


The Spartan HPC System at the University of Melbourne

Spartan: From Experimental Hybrid towards a Petascale Future



eResearchAU, October 20, 2020

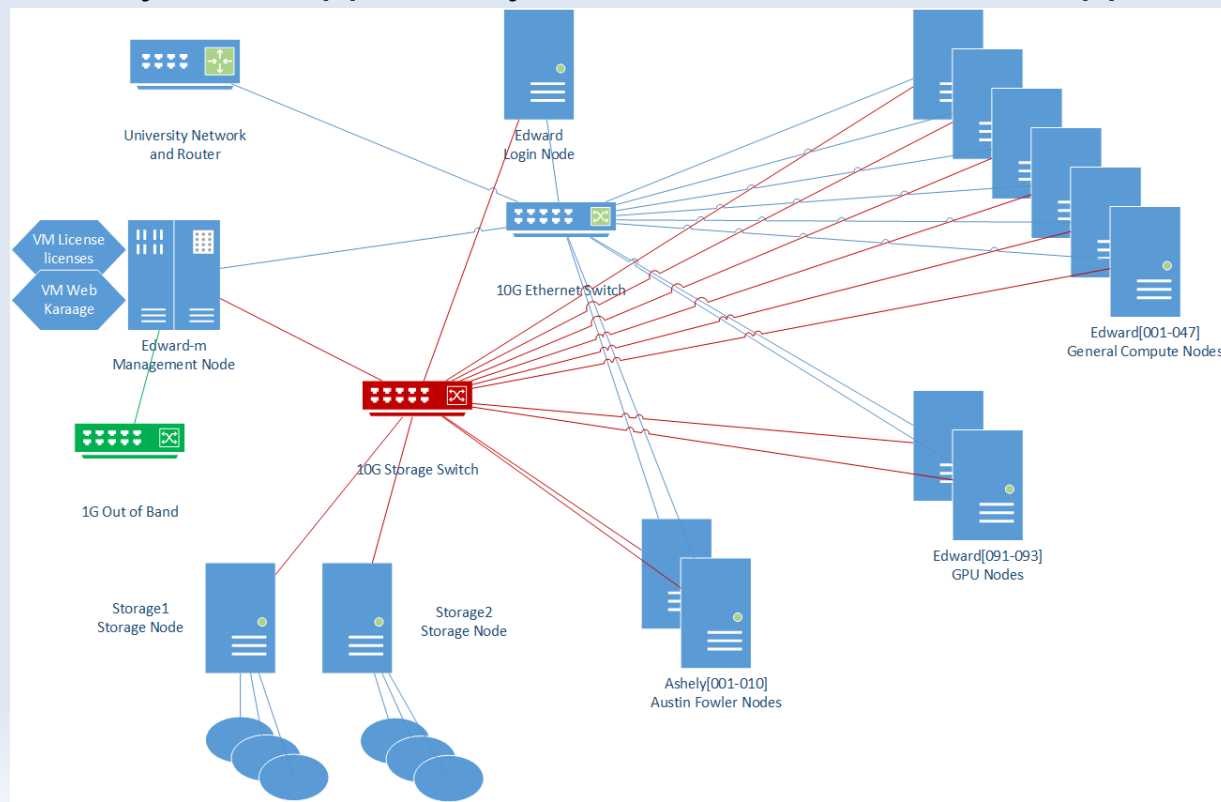
lev.lafayette@unimelb.edu.au

From Edward to Spartan

"Spartan" is the general purpose high performance computing HPC system at the University of Melbourne, first introduced in 2016.

Previous systems were "Edward" and "Alfred", named after the Kings of Wessex. Next system would have been "Æthelstan" and/or "Ælfweard"?!

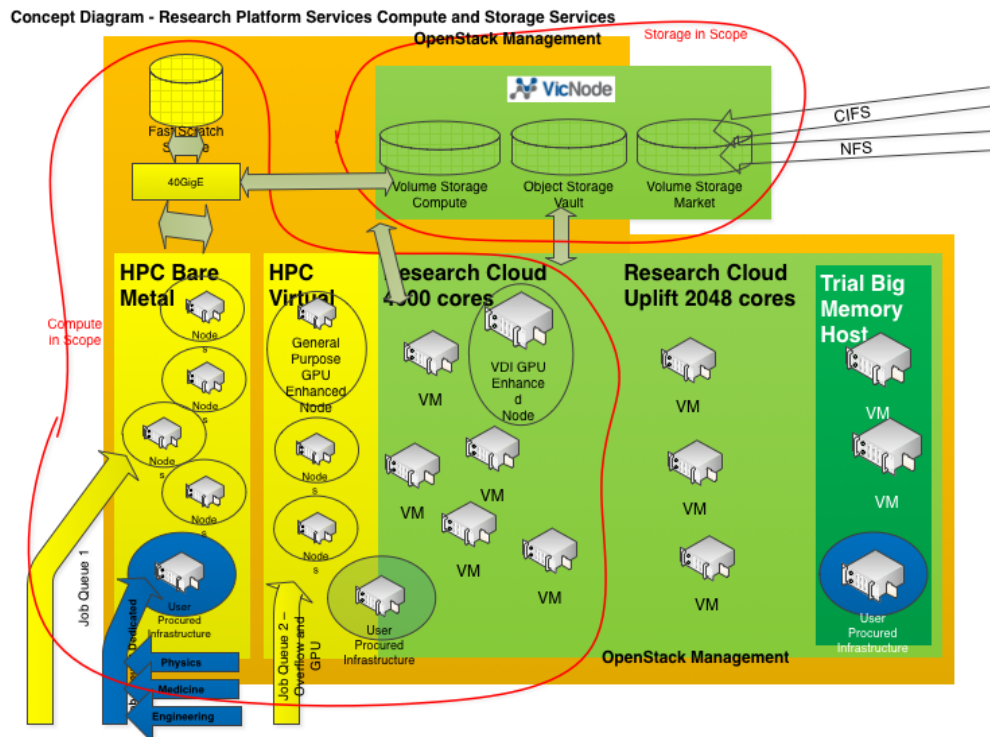
"Edward" was a fairly typical cluster, similar architecture to many others all around the world. But most of the jobs were single node, used for capacity rather than capability computing. So with a new system, opportunity cost metrics need to be applied.



This is Spartan!

Spartan was built on a very limited budget but made use of novel technologies; small traditional HPC "bare-metal" partition, larger use of NeCTAR research cloud virtual machines as compute nodes. Made use of OpenStack for deployment, RoCE interconnect technologies. CephFS for file system, Slurm for workload manager, EasyBuild and Lmod for applications.

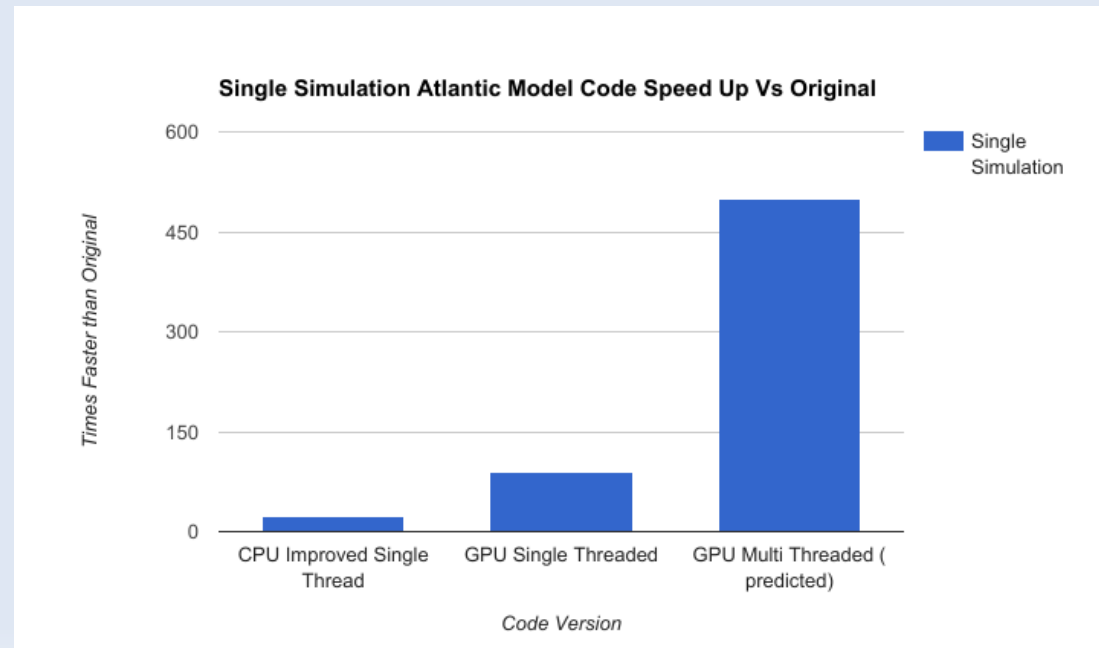

Initial test cases in March 2016, main launch of June, presentation to eResearchAU in October, followed by presentations throughout HPC centres European; Frankfurt, Stuttgart, Freiburg, CERN, Montpellier, Barcelona.



Laconic Successes to GPGPUs

Spartan completed 266% as many jobs in its first year compared to Edward's last year. LIEF grant LE170100200 from the ARC, a partnership between the University of Melbourne, La Trobe University, Deakin University, and the Royal Melbourne Institute of Technology (RMIT) led to a massive increase in computational power. From an experimental system to within the same league as those in the Top500!

Also introduced FastX desktops and Open OnDemand for Rstudio and Jupyter notebooks for interactive visualisation during this period.




Spartan 2020

Petascale Campus Initiative (PCI) has been a notable driver with professional team at Research Compute Services (RCS). Extensive training programme; 123 researcher training days enrolment from 2012-2015 to 1321+ for 2016-2020. Completed jobs has from 1,434,474 (2010 to 2016, Edward) to 20,462,064 (2016 to 2020, Spartan). User and projects from 886 and 371 (Edward) to 3,506 and 1,402 (Spartan).

Following 2020 upgrade Spartan has: 162 nodes across, peak performance of 1698 Teraflops (would be #283 in the Top500), 5,100 physical cores, file system changed from CephFS to IBM Spectrum Scale, 2.3PB storage with an additional 550TB flash, various network upgrades.

Spartan's architecture is a proven success. It provides performance, flexibility, reliability, and can continue to grow as needed, even to the extent that all components can be replaced – like the ship of Theseus. It will always have something to offer.



References

- [1] Keynes, Simon. "Rulers of the English, c.450–1066". In Michael Lapidge; John Blair; Simon Keynes; Donald Scragg (eds.). The Blackwell Encyclopedia of Anglo-Saxon England. Blackwell Publishing, p. 514, 2001
- [2] Lafayette, Lev., Tosello, Daniel., Mead, Bernard (2016). Spartan: Performance and Flexibility : An HPC-Cloud Chimera. Presentation to eResearchAustralasia 2016
- [3] Solihin, Yan. Fundamentals of Parallel Architecture. CRC Press, 2016
- [4] Lafayette, Lev., Sauter, Greg., Vu, Linh, Meade, Bernard (2016). Spartan Performance and Flexibility: An HPC-Cloud Chimera, Presentation to OpenStack Summit, Barcelona, October 27, 2016. Youtube video also available, doi.org/10.4225/49/58ead90dceaaa
- [5] Telfer, Stig. The Crossroads of Cloud and HPC: OpenStack for Scientific Research, Open Stack, 2016
- [6] Lafayette, Lev., Wilcox, Mark., Turnbull, Mitch., Treml, Eric A. Performance Improvements with GPUs for Marine Biodiversity, Presentation to EResearchAustralasia October 2018
- [7] Lafayette, Lev. Exploring Issues in Event-Based HPC Cloudbursting. Presentation to HPC Advisory Council, Fremantle, August 2018
- [8] Melbourne Data Analytics Platform, Shaping the future of data-intensive research 2019-2020, University of Melbourne
- [9] Lafayette, Lev., Software Tools Compared To User Education in High Performance Computing. Proceedings of The Higher Education Agenda conference, Gold Coast, May, 2015
- [10] Joseph, Earl., et al., Creating Economic Models Showing the Relationship Between Investments in HPC and the Resulting Financial ROI and Innovation — and How It Can Impact a Nation's Competitiveness and Innovation, IDC Special Study, 2013
- [11] Liu, J. et al., "Evaluation of HPC Application I/O on Object Storage Systems," IEEE/ACM 3rd International Workshop on Parallel Data Storage & Data Intensive Scalable Computing Systems (PDSW-DISCS), Dallas, TX, USA, 2018, pp. 24-34, 2018 doi: 10.1109/PDSW-DISCS.2018.00005.
- [12] Plutarch, (trans. John Dryden) Lycurgus, The Internet Classics Archive, Written 75 C.E.



THANKS FOR WATCHING



& LISTENING PATIENTLY