



EasyBuild site presentations

Kenneth Hoste, HPC-UGent/VSC (kenneth.hoste@ugent.be)

Alan O'Cais, JSC (a.ocais@fz-juelich.de)

Petar Forai, IMP/Vienna Biocenter (petar.forai@imp.ac.at)

Pablo Escobar, sciCORE/UniBas (pablo.escobarlopez@unibas.ch)

http://users.ugent.be/~kehoste/EasyBuild-site-pres_20150209.pdf

EasyBuild hackathon @ University of Basel
20150209



HPC-UGent: general info



<http://ugent.be/hpc>

- HPC team at central IT dept. of Ghent University (Belgium)
- 9 team members: 1 manager, ~3 user support, ~5 sysadmin
- ~1,500 user accounts, across all scientific domains
- tasks: hardware, system administration, user support/training, ...

- member of Flemish Supercomputer Centre (VSC)
virtual centre, collaboration between Flemish university associations



HPC-UGent: systems

one Tier1 cluster ('muk')

- 528 nodes (8,448 cores), Intel Sandybridge, FDR Infiniband, SL 6.6

6(+2) Tier2 clusters

- 3 batch clusters: Ethernet, SL 6.6
 - *gastly*: 42 nodes (336 cores), Intel Nehalem
 - *haunter*: 150 nodes (1,200 cores), Intel Nehalem
 - *raichu*: 64 nodes (1,024 cores), Intel SandyBridge
- 2 MPI clusters: Infiniband, SL 6.6
 - *gulpin*: 32 nodes (1,024 cores), AMD Magny Cours, 2x QDR IB
 - *delcatty*: 160 nodes (2,560 cores), Intel SandyBridge, FDR IB
- 1 ScaleMP vSMP system, (2x QDR Infiniband), SL 6.6
 - *dugtrio*: 16 nodes ($3 \times 48 + 2 \times 24$ cores), Intel Westmere
- 2 pre-production systems, Infiniband, CentOS 7
 - *golett*: 200 nodes (4,800 cores), Intel Haswell, QDR IB
 - *phanpy*: 16 nodes, Intel Haswell, dual FDR IB, 512GB RAM + SSDs



HPC-UGent: EasyBuild configuration

- one dedicated install user, shared between multiple people : (
- flat module tree, default EasyBuild module naming scheme
- Tcl/C environment modules v3.2.10 as modules tool
 - pilot setup with Lmod v5.8.6 since last week
- Tier2: one shared filesystem for software + modules, NFS mounted
 - subdirectory per cluster, e.g., /apps/gent/SL6/sandybridge
 - software + modules for Tier1 system on 'local' GPFS filesystem
- installations are done by submitting jobs on the cluster
 - `module swap cluster/foo; eb ... --job --robot`
- only install easyconfigs that have been merged into develop branch
- installed software is retained throughout lifetime of the system



'Common User Experience' VSC project

- collaboration between HPC sites @ Flemish universities
 - Ghent, Leuven, Antwerp, Brussels, Hasselt
- intention is to allow users to easily switch to a different VSC site
- important aspect: same module naming scheme
 - currently: default EasyBuild module naming scheme
- common toolchains: intel/2015a, foss/2015a
 - updated every 6 months (Dec/Jan, June/July)
 - different alternative compositions get evaluated/tested
 - other members of EasyBuild community also starting to pick these up
- minimal common divisor: toolchains, Python, Perl (bare), R (bare)
- sharing of easyconfig files, (enhancements to) easyblocks
 - via pull requests to main repository @ GitHub
- (my) intention is to also get Lmod included 'soon'



JSC: general info



<http://www.fz-juelich.de/ias/jsc/EN>

- Jülich Supercomputer Centre (JSC) has ~200 staff members
- hardware and system administration
 - division : High Performance Computing Systems with 16 staff
- user support and training
 - division : Application Support with 39 staff
- member of Gauss Centre for Supercomputing (GCS)
virtual centre, combines the three national supercomputing centres in Germany



JSC: systems

Two Tier0 systems

- JUQUEEN - BG/Q
 - 28 racks (7 rows x 4 racks) - 28,672 nodes (458,752 cores)
 - main memory: 448 TB
 - overall peak performance: 5.9 PFLOPS (Linpack: 5.0 PFLOPS)
 - I/O Nodes: 248 (27x8 + 1x32) connected to 2 CISCO switches
- JURECA - cluster architecture (replaces JUROPA)
 - projected peak performance of about 1.8 PFLOPS
 - Intel Xeon E5-2680 v3 Haswell, 24 cores per node
 - latest InfiniBand with full fat tree

Number of additional research systems

- QPACE - hot cooling, Cell processors
- JUDGE - GPU cluster
- DEEP - cluster + booster



JSC: EasyBuild configuration

- shared installation group software with ~ 25 members
 - use GID permissions to control
- hierarchical module tree, default EasyBuild module naming scheme (currently)
- Lmod as modules tool
- one shared filesystem for software + modules
 - subdirectory per cluster
- installations are done by submitting jobs on the cluster
- software is provided in stages
 - new module hierarchy or features are built in new stage
 - dedicated user for rebuilds
 - deprecated toolchains and software versions are retired
 - soft links \implies instant transition
 - previous stages can be loaded (installed software is retained)



Vienna Biocenter: general info



- Vienna Biocenter is a molecular biology cluster
- consists of 4 research institutes and ~30 biotech companies
- research institutes: IMP, IMBA, MFPL, and GMI
- ~450 scientists at IMP and IMBA, ~100 at GMI, ~700 at MFPL
- 1 person responsible for HPC infrastructure at GMI
- 2 persons responsible for HPC infrastructure at IMP/IMBA
- no collaboration on HPC level



Vienna Biocenter: systems

Not really high performance computing shops.

Almost all of workload is data analysis with R, Python, Perl, Java, etc.

GMI

- 1,600 cores Intel Sandy Bridge, IB FDR Torus, SLES11, Lustre
- 128 cores Intel Sandy Bridge, NUMALink6, SLES11, Lustre

IMP/IMBA

- 600 cores Intel SandyBridge, 1GbE and 10GbE, Debian7, 2PB Isilon



Vienna Biocenter: EasyBuild configuration

GMI

- set up the infrastructure there in 2011
- EasyBuild deployed since 1.0 in 2012
- default EasyBuild naming scheme, one install user
- Lmod in use since 2012, flat module tree
- software and modules are on NetApp NFS cluster

IMP & IMBA

- inherited the setup in Dec 2014
- Tcl/C environment modules v3.2.10
- one install user, flat module tree
- no EasyBuild yet :(
- projects to deploy EasyBuild, Lmod, and XALT on the way



sciCORE: general info

- providing HPC facilities and support for the University of Basel
- 7 team members:
 - 1 manager
 - 3 user support/development
 - 3 sysadmin
- ~350 user accounts



sciCORE: systems

- Maia cluster:
 - ~300 compute nodes
 - Intel Sandy Bridge & Ivy Bridge
 - ~4,000 cores
 - QDR infiniband
 - ~50TB BeeGFS (Fraunhofer)
 - CentOS 6.5
- BC2 cluster:
 - ~150 compute nodes
 - Intel Sandy Bridge, AMD Bulldozer
 - ~3,000 cores
 - Ethernet
 - ~2PB GPFS
 - CentOS 6.5



sciCORE: EasyBuild configuration

- one dedicated install user, shared between multiple people
- flat module tree, EasyBuild default module naming scheme
- Lmod 5.7.4 (to be upgraded soon for performance improvements)
- default toolchains: `goolf-1.4.10` and `iomkl-2015.01`
- rely on system packages for common libraries like `zlib`, `ncurses`, `libpng`, `bzip2`, ...
 - dependencies are filtered using `"eb --filter-deps"`
- BC2 cluster: one shared filesystem (GPFS) for software and modules
 - subdirectory per cpu type, e.g. `/import/bc2/apps/sandybridge`
 - every machine accesses the software in a common path (`/import/bc2/apps`) mapped with automounter
- Maia: software/modules is copied to local disk of each node
- compilation is done in login nodes (SGE support in EasyBuild?)

