



Accelerating Datacenter Workloads

FPL 2016

PK Gupta,
GM of Xeon+FPGA Products
Datacenter Group



Overview

- **Data Center and Workloads**
- Xeon+FPGA Accelerator Platform
- Deployment of FPGAs in the Datacenter
- Applications and Eco-system


Cloud Infrastructure Enables New Usage

Cloud 2015 Digital Services Economy



66%¹ of current cloud demand comes from consumer services

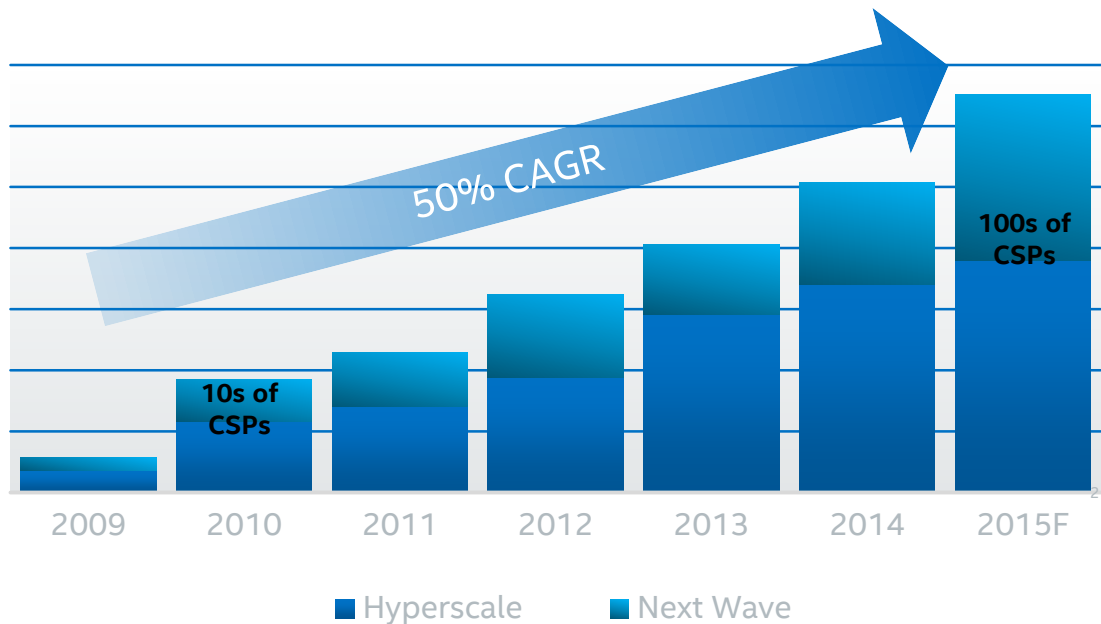
Cloud 2020 IoT, Big Data, and Enterprise



By 2020, 65-85%² of apps will be delivered via cloud infrastructure

Tremendous Cloud Growth Fueled by New Usages

Public Cloud Intel CPU Sales¹




1. Source: Intel

2. Source: Internal Intel forecast, based on available industry data, 2015

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

Cloud Adoption: Next Wave and Broad Enterprise

Enterprise Workloads
Destination
2015-2016¹

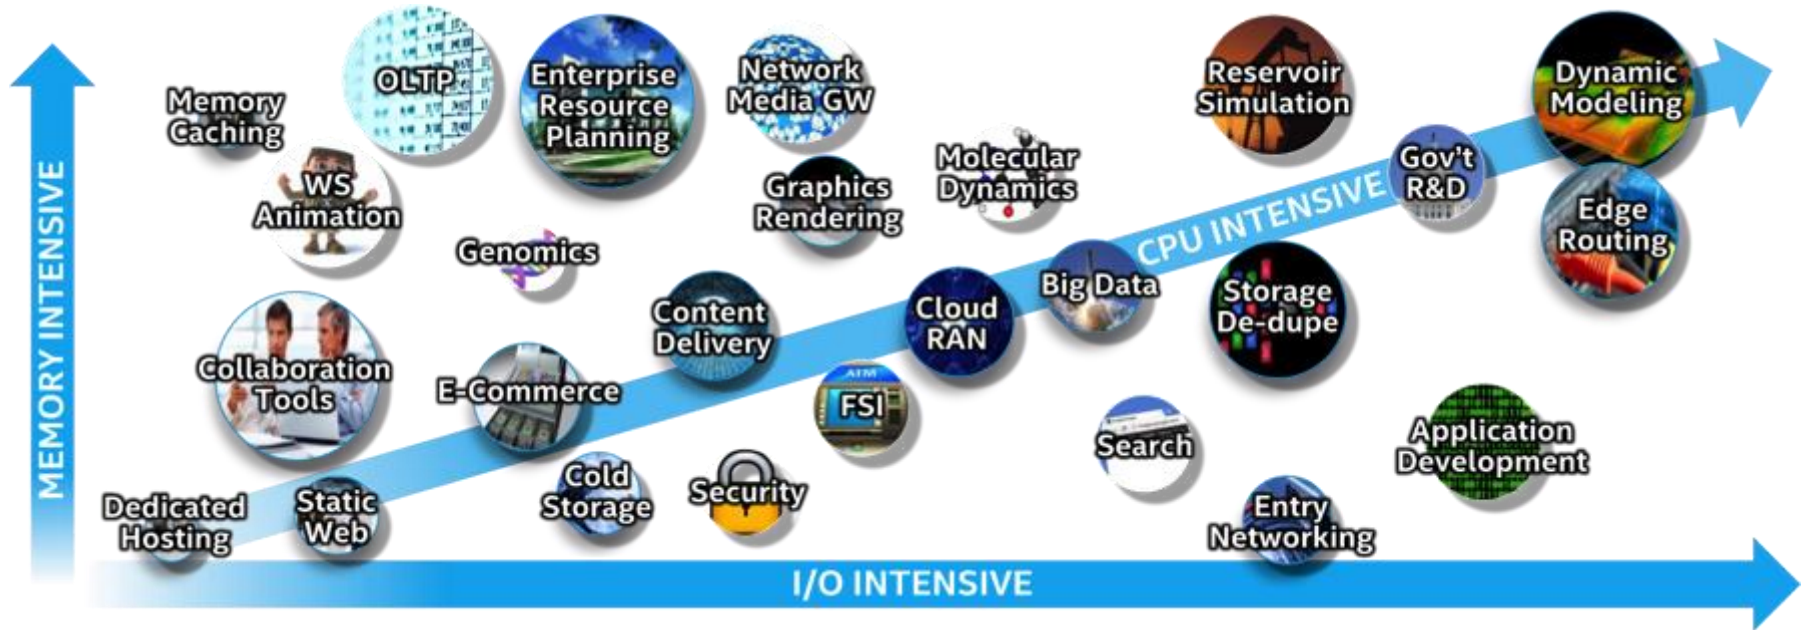


Intel® Cloud for All

Unleash Tens of Thousands of New Clouds*



Diverse Data Center Demands



Accelerators can increase performance at lower TCO for targeted workloads

Overview

- Data Center and Workloads
- Xeon+FPGA Accelerator Platform
- Deployment of FPGAs in the Datacenter
- Applications and Eco-system

Range of Acceleration Options

Highest Flexibility / Customization

Highest Performance / \$

PROCESSOR

FPGA

ASSP /
WORKLOAD SPECIFIC

General Purpose
Applications

Acceleration for
Flexible Workloads

Acceleration for
Standard Workloads

Software Flexibility

CPU socket compatible
access to FPGA
capabilities

Scalable range of FPGA
options (I/O, TDP,
Price, Mem, Features)

Built-in Standard
platform acceleration,
Highly Optimized

GENERAL
Workloads

Standard Silicon

Integrated FPGA
Solution

Discrete FPGA
Solutions

Chipset Solutions

TARGETED
Workloads

For SW
Customization

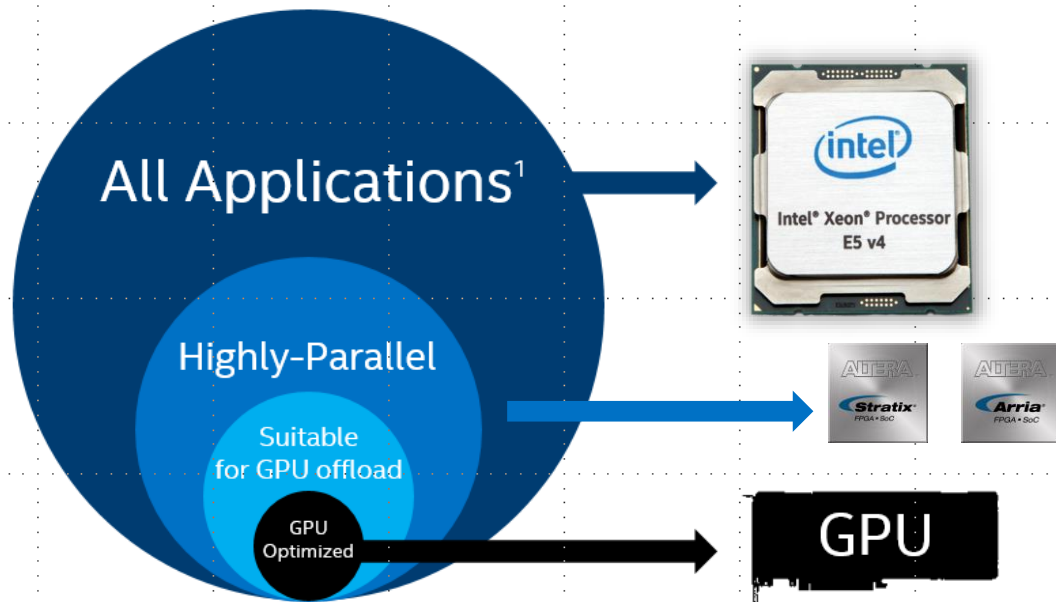
For HW
Customization

For Specific
Acceleration



Intel® Xeon with FPGA Meet More Applications


GPU accelerators in the datacenter are frequently idle, consuming space, power and capital budget*



*Based on Intel customer survey feedback of customers with significant GPU deployments

† All x86 applications that run on Intel® Xeon® processors

Intel® Xeon® + FPGA Pilot Systems



DELIVERING CUSTOM ACCELERATION

APPLICATIONS

Image Identification



Security



Firewall, VPN, Router...



XEON



FPGA

ALGORITHMS

Convolutional Neural Network

Encryption

Virtual Switching

Overview

- Data Center and Workloads
- Xeon+FPGA Accelerator Platform
- **Deployment of FPGAs in the Datacenter**
- Applications and Eco-system

Key Challenges Impacting FPGAs in the Data Centers

Increasing Velocity of Unique Workloads

Homogeneity vs. Customization

Power Consumption

Densely packed processing and acceleration

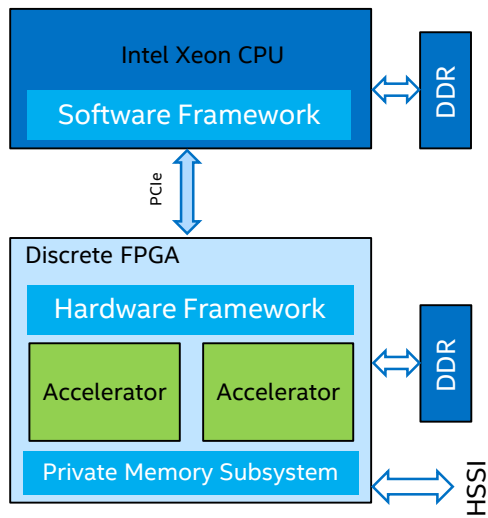
Security

Protecting sensitive data from hackers



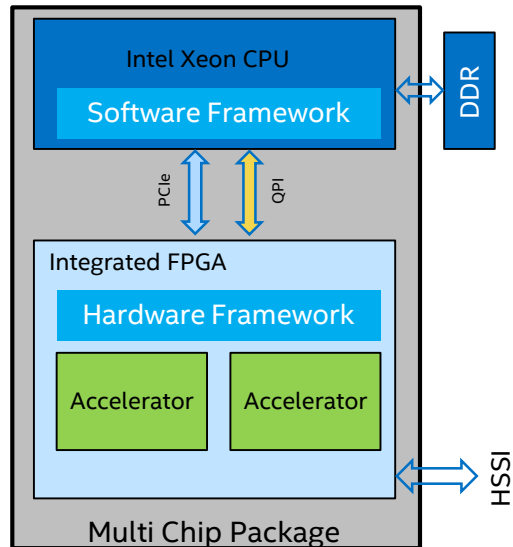
Discrete and Integrated FPGA Platforms

Discrete Platform (DCP)



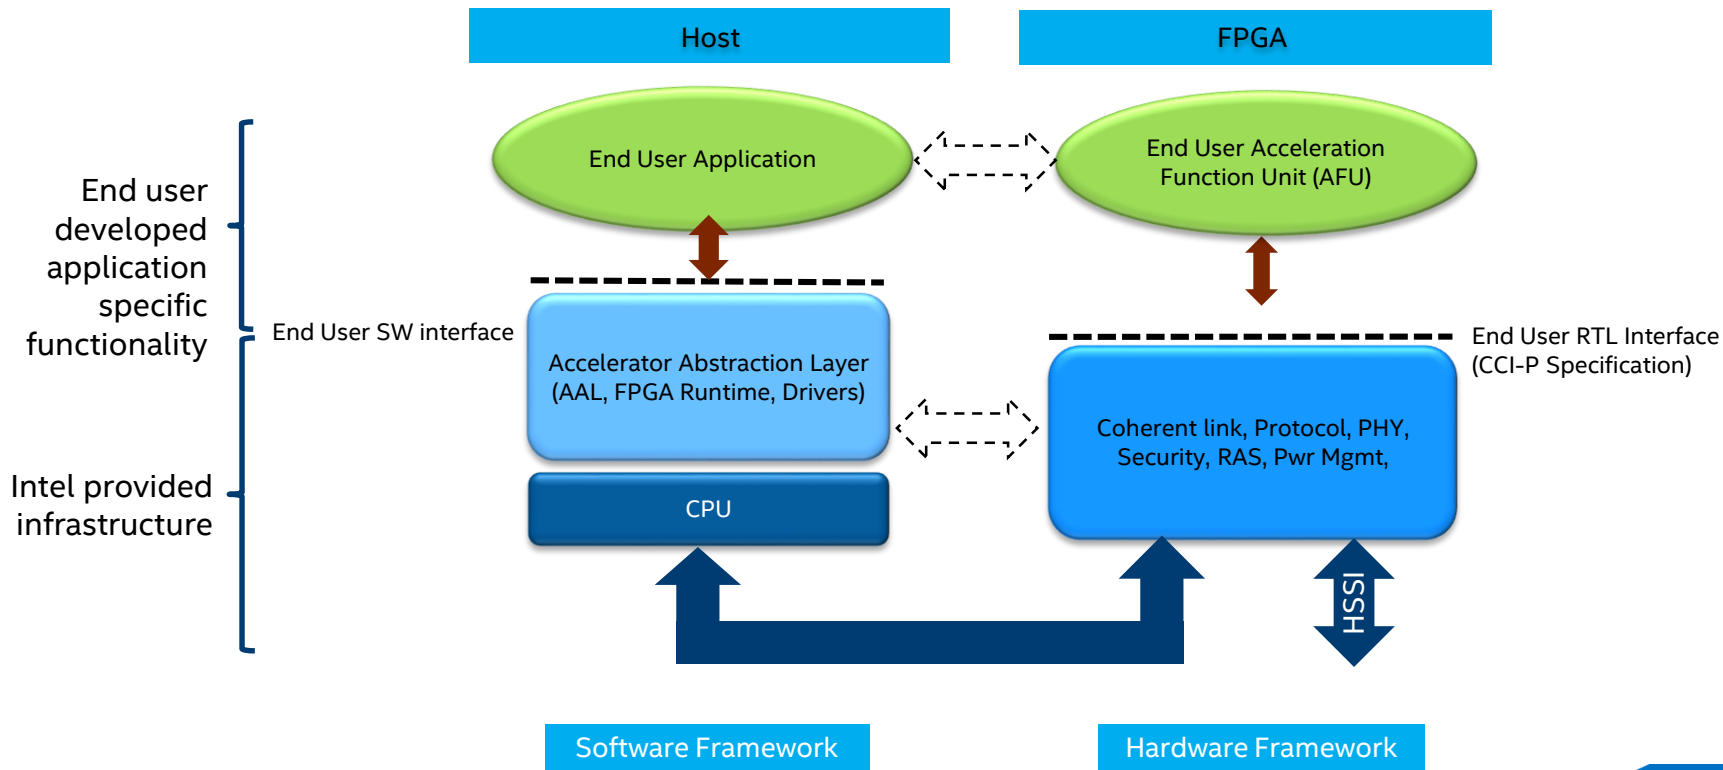
Today: Arria[®]10 PCIe Available

Intel[®] Xeon[®]+FPGA Integrated Platform (MCP)

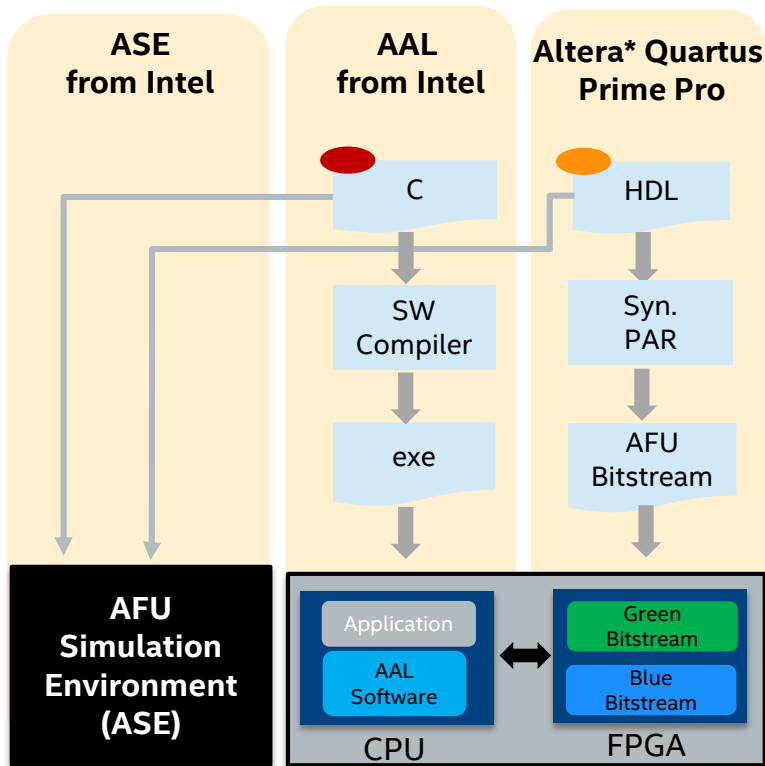


Today: BDX+FPGA MCP Pilot

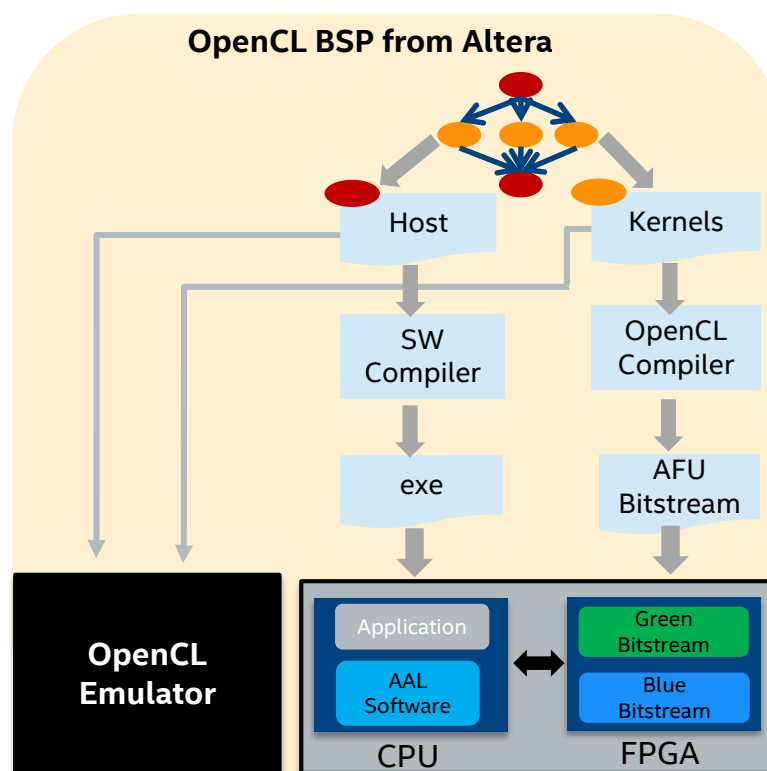
End User Programming Interfaces




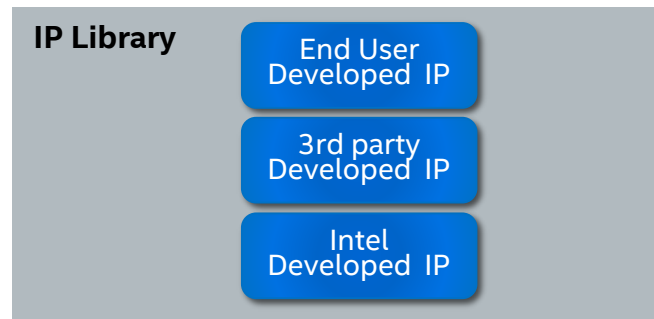
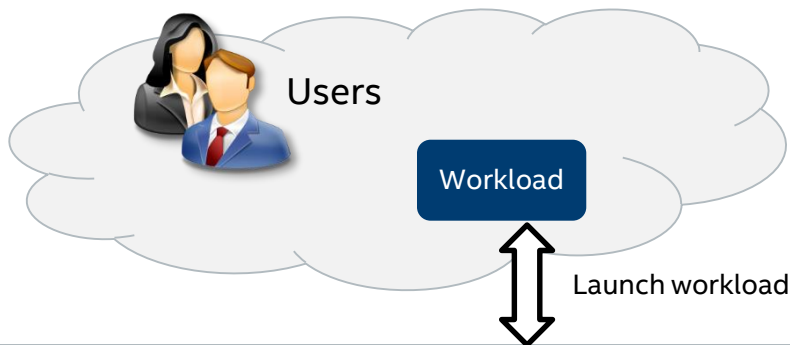
HDL Programming






OpenCL* Programming



Intel® Xeon® + FPGA in the Cloud



Software Defined Infrastructure



OpenStack* based sample solution available from Intel

Overview

- Data Center and Workloads
- Xeon+FPGA Accelerator Platform
- Deployment of FPGAs in the Datacenter
- Applications and Eco-system

Intel Machine Learning Strategy

Solutions

ADAS

Health & Life Sciences

Energy

Retail

Intel Solution Architects, Data Scientists, and Software Engineers

Trusted Analytics Platform



Open Source, ISV, SI, & Academic Developer Outreach




Optimized with Intel kernels / primitives for Deep Learning - NEW



Intel® Math Kernel and Data Analytics Acceleration Libraries

Linear Algebra, Fast Fourier Transforms, Random Number Generators, Summary Statistics, Data Fitting, ML Algorithms



+ FPGA

Support the industry **innovation** across verticals

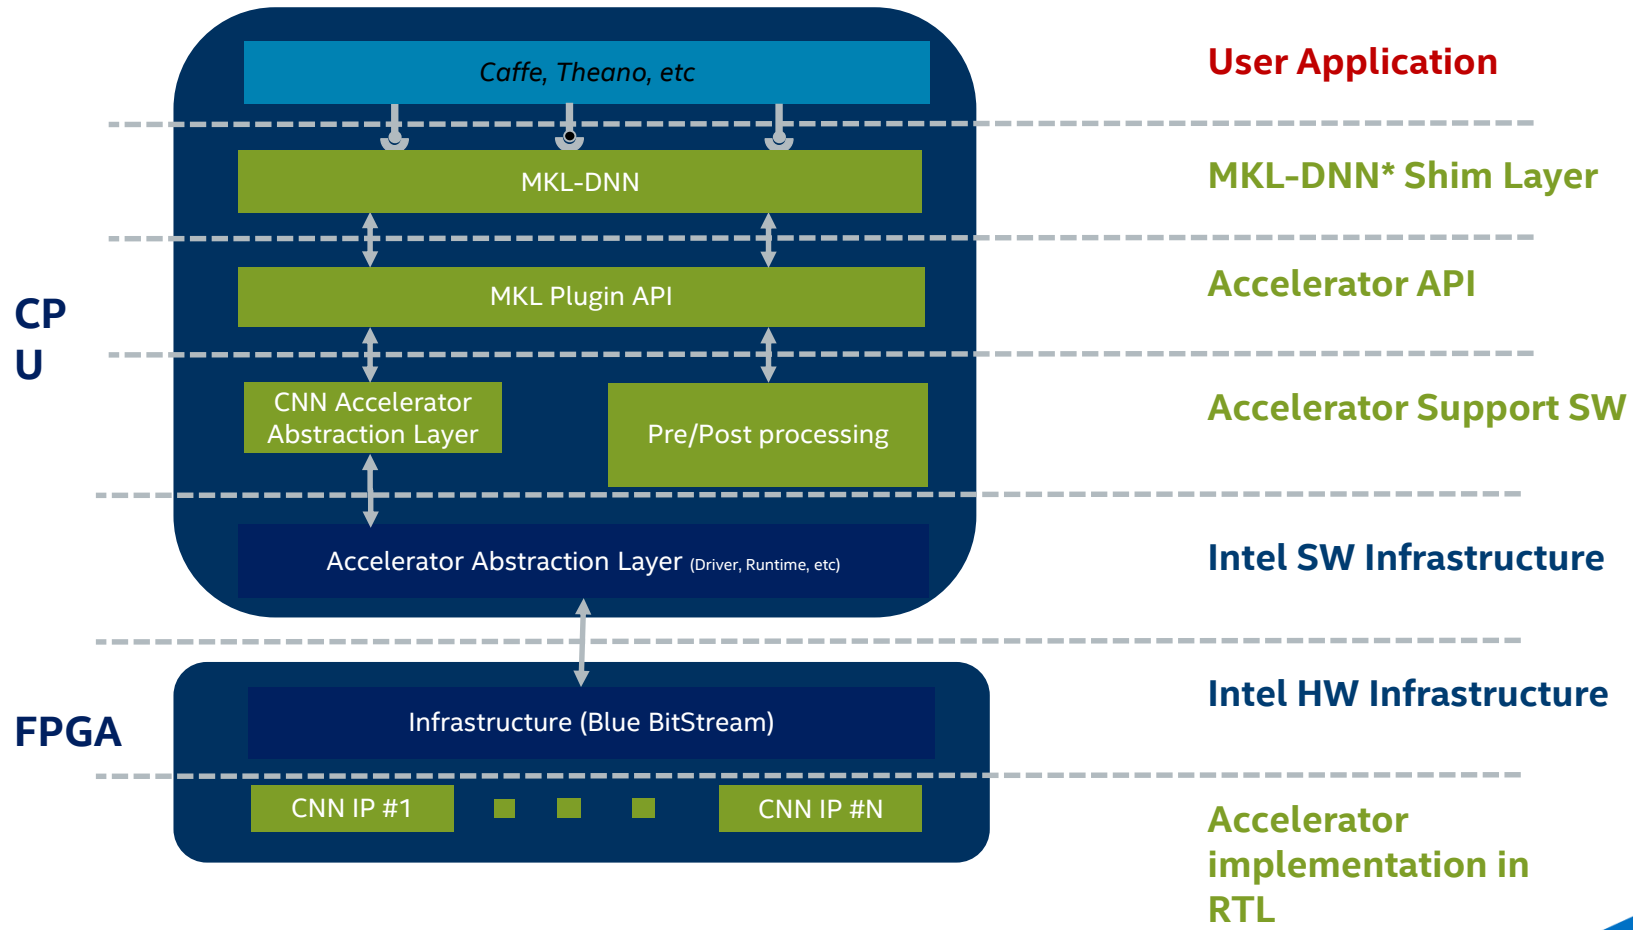
Accelerate adoption by providing tools to the ecosystem

Enable and **optimize** key industry frameworks

Extract maximum **performance** through libraries

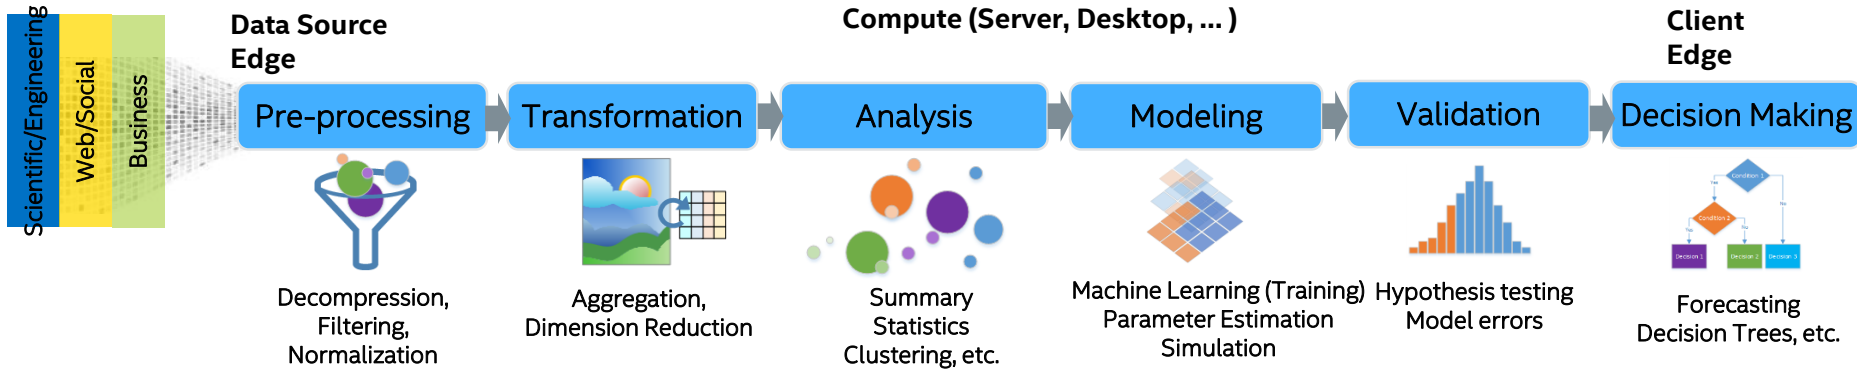
Enable optimization of **single-node and cluster** performance for Compute, Networking and Storage

Xeon+FPGA Software Stack for Machine Learning

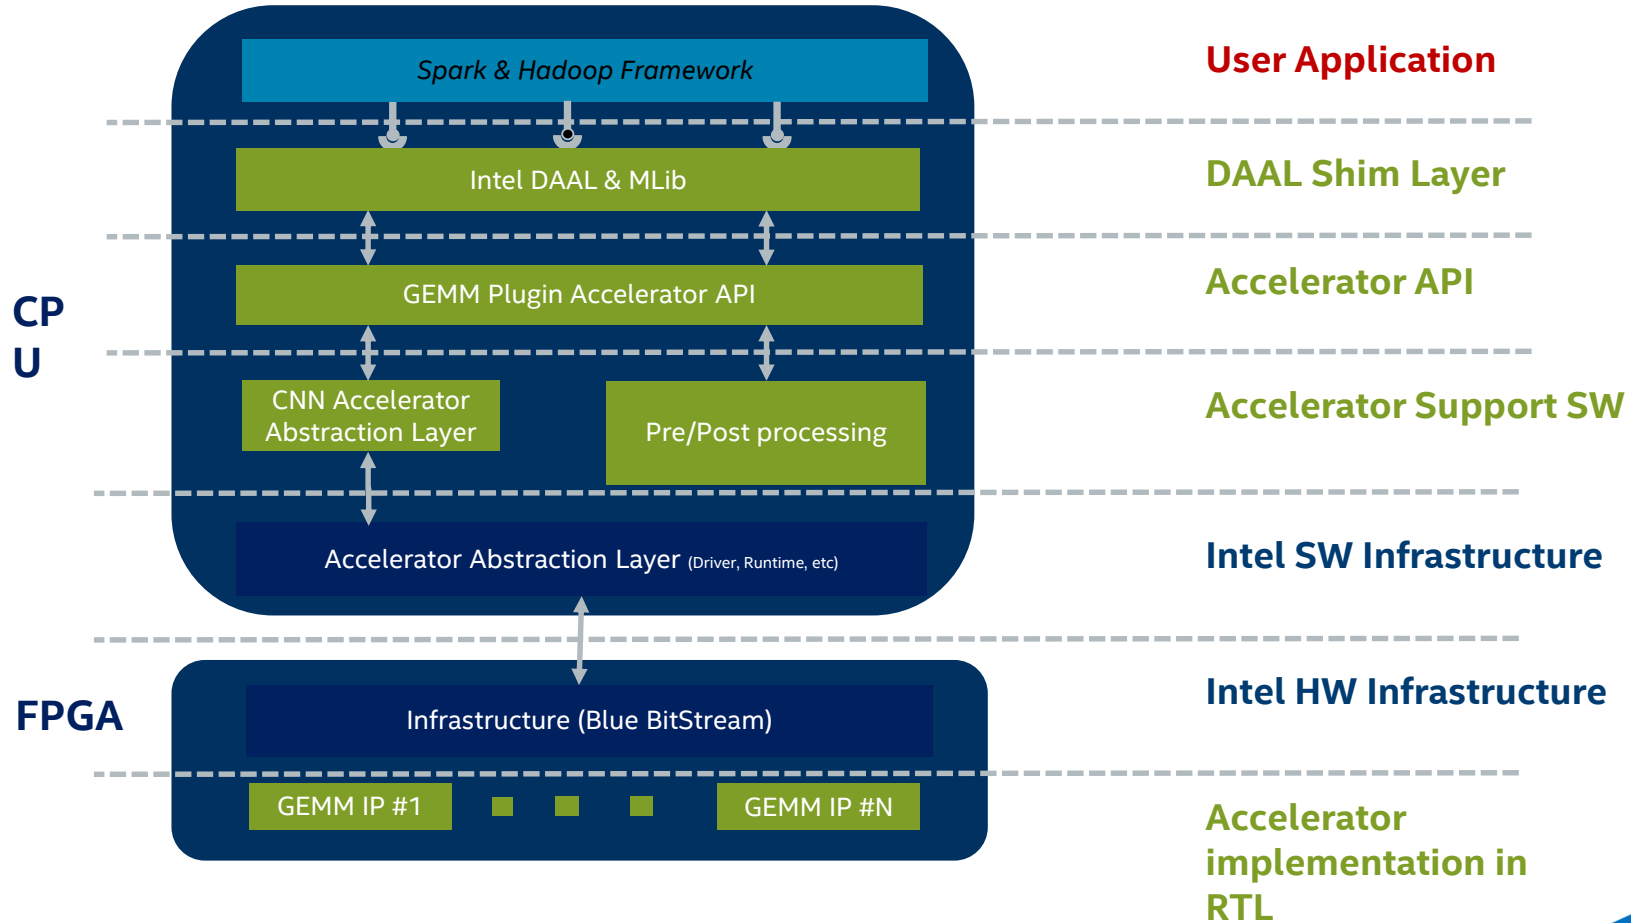


Intel® DAAL for Big Data Analytics

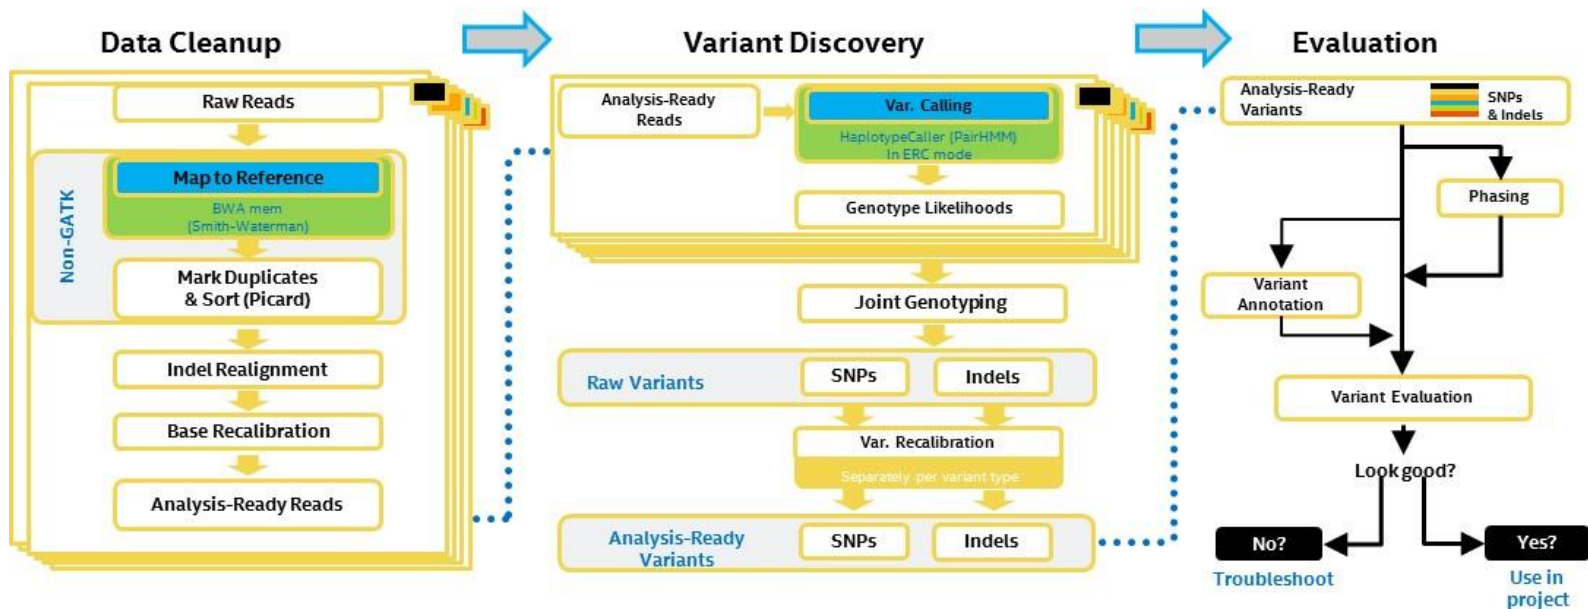
- Library of optimized building blocks covering all stages of the data analysis, from data extraction till data-driven decisions
- Targets both data centers (Intel® Xeon® and Intel® Xeon Phi™) and edge-devices (Intel® Atom)
 - Perform analysis close to data source (sensor/client/server) to optimize response latency, decrease network bandwidth utilization, and maximize security.
 - Offload data to server/cluster for complex and large-scale analytics only.



Xeon+FPGA Software Stack for SPARK/Hadoop

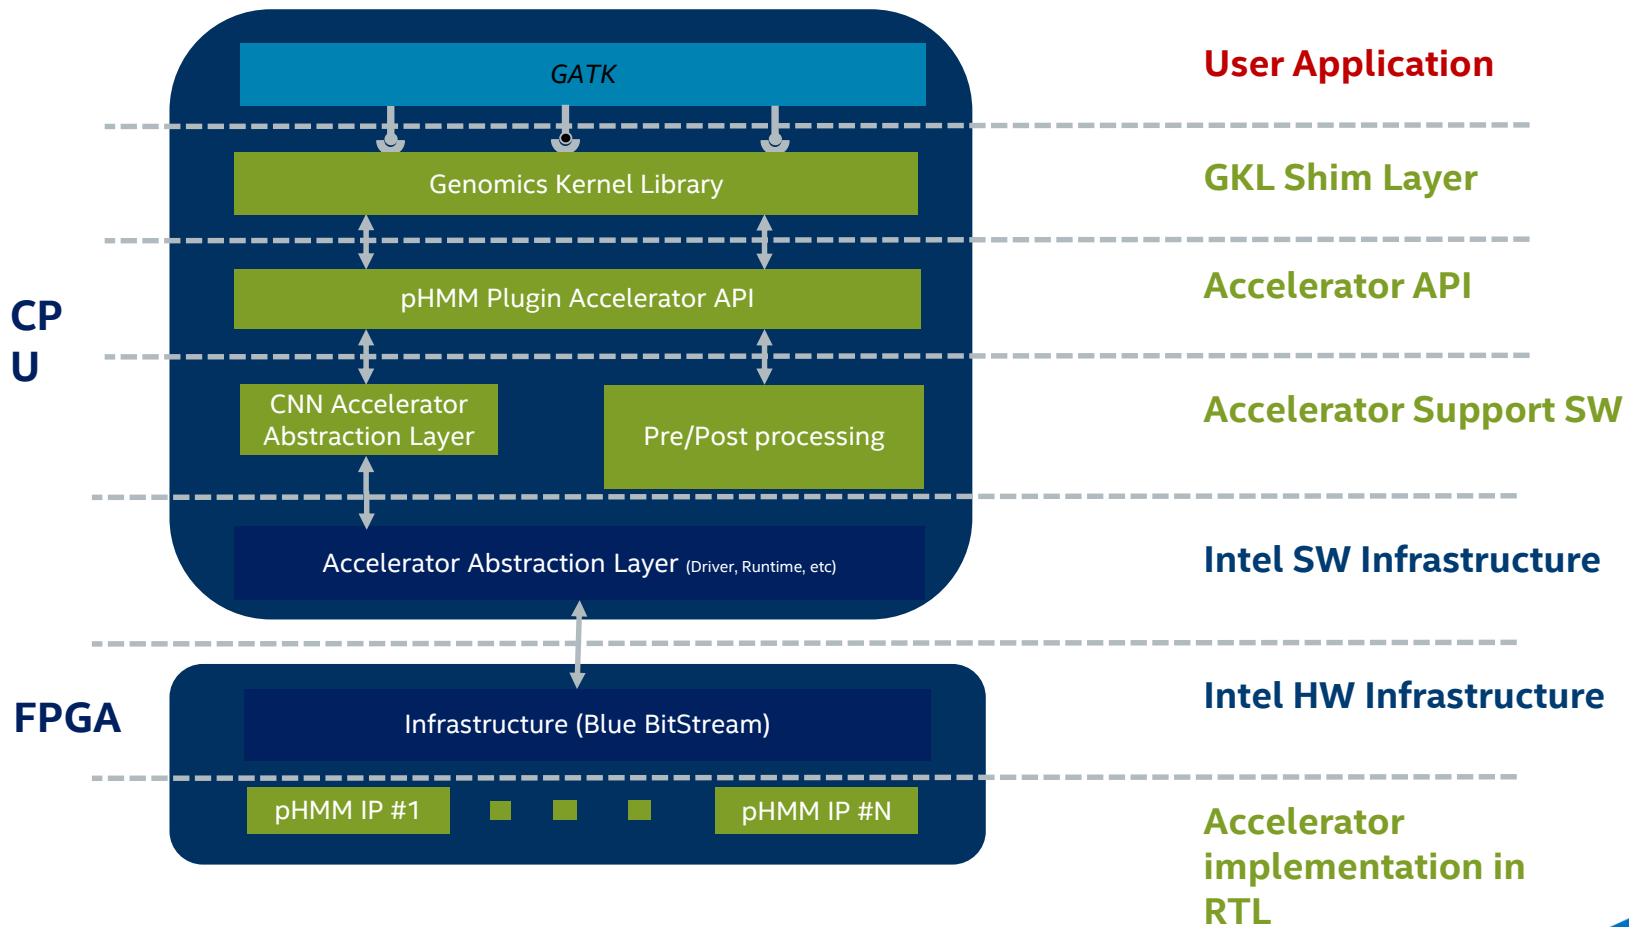


Genomics Analysis Toolkit



The Genome Analysis Toolkit or GATK is a software package developed at the Broad Institute to analyze high-throughput sequencing data.

Xeon+FPGA Software Stack for GATK



Academic Research in FPGA Usages : HARP 1

Call for Proposals 2015

Intel-Altera Heterogeneous Architecture Research Platform (HARP) Program

Intel® Corporation and Altera® Corporation are pleased to announce the Heterogeneous Architecture Research Platform (HARP) program, which will provide faculty with computer systems containing Intel microprocessors and an Altera Stratix® V FPGA module that incorporates Intel® QuickAssist Technology in order to spur research in programming tools, operating systems, and innovative applications for accelerator-based computing systems.

2015: Over 30 Intel® Xeon® E5-2600 v2 + FPGA systems shipped to universities in 2015



Global community, across continents

Winning Academic Mindshare with HARP 1

1. Academics are focusing on novel hybrid CPU—FPGA use cases

- Before: what can I offload to FPGA?
- Now: what's CPU great at? what's FPGA great at? how to collaborate?
- E.g., Genomics, Database, Graph/irregular, Sort

2. Academics are rethinking hybrid CPU – FPGA systems

- FPGA is becoming 1st class citizen, tighter integration to CPU
- What technologies needed to best take advantage of hybrid CPU-FPGA systems?
- E.g., JIT to FPGA, SPARK cloud + FPGA, OpenMP for FPGA

3. Academics are publishing on top FPGA conferences using Xeon-FPGA

- ISFPGA (Feb 2016): 1 out of 20 full papers use HARP
- FCCM (May 2016): 2 out of 18 full papers use HARP
- ASSP (July 2016) : 1 full paper based on HARP

Announcing HARP 2

- Intel® Xeon® E5-2600 v4 + Arria 10 FPGA to ship to universities in next few weeks for continued research in FPGA acceleration.
- Also, HARP 2 will be installed in clusters at sites in US and Europe.

Q & A