# Intel Scalable System Framework Fueling Your Future HPC Systems

Pawel Gepner

Intel Corporation, Pipers Way, Swindon, Wiltshire, SN3 1RJ, United Kingdom

email: `pawel.gepner@intel.com`

## 1.  Introduction

High Performance Computing is a fundamental tool used in scientific research and in industry. As the complexity of our understanding of the universe grows, the only practical way to compare observations to theory is via computational simulation, such approach involves massive usage of HPC. Similarly, for business and industry, HPC allows to use this understanding to model and simulate designs, enabling to create better products and services faster and for a lower total investment. In both science and industry, the size and diversity of data is increasing exponentially. HPC is key factor to enabling of this data to drive deeper insight and faster reactions to changing conditions.

## 2.  Problem statement

Total size of the x86 market in 2014 was 44,324M\$ and x86 HPC market spend \$9,464 M\$. One fifth of data center spending on hardware goes towards HPC. However, HPC could do much more than is currently possible, if we can overcome several growing challenges facing HPC. The challenges represent the base technical barriers the industry is facing today in compute, memory, I/O, storage and power. This is mix with different workloads hitting these barriers differently, driving separate systems with specialized hardware and software for each workload – increasing acquisition and management costs and reducing efficiencies (load balancing). Both of these approaches have driven the complexity of programming for an HPC system, and we need new methods and tools to reduce this complexity to extend HPC's reach (Cloud, holistic software stacks, new programming models). Reducing this complexity is a crucial step in enabling more people to take advantage of the aforementioned benefits HPC offers. This democratization could take many forms: cloud resources to reduce the complexity of acquiring and managing an HPC system, holistic system software stacks to reduce the complexity of utilizing the hardware or new programming models to reduce the complexity of creating HPC applications.

## 3.  Description of a problem solution

Any solution to these challenges will require three elements:

- Innovative technologies in compute, memory, fabric, storage and system software.
- Taking these technologies and integrating them tightly together (e.g. memory +fabric +power) and be performant across workloads
- Code that's modernized to take advantage of these hardware advances – whether that's community, ISV or proprietary (in house) developed code

Intel's strategic architectural direction designed to overcome these challenges is the Intel Scalable System Framework, consisting of four pillars of innovation and three guiding principles.

The four pillars of innovation are:

- Compute: consisting of Intel Xeon processors, Intel Xeon Phi coprocessors, and Intel Xeon Phi processors
- Memory and Storage: consisting of Intel 3D Xpoint technology, Intel Optane technology, Intel Solid-State Drives (SSDs), and Intel Lustre-based solutions
- Fabric: consisting of Intel True Scale Fabric, Intel Omni-Path Architecture, Intel Ethernet, and Intel Silicon Photonics Technology
- Software: consisting of Intel software tool, HPC Scalable Software Stack, and the Intel Cluster Ready Program

The three guiding principles are:

- Reliability and Resiliency
- Power Efficiency
- Price / Performance

As Intel's roadmap and technologies develop over time, these guiding principles will be used to continue to deliver breakthrough solutions, capable of meeting the increasingly complex demands of the high performance computing community.

There are four key end user benefits that arise from this:

- Breakthrough performance in four key areas
- All built on standards based programmability
- Which allows you to build a common infrastructure performant across a broad range of workloads
- readily available from most major hardware vendors (no change to how you research and purchase systems)

The true end user value can be seen in four main areas. First, we're developing innovative technologies that will deliver breakthrough performance across the range of needs of an HPC system. Including industry leading CPU performance, fast access to data, high-bandwidth, low latency fabric and software to ease the deployment and management of your system. Second, all of this technology is built on the standards-based programmability that IA is known for. While optimizations are often necessary to take advantage of new features and advances in hardware, with Intel Scalable System Framework you can carry forward and build on your previous modernization efforts and your current application bases. All without learning a new programming language or having to completely rewrite legacy code. These enable you to build a common infrastructure to serve your entire application base, whether that's modeling and simulation, big data analytics, machine learning or visualization. While individual systems can be built to provide optimal performance for a specific workload (more memory, storage, fabric or compute, or all of the above) it is still a single system architecture which you can load balance across. Reducing acquisition and management costs while increasing efficiency. Finally, Intel Scalable System Framework, as with the rest of Intel's portfolio, will be available from most major hardware vendors, building on the segments largest software and hardware ecosystem.

There are two parts to Intel Scalable System Framework:

First, we have a portfolio of innovative technologies, including compute, memory and storage, fabric, and software. These technologies provide much of the breakthrough performance from above. By acquiring and developing these technologies in house, we've enabled integration on previously unseen levels.

This integration is encompassed in two main types; component integration and system level integration. Again, it is this system level integration that enables the creation of systems that overcome multiple walls and are performant across workloads. Component integration brings key components much closer together, providing increases in bandwidth and density while decreasing latency, power and cost. Component integration brings system components onto the CPU package or die. We've seen this in the past with integrating cores, the memory controller, I/O (PCIe bus) and graphics, and we'll continue to see this with the integration of fabric and memory (Knights Landing). This provides increased bandwidth, increased density, reduced latency, reduced power and ultimately reduced cost. System level integration brings the system as a whole closer together via two main methods:

- High bandwidth on-package memory brings local memory closer to compute, Intel® DIMMs based on 3D XPoint™ Technology bring storage onto the memory bus, Intel® Optane™ Technology SSDs bring storage into a higher bandwidth, lower latency form factor on the compute or I/O nodes. All of which brings data much closer compute
- Intel® Omni-Path Architecture integrates the nodes together much more tightly, increasing bandwidth and reducing latency between nodes, greatly reducing the penalty of spanning a job across multiple nodes

Both HPC and BD use these capabilities, but their usage is weighted differently. For example, HPC emphasizes high bandwidth configurable memory. Big Data uses in package memory, but focuses on configurable memory and local application storage.

## 4. Conclusions

To take full advantage of the hardware advances however, in Intel Scalable System Framework, we need to modernize the code running on these systems. Specifically, we need to take advantage of three key advances: Parallelism in the CPU, specifically thread, vector and data level parallelism, the ability to keep more data closer to the CPU with High bandwidth in package memory and Intel 3D XPoint Technology NVM, and the ability to more tightly integrate nodes with the increased bandwidth and reduced latency of Intel Omni-Path Architecture.