

# Agenda of the break-out session on Software

---

## Yesterday

- ❑ Worked on tables listing software that are supposed to contribute to future Exascale stacks.
- ❑ Start from the tables of software provided by all regions for Exascale: Japan, Europe, USA, China, Russia
- ❑ Review of the components by Region and ask leaders to update their table
- ❑ Refined the definition (interregional instead of international)
- ❑ Long discussion about the classification
- ❑ Asked leader to update their table over the night with the new code and with mentioning already known collaborators

## Today

- ❑ Improved (more details) tables adding collaborators and funding body
- ❑ → built a table that could be used as an instrument by funding agencies

# Color Codes and Conventions

---

- In Column 1
  - Orange = Developers provide support
  - Green = Developers collaborate with vendor for vendor support
- Everywhere: Underline anything that is NOT open source
- Classes of Interregional Relations
  - C0 = Regional now but open to interregional
  - C1 = Cooperation (loose sharing, no interregional fund)
  - C2 = Collaboration (interregional funded joint effort)
  - C3 = Co-development with concrete software deliverable
  - C? = don't know at this point

# Stack Components: Japan (updated April 13)

Group provides in production quality	Experimental versions only	No major research effort—rely on vendors or other countries
<ul style="list-style-type: none"> <li>- Programming language <i>XScalableMP</i> (UTsukuba-Riken, Fujitsu) (C2, INRIA, JST-ANR), and accelerator extension XMP-dev (U. Tsukuba) (C2, INRIA, JST-ANR)</li> <li>- OS kernel for many core architecture MIC, node level scheduling, including PGAS style (UTokyo-Riken, Fujitsu, NEC, Hitachi, PCCC) (C0)</li> <li>- FT-MPI (C2, Riken, ORNL, JST)</li> <li>- File system-Lustre Enhancement (Fujitsu) (C3, OpenSFS* &amp; WhamCloud*, K-funding)</li> <li>- File system-GFARM (UTsukuba) : (C2, INRIA, JST-ANR)</li> <li>- Communication libraries (below MPI layer) (UTokyo, Fujitsu, NEC, PCCC) (C0)</li> <li>- Scripting language for parallel workflow-Xcrypt (Kyoto U., Fujitsu) (C0)</li> <li>- *usa, with international partners</li> </ul>	<ul style="list-style-type: none"> <li>- FT frameworks &amp; libraries-FTI, fault detectors(Tokyo Tech – UTokyo) (C2, INRIA-NCSA, JST-ANR; G8)(C1, LLNL)</li> <li>- Domain specific languages, <i>Physis</i> (Riken – TokyoTech) (C0)</li> <li>- Numerical libraries: z-Pares, Eigen-K (UTsukuba, AICS, etc.) (C0),</li> <li>- Autotuning functions (UTokyo, UTsukuba, KyotoU, etc.) (C2, INRIA-Saclay, ENSEEIHT-IRIT, Uversailles, JST-ANR)</li> <li>- “App Framework with Automatic Tuning” ppOpen-HPC (UTokyo, KyotoU, JAMSTEC) (C1, LBL)</li> <li>- “Application Framework” Framework (U-Tokyo)(C0)</li> <li>- Development Environment (Tohoku-U, U-Tsukuba, U-Tokyo) (C1, GWTH Aachen, UIUC, HLRS)</li> <li>- Communication library (space efficient, dynamic optim.) (KyushuU, Fujitsu, ISIT-Kyushu) (C0)</li> <li>- Power management and scheduling framework-scheduling strategies limiting peak power (Tokyo Tech.) (C2, NVIDIA, HP, GRC, etc., NVIDIA COE)</li> <li>- Data Analytics Framework (ChuoU-Tokyo Tech.) (C1)</li> </ul>	<p><u>Debuggers</u></p> <p>Performance tools</p> <p>Batch / Machine-wide scheduler</p> <p><u>Node level compilers for many-core</u></p>

# Stack Components: Europe (updated – April 13)

Group provides in production quality	Experimental versions only	No major research effort—rely on vendors or other countries
<p>Programming models/ languages <b>MPI/OmpSs</b> (C0), <u>in the long run may move to OpenMP 4.0</u> (C3, multiple)</p> <p>Compiler for heterogeneous systems <b>HMPP</b> (C2, CAPS, ORNL, DOE) <b>OpenACC</b> (C3, multiple)</p> <p>Node-level runtime <b>NANOS++</b> (C0), <b>StarPU</b> (C1, INRIA, UTK) (C2, Tsukuba, TiTech, JST-ANR)</p> <p>Performance Tools <b>Paraver</b> (C0), <b>Scalasca</b> (C1, Julich, UTK) (C2, JSC, Moscow, EU FP7) , <b>Score-P</b> (C3, multiple, Uoregon, DOE), <b>Vampir</b> (C2, TUD, ORNL, DOE), <b>OPT</b> (C1), <b>ThreadSpotter</b> (C1)</p> <p>Debuggers <b>DDT</b> (C2, Allinea, ANL, LLNL, ORNL, DOE)</p> <p>MPI Correctness <b>MUST</b> (C3, TUD, LLNL, DOE)</p> <p>Commit to contribute to MPI implementations <b>MPICH</b>, <b>Open MPI</b> (C1, multiple) <b>MPI Madeleine</b> (C2, INRIA, U. Tokyo, ANR-JST)</p>	<p>Communication libraries (below MPI) (C1, ?) , 1-sided communication library <b>GPI</b> (C1, ?)</p> <p>Fault tolerant libraries FTI, MPI (C1, INRIA, TiTech)</p> <p>Performance Tools Periscope (C0, TUM), Task-based Debugging Ayudame/Temanejo (C0, HLRS)</p> <p>Domain specific languages (for autotuning) (C?)</p> <p>Batch scheduler research MOAB (C1, ParTec, Cluster Resources)</p> <p>Power measurement management eeClust and others (C1, eeClust, ??)</p> <p>I/O middleware DAMARIS (C2, INRIA, UIUC, Joint Lab), SIONlib (C1, JSC, ???)</p>	<p>Node level compilers (vendor compilers)</p> <p>PGAS languages</p> <p>Scalable RAS system</p>

# Stack Components: Europe (updated – April 13)

Group provides in production quality	Experimental versions only	No major research effort—rely on vendors or other countries
<p>File systems <b>EOFS Lustre</b> (C3, multiple, NEC, Netapp, Whamcloud)</p> <p><b>Numerical libraries</b> (C1, multiple, UTK and others)</p> <p>Numerical libraries <b>NAG</b> (C1)</p> <p>OS kernel + cluster env <b>BullX</b> (C1)</p> <p>Log analysis <b>HELO</b> (C2, INRIA, UIUC, Joint Lab)</p>	<p>OAR batch scheduler (C0)</p>	

# Stack Components: USA (updated April 13)

Group provides in production quality	Experimental versions only	No major research effort—rely on vendors or other countries
<p>Programming models: <i>CAF Fortran Standard Committee</i> (C2, multiple), <i>OpenMP Standard Committee</i> (C3, multiple), <i>OpenAcc Consortium</i> (C3, multiple), <i>OpenCL Consortium</i> (C2, multiple)</p> <p><i>UPC LBNL</i> (C0)</p> <p><i>OS Kernel ANL SNL</i> (C1, ANL, U. Tokyo)</p> <p>FT Framework: <i>Teal IBM ANL LLNL NCSA</i> (C0)</p> <p><i>Integrated System Console NCSA</i>(C2, INRIA, UIUC, Joint-lab)</p> <p>FT env: <i>SCR</i> (C1, LLNL IBM, Titech)</p> <p>Runtime: Charm++ <i>UIUC</i> (C0)</p> <p>Communication library: <i>MPICH</i> (C1, ANL Cray IBM Intel Microsoft OSU UBC INRIA), <i>Open MPI</i> (C1, UTK, ORNL Bull)</p> <p>Tool infrastructure: <i>MRNet</i> (C0, Cray UWI ORNL UNM)</p> <p>Debugging: <i>STAT</i> (C0, LLNL Cray) <i>TotalView</i> (C0) <i>Valgrind/Memcheck</i> (C3, OpenWorks, LLNL IBM, DOE)</p>	<p>Performance tools: mpiP (C0), Fault tolerance backplane: CIFTS (C0, ANL ORNL)</p> <p>Compilers: Open64 (C1, UHous ORNL NVIDIA, CAS, Tsinghua)</p> <p>Compiler framework: Rose (C0)</p> <p>Compiler support for heterogeneous processing (C0, UHous NVIDIA)</p> <p>Gmac, dl, tc (C0, UIUC)</p> <p>Debugging and validation: MUST (C0)</p> <p>Runtime: Unistack/ Plum (C0)</p> <p>CCI (C1, INRIA, ANL, SNL, UTK)</p> <p>PowerMgmt Layer (C?)</p> <p>Low-level threading lib: Qthreads (C0)</p> <p>Task scheduler: ExM (C0)</p> <p>File system: PVFS (C0, multiple USA, UC3M, Hamburg), PLFS (C0)</p> <p>FT API (C0, Cray UTAustin)</p> <p>Data analytics libraries: Xanalytics (C0)</p> <p>Data model storage library: Damsel (C0)</p> <p>PnetCDF (C1, NorthW, ANL, U. Tokyo, RIKEN)</p> <p>Composition frameworks: COMPOSE-HPC (C0)</p>	<p>Commercial compilers</p>

# Stack Components: USA (updated April 13)

Group provides in production quality	Experimental versions only	No major research effort—rely on vendors or other countries
<p>Performance tool: <b>PAPI</b> (C1, <i>UTK IBM</i>), <b>OpenSpeedShop</b> (C1, <i>Krell</i>), <b>TAU</b> (C1,, <i>U. Or, Juelich</i>), <b>HPCToolkit</b> (C0, <i>Rice</i> ) Dyninst (<i>C1, U. Maryland, UWI, IBM, UAB, multiple</i>)</p> <p>File systems: <b>Lustre</b> (C3, <i>WhamCloud Cray LLNL ORNL, OpenSFS Xyratex, DDN, DOE and commercial</i> )</p> <p><b>I/O delegation</b> (C0 <i>Northwest., ANL</i>)</p> <p>I/O middleware: <b>IOFSL</b> (C1, <i>ANL LANL ORNL SNL, U. Tokyo</i>)</p> <p>Visualization: <b>VTK, Visit</b> (C0)</p> <p>I/O libraries: <b>HDF5, pNetCDF</b> (C0/C1)</p> <p>Numerical libraries: <b>MAGMA/ PLASMA</b> (C1, <i>INRIA, KAUST, UCo UCB</i>), <b>PETSc</b> (C0/C1, <i>ANL</i>), <b>Trilinos</b> (C1, <i>SNL</i>), <b>SuperLU</b> (C1, <i>UCB, INRIA</i>), <b>hypre</b> (C0, <i>LLNL</i>)</p> <p>Resource managers: <b>SLURM</b> (C2/C3 <i>SchedMD LLNL Intel BULL HP Ubarcelona</i>), <b>MOAB</b>. (C0, <i>Cluster Resources Cray IBM HP APPRO SGI etc</i>)</p>	<p>New programming models and approaches: Sketch, DSLs (C?), Chapel (C3, <i>CRAY LLNL, SNL, BSC, U. Malaga</i>), X10 (C2, <i>IBM, Titech, JST</i>)</p> <p>Benchmarks and mini-apps: SHOC (C0), SPP (C0, <i>NCSA</i>)</p> <p>DAGUE/QUARK –runtime DAG execution (C0)</p>	

# Stack Components: China (updated - April 13)

Group provides in production quality	Experimental versions only	No major research effort—rely on vendors or other countries
<p>Operating system: <b>Kylin OS</b> (compatible with Linux, optimizations for HPC) (C0)</p> <p>File systems: <b>Lustre enhancement</b> (C1 NUDT, Whamcloud)</p> <p>Parallel Programming framework (structured mesh): <b>JASMIN</b> (C0)</p> <p>Light-weight fault-tolerant MPI: <b>NR-MPI</b> (C0)</p> <p>HPC virtual environment : <b>HPVZ</b> (C0)</p> <p>Resource management: <b>Slurm enhancement</b> (C1, NUDT, SchedMD)</p>	<p>Resilience computing framework (C2, NUDT, IIT Chicago, UIUC; NSFC)</p> <p>Hybrid tiered file system (C0)</p> <p>Heterogeneous programming runtime system (load balance &amp; pipeline) (C0)</p> <p>Domain specific programming framework (unstructured mesh and others) (C0)</p> <p>Autonomic resource management and scheduler (C0)</p> <p>Integrated development tools (performance analysis, debugger, workflow) (C0)</p>	<p>Compilers for commercial processors</p>

# Stack Components: Russia

Group provides in production quality	Experimental versions only	No major research effort—rely on vendors or other countries
<p>Numerical packages (<a href="#">FlowVision (TESIS)</a> (C2), <a href="#">FireFly</a> (C2), ...)</p> <p>OS Kernel (C3)</p> <p><a href="#">Software suit for multilevel programming of reconfigurable FPGA-computers (SFU)</a> (C1)</p> <p><a href="#">Programming language COLAMO (SFU)</a> (C1)</p> <p>Batch scheduler Cleo (MSU) (C2)</p> <p>Batch scheduler SUPPZ (JSCC) (C1)</p>	<p>New programming models and languages (MSU, NUDA) (C1)</p> <p>Communication libraries (below MPI, DISLIB) (C1)</p> <p>Fault tolerant libraries (T-PI) (C2)</p> <p>Commit to contribute to MPI implementations (T-PI) (C?)</p> <p>Performance tools (MSU, LAPTA-HOPSA) (C3)</p> <p>Power measurement management (T-PI, RSC) (C?)</p> <p>I/O middleware (T-PI) (C?)</p> <p>Convergent cloud/HPC platform (RSC) (C1)</p> <p>Numerical libraries on GPU (SpVM): Sparse matrix-vector multiply, autotuning (ISP) (C1)</p> <p>Numerical packages: OpenFOAM on GPU (ISP) (C2)</p> <p>Compiler framework: automatic OpenCL generation in GCC (ISP) (C?)</p> <p>Tools and performance: Java+MPI bindings (ISP) (C?),</p> <p>HPC Virtualization: Palacios VMM improvements (ISP) (C1)</p>	<p>Debuggers, Compilers</p>

## Specific Areas Where Tight Collaboration would be useful

---

- Programming models (system level)
- Node level paradigms (manycores, accelerators, etc.)
- Performance and correctness tools
- Fault tolerance (system and application)
- Power efficiency
- Storage and file system
- Flexibility and adaptability

# Specific Areas Where Tight Collaboration would be useful

---

## □ Observations:

- Many P2P funding, almost no multi-region funding (>2)
- No sufficient multi-region funding (>2) and no instrument
- Very few US-to-Europe (and vice versa) funding compared to Japan-to-Europe (and vice versa)

## □ Planning to improve the situation

- Consolidate the list along the topic area and make suggestions
- Identify missing point: Comparison between the roadmap and the tables (gap analysis)
- Providing Evidences of the benefit and success of these collaborations