



arm

The Arm ecosystem for HPC and server

Chris Goodyer

Going Arm : Arm HPC User Group

13th November 2017

Copyright © 2017 Arm Limited

Contents

Overview of Arm

- HPC engagements

Arm partner information

- Latest deployment information

Arm Software Ecosystem


- Software stack enabling
- Our priorities on libraries and applications,

More events this week



Arm's mission

- Deploy energy-efficient Arm-based technology, wherever computing happens...

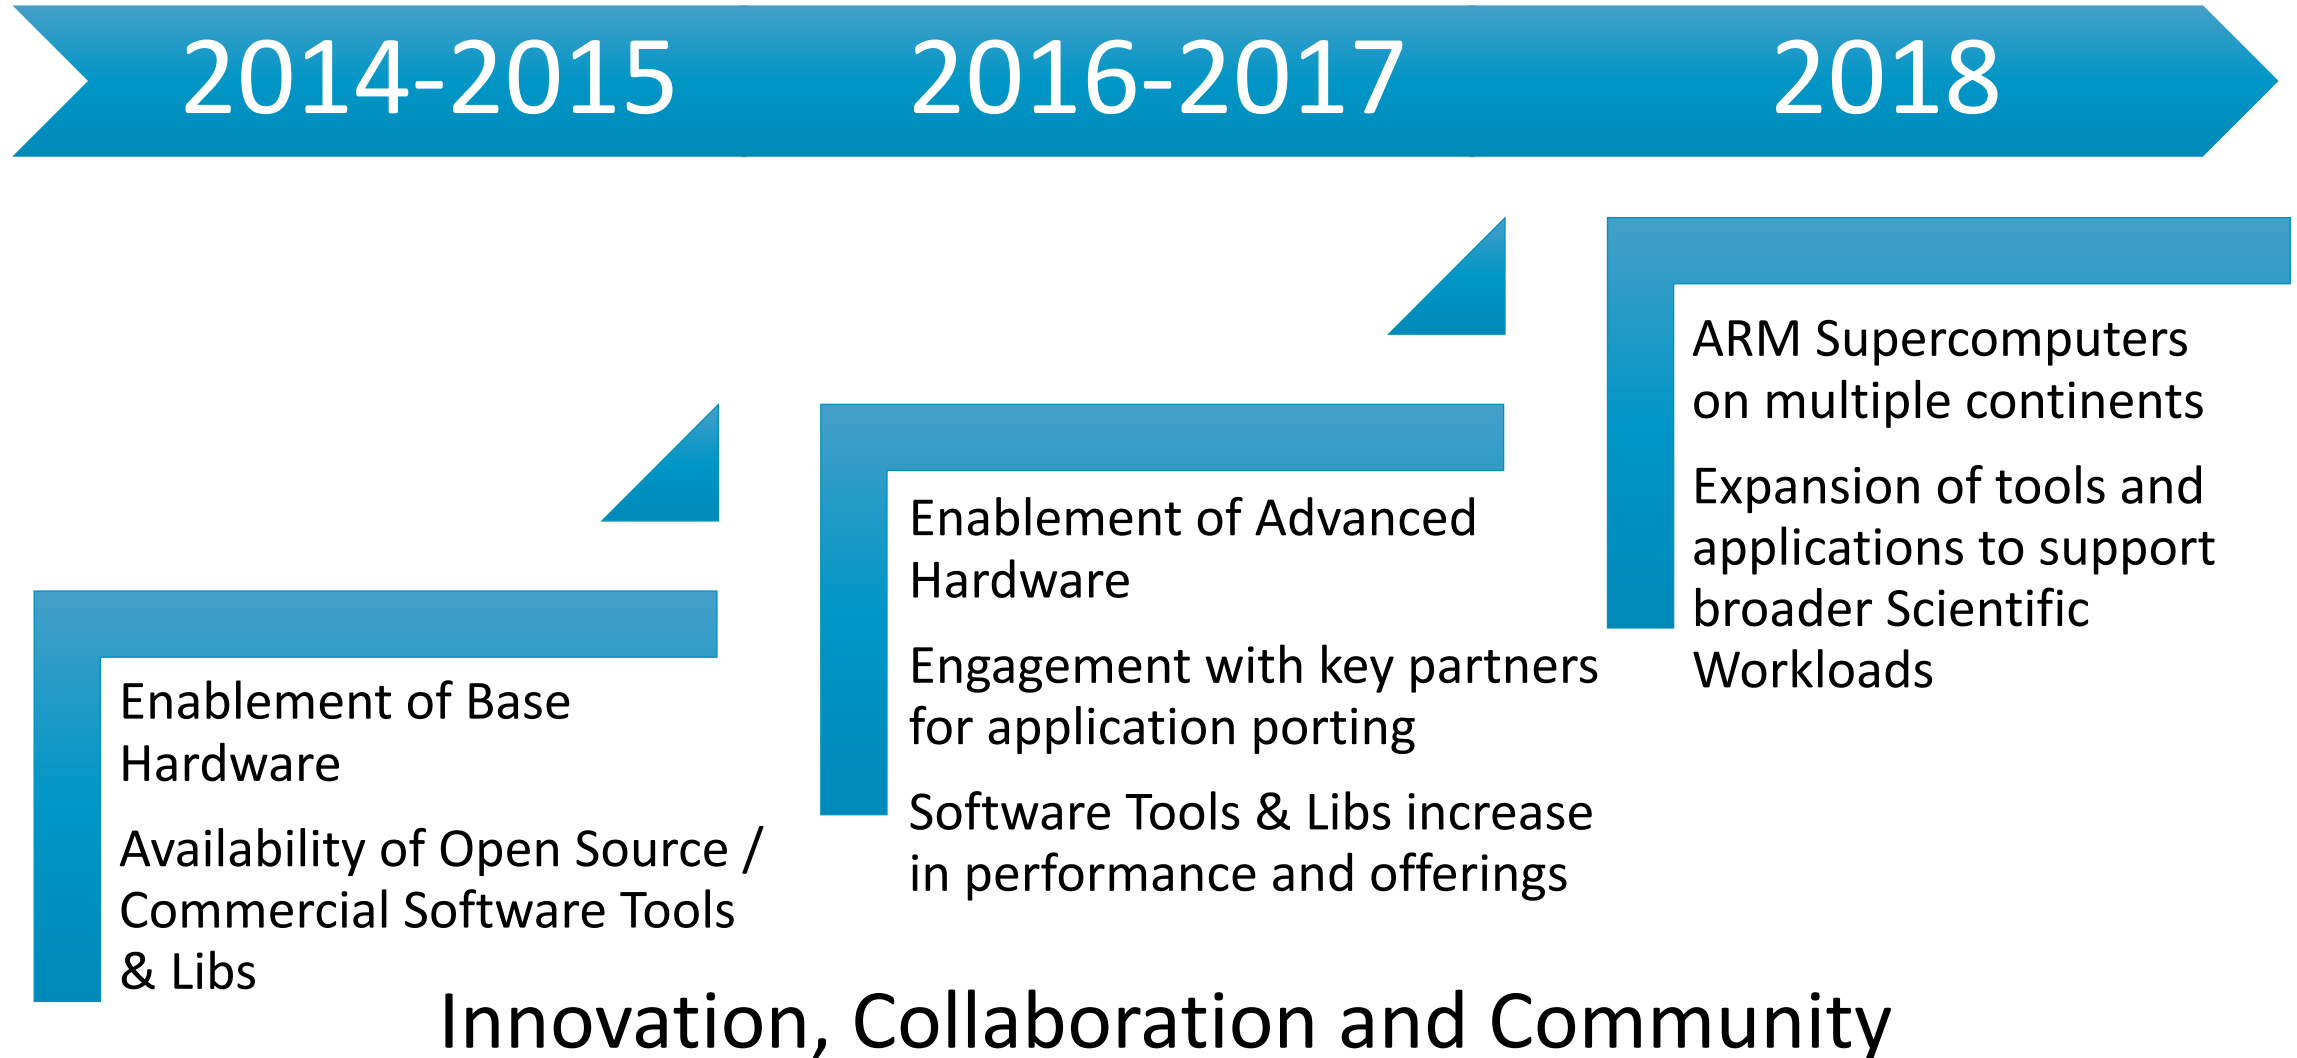


Leading in wearables and
the Internet of Things


~85% share of
laptops, tablets,
and smartphones


Partnering to deliver
data center efficiency

Roadmap – from two years ago



HPC deployments happening...






January 17th 2017

Announcing the GW4 Tier 2 HPC service, 'Isambard':
named after Isambard Kingdom Brunel

System specs:

- Cray CS-400 system
- **10,000+** ARMv8 cores
- HPC optimised software stack
- Technology comparison:
 - x86, KNL, Pascal
- To be installed March-Dec 2017
- £4.7m total project cost over 3 years



I.K.Brunel 1804-1859

Simon McIntosh Smith, simonm@cs.bris.ac.uk,
@simonmcs

5


bristol.ac.uk

SANDIA REPORT
SAND2017-6783 O
Unlimited Release
Printed June 2017

Vanguard Request for Information (RFI)


The Mont-Blanc prototype ecosystem

Prototypes are critical to accelerate software development
System software stack + applications



Mont-Blanc prototype

- 1080 compute cards
- Dual Cortex-A15
- GPU Mali-T804
- 4 GB LPDDR3



Mont-Blanc 3

- Bull Sequana™
- Cavium ThunderX2™

Mini-clusters

- Arndale

PRACE prototype

- Tibidabo
- Carma
- Pedrafor

2011

BSC KEEPS ITS HPC OPTIONS OPEN WITH MARENOSTRUM 4

December 1, 2016 Timothy Prickett Morgan



When it comes to supercomputing, you don't only have to strike while the iron is hot, you have to spend while the money is available. And that fact is what often determines the technologies that HPC centers deploy as they expand the processing and storage capacity of their systems.

A good case in point is the MareNostrum 4 hybrid cluster that

Program Collaborators

Drive open HPC ARM SW stack to be as robust as x86.

Platform (HPE)



- Develop & Deploy a stable ARM64 Platform
- Coordinate & Facilitate The Collaboration
- Sponsor Forums for Discussion / Communication
- Honest Broker

Partners



- Port/create, develop, optimize, validate & support key software components
- Collaborate with HPE and customers.
- Participate in Forums and discussion

Customers



- Provide real-world exposure to ARM platform.
 - Select and port/develop workloads/projects
 - Provide feedback and insight into results
 - Identify, implement & distribute improvements
 - Cooperate with partners to resolve problems.
- Collaborate
 - Port/develop as required.
 - Measure performance & optimize.
 - Communicate inside and outside the program



GoingARM Workshop, ARM Research Summit September 2017



Arm in HPC worldwide

- **United States**

Arm has been a participant in two Department of Energy funded Exascale projects: **Data Movement Dominates** and **Fast Forward**. Arm is now involved in at least one **Path Forward** program which includes \$270M funding to develop Exascale technologies.

- **European Union**

Through FP7 and Horizon 2020, Arm has been involved in several funded pre-Exascale projects including the **Mont Blanc** program which deployed one of the first Arm prototype HPC systems. Current projects include **Mont-Blanc 3**, **MB-2020** and **EuroEXA**.



- **Japan**

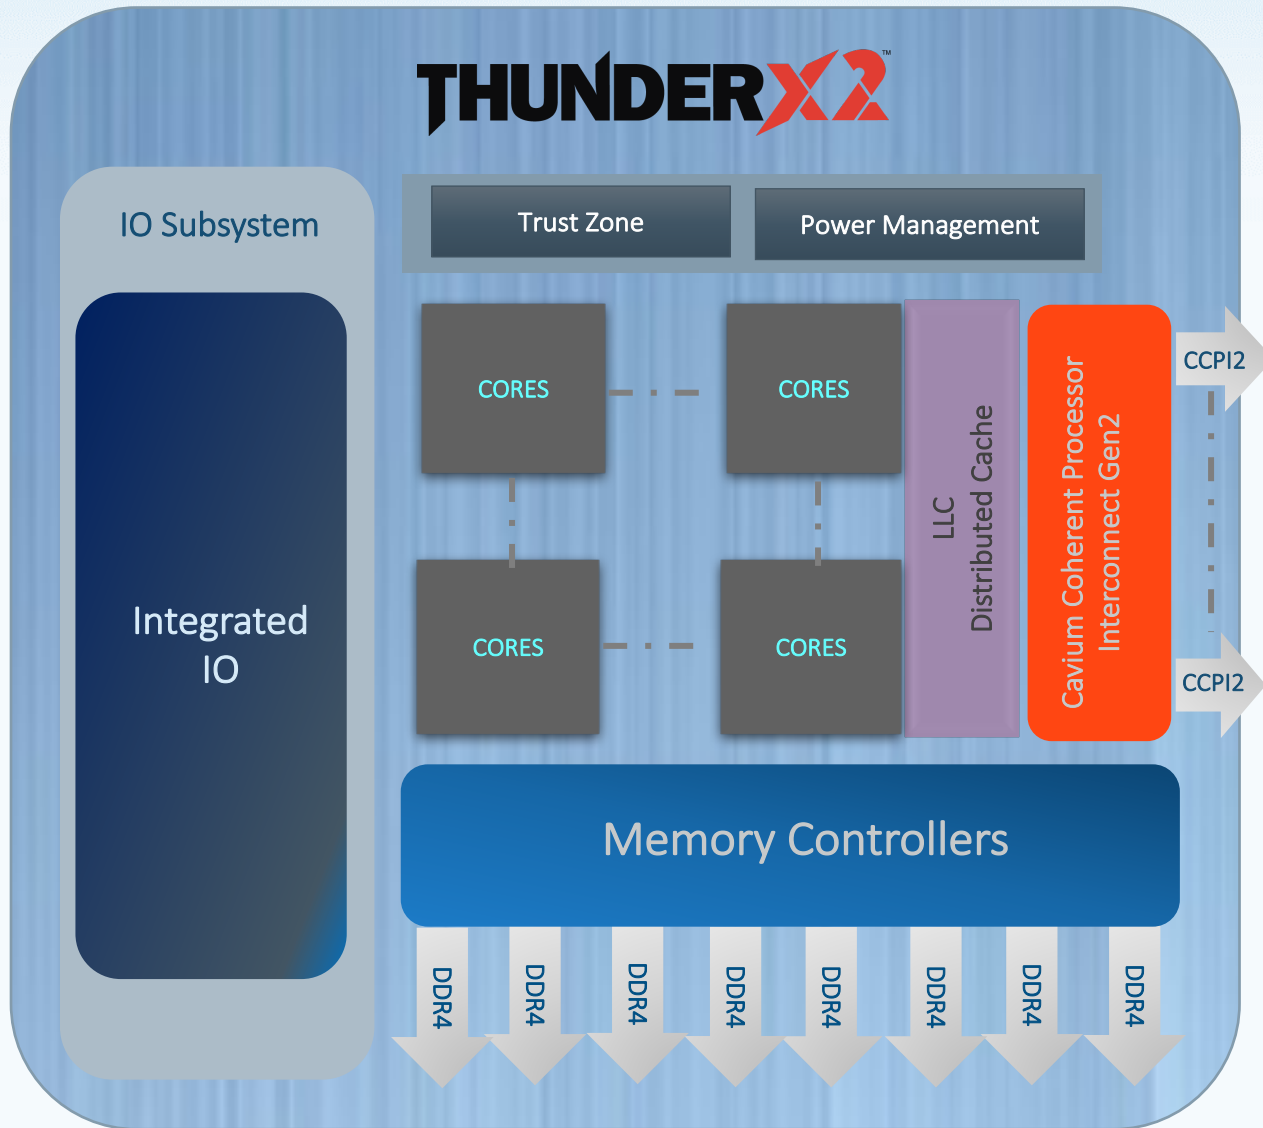
At ISC2016, Fujitsu and RIKEN announced that the **Post-K** system targeted at Exascale will be based on Armv8 and custom extensions to be described at HotChips 2016

- **China**

James Lin, vice director for the Center of HPC at Shanghai Jiao Tong University claims China will build three **pre-Exascale** prototypes to select the architecture for their Exascale system. At least one of these three are Armv8 based.

Hardware availability

Cavium CN99XX - 1st member of **THUNDERX2** Family




- 24/28/32 Custom Armv8 cores
- Fully Out-Of-Order (OOO) Execution
- 1S and 2S Configuration
- Up to 8 DDR4 Memory Controllers
- Up to 16 DIMMs per Socket
- Server Class RAS features
- Server class virtualization
- Integrated IOs
- Extensive Power Management

2nd gen Arm server SoC

Delivers **2-3X** higher performance

Qualcomm Centriq™ 2400 Processor

Accelerating Innovation in the Data Center



Performance

48-core SoC design with optimized performance for throughput data center workloads



Power

World's first and only 10nm processor for server-class performance and leading power efficiency



Total Cost of Ownership

Significant hardware CapEx and OpEx savings



Ecosystem

Robust Arm64-based ecosystem across infrastructure software and core operating system components



Security


Reliable security for the modern datacenter founded on immutable root of trust and TrustZone architecture

arm Supercomputing

...Made Possible by Cray

CRAY CATAPULTS ARM-BASED PROCESSORS INTO SUPERCOMPUTING
Cray Adds Arm Processors with Complete Software Stack to the Cray XC50 Supercomputer

Seattle, WA – November 13, 2017 – Global supercomputer leader Cray Inc. (Nasdaq: CRAY) today announced the Company is creating the world's first production-ready, Arm®-based supercomputer with the addition of Cavium (Nasdaq: CAVM) ThunderX2™ processors, based...



First purpose built Arm HPC system by the #1 HPC vendor

HPE Apollo 70 System overview



HPE Driving Multi-vendor effort to Accelerate Arm Adoption for HPC

arm



World-class HPE engineering and innovation

- Leverages proven HPE Apollo 2000 system arch
- **Management support** for HPE provisioning, cluster management and performance software
- **Tier 1 vendor services and support** (HPE Pointnext)

Leading Arm technology

- **Cavium ThunderX2** with 32 64-bit Arm v8.1 custom cores for today's HPC workloads
- **Greater density and power-efficiency** with highly integrated SoCs
- Brings **real change and choice** to HPC market

Density and scalability for HPC clusters

- **Dual Socket Nodes** – Up to 64 cores per node
- **Maximum density** - Up to 80 servers / rack
- **High speed interconnects** 100 Gb/s InfiniBand EDR
- **Architected to scale** to multi-rack levels

The Mont-Blanc prototype ecosystem

Prototypes are critical to accelerate software development
System software stack + applications

MONT-BLANC

PRACE prototypes

- Tibidabo
- Carma
- Pedraforca

Mini-clusters

- Arndale
- Odroid XU
- Odroid XU-3
- NVIDIA Jetson



Mont-Blanc prototype

- 1080 compute cards
- Dual Cortex-A15
- GPU Mali-T604
- 4 GB LPDDR3
- Up to 64 GB local storage
- USB-to-Eth network
- Fine grained power monitoring system
- Installed between Jan and May 2015
- Operational since May 2015 @ BSC



Bull
atos technologies

Mont-Blanc 3

- Bull Sequana™
- Cavium ThunderX2™



2011

2012

2013

2014

2015

2017

Software update

Open Source Compiler Highlights in 2017

Performance Improvements

- **GCC** SPEC CPU2006:
 - INT/FP+4% comparing to GCC-7
 - Selected Glibc scalar math functions significantly optimized (wrf +60%)
 - GCC loop vectorizer enhanced (hmmer +30%)
- **Glibc** single-thread optimization up to 25%-150% improvement in benchmark
- **LLVM**: About 1% performance improvement for Armv7-A

Architecture enablement

- Continuously adding new architecture support to **GCC** and **LLVM**
 - Include **Cavium ThunderX2** and **Qualcomm Falkor**
- Armv8.3-A released in GCC-7
- Armv8.4-A and SVE are targeted to be included in the GCC-8 release

Work in progress

- Resolving **GCC** vectorizer regression
- Vect-math library for Arm/AArch64, enable more vectorization opportunities
- Further loop optimizations in **GCC**
- **LLVM** GlobalSel pass enablement
- **GDB** Fortran enhancement


HPE helping Arm HPC software ecosystem mature



Engagements in HPE Comanche programme have accelerated adoption

- We have been able to assess the state of fundamental software stacks, such as MPI and NUMA capabilities ahead of Apollo 70 deployments next year
 - Collaborative work here especially great with all partners focusing on interoperability issues
 - Examples include fixing bugs with kernels, MPI drivers and OpenMP thread placement
- Engagements with collaborating labs (Argonne, LANL, LLNL, Oak Ridge, Sandia) on porting stacks has been very useful in giving confidence for future deployments
- HPC application porting and tuning work happening with recipes and improvements being made public at

<https://gitlab.com/arm-hpc/>



See the HPE Comanche rack in booth #494



Arm HPC ecosystem

Porting to Arm

Arm is engaging directly with partners and HPC scientific code developers to support porting and optimisation of common HPC libraries, tools and applications.

Initial focus on successfully building with both **Arm** and **GCC** compilers across a broad front.

Often only modest changes to environment variables, build scripts and architecture files are needed

Degree of commonality between codes



Community building

Outside the people we collaborate with, various complementary Arm HPC communities already exist:

GoingArm

Google



- Arm HPC User Group (SC) and GoingArm (ISC/ArmRS)
- Arm HPC Google Group (<https://groups.google.com/forum/arm-hpc>)
- Arm HPC GitLab pages (<https://gitlab.com/arm-hpc/>)

Encouraging our partners to use GitLab is a priority

Our app work is **engaging with code owners and users** to **get suitable test cases**, to **get Arm support built in**, and including helping them make AArch64 testing part of their development processes

Wiki

<https://gitlab.com/arm-hpc/packages/wikis/home>

Dynamic list of common HPC applications

Provides focus for porting progress

Community driven.

Maintained by Arm, but anyone can join and contribute.

Allows developers to share recipes, and learn from progress on other applications

Provides a mechanism for tracking status of applications and package sets (e.g. OpenHPC packages, Mantevo, etc.)

Up-to-date summary of package status

Package	External URL	Last Wiki Update	BuildMaturity	CompilesARMCompiler	CompilesGCC	NEONOptimized
EPOCH	http://www.ccpo.ac.uk	19/10/17 22:10:20	NeedsPatch	Yes	Yes	
SDF	https://github.com/keithbennett/SDF	20/10/17 00:13:45	NeedsPatch	Yes	Yes	
VPIC	https://github.com/lanl/vpic	19/10/17 22:10:20		Yes	Yes	
adios	http://www.oed.ornl.gov/center-projects/adios/	17/07/17 23:33:11		Yes	Yes	
arpack	http://www.caam.rice.edu/software/ARPACK/	17/07/17 23:33:11		Yes	Yes	
autoconf	http://www.gnu.org/software/autoconf/autoconf.html	01/08/17 21:48:30		Yes	Yes	
automake	http://www.gnu.org/software/automake	18/07/17 13:41:43		Yes	Yes	
bookleaf	https://uk-mac.github.io/BookLeaf/	17/07/17 23:33:11		Yes	Yes	
boost	http://www.boost.org	18/07/17 11:29:00		Yes	Yes	
ccs-qcd	https://github.com/iber-miniapp/ccs-qcd	01/08/17 21:43:40	Yes			
cloverleaf	http://uk-mac.github.io/CloverLeaf/	19/10/17 22:10:20	Upstream	Yes	Yes	
cloverleaf3d	http://uk-mac.github.io/CloverLeaf3D/	24/07/17 21:41:31	Upstream	Yes	Yes	
comjd	http://eematex.github.io/CoMD	24/07/17 21:36:31	Upstream	Yes	Yes	
dgemm	http://www.nersc.gov/research-and-development/apex/apex-bench	24/07/17 21:36:31	NeedsPatch	Yes	Yes	
fftw	http://www.cisl.isi.u-tokyo.ac.jp/riss/english/project/fftw/	24/07/17 21:36:31	Yes			
fftw	https://github.com/EFTW/fftw3	17/07/17 23:33:11	Yes	Yes	Yes	Yes
gnu-scientific-library	http://www.gnu.org/software/gsl/	18/07/17 07:08:24		Yes	Yes	
gromacs	http://www.gromacs.org/	24/07/17 21:36:31		Yes	Yes	
hdf5	http://www.hdfgroup.org	17/07/17 23:33:11		Yes	Yes	
hpc-challenge	http://icl.cs.utk.edu/hpc/index.html	24/07/17 21:36:31	Upstream	Yes	Yes	
hpcga	http://mantevo.org/downloads/HPCGA-1.0.html	24/07/17 21:36:31	Upstream	Yes	Yes	
hpcg	http://www.nersc.gov/research-and-development/apex/apex-bench	24/07/17 21:48:22	Upstream	Yes	Yes	
hydra	https://computation.llnl.gov/project/linear_solvers/software.php	17/07/17 23:33:11		Yes	Yes	
imh		17/07/17 23:33:11		Yes	Yes	
lammps	http://lammps.sandia.gov/	19/10/17 22:10:20		Yes	Yes	
libtool		18/07/17 13:41:43		Yes	Yes	
julesh	https://codesign.llnl.gov/julesh.php	17/07/17 23:33:11		Yes		
metis		17/07/17 23:33:11	Yes	Yes	Yes	
miniAero	http://mantevo.org/downloads/miniAero_1.0.html	24/07/17 21:36:31	NeedsPatch	Yes	Yes	
miniAMR	http://mantevo.org/downloads/miniAMR_1.0.html	24/07/17 21:36:31	Upstream	Yes	Yes	
miniFE	http://www.nersc.gov/users/computational-systems/cori/nersc-8	24/07/17 21:36:31	Upstream	Yes	Yes	
miniGhost	http://www.nersc.gov/users/computational-systems/cori/nersc-8	24/07/17 21:36:31	Upstream	Yes	Yes	
miniMD	http://mantevo.org/downloads/miniMD_1.2.html	24/07/17 21:36:31	NeedsPatch	Yes	Yes	
miniXyce	http://mantevo.org/downloads/miniXyce_1.0.html	24/07/17 21:36:31	Upstream	Yes	Yes	
mpich		19/10/17 22:10:20	NeedsPatch	Yes	Yes	
numa		17/07/17 23:33:11		Yes	Yes	
mvapich-2	http://mvapich.cse.ohio-state.edu	21/08/17 13:26:19	Upstream	Yes	Yes	
namd	http://www.ks.uiuc.edu/Research/namd/	24/07/17 21:36:31	NeedsPatch	Yes	Yes	

Open source libraries for helping increase performance

Arm Optimized Routines

<https://github.com/ARM-software/optimized-routines>

These routines provide high performing versions of many math.h functions

- Algorithmically better performance than standard library calls
- No loss of accuracy

SLEEF library

<https://github.com/shibatch/sleef/>

Vectorized math.h functions

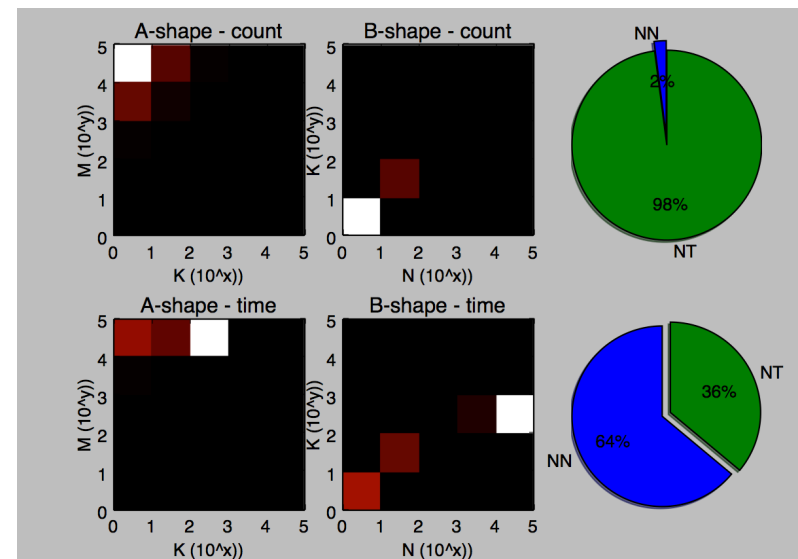
- Provided as an option for use in **Arm Compiler**

Perf-lib-tools

<https://github.com/ARM-software/perf-lib-tools>

Understanding an application's needs for BLAS, LAPACK and FFT calls

- Used in conjunction with **Arm Performance Libraries** can generate logging info to help profile applications for specific case breakdowns



Example visualization:
DGEMM cases called

arm



– Easy HPC stack deployment on Arm

OpenHPC is a community effort to provide a common, verified set of open source packages for HPC deployments

Arm and partners actively involved:

- Arm is a silver member of OpenHPC
- Linaro is on Technical Steering Committee
- Arm-based machines in the OpenHPC build infrastructure

Status: 1.3.3 release out now

- Packages built on Armv8-A for CentOS and SUSE

Functional Areas	Components include
Base OS	CentOS 7.4, SLES 12 SP3
Administrative Tools	Conman, Ganglia, Lmod, LosF, Nagios, pdsh, pdsh-mod-slurm, prun, EasyBuild, ClusterShell, mrsh, Genders, Shine, test-suite
Provisioning	Warewulf
Resource Mgmt.	SLURM, Munge
I/O Services	Lustre client (community version)
Numerical/Scientific Libraries	Boost, GSL, FFTW, Metis, PETSc, Trilinos, Hypre, SuperLU, SuperLU_Dist, Mumps, OpenBLAS, Scalapack, SLEPc, PLASMA, ptScotch
I/O Libraries	HDF5 (pHDF5), NetCDF (including C++ and Fortran interfaces), Adios
Compiler Families	GNU (gcc, g++, gfortran), LLVM
MPI Families	OpenMPI, MPICH
Development Tools	Autotools (autoconf, automake, libtool), Cmake, Valgrind, R, SciPy/NumPy, hwloc
Performance Tools	PAPI, IMB, pdttoolkit, TAU, Scalasca, Score-P, SIONLib

SC17 events where you can learn more



- **SC17 Exhibitor Forum Presentations**

- Fujitsu: *Tuesday 1:30 pm, SC room 501-502*
- Cavium: *Tuesday 2:30 pm, SC room 501-502*
- Arm: *Wednesday 2:00 pm, SC room 501-502*

- **Panel: *The Arm Software Ecosystem: Are We There Yet?***

- 1:30-3 pm, Tuesday: *SC room 201-203*
- Speakers include: nVidia, CERN, U. Bristol

- **BoF : *Arm User Experience: Testbeds and deployment at HPC centers***

- 5:15 pm, Tuesday: *SC room 701*
- Speakers include: Oak Ridge, RIKEN and Bristol

- **Arm SVE Users' Meeting : *Vectorization for all***

- 2-4 pm, Thursday: *Aspen Ballroom, Grand Hyatt Hotel*

Also look out on the exhibition floor for booths from Arm partners including:

- Arm: #1781 + #610
 - Demos from Arm HPC tools, Cavium, Fujitsu, Qualcomm, Singularity
- Cavium: #349
- Cray: #625
- Fujitsu: #1125
- HPE: #494 - Comanche

Arm HPC Workshop

BY RIKEN AICS AND LINARO

12-13TH DEC 2017 –
AKIHABARA HALL, TOKYO

linaro.co/hpc2017

Arm HPC Workshop by RIKEN AICS and Linaro

How does the Arm-Powered supercomputing future look and how can you prepare for it? The Arm HPC Workshop will bring together the leading Arm vendors, end users and the open source development community in Tokyo, to discuss the latest products, developments and open source software support.

Topics of focus are, but not limited to: Compilers including GCC, LLVM, C++, Fortran, optimisations, benchmarking and general support, OS and Runtime, Tools, Math Libraries, Orchestration, Machine Learning

For End Users: In the Arm HPC Workshop you will hear from Arm Members and their Partners, information about new trends, technologies and products for the planning and operation of an Arm-Powered supercomputer.

For Linaro Members and Partners: The Arm HPC Workshop is a two day conference to provide training and discussion panels for Arm-Powered solutions and offerings for all aspects of HPC including Server, Networking, Storage and Development.

You can expect:

- Exciting best practices and technology outlooks with peers from the Arm ecosystem
- Unique opportunities to network with current Arm users, vendors and the developer community

Sponsored by



Arm Data Center Day

- Cost (Free)
- Sponsorship options available
- Session slots available (Submit your sessions on linaro.co/hpc2017)

Register here

linaro.co/hpc2017



arm

The Arm trademarks featured in this presentation are registered trademarks or trademarks of Arm Limited (or its subsidiaries) in the US and/or elsewhere. All rights reserved. All other marks featured may be trademarks of their respective owners.

www.arm.com/company/policies/trademarks