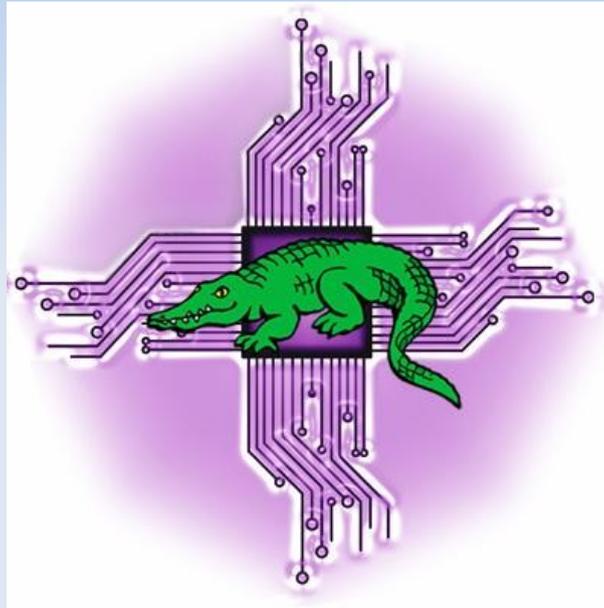


# High Performance Computing in Europe: A Selection



**Linux Users of Victoria, December 6, 2016**

[lev@levlafayette.com](mailto:lev@levlafayette.com)

# Supercomputers and Metrics

- **"Supercomputer" means any single computer system that has exceptional processing power for its time.**
- **High Performance Computer (HPC) is any computer system whose architecture allows for above average performance.**
- **Clustered computing is when two or more computers serve a single resource; this improves performance and provides redundancy in case of failure system. Typically commodity systems with a high-speed local network.**
- **One popular metric (LINPACK) is the number of floating point operations per second (FLOPS) such a system can carry out (<http://top500.org>).**
- **The HPC Challenge, in comparison, uses seven tests which cover floating point calculation speed, matrix calculations, sustainable memory bandwidth, paired processor communications, random memory updates, discrete Fourier transforms, and communication bandwidth and latency.**

# The Dominance of Linux

- In November 2016, Linux systems made up 498 of the top 500 (99.6%) systems, various forms of UNIX another 2 (0.4%).
- If one goes several years to June 2009 the distribution was 88% Linux, 5.8% mixed, 4.8% UNIX, and 0.8% MS-Windows.  
In June 2004 it was 58.2% Linux, 36.4% UNIX, 2% BSD-based, and a handy 3.2% Not Available.
- One has to go back into the last century (e.g., June 1999) to where Linux was not dominant – and then it was the various proprietary UNIXes.



(image from HPC Wire of the planned Aurora 200 petaflop system)

# Why Linux?

- **GNU/Linux scales and does so with stability and efficiency. There's an explicit orientation to ensuring that the parts interact smoothly. There are extensive modules to a minimal base kernel. The code base has been well maintained and ensures flexibility with forking for different types of hardware.**
- **Critical software such as the Message Parsing Interface (MPI) and nearly all scientific programs are designed to work with GNU/Linux in the first instance and many of those come from a very long lead time in development. Linux (1991) derives from Unix (1971) and many scientific applications can trace their lineage back to Unix systems.**
- **The operating system and many applications are provided as "free and open source", Optimisation (e.g., configuring against particular compilers and processors) leads to significant performance improvements in some computational heavy applications.**
- **GNU/Linux is built on the command line first which provides a great deal more power and is very resource efficient.**

# Spartan and Europe

- **The University of Melbourne was accepted to present at OpenStack Barcelona on their new HPC-Cloud hybrid system, Spartan (Spartan: An HPC-Cloud Hybrid. Presentation to Linux Users of Victoria, September 6, 2016, <http://levlafayette.com/files/2016spartan-luv.pdf>). An opportunity presented itself to visit other HPC centres and review their architecture and training programmes.**
- **A visit was conducted of some of Europe's major HPC centres including, Goethe University Frankfurt, University of Stuttgart, Albert-Ludwigs-University Freiburg, the European Organization for Nuclear Research (CERN), Centre Informatique National de l'Enseignement Supérieur in Montpellier, and the Centro Nacional de Supercomputación, Barcelona - as well as the presentation to the OpenStack Summit.**



# Goethe University Frankfurt

- **The Center for Scientific Computing (CSC) based at the Riedberg campus of the Goethe University Frankfurt currently operates two Linux-based computer clusters FUCHS, and LOEWE-CSC.**
- **FUCHS has 14 air-cooled and 5 water-cooled racks using AMD Opteron (Istanbul and Magny-Cours) with 39956 cores total, mixed 4X DDR-QDR InfiniBand fabric and, a parallel scratch file system with an aggregated bandwidth of 6 GB/s and a capacity of 600 TB. 41 TFlops peak performance.**
- **LOEWE was installed in 2010 consists of 825 compute nodes (169344 cores total) in 36 water-cooled racks, across two server rooms, with QDR and FDR InfiniBand interconnects, and a parallel scratch filesystem with a capacity of 764 TB and an aggregated bandwidth of 10 GB/s. CPUs performance: 226 TFlop/s (peak), GPUs performance: 597 TFlop/s (peak)**
- **CSC Frankfurt offers every second month training course in using LOEWE and FUCHS. These courses cover introductory UNIX, shell scripting, cluster computing and Slurm, and Python.**

# Goethe University Frankfurt



(LOWE: image from CSC Frankfurt University)

# Stuttgart University

- **The High Performance Computing Center of Stuttgart University (HLRS) is Germany's first national HPC centre, established in 1996. In addition to being the home of Hazel Hen, the world's #9 machine (185,088 cores, Intel Xeon CPU E5-2680s, 7420 TFlops, ~10 PB of disk), there is also a 64 nodes Cray Urika-GX, A 64-node NEC SX-ACE system for vector processing, and the 751 node Laki and Laki 2 heterogeneous NEC systems.**
- **The HLRS also has a Department of Philosophy of Science & Technology of Computer Simulation. It is an interdisciplinary research group of postgraduates in philosophy and scientific disciplines who are concerned with the human relationship with computer simulations especially with regard to policy, trust, and expectations.**
- **The HLRS has an extensive programme in education and training for high performance computing; this includes Fortran, Computational Fluid Dynamics, Linear Solvers and Parallelisation, Optimization Workshops, and a week-long course in parallel programming taught by a leading member of the MPI forum.**

# Stuttgart University



# Stuttgart University



# Freiburg University

- **The Albert-Ludwigs-Universität Freiburg is home for an HPC centre which specializes in Neuroscience, Elementary Particle Physics and Microsystems Engineering giving the name to the system: NEMO. NEMO is a CentOS 7 cluster with has 756 compute nodes (plus a few others) using Intel Broadwell for a total of 526 TFlops of performance and OmniPath interconnect. They run training courses in Introduction to Linux and HPC.**
- **The most interesting thing about NEMO is that they make use of a HPC/Cloud hybrid - like the University of Melbourne - but from a different approach. At Freiburg, they allow user defined cloud virtual machines to run on compute nodes with workflows which use MOAB/TORQUE and OpenStack - either as job submission with or without pbs\_mom on the VM or via the OpenStack dashboard. Comparing this approach with that of the University of Melbourne has led to a paper being written by both institutions for the International Supercomputing Conference in Frankfurt, June, 2017.**

# Freiburg University

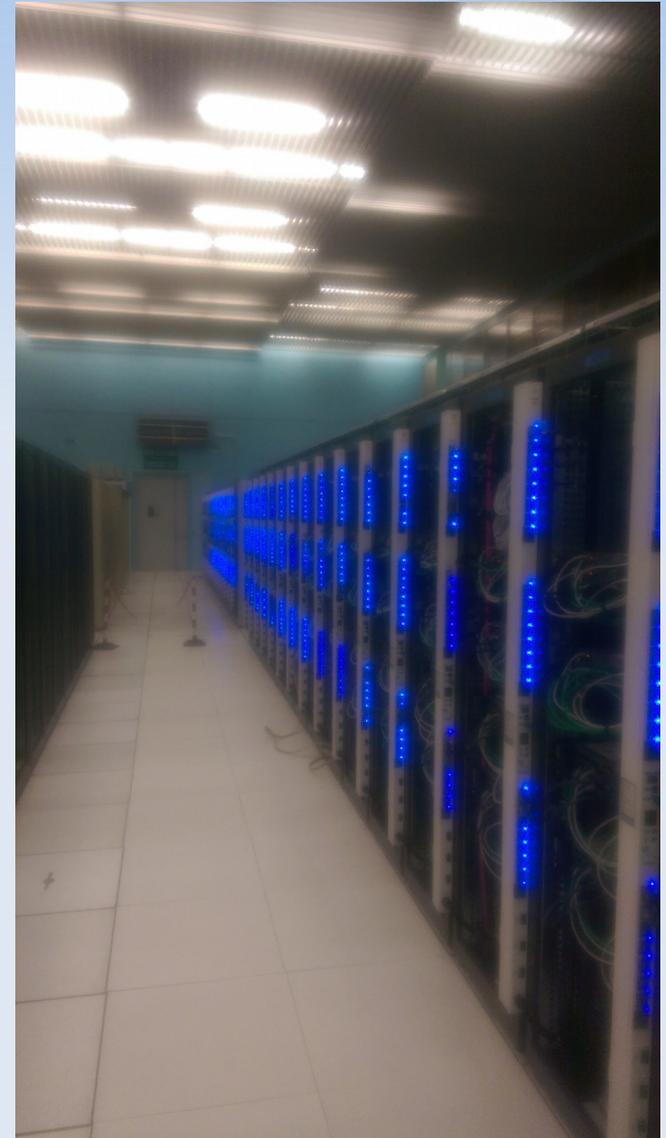


# CERN

- **Large data centre (Meyrin data centre), two 100GBe connections with Wigner RCP in Budapest. Heavy use of cloud compute operations using OpenStack. 190K cores in production under OpenStack, >90% of computer resources are virtualised. An additional >100K cores to be added in the next six months. Expecting to record 400PB/year by 2023 with the High Luminosity LHC upgrade (Facebook is currently about 180PB, Google's search DB about 100PB).**
- **For computer services, universal resource provisioning layer for bare metal, containers and VMs. Opensource HTCondor as the single end user interface, moving from proprietary LSF. 113k CPU cores and increasing for high-throughput batch service with fair-share policies. HPC is used for MPI applications (Theory Lattice QCD, accelerator physics, beam simulation, plasma simulations, CFD). Currently lxbatch resources (SLC6) running MPI. Planning New Theory QCD Infiniband cluster (72 Quanta E4 16 core / 64Gb, Infiniband), Dedicated (recent) clusters for TE plasma simulations (16 core Quanta) and Windows Engineering HPC service (Ansys, Comsol etc, IT-CDA)**

# CERN

- **Scientific Linux 5 and 6 are still supported; SL 6.8 released May 2016 and SL5 security updates to 2019. CERN Centos 7 has been introduced since April 2016; some additional CERN CentOS software. About 800 licenses for RHEL. Staged updates and internal snapshots for CERN Centos with a workflow integrated with ticketing / gitlab and QA testing of repositories enabled before production.**
- **<http://linux.web.cern.ch/linux/centos7>**



# Centre Informatique National de l'Enseignement Supérieur

- CINES is a French national facility based in Montpellier, employs c60 people. Created in 1999. Two main systems; Cristal for pre/post processing (13.1 TFLOPS, 224 core plus GPGPUs), and Occigen (2.1 Pflops, 50,554 Intel Haswell, Infiniband, Lustre for scratch and Panasas for home). Expanding to include another 1260 nodes and 35280 cores with Intel Broadwell. Project applications are peer-reviewed; 470 millions CPU hours produced in 2015.
- Use of Slurm for workload management. Users are required to submit time, nb\_nodes or nb\_tasks. Usual submission information (e.g., partition) is not set by users but rather by job\_submit plugin. Emphasis is orientated towards large (2400-16800 cores), short (0 to 24 hour) jobs. System is highly available (98.7% in February 2016) and busy (91.6% utilisation).
- Development of virtual quotas; in-house development for automatic management of storage spaces. Check usage rate (number of files, volume, age, etc). Integrated with SLURM; in case of overflow running jobs can complete, new job submissions are rejected. Future implementation with check usage quality and automated file deletion.

# Centre Informatique National de l'Enseignement Supérieur



# Centro Nacional de Supercomputacion

- **The Centro Nacional de Supercomputacion is located at Barcelona and is home to MareNostrum (1.1 Pflops, 48,896 Intel Sandy Bridge processors in 3,056 nodes, including 84 Xeon Phi 5110P in 42 nodes, and 2 PB of storage; 29th system in the top500 in June 2013). The also have Mino Tauro, a heterogeneous GPGPU cluster.**
- **The BSC/CNS has an extensive PhD and Masters programs (with the Polytechnic University of Catalonia), internship, and diverse training programme with PRACE, including programming, performance analysis, data analytics, and HPC systems administration.**
- **MareNostrum is not the most powerful system in the world, but it is the most beautiful. It is housed in the Chapel Torre Girona, a 19th century deconsecrated church.**

# Research Compute in Europe

- In relative terms the European countries visited take research computing far more seriously than we do in Australia. Top500 systems as of Nov 2016 AU 3 for RMax 2,791,709GFlops, FR 4 for 25,398,803, CH 4 for 12,273,082, DE 31 for 36,501,435 (ES is an exception, 1 for 925,058 GFlops).
- They do have issues regarding funding of the facilities with multinational, state, regional, and local contributors.

*(photo by Martin Paulo)*



# Thanks

- **Special thanks are given to the members of the various European facilities who took their time to accommodate my visit and provide tours of their facilities.**
- **This includes (my deepest apologies for names I've overlooked!):  
Hans-Christian Jankowiak, David Palao, and Anja Gerbes at CSC Frankfurt University.  
Andreas Kaminski, Alena Wackerbarth, Nico Formanek, Michael Herrmann, Bastian Koller, Rolf Rabenseifner, Ludger Benighaus, Jutta Oexle, Martin Bernreuther, Heinz Pöhlmann, Thomas Beisel, and Bernd Krischok at the HLRS Stuttgart University  
Bernd Wiebelt, Michael Janczyk, Dirk von Suchodoletz, and Konrad Meier at Freiburg University  
Tim Bell, Gavin McCance, Caroline Lindqvist, Philippe Ganz, Nils Høimyr at CERN  
G rard Gil, Jean-Christophe Penalva, Johanne Charpentier, Nicole Audiffren at CINES  
Vassil Alexandrov, Maria-Ribera Sancho, Fabrizio Gagliardi, Javier A. Espinosa Oviedo at the Barcelona Supercomputing Centre**

**THANKS FOR WATCHING**



**& LISTENING PATIENTLY**