HPC Training Generates HPC Usage

by Lev Lafayette, HPC Service Team Leader, University of Melbourne
lev.lafayette@unimelb.edu.au

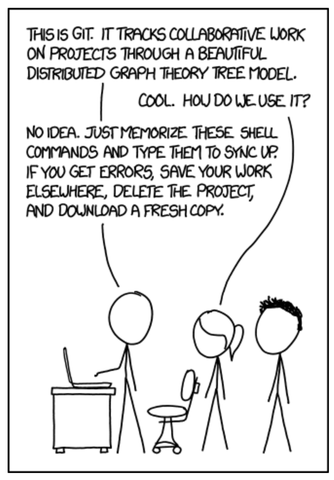

The Problem with HPC and Training

Dataset and problem complexity continues to grow at a rate faster that the capability of personal computing devices. Despite clock speed, parallel applications, multithreading, multicore, GPUs, etc., there are fundamental physical limits to making computers transfer and process data (Markov, 2014). As a result, increasing numbers of researchers look at High Performance Computing (HPC) as a preferred solution for their computational needs where performance is important (Fox, 2015). HPC is the best tool for processing large and complex datasets.

For many researchers, the HPC environment is significantly different; they don't have the core competencies to use the system effectively due to a lack of experience. Usage skills are rarely taught within the curriculum. HPC centre providers usually don't have teaching and training as part of their duties and, when they do add it, they don't have people who are HPC engineers, sysadmins, *and* educators.

Evidence of HPC Success

There's good evidence that the provision of HPC systems correlates with research output (Apon et al, 2010).

There's even better evidence that HPC generates excellent return-on-investment (Joseph, et al, 2013); about $44 in new income or cost-savings per $1 invested. Although this is almost entirely as externalities as an infrastructure investment.

Observed Experience of HPC Training

Change in cluster at VPAC/V3 Alliance saw significant increase in usage measured in CPUs hours from both partner institutions; but the one which make use of HPC training improved *much* more (Lafayette, 2015).

Usage (CPU Hours) to December 31st 2014

Year	RMIT	La Trobe	HPC System
2012	1,729,837h	1,719,554h	Tango
2013	8,108,695h	3,301,052h	Trifid
2014	9,760,919h	4,964,297h	Trifid

Course Enrolments from January 2013 to December 31st 2014: RMIT enrolments 229 La Trobe enrolments 29

The University of Melbourne had a general purpose cluster called Edward from 2010 to 2015, which had 123 researcher training days enrolment from 2012-2015. With the introduction of the new system Spartan, there was 1321 researcher training days for 2016-2020. Edward completed a total of 1,434,474 jobs from 886 users across 371 projects during its existence, including 375,000 in its final year. Spartan completed over 1,000.000 jobs in its *first* year. It's now up to 56,063,814 jobs at the time of writing from approximately 2,500 researcher training days with 6,217 users across 2,143 projects.

In 2023, unique identifiers for 263 users who received HPC training the previous year. Researchers may enrol in multiple courses and may return for revision; all these users are counted once only. From the unique users a total of 212 usernames could be determined from email addresses. There were 97 users who established an account but did not use Spartan (compute hours = 0). Of the remaining 115 users the total of job hours was determined from trained users was 6,280,454, after they received training.

The total allocated hours of cluster utilisation for the year was 11,597,951; at least 54.14% of cluster utilisation measured by job submission was conducted by users after receiving training.

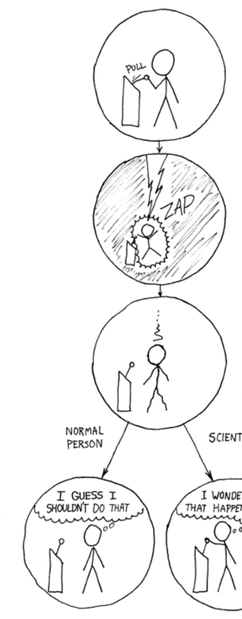

HPC Training at the UniMelb

Almost every month of the year, HPC training workshops are run by Research Computing Services. These follow the same format; 10am to 3pm, originally face-to-face, online during the COVID years, and with blended content coming this year. Class sizes are deliberately kept quite modest to allow participation from researchers, which is actively encouraged (the training workshops are about solving the researcher's problems). Formative spot questions are suggested throughout the presentation.

All attendees are provided with a copy of the video and chat from the session and access to the training content. There are 140 application example directories for job submission, with major applications have multiple job submission examples (e.g., Python, R, Octave, MATLAB, Ansys).

The most popular workshops are Introduction to Linux and HPC, Advanced Linux and Shell Scripting for HPC. Others include High Performance and Parallel Python, Parallel Processing and Programming (OpenMPI, MPI), Regular Expressions on HPC, HPC and Databases, HPC Mathematical Applications and Programming, and GPGPU Applications and Programming. Specific disciplinary groups receive tailored workshops with examples from their discipline (e.g., Bioinformatics, Neuroscience, Mechanical Engineering).

Course content is aligned with the examinable content of the HPC Certification Forum (Kunkel, et al, 2020).

Further, we have an "Spartan Champions" programme, for skilled users to help train other researchers in their team – they receive access to a high-priority queue and priority access to help-desk support.

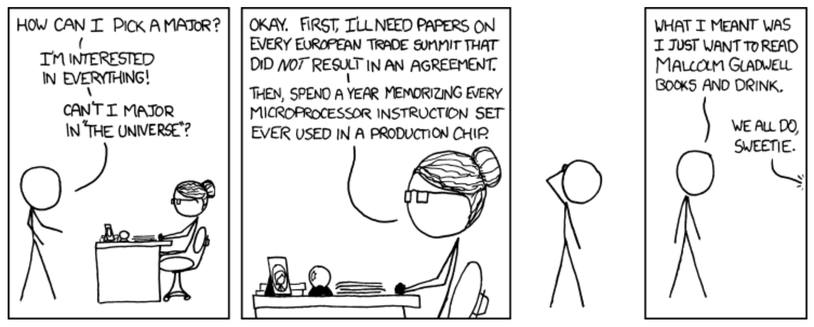

Training Approach

Learners benefit when content is taught in a structured manner (i.e., the content has related learning outcomes and is embedded in a conceptual structure) and with scaffolding (new content has a competency dependency on previous content). Content should included grounded theory and practical examples from which the learner can understand (rather than just repeat) and elaborate to new examples; this is superior to providing software tools and interfaces that attempt to by-pass comprehension. Content should be interesting, amusing, and motivational at the point of delivery and relatively terse for future asynchronous reference. Update reference material regularly! Use of formative assessment, MCQs, spot questions etc to provide explanations to answers. Summative assessment will test concepts.

Note the continuum between pedagogy and andragogy with differences in independence, experience, and time orientation. Adult learners are more prone to increasing intrinsic motivations, rather than extrinsic ones. Extrinsic motivations include fear of punishment, loss or gain of social status and/or wealth, whereas intrinsic motivations include self-identity or value for its own sake. Adult learner are social equals with their own experiences and problems that they bring to the learning environment; their learning is best achieved when the content and teaching relates directly to their experience and the challenges they face. Collaborative and cooperative learning among the learners should be encouraged to help form an organic community of practise. Always seek feedback mechanisms and invite ideas for new content from learners.

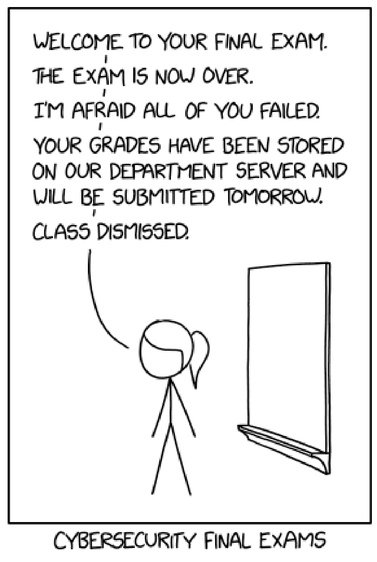

Further Projects for 2024-2025

Educators, being physical people, don't scale. They are subject to "cost disease of the service sector", and especially so when there is increasing numbers of researchers who want to learn the environment. Definite need to further develop and expand "Spartan Champions" programme.

Rather than increasing the number of training sessions to cope with the increased demand, stronger orientation this year in producing smaller and finer-grained structured and dependent components in a more permanent form (written and video) *prior* to researchers taking workshops; easier to update, converts workshops into a more interactive experience. Further tighten curriculum alignment with international HPC Certification Forum.

More Reading!

Apon, Amy, et. al., High Performance Computing Instrumentation and Research Productivity in U.S. Universities, Journal of Information Technology Impact, Vol 10, No 2, pp87-98, (2010)

Fox, Geoffrey, et al. "Big data, simulations and hpc convergence." Big Data Benchmarking. Springer, Cham, 3-17. (2015)

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, October (2013)

Kunkel, Julian, et al. "One year HPC certification forum in retrospective." Journal of Computational Science 11.1 (2020).

Lafayette, Lev. "Software Tools Compared To User Education in High Performance Computing", The Higher Education Technology Agenda Conference Proceedings (2015)

Markov, Igor L. "Limits on fundamental limits to computation." Nature 512.7513: 147. (2014)

Cartoons from Randall Monroe of XKCD (xkcd.org), of course!