

# SPARTAN

Performance & Flexibility

## An HPC-Cloud Chimera

eResearch Australasia, Melbourne, October 11, 2016

Lev Lafayette, University of Melbourne  
Daniel Tosello, University of Melbourne  
Bernard Meade, University of Melbourne

# HPC, Cloud, and Edward

- HPC systems provide excellent performance, especially for multinode computational tasks.
- Cloud architectures provide flexibility and cost-efficiency - however they don't have the power of HPC.
- University of Melbourne's general purpose HPC system (Edward) was due for retirement.
- Also provided specialist queues for particular departments.
- Edward supported 886 users and 371 projects.

# Why change?

- Debate about the value proposition of HPC
- Profiling of Edward users indicated that usage was very heavily biased towards lots of single node and low memory tasks.
- Demand for more cores, more RAM, faster storage and interconnect is **actually** a demand for shorter queues
- Access to national peak facilities is not simple and many researchers choose to DIY with grant money
- Substantial investment in the Research Cloud

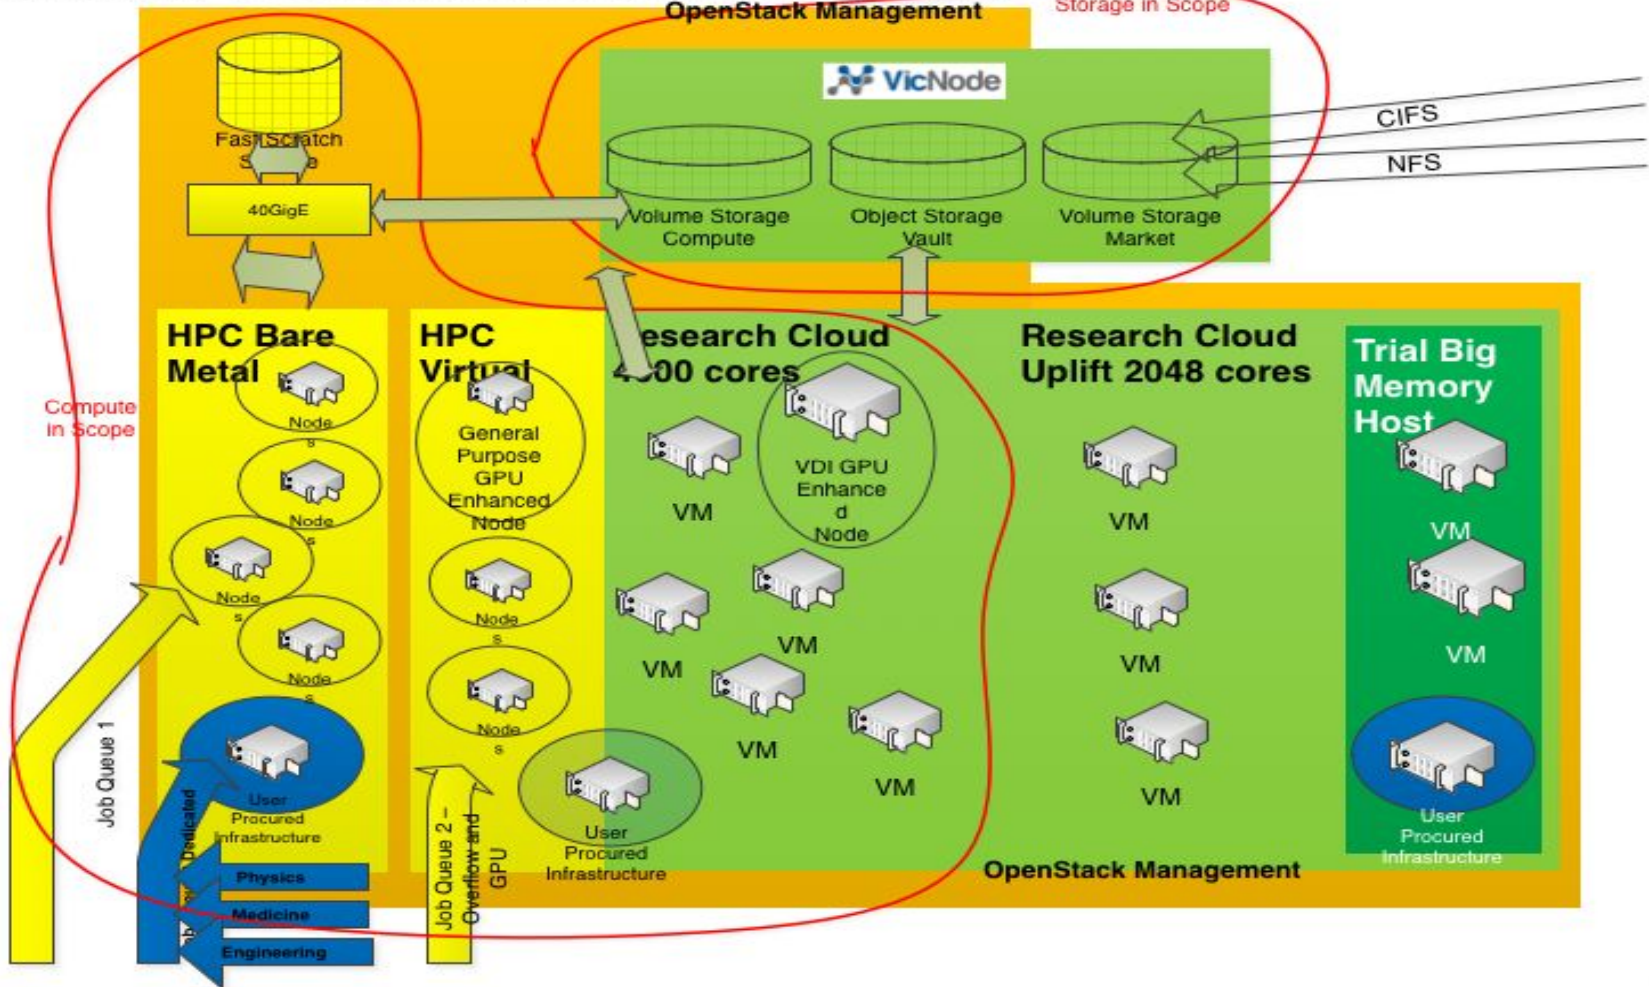
# The Spartan Vision

- A mix of bare-metal and virtualized resources
- Operated as a service in the Research Cloud
- Dynamically scalable depending on demand and availability
- Investment can match the usage demand
- Greatly improves the Research Cloud vCPU utilisation
- No longer need forklift upgrades
- Grant money/departmental investment can be used to expand the required aspects of the service

# Concept Diagram - Research Platform Services Compute and Storage Services

## OpenStack Management

Storage in Scope



# The Plan

*".. no plan of operations extends with any certainty beyond the first contact with the main hostile force."* (Helmuth von Moltke the Elder, 'On Strategy', 1871)

- Requirements gathering - 2014
- Partnering with Victorian Life Sciences Computation Initiative (VLSCI)
- Input from other research institutions
- External architect contracted for solution design

# The Implementation

- Extensive training and transition programme.
- Integration with campus Active Directory was changed to LDAP instead.
- Noisy neighbour effects caused issues. Turn off overcommit on cloud nodes.
- All vHPC converted to cloud-burst using Slurm's power-saving feature.
- Connections to VicNode Aspera Shares and specific I/O nodes for data transfer.

# Spartan Hardware

- Spartan-m and Spartan: 16 core Intel Xeon E312xx @ 2294MHz, 64 GB RAM
- Spartan Bare Metal: 12 core Intel(R) Xeon(R) CPU E5-2643 v3 @ 3.40GHz, 294 GB RAM, 252 cores total
- Spartan Cloud: 8 core Intel Xeon E312xx @ 2294MHz, 64 GB RAM (1024 cores total)
- The interconnect 10gb Ethernet (cloud), 25gb/56gb Ethernet (physical) with Mellanox SN2700 and SN2100 leaf switches with 100G connections to core.



# Scheduler and Applications

- SLURM Workload Manager for resource manager and job scheduler.
- EasyBuild is used for software installations; builds from source codes with specified toolchains and dependencies.
- Automatically generated LMod environment module files.

# Storage, OS, Management

- There is Fast CEPH storage from the spine, and multiple 56gb scratch storage nodes from the leaf.
- Using Linux version 3.10.x and Red Hat Enterprise Linux Server release 7.2
- Consistency in operating system and libraries in the management, login, i/o, physical, and cloud compute nodes maintained through Puppet reviewed through Gerrit.

# The Future

- Large separate partition for proteomics project.
- Closure and transition of Edward users (Dec 2016 deadline).
- Separate partitions for different hardware and separate login nodes for special groups (e.g., classes).

**Prediction: Spartan's architecture will be the model for high throughput general purpose computing in the future.**