimatter ?

We build instruments large detectors to record the collisions of matter and antimatter

SLAC National Accelerator Lab

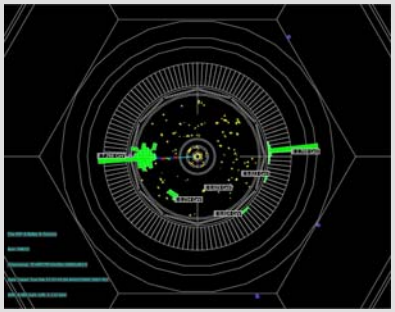

Record billions of particle collisions

"Events"

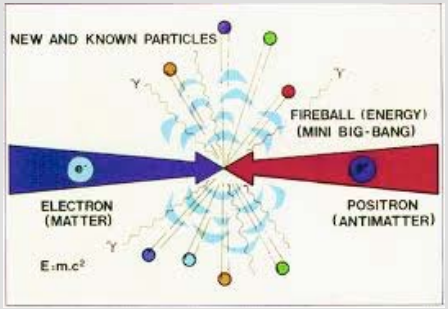

Computing solutions

Users build their analysis code and submit many batch jobs

BaBar experiment uses multiple and independent facilities

Newer generation experiments (LHC) use *grid* technologies to construct an integrated environment using many sites around the world

Role of clouds in research computing

Parallel applications require large, dedicated facilities
(High-performance computing HPC environment)

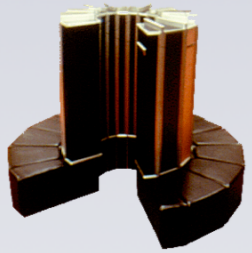

Large-scale, data intensive, embarrassingly parallel applications well suited for the Grid
(Tight integration of the application and systems)

Commercial and science clouds provide SaaS and IaaS research computing solutions

SaaS (Software-as-a-Service)
IaaS (Infrastructure-as-a-Service)

Complex research environments

BaBar is a Complex Application

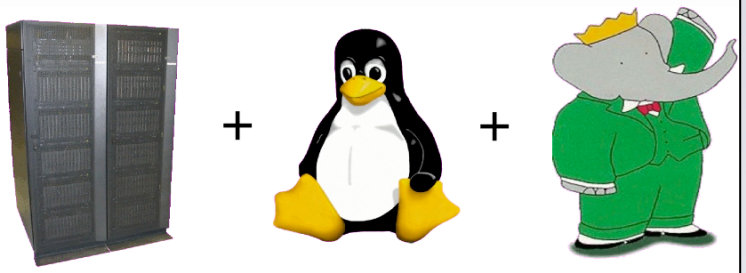

Issues:

- 9.5 million lines of C++ & Fortran
- Requires particular OS, compiler, & library levels
- Certification of environment required
- Not easy to share environment
- Sysadmins almost need to become application specialists

How do we analyze the BaBar in the coming few years?

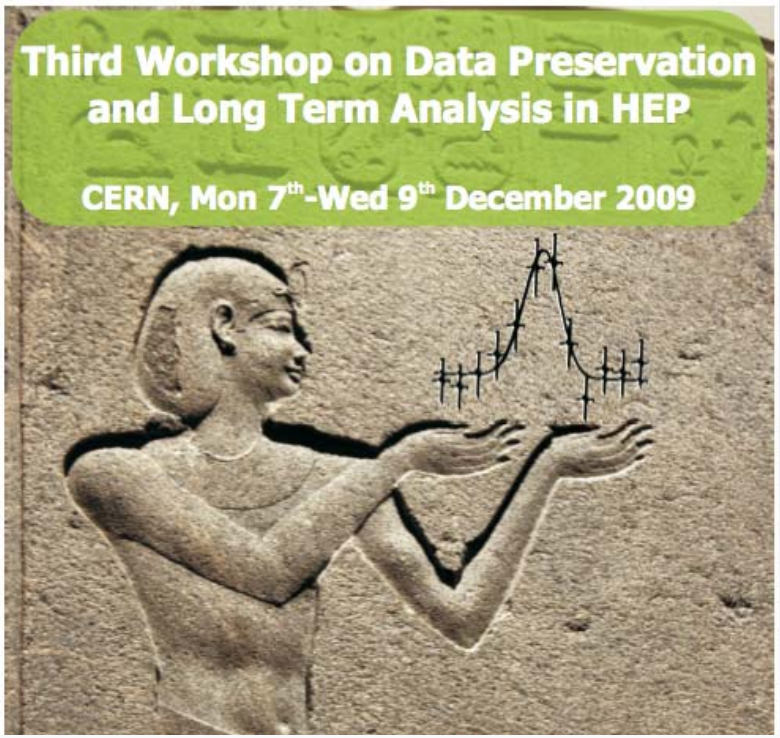

Data Preservation:

We need to archive the data and the software for many (>10) years

Distributed compute cloud

Sophisticated user communities in physical sciences
- Non-GUI users
- Batch computing environments

Complex software packages and demanding system requirements
- Specific OS system
- Specific application libraries and compilers

Medium-scale data sets (100s TBs)
- Data accessed (on-demand) from remote or local repositories

Distributed compute cloud
System to boot user-customized VMs on any number of science or commercial clouds in a familiar batch computing environment

Often referred to as *Sky Computing* or *Grid of Clouds*

Components

Application encapsulation
Image replication
eg Xen, KVM

Clouds
WS interface
eg Nimbus, OpenStack, EC2

Dynamic resources
eg Condor, SGE

Job Scheduler

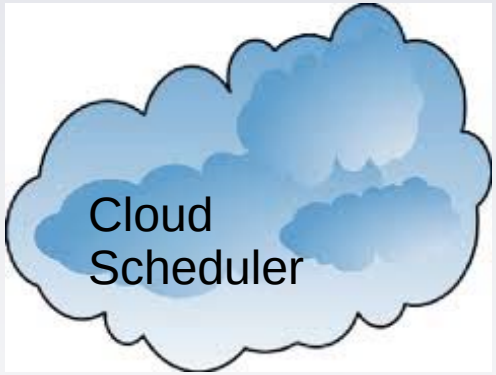

Managing multiple clouds
eg Cloud Scheduler

The Interactive System

User saves the modified environment as a new image

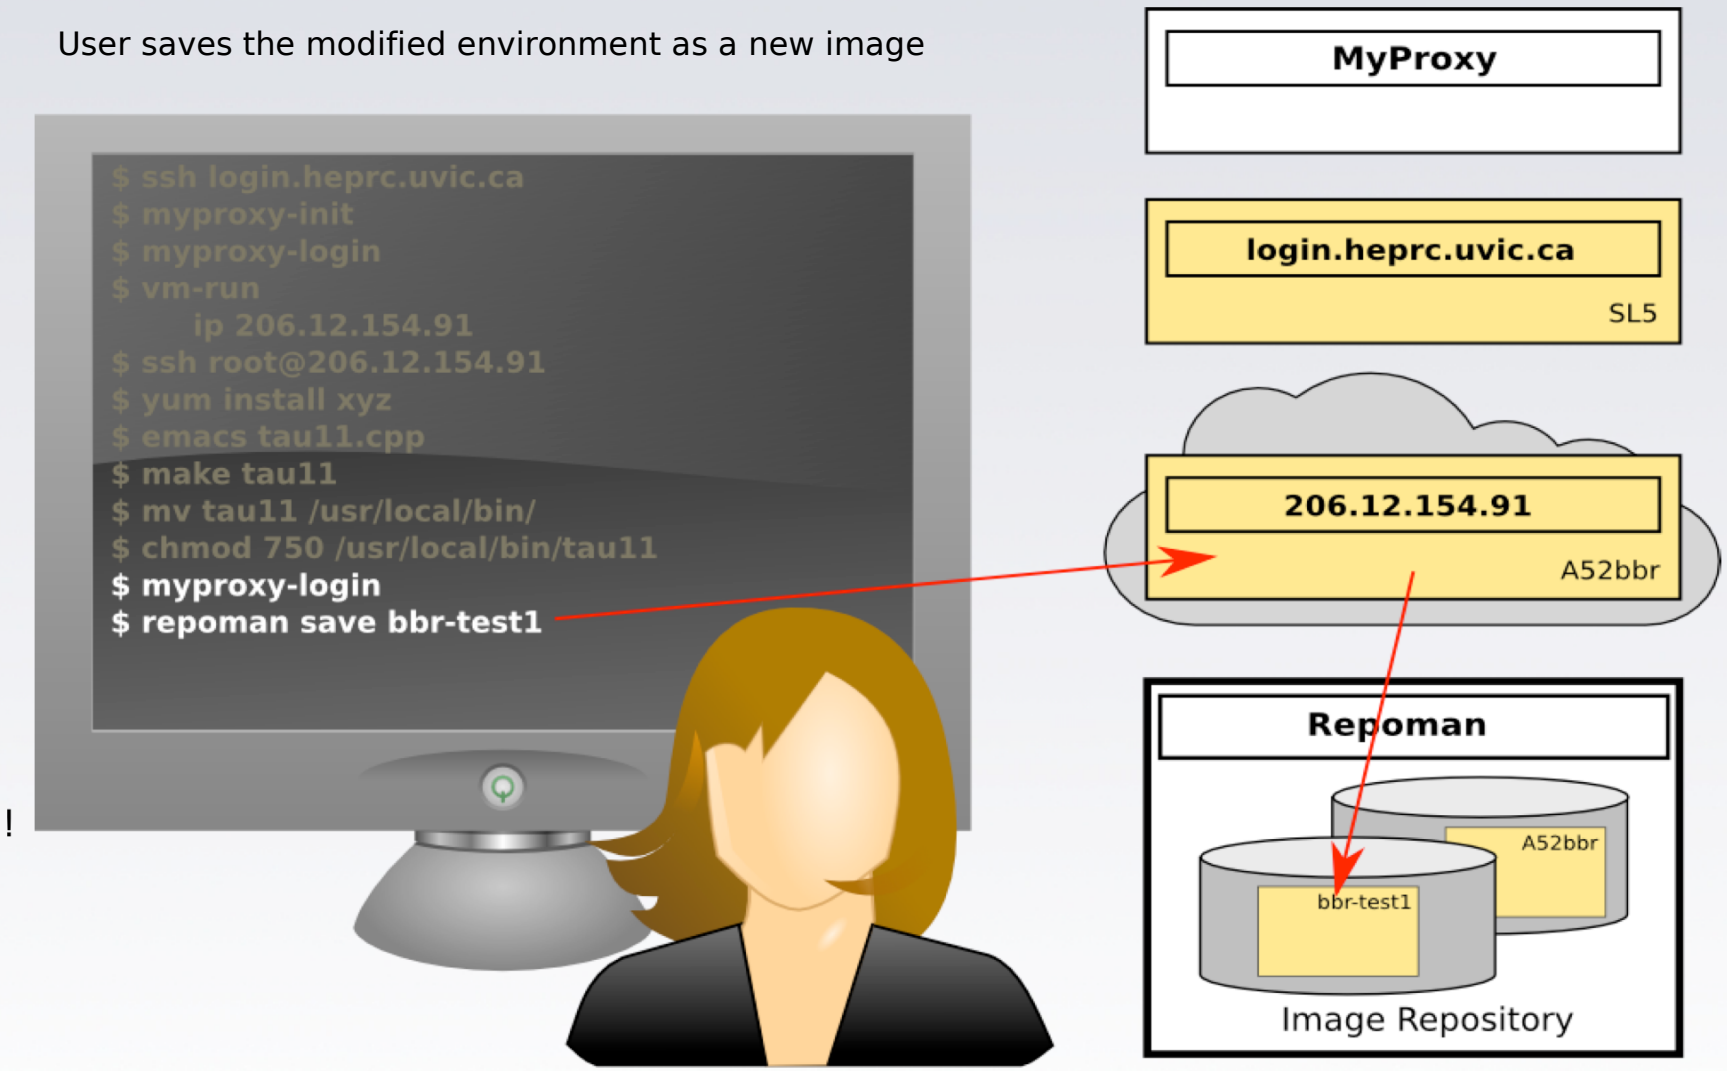

!

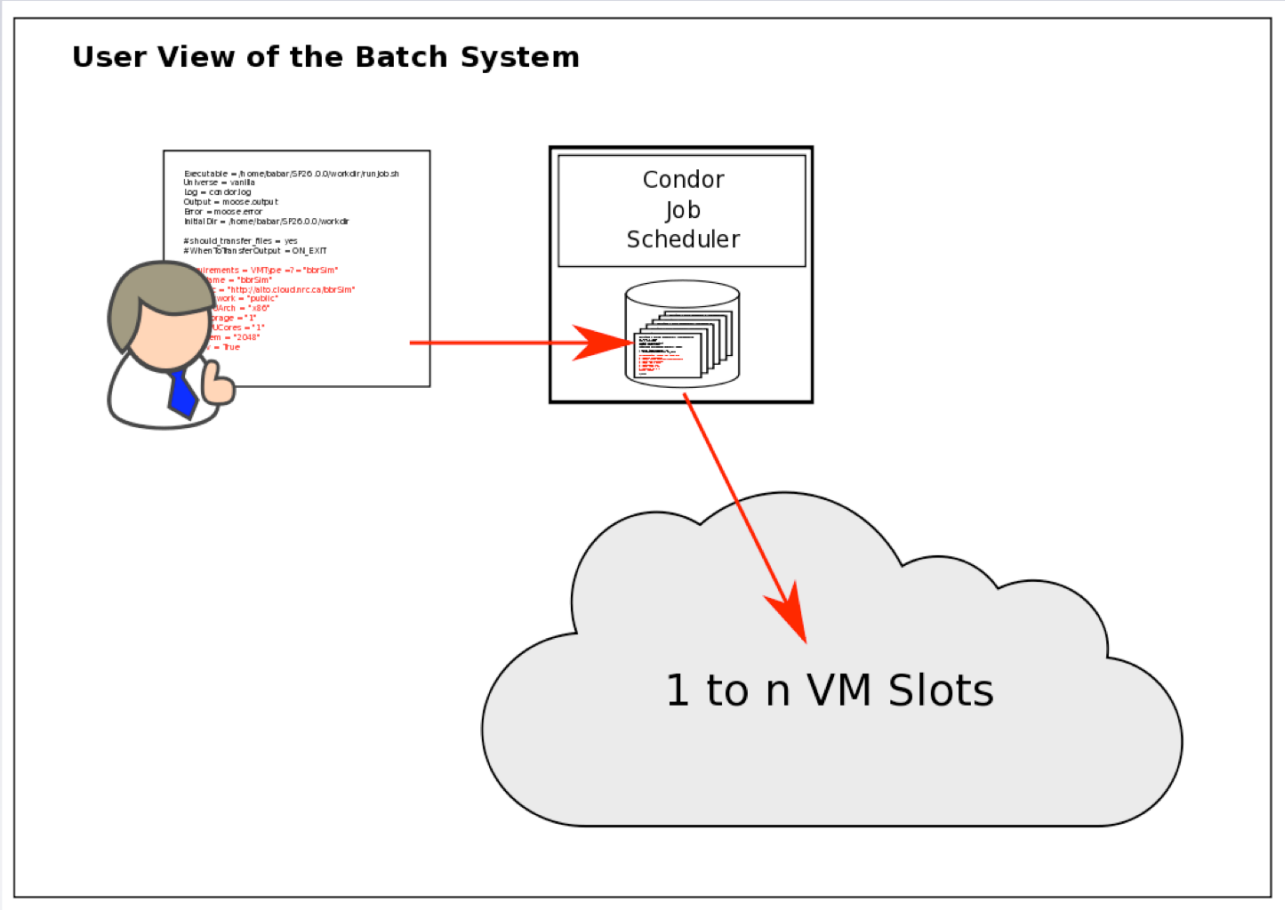

User view of the system is the same as a standard batch environment

Job script contains a link to the user's VM required for the job

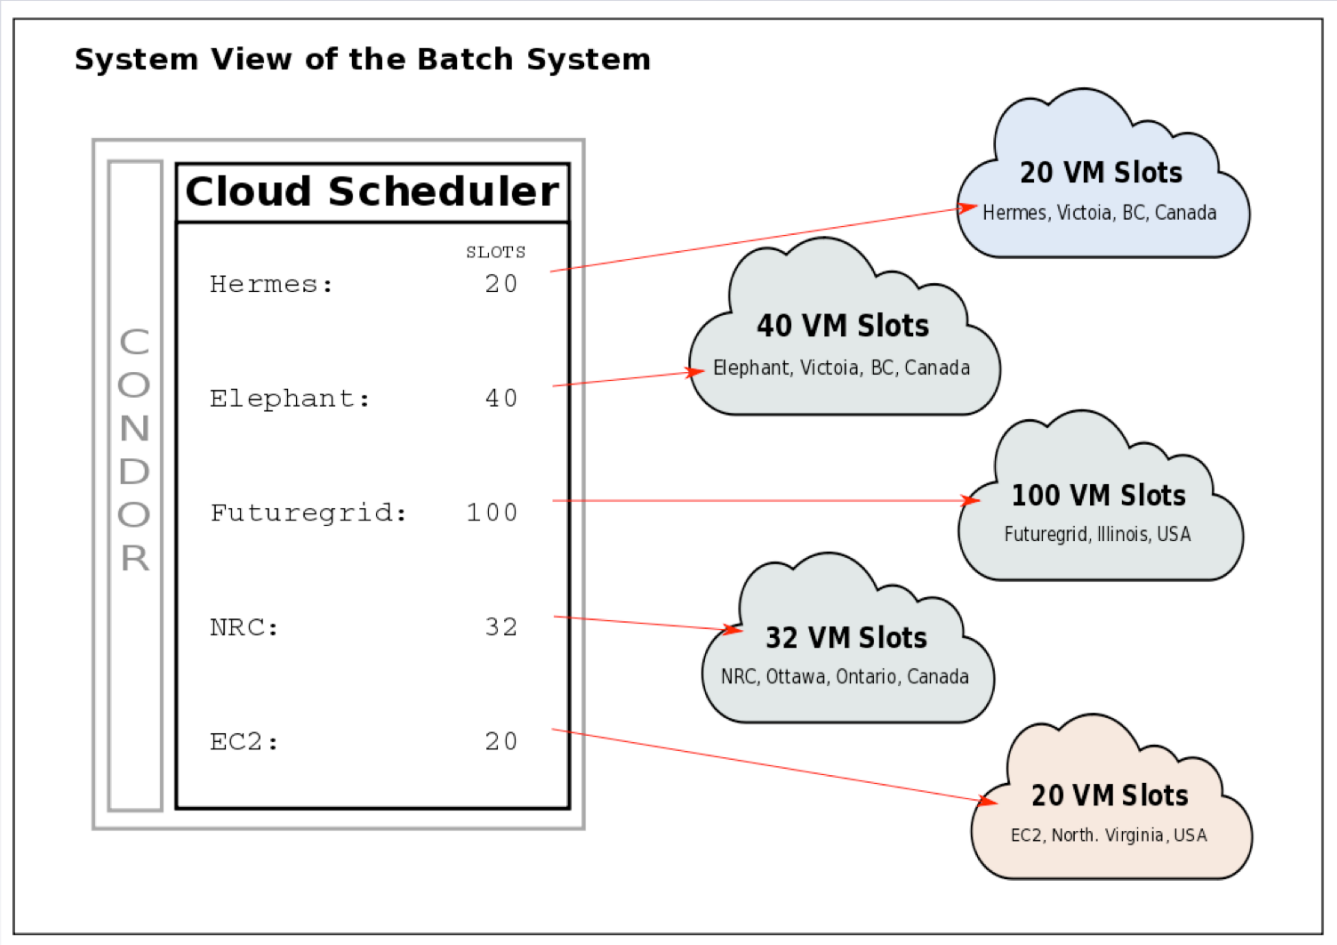

CS looks at the job queue and sends a request to the next available cloud to boot the User-VM

Simulation Production

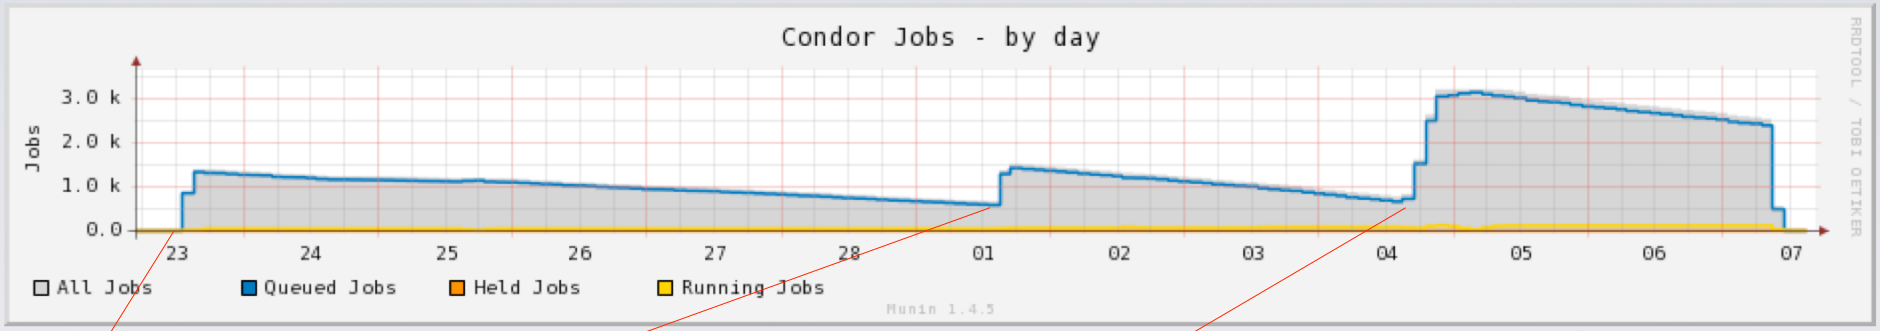

~1300 jobs
50 slots
2 time zones

~850 more jobs, 20 more slots

15 more slots, 3rd time zone

2500 more Jobs, 35 more slots

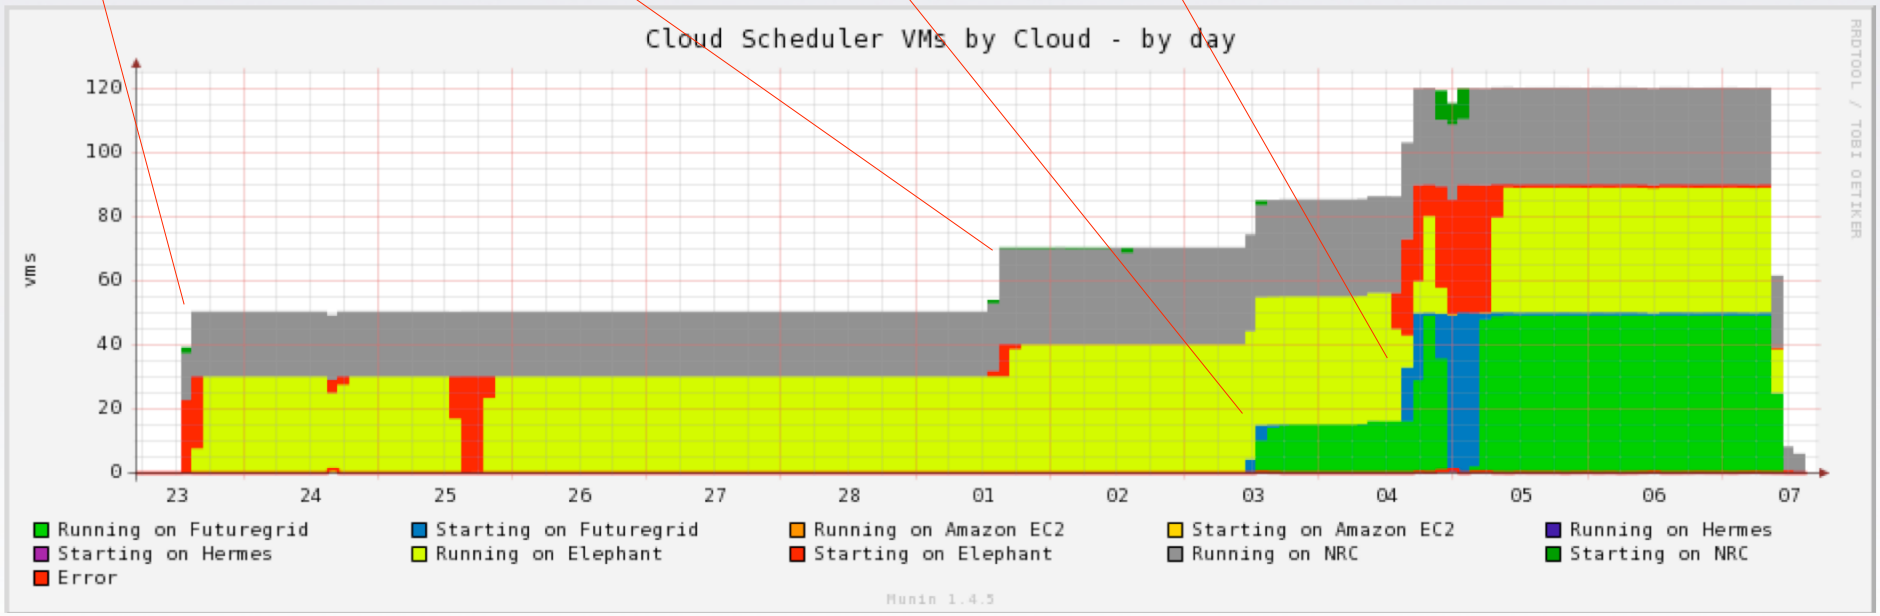

Astronomy applications

CANFAR Project
Canadian Advanced Network for Astronomical Research
UIVC, UBC, NRC-HIA
CANARIE-funded project

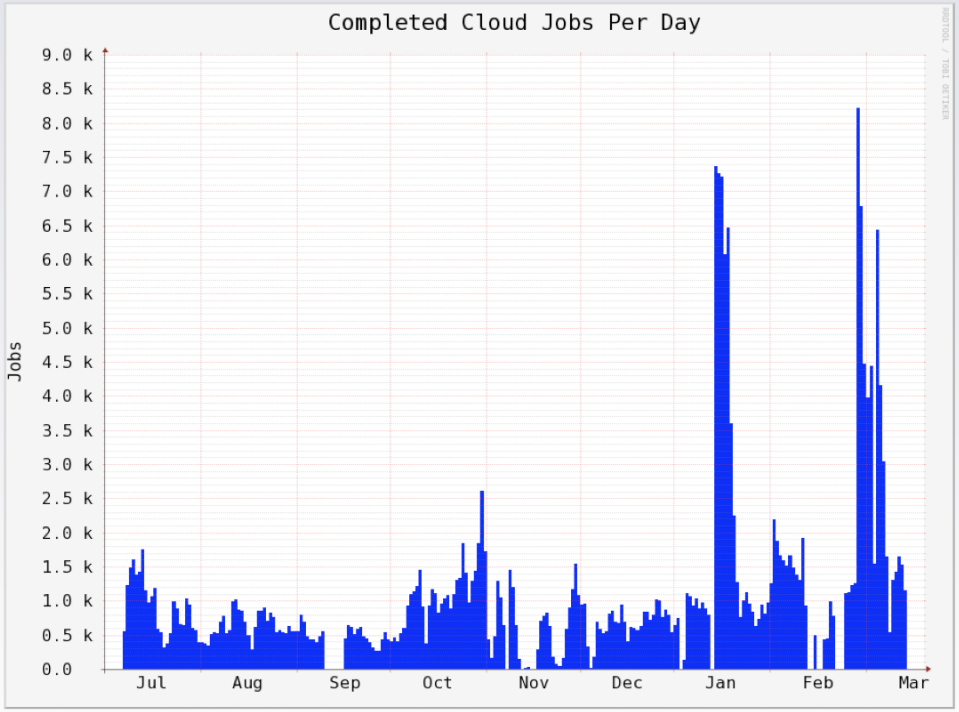

Distributed cloud used to process survey data

In production for 8 months using different IaaS cloud resources

Compute Canada cloud site at UVIC

Enabling system for user analysis as well as production jobs

Summary

- We have established a distributed cloud for research applications
  - Focus is on applications in physical sciences with large high-throughput (HTC) workloads and a knowledgeable user community
  - Fault-tolerant system using multiple-IaaS (commercial or science) cloud resources
  - Based on open-source components with two new in-house elements
  - Easily scales for low-IO applications
  - We are currently studying the scaling to high-IO applications where the data located at a few repositories

Support provided by CANARIE, NSERC, NRC, Amazon, Google, FutureGrid (NSF)