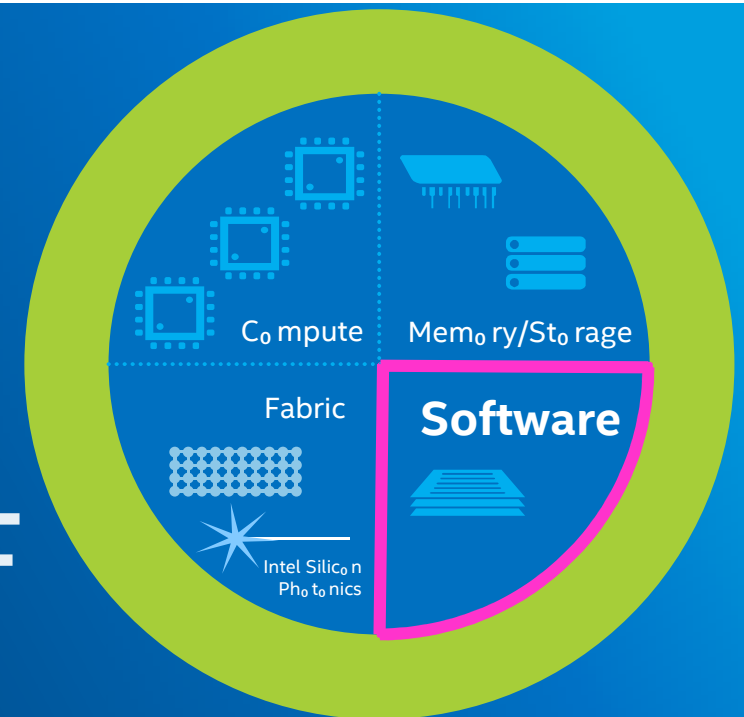


# OpenHPC and SSF



Dr. Robert W. Wisniewski, Chief Software Architect Extreme Scale Computing, Intel

SOS 20 March 24, 2016

# Legal Disclaimer

Results have been estimated based on internal Intel analysis and are provided for informational purposes only. Any difference in system hardware or software design or configuration may affect actual performance.

Intel, processors, chipsets, and desktop boards may contain design defects or errors known as errata, which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Intel, Intel Xeon, Intel Core microarchitecture, and the Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at [intel.com](http://intel.com), or from the OEM or retailer.

No computer system can be absolutely secure.

Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase. For more complete information about performance and benchmark results, visit <http://www.intel.com/performance>.

Cost reduction scenarios described are intended as examples of how a given Intel-based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest forecast, schedule, specifications and roadmaps.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

Intel does not control or audit third-party benchmark data or the web sites referenced in this document. You should visit the referenced web site and confirm whether referenced data are accurate.

Intel, the Intel logo, Xeon, Xeon Phi, Core, and others are trademarks of Intel Corporation in the U.S. and/or other countries. \*Other names and brands may be claimed as the property of others.


\*Other names and brands may be claimed as the property of others

Copyright © 2016, Intel Corporation. All rights reserved.



# INTEL'S SCALABLE SYSTEM FRAMEWORK

A design foundation enabling a wide range of highly workload-optimized solutions



Small Clusters Through Supercomputers

Compute and Data-Centric Computing

Standards-Based Programmability

On-Premise and Cloud-Based

Intel® Xeon® Processors

Intel® Xeon Phi™  
Coprorocessors

Intel® Xeon Phi™ Processors

Intel® True Scale Fabric

Intel® Omni-Path  
Architecture

Intel® Ethernet

Intel® SSDs

Intel® Lustre-based Solutions

Intel® Silicon Photonics  
Technology

Intel® Software Tools


HPC Scalable Software Stack

Intel® Cluster Ready Program

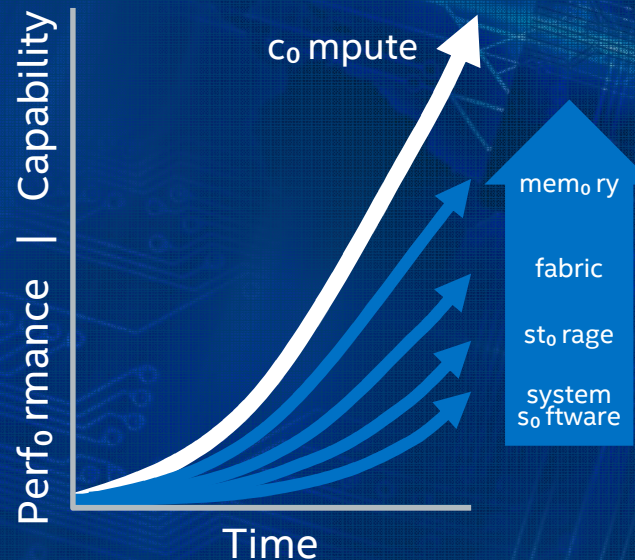


# A SYSTEMS APPROACH FOR INNOVATION


Continue the pace of  
**FASTER COMPUTE**



Develop new technologies to  
**REMOVE BOTTLENECKS**



Achieve full potential via  
**TIGHTER INTEGRATION**


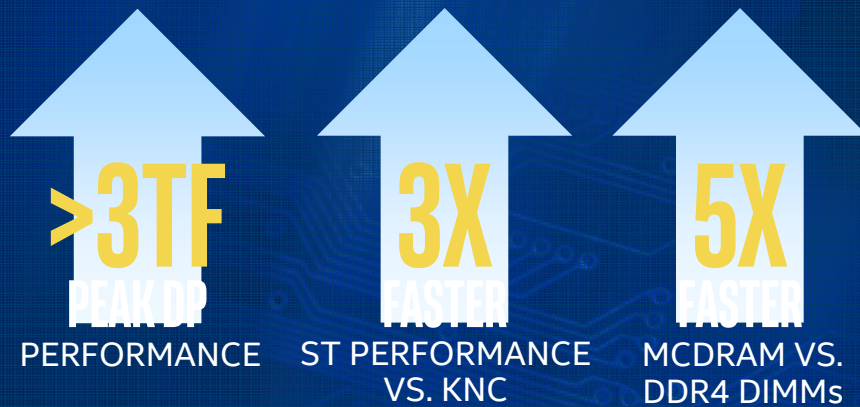



Better power, performance,  
density, scaling & cost



# KNIGHT'S LANDING

## NEXT-GEN INTEL® XEON PHI™ PROCESSOR



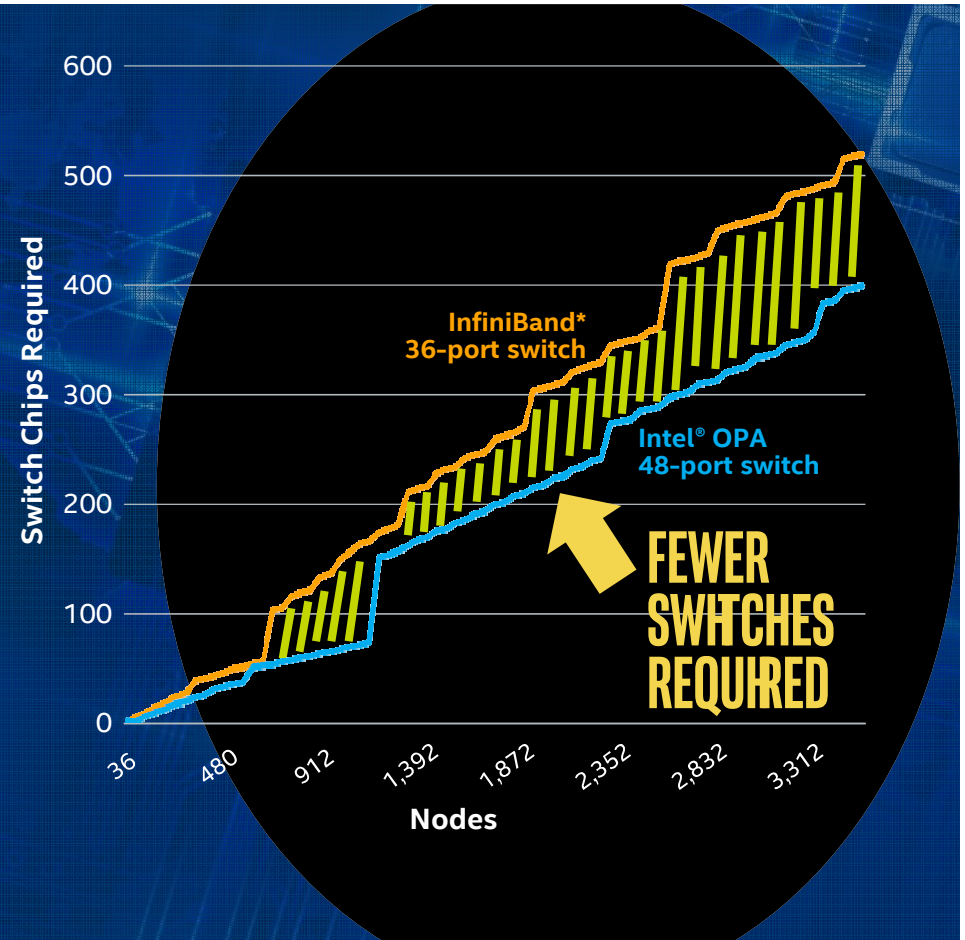
# OMNI-PATH™ ARCHITECTURE



**100**  
GBIT/S  
PER  
PORT

**1.3X**  
HIGHER  
SWITCH  
DENSITY

**2.3X**  
GREATER  
FABRIC  
SCALABILITY



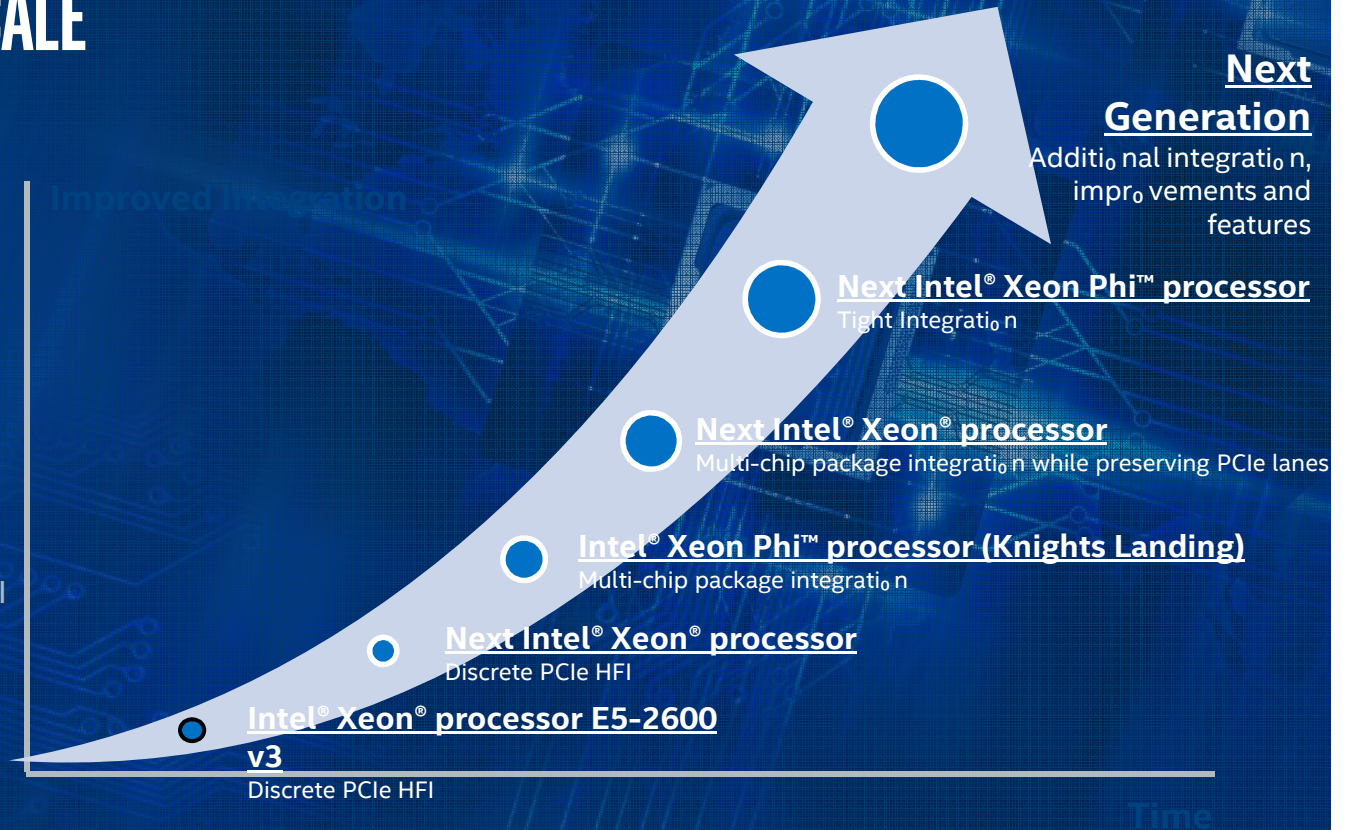
Number of switch chips required, switch density, and fabric scalability are based on a full bisectonal bandwidth (FBB) Fat-Tree configuration, using a 48-port switch for Intel® Omni-Path Architecture and 36-port switch ASIC for either Mellanox or Intel® True Scale Fabric. \*Other names and brands may be claimed as the property of others. 2.3X fabric scalability based on a 27,648-node cluster configured with the Intel® Omni-Path Architecture using 48-port switch ASICs, as compared with a 36-port switch chip that can support up to 11,664 nodes.



# PROCESSOR WITH THE INTEL® OMNI-PATH ARCHITECTURE INTEGRATION IS REQUIRED FOR EXASCALE


## KEY VALUE VECTORS

- ✓ **Performance**
  - >30% reduction in delivered MPI message latency
  - >2x messaging rate
  - Bandwidth no longer package and PCIe limited
- ✓ **Density**
  - >2x nodes/rack in Aurora form Integration
- ✓ **Cost**
  - 4-6x savings for single/dual rail HFI
- ✓ **Power**
  - >4x reduction for Aurora
- ✓ **Reliability**
  - Large FIT rate improvement by completely eliminating parts



# INTEL'S SCALABLE SYSTEM FRAMEWORK

A design foundation enabling a wide range of highly workload-optimized solutions



Small Clusters Through Supercomputers

Compute and Data-Centric Computing

Standards-Based Programmability

On-Premise and Cloud-Based

Intel® Xeon® Processors

Intel® Xeon Phi™  
Coproprocessors

Intel® Xeon Phi™ Processors

Intel® True Scale Fabric

Intel® Omni-Path  
Architecture

Intel® Ethernet

Intel® SSDs

Intel® Lustre-based Solutions

Intel® Silicon Photonics  
Technology

Intel® Software Tools

HPC Scalable Software Stack

Intel® Cluster Ready Program



## Goals for the HPC Software Community

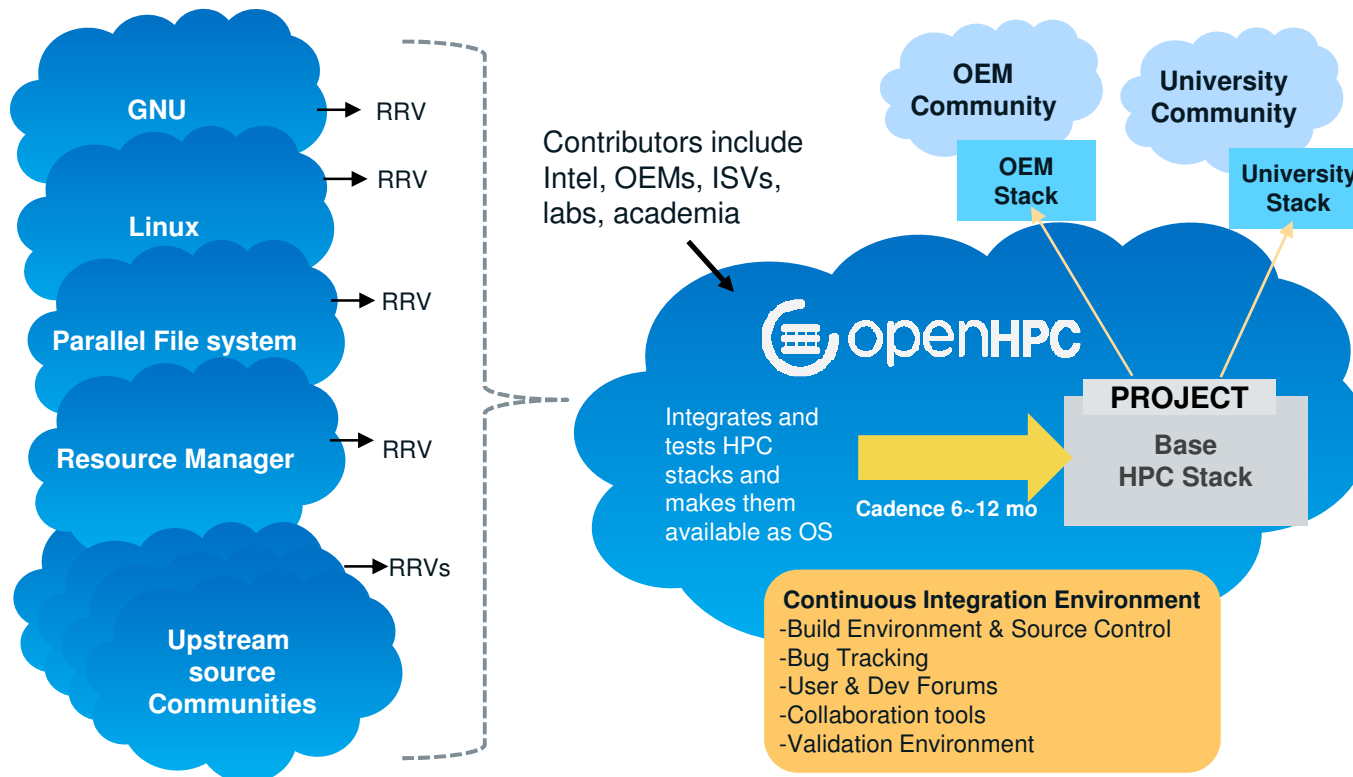
Community name: **OpenHPC**

Web Address: **[www.openhpc.community](http://www.openhpc.community)**

- ❖ Provide a common platform to the HPC community that works across multiple segments and on which end-users can collaborate and innovate
- ❖ Simplify the installation, configuration, and maintenance of a common HPC software stack
- ❖ Facilitate contributions and input across community
- ❖ Enable developers to focus on their differentiation and unique area, rather than having to spend effort on developing, testing, and maintaining a whole stack
- ❖ Deliver integrated hardware and software innovations to ease the path to extreme scale



# OpenHPC and Stack Curation Framework







# Early OpenHPC Members (as of January 2016)

## OEMs



## Users



## ISV-OSV



# OpenHPC High-Level Component List (potential future add)

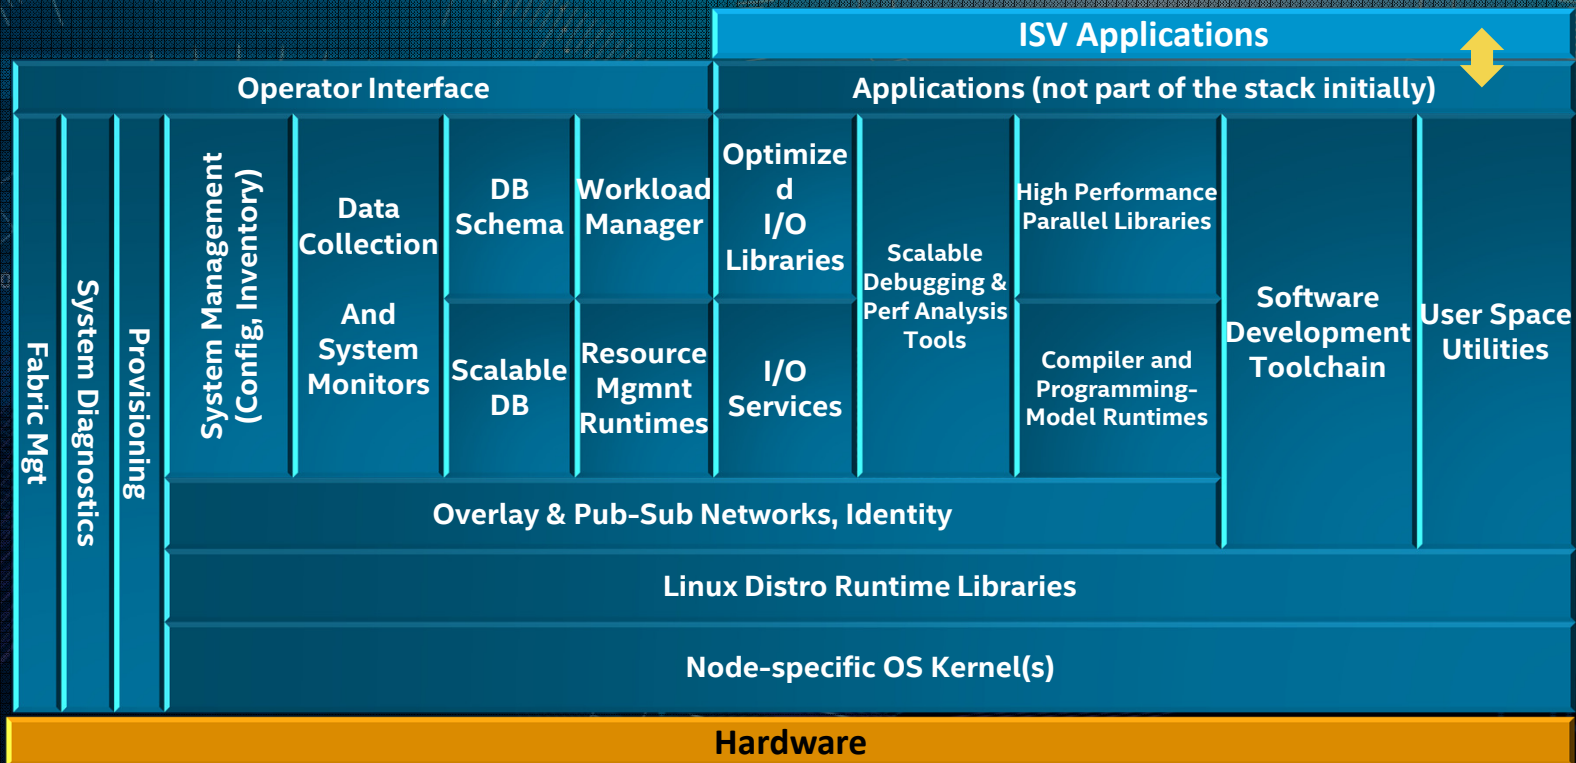
| Functional Areas               | Components  | Adds Contributed by                   |
|--------------------------------|---|---------------------------------------|
| Base OS                        | RHEL/CentOS 7.1, <b>McKernel, Kitten, mOS</b>   | <b>RIKEN, Sandia, Intel®</b>          |
| Admin Tools                    | Conman, Ganglia, Intel® Cluster Checker**, Lmod, LosF, Nagios, pdsh, prun, EasyBuild, <b>ORCM</b>   | <b>Intel</b>                          |
| Provisioning                   | Warewulf, <b>xCAT</b>   | <b>community</b>                      |
| Resource Mgmt.                 | SLURM, Munge, <b>ParaStation management, PMIx, PBS Pro</b>  | <b>ParTec, community, Altair</b>      |
| Cross Cutting                  | <b>OpenStack HPC suitable components</b>  | <b>Cray</b>                           |
| Runtimes                       | OpenMP, <b>OmpSs, OCR, HPX-5</b>  | <b>BSC, Intel, Indiana University</b> |
| I/O Services                   | Lustre client   |                                       |
| Numerical and Scientific Lib's | Boost, GSL, FFTW, Metis, PETSc, Trilinos, Hypre, SuperLU, Mumps, Intel MKL**  |                                       |
| I/O Libraries                  | HDF5 (pHDF5), NetCDF (including C++ and Fortran interfaces), Adios  |                                       |
| Compiler Families              | GNU (gcc, g++, gfortran), Intel Parallel Studio XE (icc, icpc, ifort)**   |                                       |
| MPI Families                   | MVAPICH2, Intel MPI**, OpenMPI, <b>MPICH, ParaStation MPI</b>   | <b>Argonne, ParTec</b>                |
| Dev. Tools                     | Autotools (autoconf, automake, libtool), Valgrind, R, SciPy/NumPy, Intel Inspector**  |                                       |
| Performance Tools              | PAPI, Intel IMB, mpiP, pdtollkit, TAU, Intel Advisor**, Intel Trace Analyzer and Collector**, Intel Vtune Amplifier**, <b>Paraver, Scalasca</b> | <b>BSC, Jülich</b>                    |

\*\* Additional license required




# OpenHPC Stack Component View

- ❑ Intra-stack APIs to allow for customization/differentiation
- ❑ External APIs to develop on and around the stack



## Conclusions

- SSF is leading the way to tighter integration and therefore quicker and more effective delivery of new technology
- Intel will contribute to the open community work needed to scale to extreme scale
-  is live - check out [www.openhpc.community](http://www.openhpc.community)