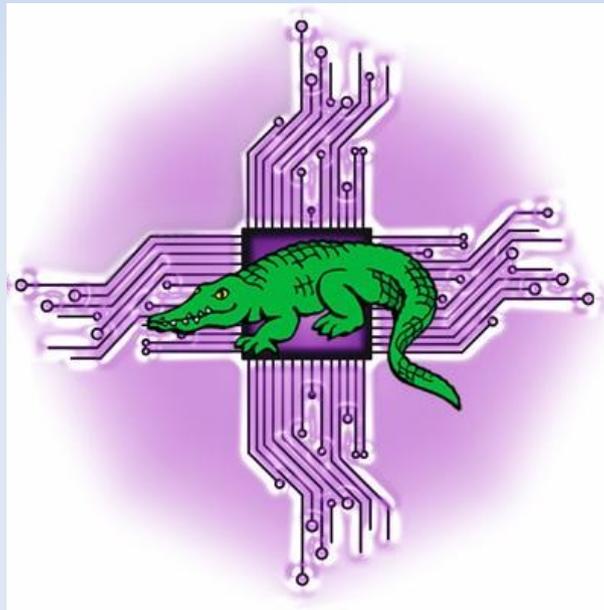


# A Laconic HPC with an Orgone Accumulator

**Presentation to Multicore World 2016**



**Wellington, February 15-17, 2016**

<http://levlafayette.com>

# Edward - University of Melbourne Cluster - System

- **Installed and operational since 2011. Named after King of Wessex (prior cluster was Alfred)**
- **Hardware : 48 compute nodes. 2x AMD Opteron 6128 (16 cores per node), 1TB localdisk, 32 GB memory**
- **2 Storage Nodes. IBM x3650 (for user and project data, 45TB each)**
- **Interconnect and Filesystem : 10GbE with XFS, NFS**
- **Deployment, Resource Manager and Job Scheduler: xCAT, Torque, Moab**
- **Authentication and Accounts: LDAP, Kerberos, Karaage**
- **Operating System : Scientific Linux 6.**
- **Main Application Software: Matlab(TM), R, OpenMPI, FDS, FSL, Comsol.**



# Edward - University of Melbourne Cluster - Users and Usage

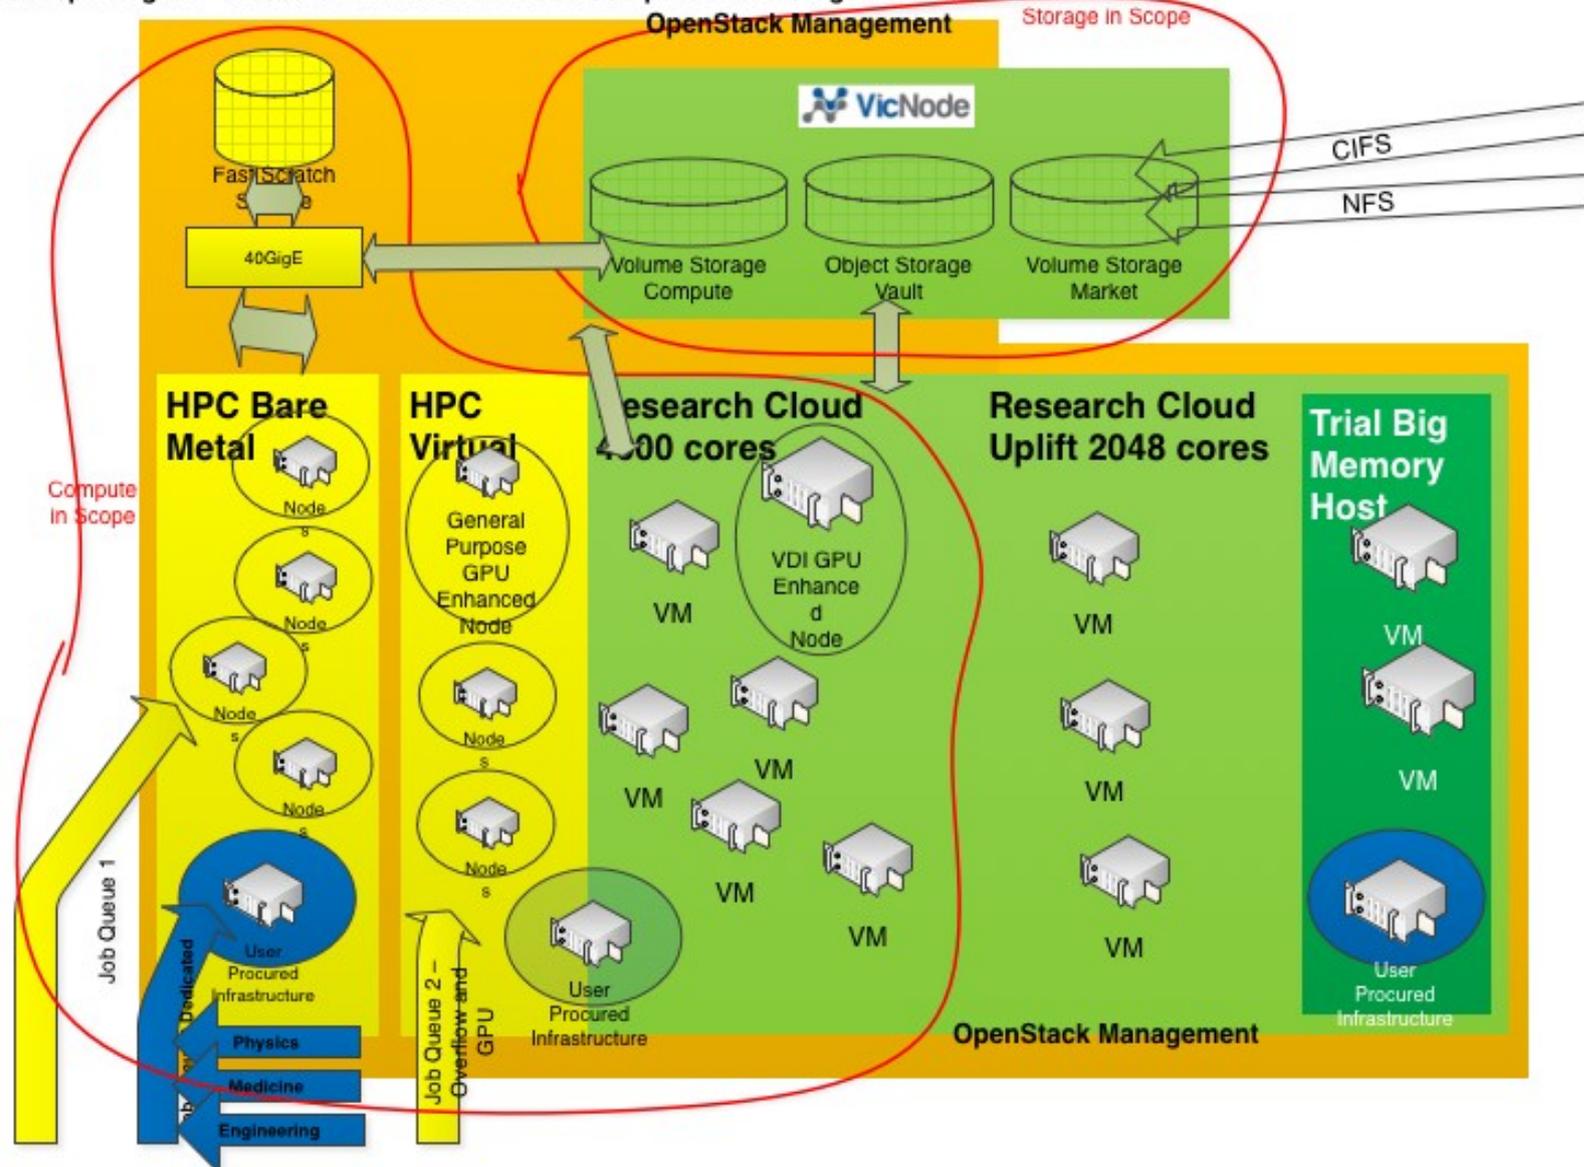
- As of February 9, 2015: 816 active users: (including 25 from Victoria University) and 333 projects.
- From January 1st 2015 to December 31st 2015: 4,930,340 compute hours from 375,458 jobs; 74.84% utilisation (and increasing).
- Major projects in the past year; physarum optimisation, fluid flow with deforming interfaces in microchannels, bushfire dynamics (three different projects), path analysis for genomic datasets, efficiency in luminescent solar concentrators ....
- Users have not had been engaged in an extensive HPC training program; approximately only 38 student/days training from 2012-2014. Compares strangely to other extensive ResBaz activities.
- As a result, single-core and low memory jobs dominate; 76.35% of jobs from Feb 9 2015 to Feb 9 2016 were single core, and 96.83% used 1-4GB of memory.

# Moving Towards A New System

- A detailed review was conducted last year looking at the infrastructure of the Melbourne Research Cloud, High Performance Computing, and Research Data Storage Services.
- University desired a 'more unified experience to access compute services'
- University was already using Openstack cloud software as part of the NeCTAR Research Cloud, with storage provided by NetApp.
- Recommended solution, based on technology and usage, is to make use of existing NeCTAR Research cloud with an expansion of general cloud compute provisioning and use of a smaller "true HPC" system on bare metal nodes.
- Infrastructure deployed across two data centres with HPC at one. The replacement RDSS storage service will use the same Netapp shared storage solution.

# Moving Towards A New System

Concept Diagram - Research Platform Services Compute and Storage Services



# The Spartan HPC ....

- The HPC will be called "Spartan" (not Æthelstan or Ælfweard!).
- The 'bare metal' HPC component really will be laconic. Think of Sparta's citizenship structure - the few *Spartiate* citizens are bare metal HPC, the more numerous *Perioeci* free inhabitants are vHPC nodes, and the many *Helot* slaves are elastic compute nodes.
- Real HPC is a mere c200 cores, 16 GB per core. 2 socket Intel E5-2643 v3 CPU with 6-core per socket, 192GB memory, 2x 1.2TB SAS drives, 2x 40GbE network



# .... with Cloudbursting

- Management and login nodes will use vHPC.
- Slurm will be implemented for job management and integrated with Openstack with Cloudbursting ("elastic computing"). Burst capability into the Research Cloud when resources are not available in the HPC cluster (vHPC) (1:1 subscription of CPU for vHPC)
- Single authentication with Openstack and HPC Job Management
- 4000 cores with 4+GB per virtual core for the Research Cloud Compute environment.
- Server hardware is 2 socket Intel E5-2699 v3 CPU with 18-core per socket 384 GB memory.



# Implementation and Future

- Initial live tests being conducted by 250 students in March entirely on vHPC system.
- The decision to implement this model is driven by ease of management, integration of existing services, cost, and current user profile. It is not primarily driven by performance (a "high batch computer" rather than a "high performance computer").
- Widespread existing performance metrics suggest that cloud computing performs poorly compared to HPC in MPI and distributed computing but is comparable with single-core or multi-core shared memory nodes with minimal overhead.
- Existing minimal training courses will be expanded; plan to have three days of courses every second month for researchers, ranging from basic Linux/HPC to parallel programming.
- This will likely increase the use of shared and distributed multicore job submissions which will put pressure on small 'bare metal' HPC nodes. Sparta may need more citizens!

**THANKS FOR WATCHING**



**& LISTENING PATIENTLY**