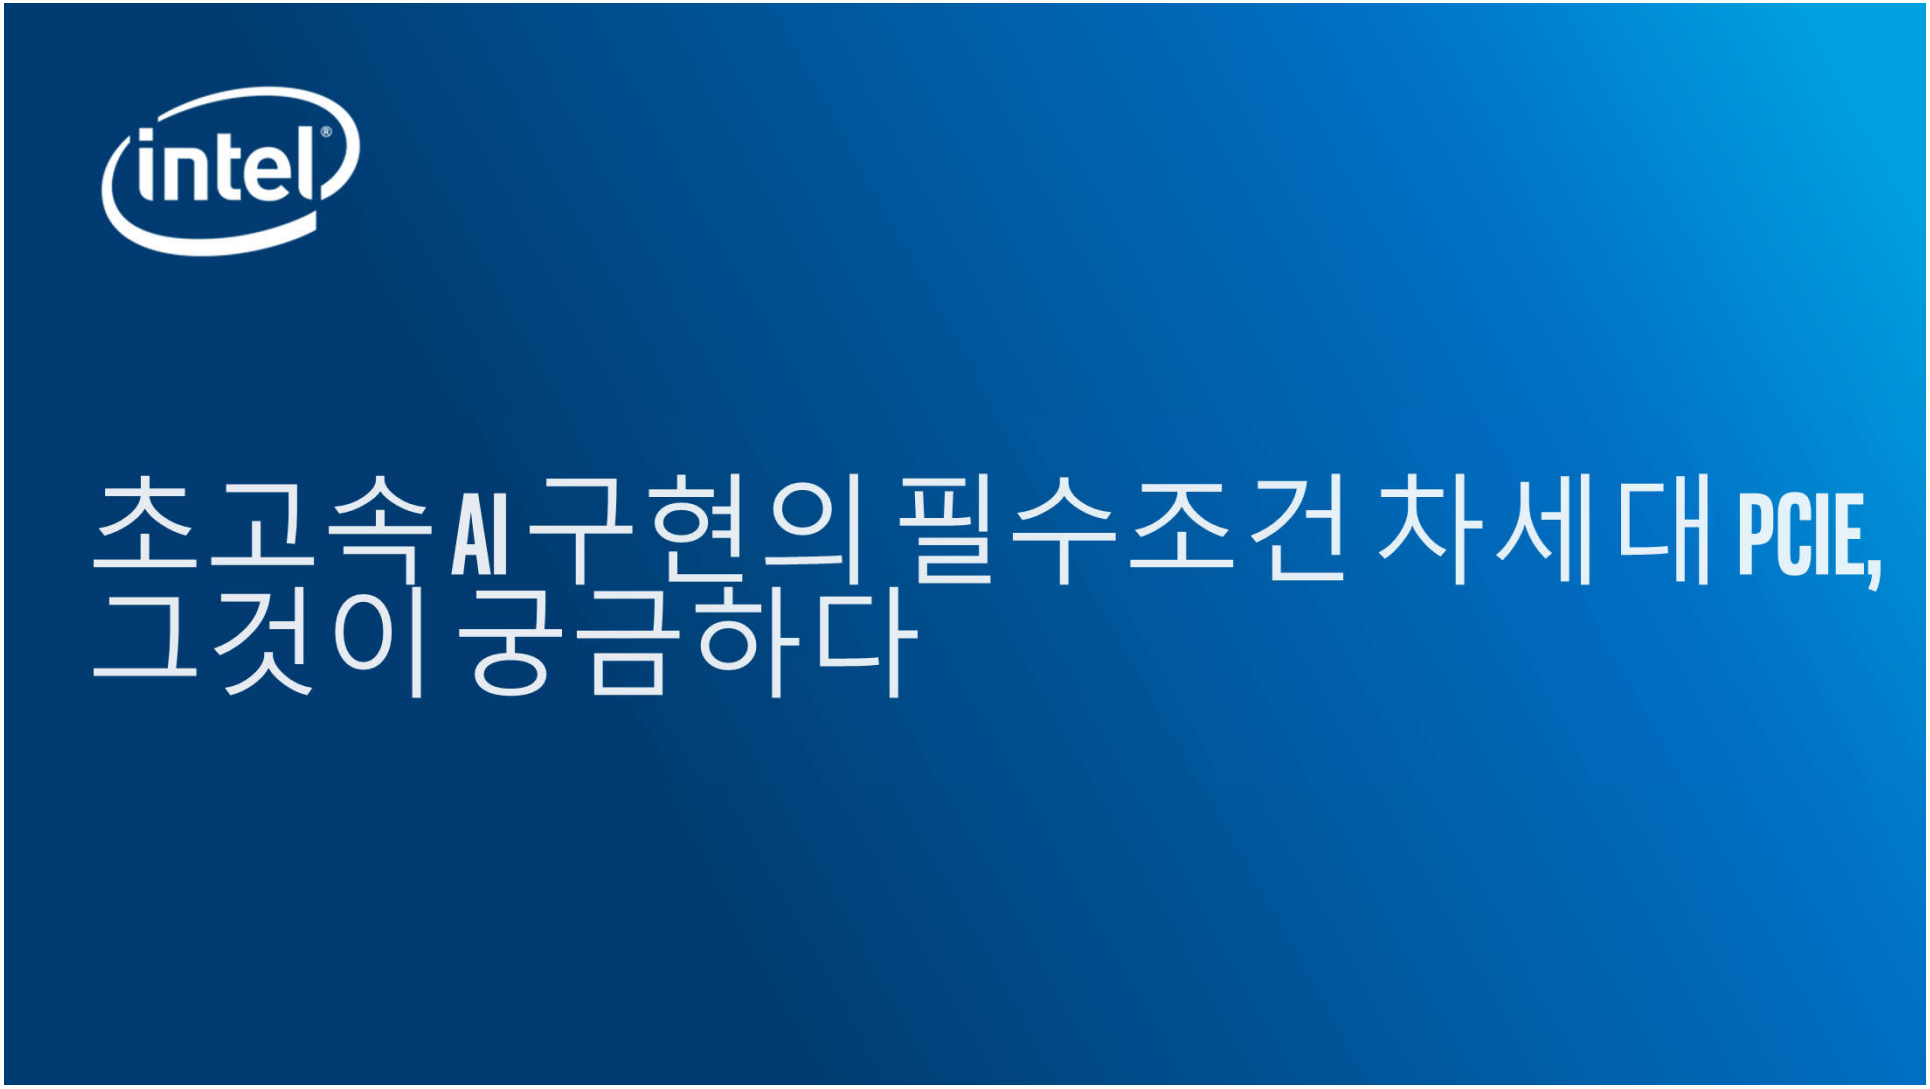
초고속AI구현의필수조건차세대PCIE,
그것이궁금하다

Agenda

- Intel® FPGA Connectivity Options
  - PCI Express (PCIe), Ethernet, PAM-4 Transceivers
- Intel® FPGA PCIe Multi-Generational Support
- Intel® FPGA PCIe Gen4 and Gen5 Features
  - Hard IP, Transaction Layer Bypass, Virtualization Intellectual Property (IP)
- Intel® PCIe Gen5 PHY Test Chip and Interoperability Enablement
- Intel® FPGA PCIe Development Kits
- Accelerator with Intel FPGA & Intel PAC
  - Intel® programmable Acceleration Cards : Accelerate 5g
  - FPGA based Data Center : Workload accelerators
- Summary

Intel® FPGA Connectivity Features to Move Faster

TODAY

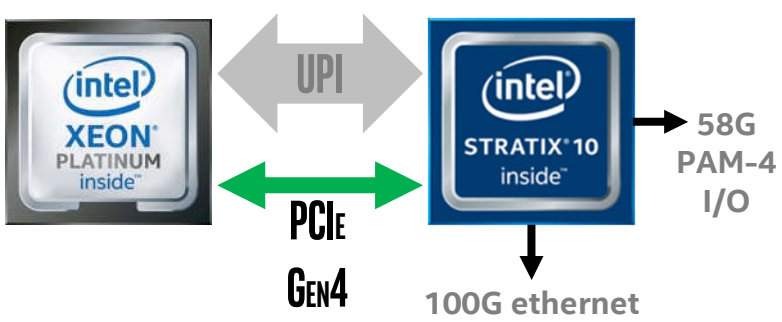

- 1st FPGA listed on PCI-SIG integrators list as PCIe Gen.4 compliant (x16 lanes, end point or root complex)
- New PCIe virtualization features and IP supported
- 58G PAM-4 and 100G Ethernet in volume production (since Feb. 2018)

COMING

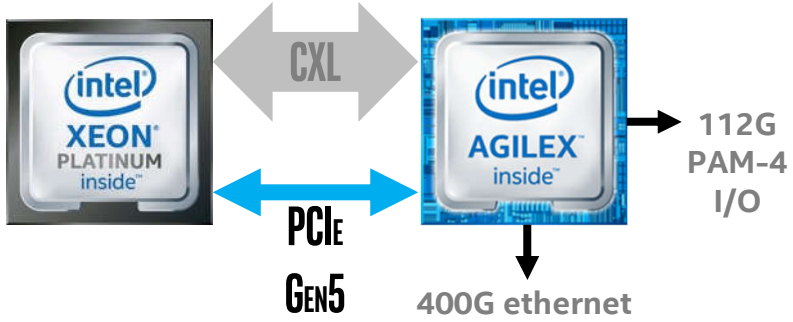

- PCIe Gen5 PHY test chip – demonstration at PCI-SIG Developer Conference (see YouTube video)
- 112G PAM-4 and 100G Ethernet test chip demo. (since OFC 2019)
- PCIe Gen6 will use PAM-4 and forward error correction FEC technologies

Intel® FPGA PCI Express Support: Multi-Generational New Features

Gen3 x8 Hard IP

- Gen3 x8
- ***SR-IOV 4PF/2048VF***
- ***Design examples***
- ***Signal Tap***

Stratix® V FPGAs

Intel® Arria® 10 FPGAs

- ***Gen3 x16***
- SR-IOV 4PF/2048VF
- Design examples
- ***Signal Tap II***
- ***Link Inspector***

Intel® Stratix® 10 GX/SX/MX/TX FPGAs

- ***Gen4 x16 with bifurcation***
- ***TLP bypass for PCIe switches***
- ***SR-IOV 8PF/2048VF, SVM, VirtIO***
- ***Scalable IOV***
- Design examples
- Signal Tap, Link Inspector

Intel® Stratix® 10 DX FPGAs

- ***Gen5 x16 with bifurcation***
- TLP bypass for PCIe switches
- SR-IOV, shared virtual memory (SVM), VirtIO
- Scalable IOV
- Design examples
- Debug tools

Intel® Agilex™ FPGAs

* Bold italic items are new features compared to prior generation.

Intel® Stratix® 10 FPGAs – Overview

Embedded Multi-die Interconnect Bridge (**EMIB**) packaging innovations

AIB electrical I/F = Now open sourced

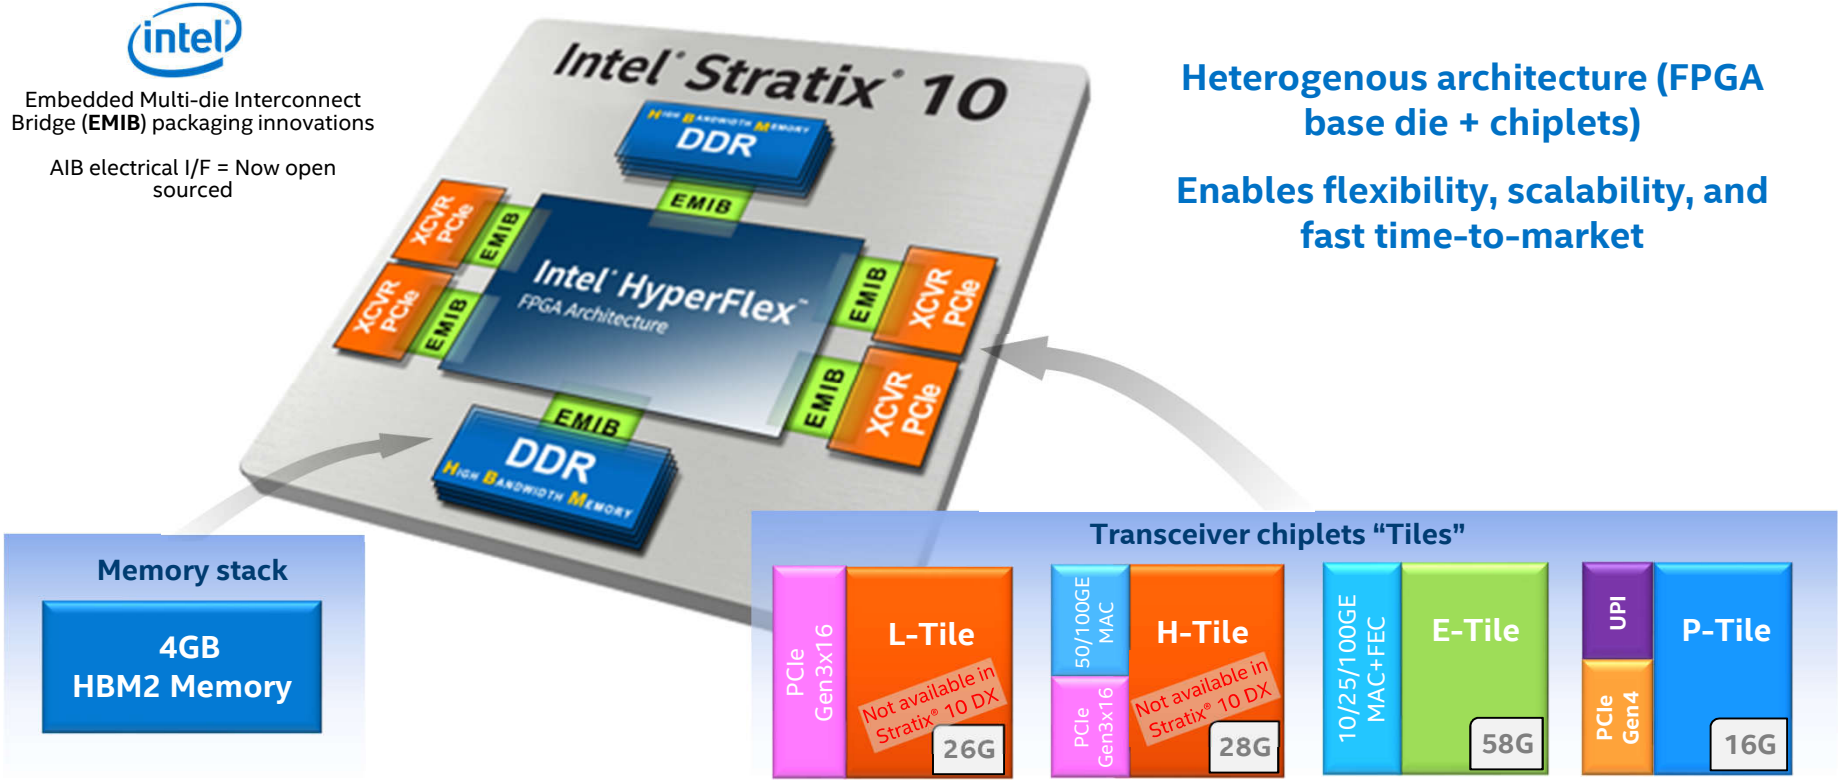

Heterogenous architecture (FPGA base die + chiplets)

Enables flexibility, scalability, and fast time-to-market

Intel® Stratix® 10 Portfolio Transceiver Tile Comparison

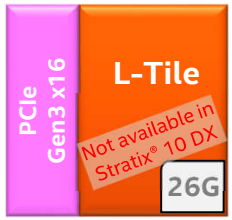

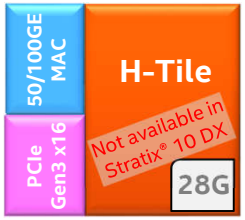

Transceiver Tile Variants	L-Tile (GX/SX)	H-Tile (GX/SX/TX/MX)	E-Tile (TX/MX/DX)	P-Tile (DX)
Max Chip-to-Chip	8x26.6G NRZ	16x28.3G NRZ	57.8G PAM-4 30G NRZ	16x 16G NRZ (PCIe Gen4), or 20x 11.2G NRZ (UPI)
Max Backplane	12.5G NRZ up to 25dB (>12.5G - 26G backplane use re-timers)	30dB @ 28.3G NRZ	30dB w/ FEC @ 57.8G PAM-4 35dB w/ FEC @ 30G NRZ 30dB w/o FEC @ 30G NRZ	PCIe Gen3/4 specification compliant
Hard IP	PCIe Gen3, x16	PCIe Gen3, x16 with 4 PF/2K VF SR-IOV 50/100GE MAC	10/25/100GE MAC, PHY and RS-FEC, KP-FEC	PCIe Gen4, x16 with 8 PF/2K VF SR-IOV, EP/RP. UPI PHY and Link Layer
Power (same rate)	1.0	1.0	~ 0.6	~ 0.7
Pin Migratable	Yes		No	No

Production device capabilities shown above.
Some features not available in Engineering Sample (ES) silicon. Refer to silicon and software/IP rollout schedules for feature enablement details.
Not all channels support the maximum data rate specification. Refer to transceiver tile descriptions and ordering information for channel count details.

Intel® Stratix® 10 Portfolio

Up to 512 GB/s HBM2 DRAM Bandwidth

Dual Mode 58G PAM-4 30G NRZ Transceivers

Quad Core 64-bit ARM A53 HPS up to 1.5 GHz

HyperFlex Architecture
Up to 10 TFLOPs DSP
Secure Device Manager
Monolithic Fabric
28G NRZ Transceivers (GX/SX/TX/MX)
PCIe Gen4 or UPI (DX)

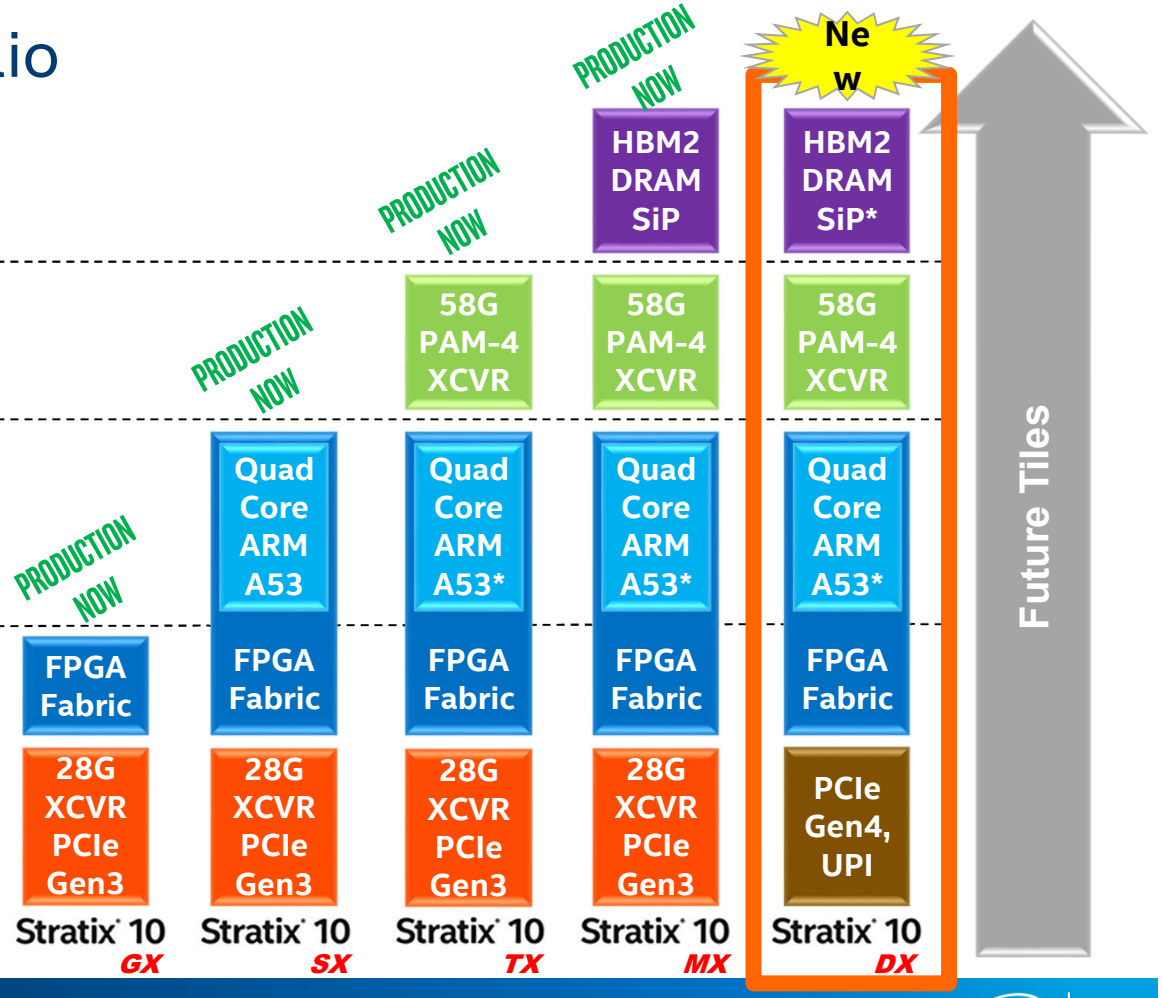

* Not all features are available in every device.

PCIe Gen4 IP Advantages (Intel® P-Tile Implementation)

https://pcisig.com/developers/integrators-list

PCIE 4.0

- PCI-SIG CERTIFIED
- BASE SPEC REV1.0
- MULTIPLE MODE SUPPORT (EP/RP/BIFURCATION)
- VALIDATION SHOWS EXCELLENT RESULTS

PCIe Gen4 TX Eye, 16G, PRBS15

ADVANCED CLOUD FEATURES

- VIRTIO SUPPORTED
- BEST IN CLASS[1] SR-IOV AND SCALABLE IOV
- SHARED VIRTUAL MEMORY (SVM) FOR ZERO COPY

SWITCH IP

- TLP-BYPASS TO ENABLE PCIE SWITCH APPLICATIONS
- COUPLED WITH VIRTIO AND SMARTNIC IP
- OPTIMIZED FOR SMALLEST ADAPTIVE LOGIC MODULE (ALM) COUNT
- 75% LESS THAN 3RD PARTY[2]

1SD280 P-Tile 4x locations

(top down view)

1. Compared to Xilinx FPGA offerings.
2. Compared to 3RD party vendor implementation of PCIe Soft IP switch, for FPGAs, using PCIe PIPE interface. See pages 10-11 for more details.

Intel® P-Tile: PCI-SIG Gen4 Compliant

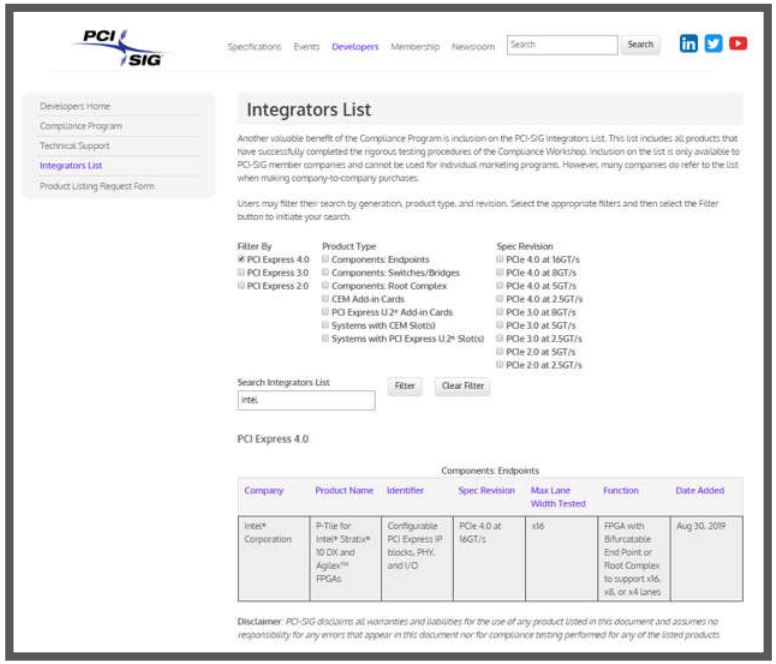

- Attended compliance workshop #110 (Aug. 2019)
- 100% compliant = Passed 12 of 12 vendor tests[1]
  - Includes CPU, ASSP, and test equipment suppliers
- PCI-SIG integrators list

Intel® Stratix® 10 DX Devices

- Same P-Tile used in Intel® Agilex™ FPGAs

Intel® Agilex™ FPGAs (F-series devices)

1. Contact the Intel® FPGA team for compliance test result details.

EMIB = Embedded Multi-die Interconnect Bridge

Intel® P-Tile: PCI Express Hard IP

Hard IP has pass PCI-SIG compliance testing:

- PCIe Gen4 Base Spec. Rev. 4.0, Ver1.0
- PCIe Gen3 Base Spec. Rev. 3.1

PCIe x16 lanes support (Gen3 and Gen4) per Tile

- Both root-port (RP) and end-point (EP) modes supported

Bifurcation Support (Tile can be split into independent links)

- EP: Two Gen4 x8 or two Gen3 x8
- RP: Four Gen4 x4 or four Gen3 x4

Mode control: Intel® Quartus® Prime Design S/W (at compile time)

Hard IP Bypass Mode

- Allows realization of PCIe switch and other functionalities in the fabric core

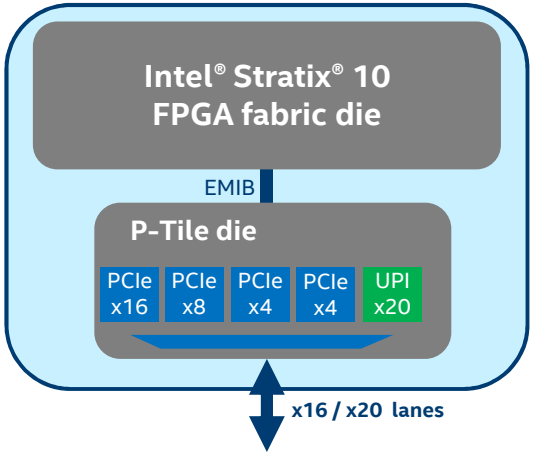

Intel® Stratix® 10 DX FPGA Development Kit (4x P-Tiles, 1x E-Tile)

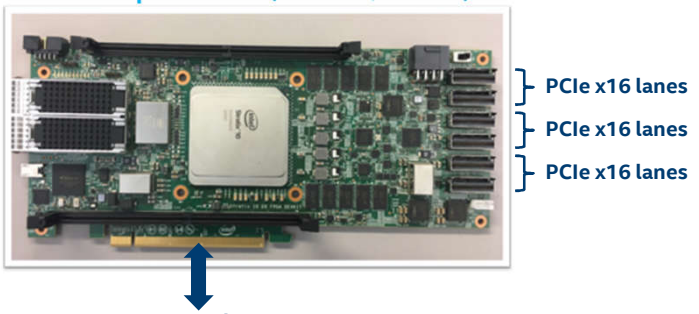

Intel® P-Tile: Hard IP Bypass Mode

Bypasses Transaction Layer (TL)

- Hard IP used as a 'TL packet (TLP) data transporting' link
- Malformed TLPs dropped with error notification
- In this case, both configuration registers (e.g. BARs) and TLP need to be implemented in FPGA user logic

Incoming TLPs 'handed over' to User Logic

- Hand-over to/from user logic using Avalon® -ST hard IP for PCIe

'Lite'-weight TL block needed (user design or 3rd party)

- Some configuration registers still used by Hard IP
  - E.g. Registers related to link state
- User logic can access these registers with Avalon® Memory-Mapped IP for PCIe

Applications: PCIe switch (SWUP/SWDN), VirtIO

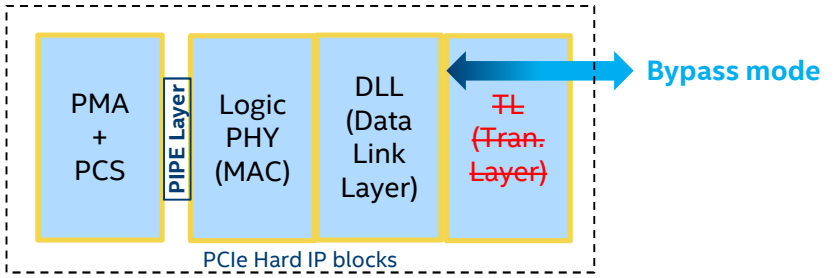

Example of PCIe Hard IP in Bypass mode

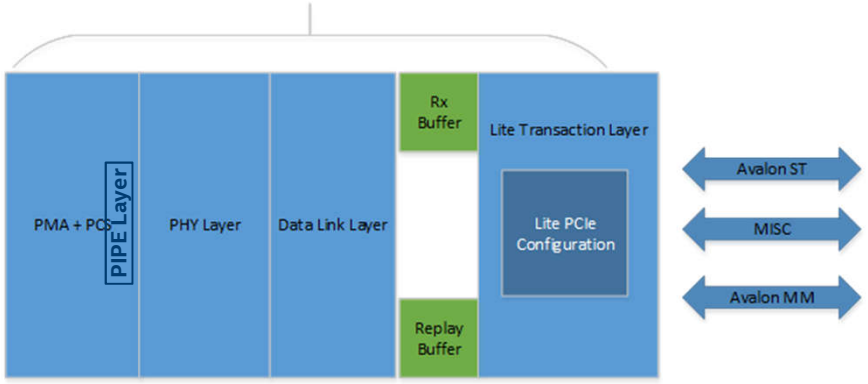

Switch Port IP logic size:
- With PIPE attach: ~43K ALMs

Implementing Soft PCIe Switch IP (PCIe PIPE)

SW-UP Port
- Address based routing:
  - Admit only if *TLAddr* within UP windows
  - Standard error detection (malformed, ECRC etc.)
- ID-based routing
  - Forward TLP if bus # >= Secondary bus #
  - Handle *CFG0* internally

Switch Fabric
- Address based routing:
  - Route TLP to the corresponding DN when inside its windows
  - Error detection and generation (Generate CPL/Error message)
- ID-based routing
  - If bus number = UP secondary bus number
    Program DN config space register and send CPL (SC or UR)
  - If bus number = DN secondary bus number
    Convert *Type0* and send to DN port
  - Else forward it
- Arbiter Mux on outgoing TLPs

SW-DN Port
- No *checkinf* is done on transmitter (TX)
- Receiver (RX) checks that TLP is outside its window

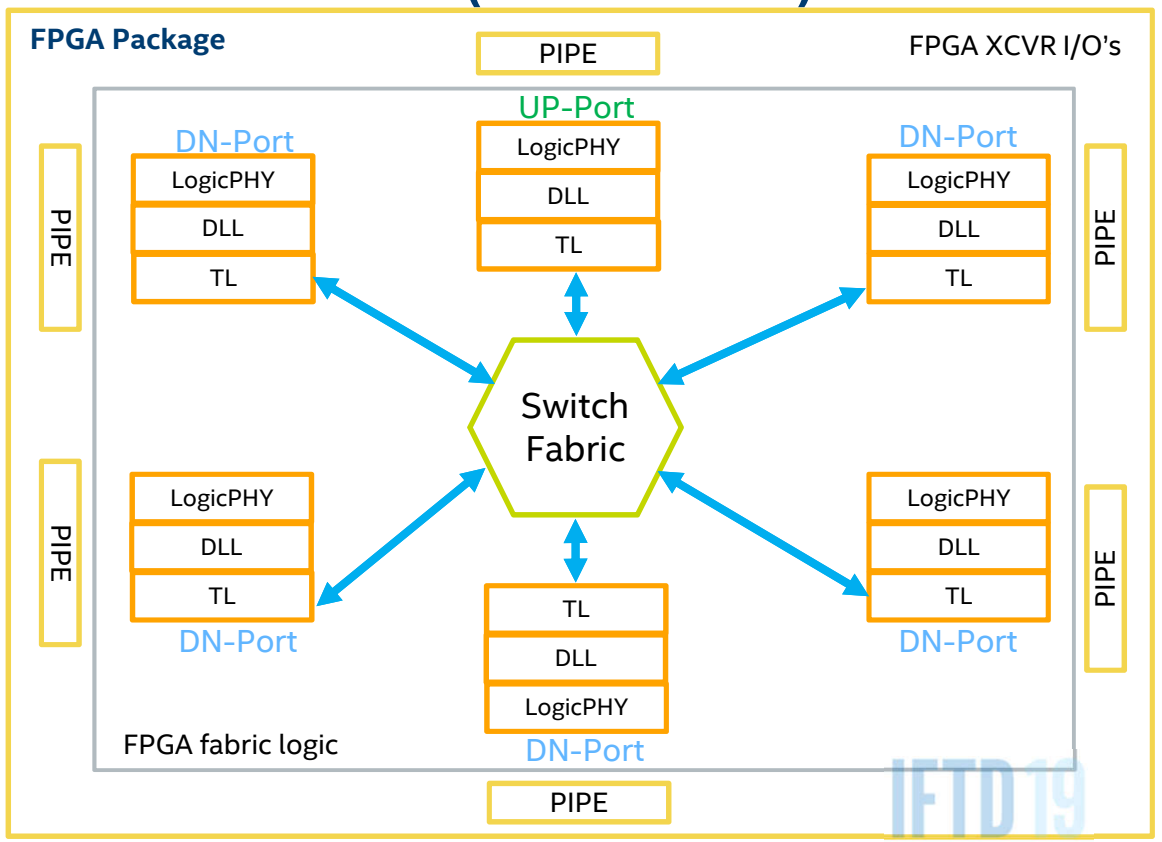

Implementing Soft PCIe Switch IP (P-Tile)

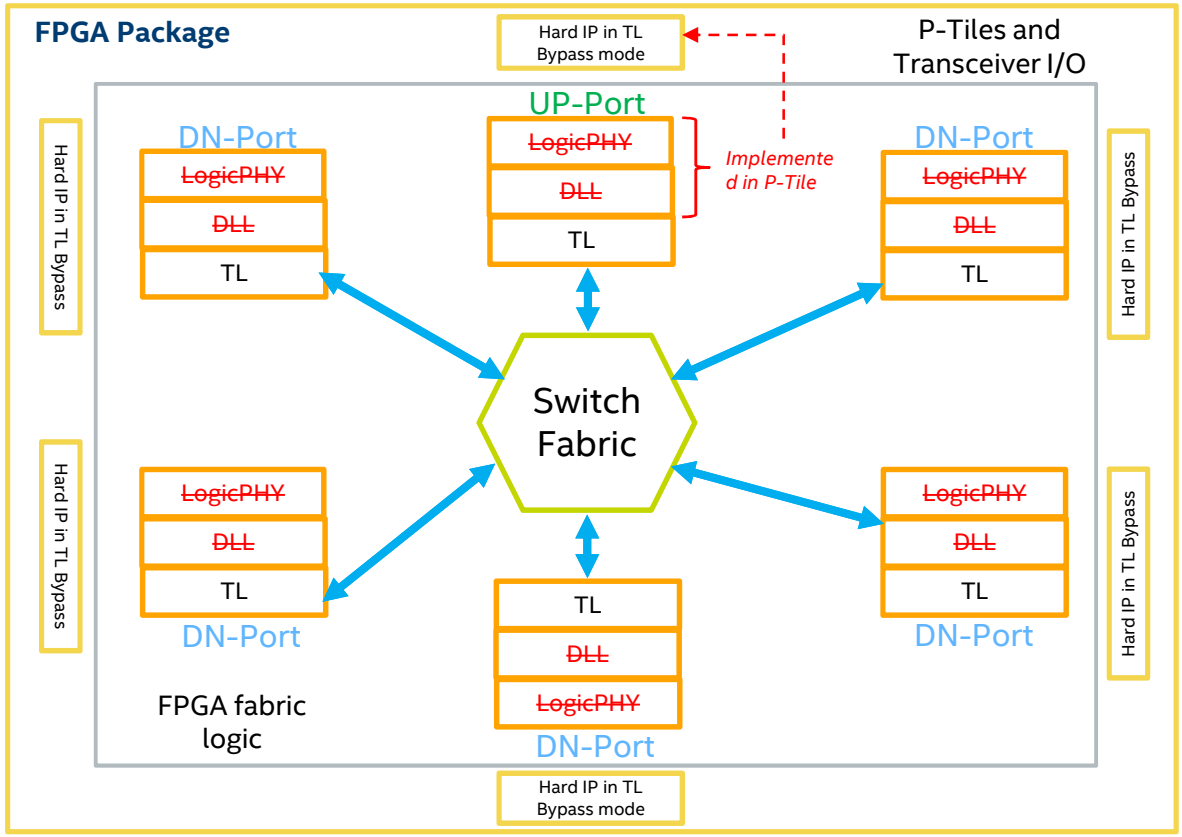

Switch Port IP logic size:

- With PIPE attach: ~43K ALMs
- With TL bypass: ~6K ALMs

Result = Significantly less logic and lower power

1. Based on Intel® estimates for P-Tile vs. 3RD party vendor implementation of PCIe Soft IP switch, for FPGAs, using PCIe PIPE interface = 43K ALMs reduced to 6K ALMs.

Intel® P-Tile: PCIe Functionality Accessed via Avalon® Streaming (Avalon-ST) Hard IP for PCIe

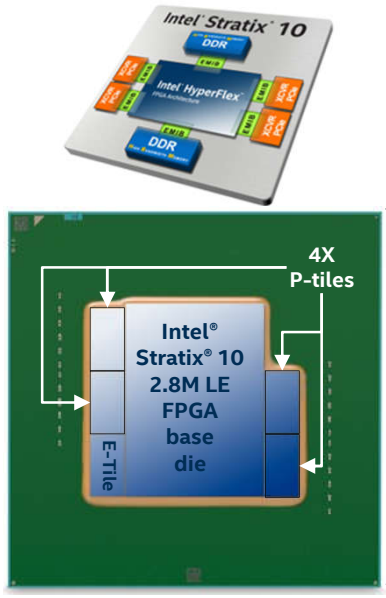

Actual picture of Intel® Stratix® 10 DX device (de-capped), used for illustration purposes only.
(Top Down view)

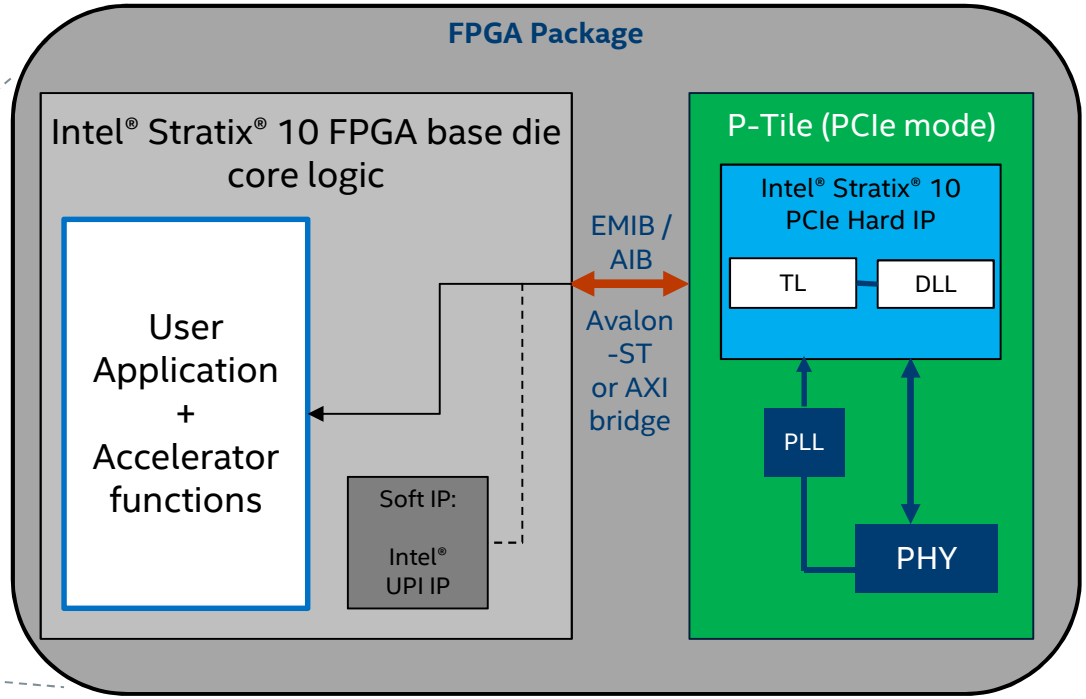

Note: Intel® AIB = Advanced Interface Bus. A die-to-die interface used with Intel's Embedded Multi-die interconnect Bridge (EMIB) packaging technology to create a physical and electrical interface between an FPGA die and chiplet/tile die. Intel® AIB is now available via royalty-free license: http://github.com/intel/aib-phy-hardware

Intel® P-Tile: PCIe Functionality Accessed via Avalon® Memory-Mapped (Avalon-MM) IP for PCIe

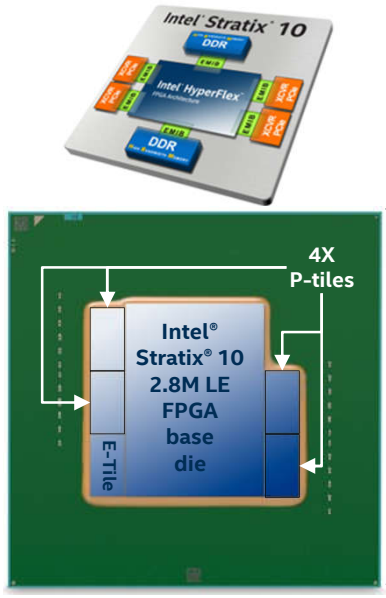

Actual picture of Intel® Stratix® 10 DX device (de-capped), used for illustration purposes only.

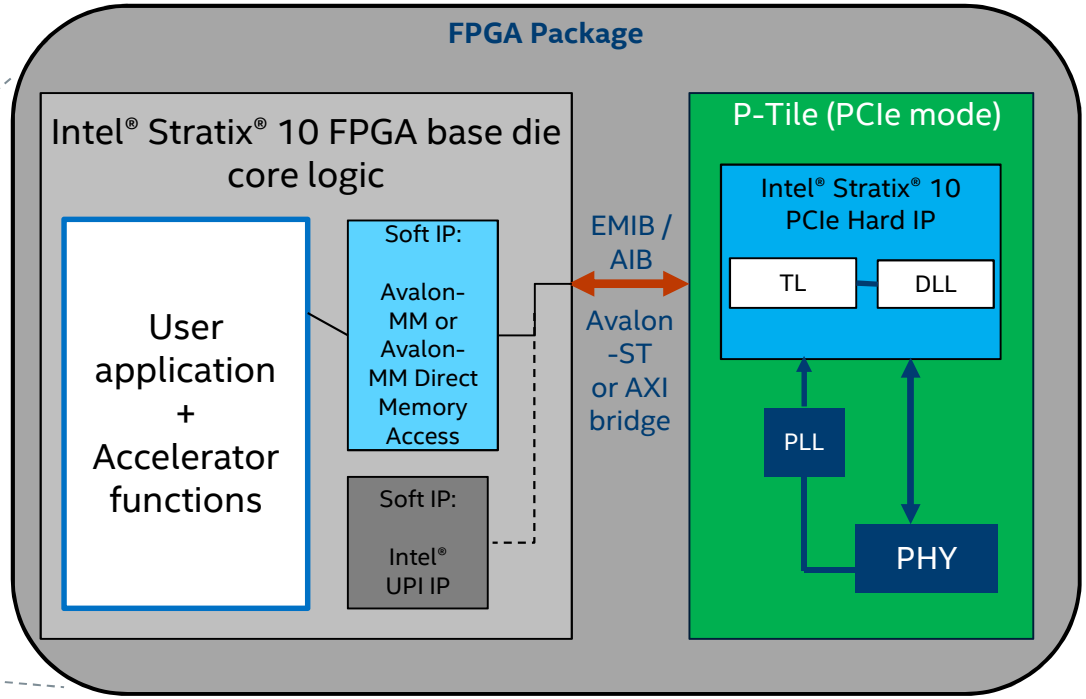

*AIB = Advanced Interconnect Bus. A die-to-die interface used with EMIB packaging technology to create physical & electrical interface between FPGA die and chiplet/tile die. AIB is now available via royalty-free license: http://github.com/intel/aib-phy-hardware

Intel® P-Tile: PCIe Advanced Features for Virtualization or Accelerators

Virtualization Trend

- To increase performance and maximize physical compute resources (CPU, network interface card (NIC), storage, etc.) in a virtual architecture (cloud, virtual machines (VMs), containers, dis-aggregated servers, etc.)

P-Tile PCIe* can support multiple virtualization options and/or accelerator functions

1. Single Root I/O Virtualization (SR-IOV)
2. VirtIO using PCIe PFs / VFs
   - Linux Para-virtualization with SR-IOV
3. Scalable I/O virtualization (Scalable IOV)
   - Allows for dynamic virtual device creation
   - More granular than SR-IOV VFs
4. Shared virtual memory (SVM)

Selection based on requirements for:
- Latency
- # of VMs/containers
- North/South vs. East/West traffic patterns
- NIC hardware
- Etc.

Used for FPGA accelerators
- Zero-copy functionality, etc.

Single Root I/O Virtualization

Contact Intel® FPGA marketing for more details on SR-IOV network functions virtualization (NFV) solutions or Intel programmable acceleration cards (Intel PACs).

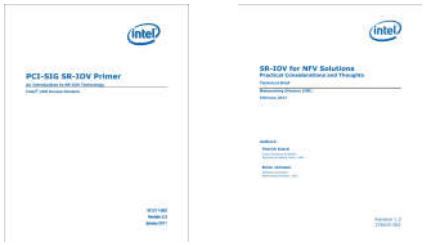

https://www.intel.com/content/dam/doc/application-note/pci-sig-sr-iov-primer-sr-iov-technology-paper.pdf

https://www.intel.com/content/dam/www/public/us/en/documents/technology-briefs/sr-iov-nfv-tech-brief.pdf

Intel® P-Tile has *Best in Class*[1] FPGA Single Root I/O Virtualization (SR-IOV) features

- Eight Physical and 2,048 virtual functions
  - Implemented in hard-IP
- Customer benefits include:
  - Supports more VMs[1]
  - Uses significantly less FPGA user logic (vs. competition)[2]
- SR-IOV generally has lower latency and uses less CPU resources vs. other virtualization options

2X – 8X more than FPGA competitor[1]

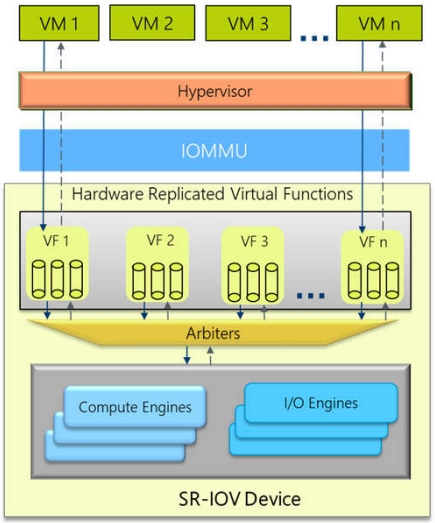

1. Compared to Xilinx FPGA datasheet specifications.
2. Comparison of PCIe features implemented in Hard vs. Soft IP, versus Xilinx FPGAs.

VirtIO

Contact Intel® FPGA Marketing for more details on Intel FPGA PACs with VirtIO Solutions

Virtualization solution completely implemented in software

Part of standard Linux library *libvirt* (most Linux versions)

- Customer benefits include:
  - Runs on any hardware
  - Easier software coding
  - Easier VM-to-VM migration

Contact Intel® FPGA team for more details on VirtIO Solutions

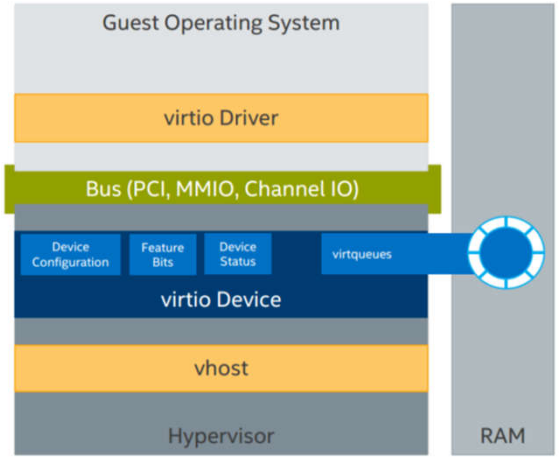

https://wiki.libvirt.org/page/Virtio

Sources: https://www.juniper.net/documentation/en_US/junos/topics/concept/disaggregated-junos-virtio-sr-iov.html and https://www.stratoscale.com/blog/compute/hypervisor-comparison-of-io-virtualization-models/

Intel® Scalable I/O Virtualization (IOV)

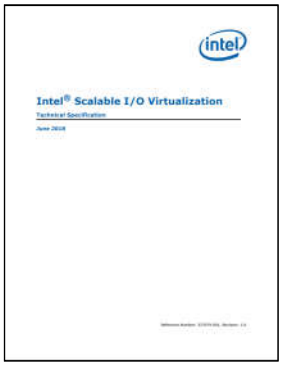

https://software.intel.com/sites/default/files/managed/cc/0e/intel-scalable-io-virtualization-technical-specification.pdf

New I/O virtualization architecture → For dynamic creation of virtual devices

- Lighter weight and more scalable compared to SR-IOV
- Can co-exist with existing I/O virtualization solutions
- Utilizes existing PCIe capabilities e.g Process Address space IP (PASID TLP Prefix) which "… enables significantly more number of domains to be supported…"[2]
  - ***supported in Intel's P-Tile***

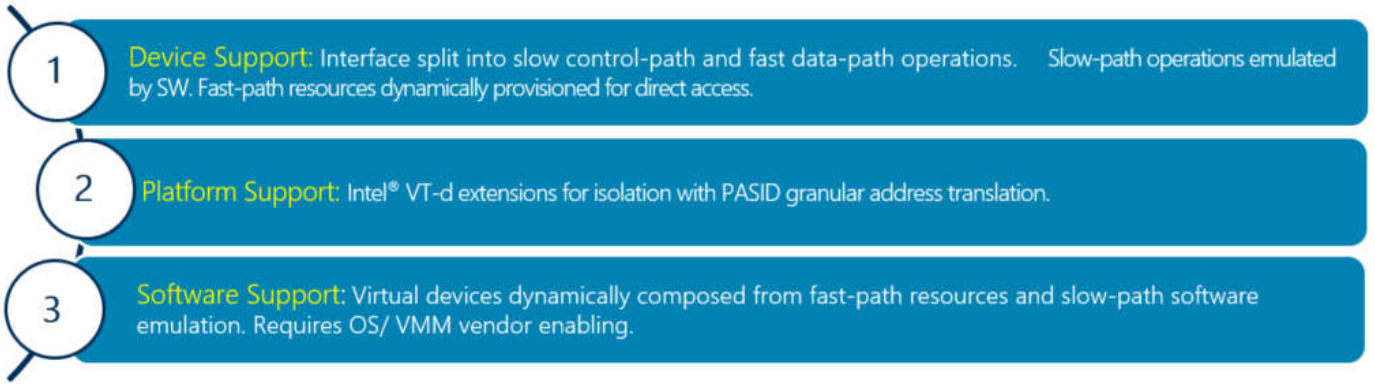

Warning

CPU & software support required to implement this feature.

1. Implementation of Intel® Scalable IOV requires compliant End Points devices (i.e. FPGA) combined with compliant Host System (i.e. CPU) and Operating Systems.
2. Intel® Scalable I/O Virtualization specification, section 2.2.3.

SR-IOV vs Intel® Scalable IOV: Hardware View

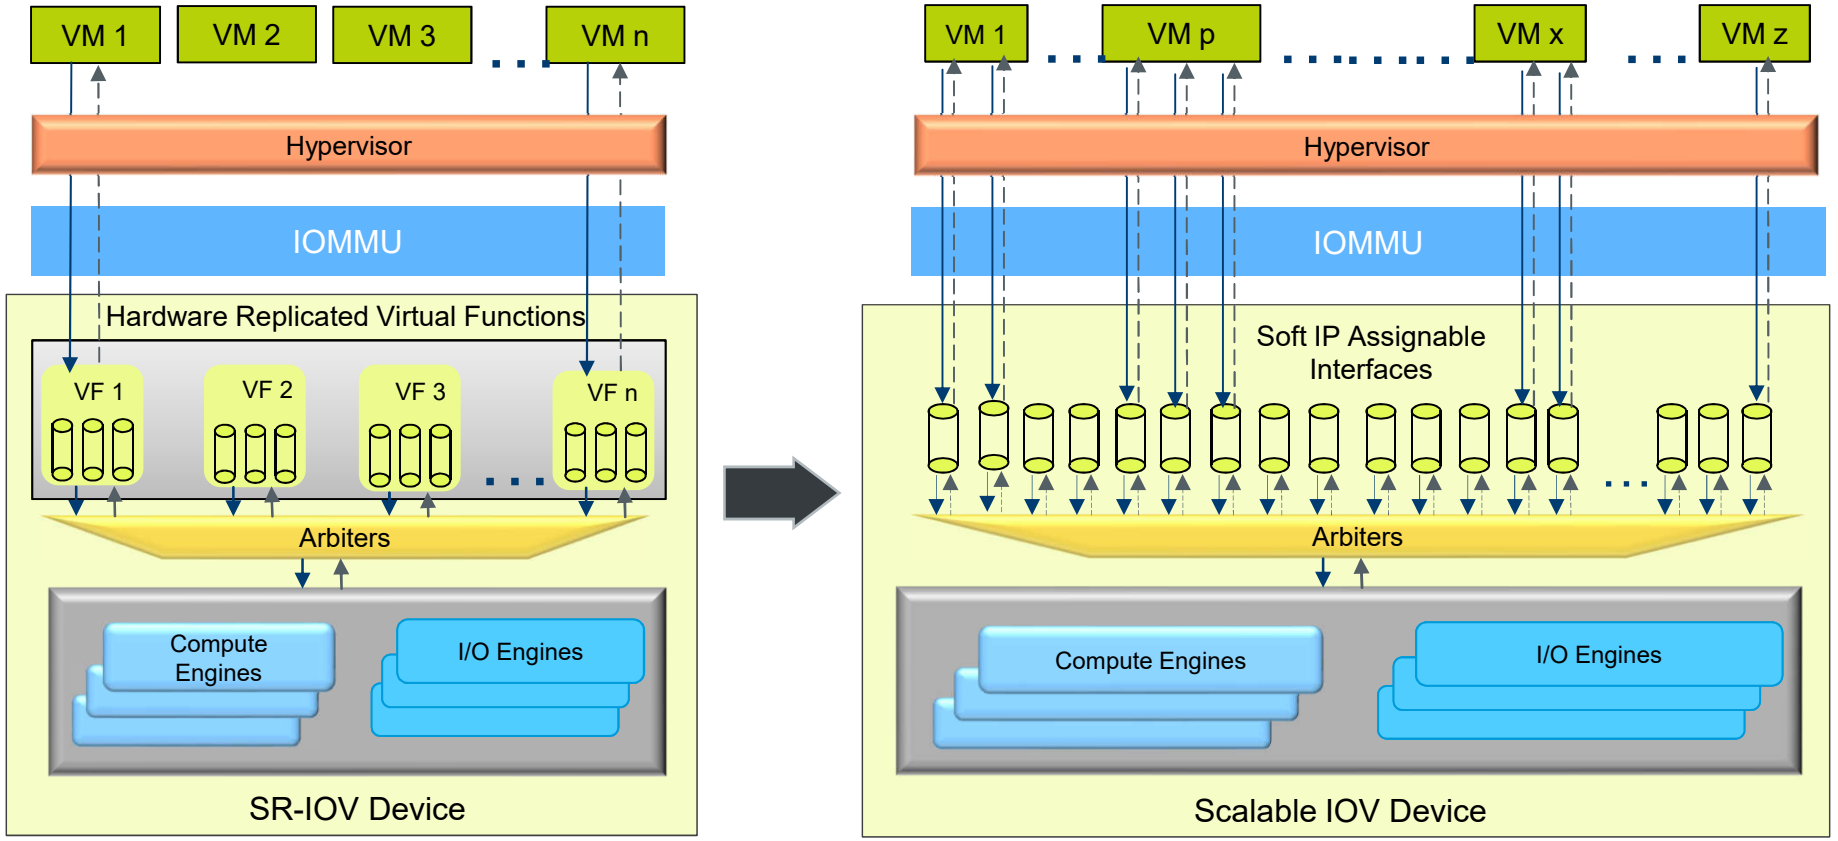

Shared Virtual Memory (SVM)

Simplified Zero Copy design example available on Github. Contact Intel® FPGA Marketing for .PPT overview of design or POC demo(s).
https://github.com/OPAE/intel-fpga-bbb/wiki/MPF-VTP-Virtual-to-Physical

For Fixed-function Accelerators

- Customer benefits:
  - Higher performance[1]
  - Enables additional functions zero-copy, memory pining, and support for large or irregular working sets

For *Programmable* Accelerators (such as OpenCL™ SVM)

- Customer benefit: Low-latency cache coherency
- Specified with memory consistency models (coarse-grained vs. fine-grained coherency, synchronization primitives via atomics)

Programmer's Perspective

- Seamless sharing of pointer rich data structures between host and task offloaded to a device (accelerator)
- Without SVM, a programmer or toolchain is required to marshal data between host and device view of memory

1. Based on internal Intel® FPGA estimates.

Intel® Agilex™ FPGAs: Product Families / Variants

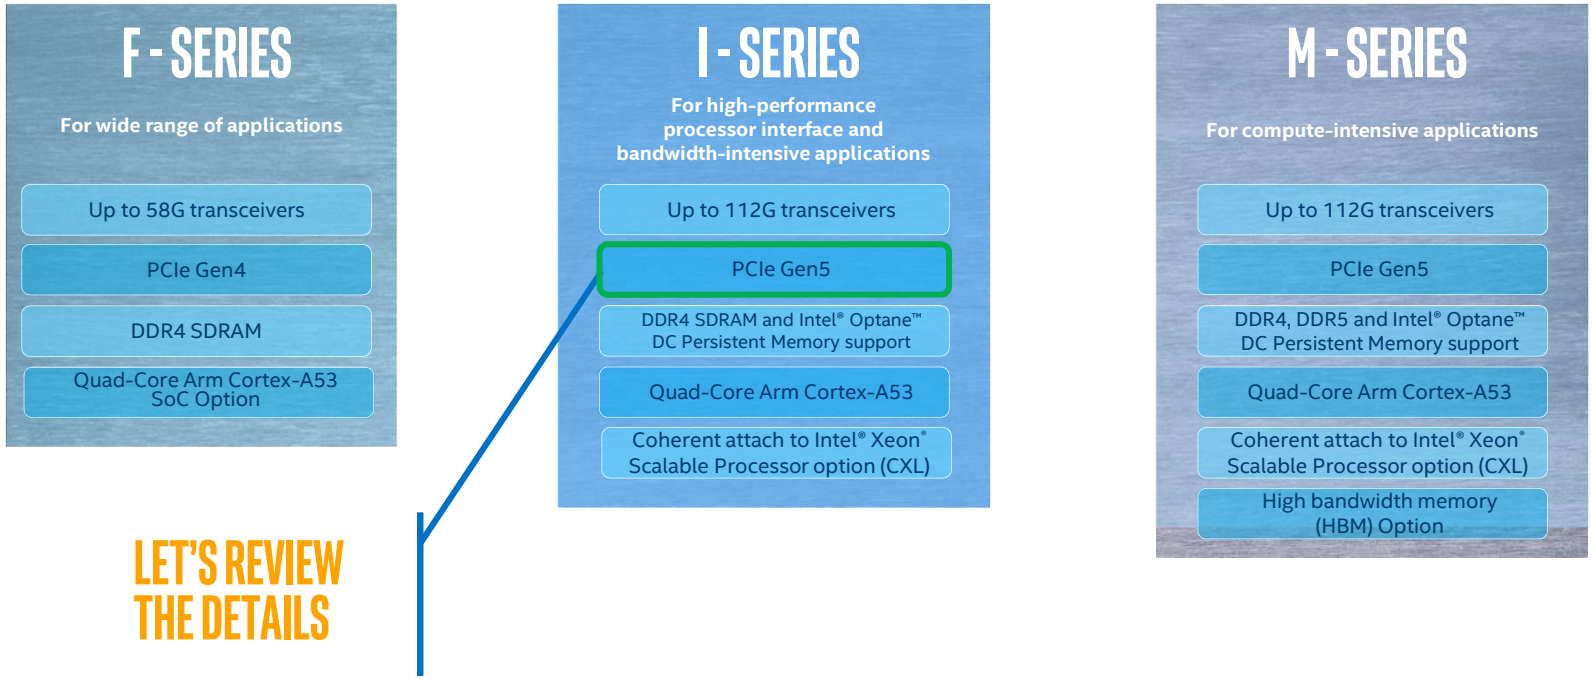

Intel® Agilex™ I-Series FPGAs: 6 Package + Tile Options

FPGA Logic Densities (K LEs)	Tile Configurations	Package Details
2,200	1x F-Tile + 3x R-Tiles	57.5 x 49 mm², hex, 1.0 mm pitch, R2979A
2,692		
2,200	3x F-Tiles + 1x R-Tile	60 x 59 mm², hex, 1.0 mm pitch, R3803A
2,692		
2,200	4x F-Tiles	59 x 53 mm², hex, 1.0 mm pitch, R3343A
2,692		

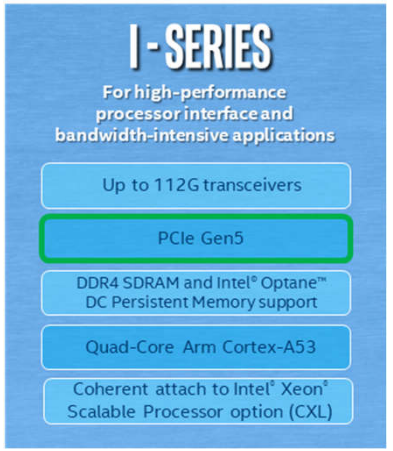

R2979A Package

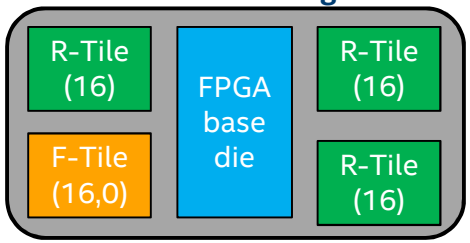

R3803A Package

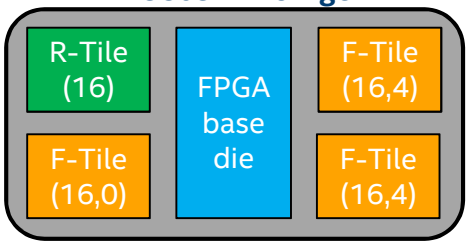

R3343A Package

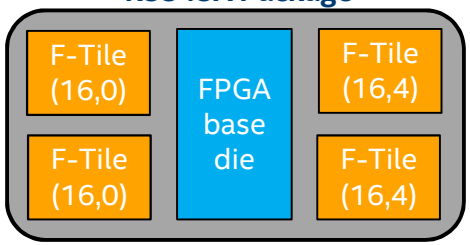

R- and F-Tile locations ***(bottoms up views)***

Notes: R-Tile supports PCIe Gen5 while F-Tile only support PCIe Gen4.
R-Tile (xx) refers to the number of 32G transceiver channels available. F-Tile (xx, y) refers to the number of 58G/112G transceiver channels available.

Intel® Agilex™ FPGA: PCIe Gen5 Support

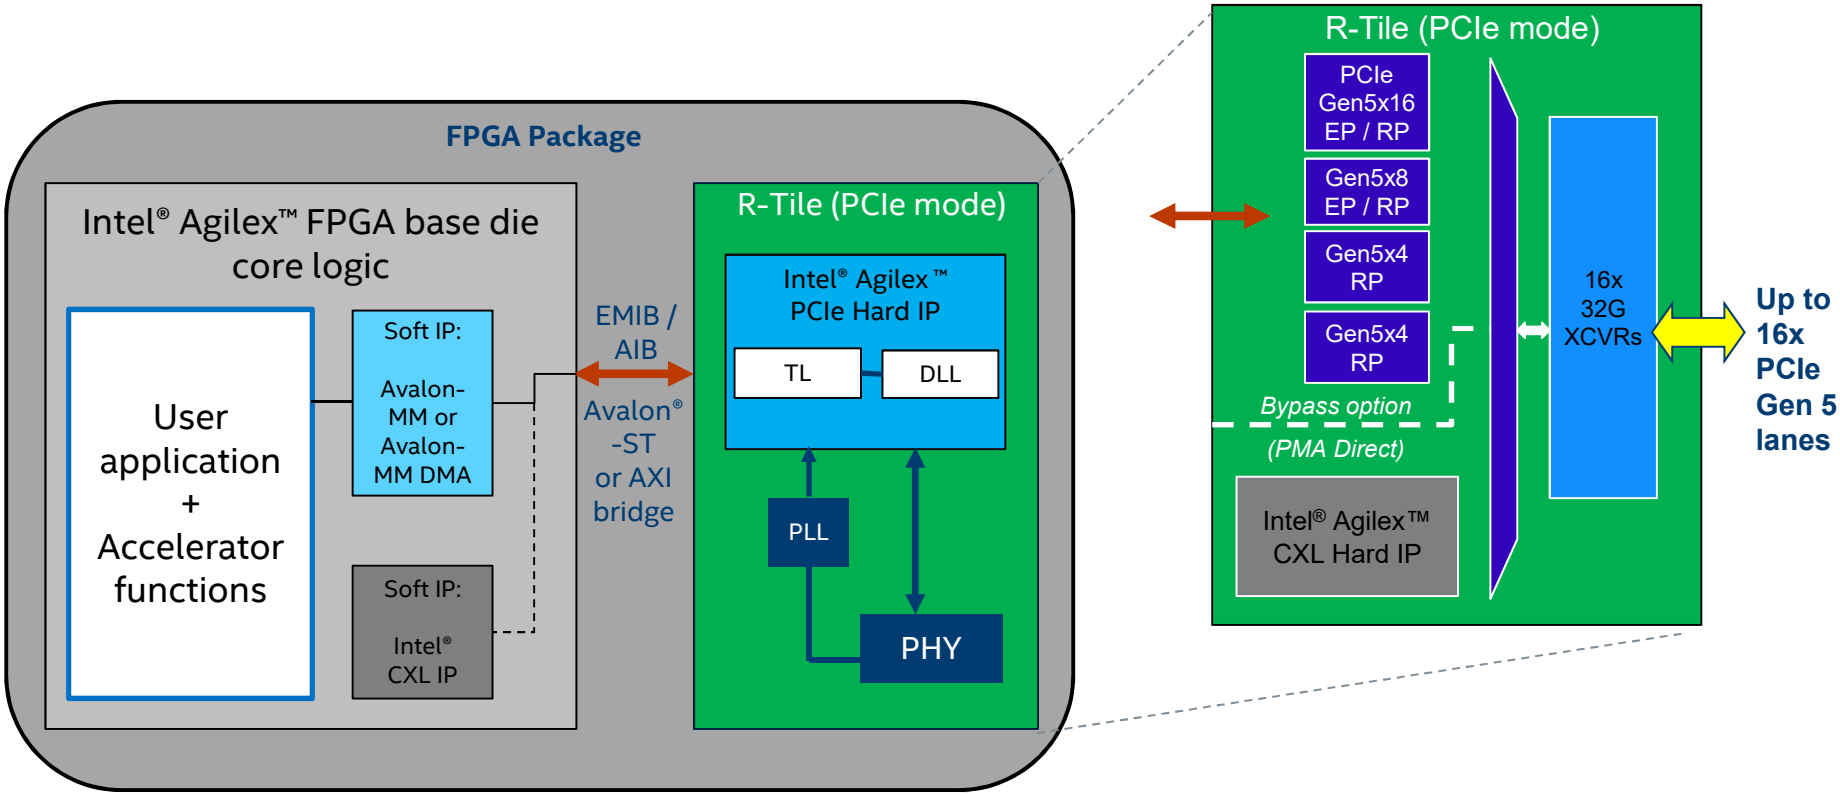

Intel's PCIe Gen5 Interoperability (PHY Test Chip)

Early testing with selected IHVs

- Interop testing with multiple vendors
  - Goal = Identify issues in early independent hardware vendors (IHV) designs and provide feed back into next generation Intel® Xeon® CPU R&D
  - Preliminary testing with IP providers, retimers, and test silicon from key providers – **Done** ✓
  - Demonstrated PCIe 5.0 interoperability with Synopsys at PCI-SIG Developers Conf. (June 2019)
- Intel using 'Ivey Creek' for early testing with IHV's and initial CPU validation
- Next phase = PCIe 5.0 testing with CPU + FPGA samples (contact Sales for schedules)

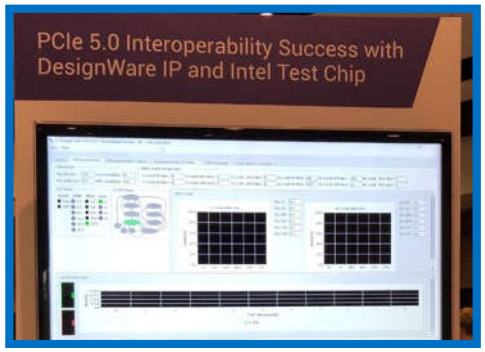

https://youtu.be/fR00tPwZFQg

Testing Performed
Link Training
• Link Train to L0
• Speed Change Loop
• Recovery via link retrain Loop
• Link Disable/Enable Loop
• TxEq Redo
• Lane Margining (Gen4)
Reset Tests
Hot Reset, Secondary Bus Reset, SBR Loop
Power Management
• L1 PM
• L1 ASPM

Subset of interoperability tests

New Intel® FPGAs – Start Designing PCIe Gen4 Today

P/N: DK-DEV-1SDX-P-CES

- 4X ports of PCIe Gen4 x16 with Stratix 1SD280 base die
- 2X DIMM slots (DDR4 or Intel® Optane™ DC persistent memory)
- 2X QSFP (10GbE / 25GbE / 100 GbE)

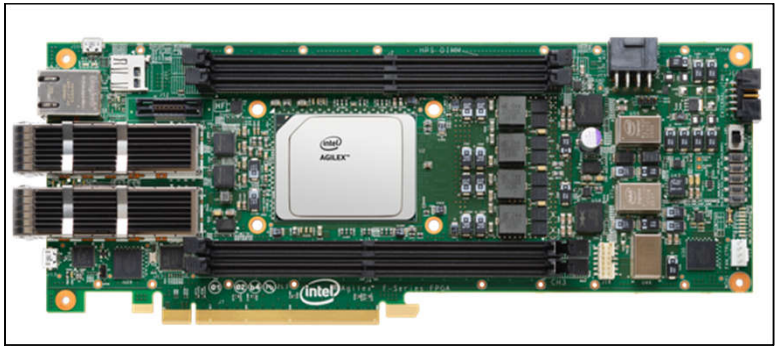

P/N: DK-DEV-AGF014E2ES

- 1X port PCIe Gen4 x16 & Quad-core Arm A53 subsystem with Intel Agilex AGF014 FPGAs base die
- 4X DIMM slots (DDR4 or Intel® Optane™ DC persistent memory)
- 2X QSFP-DD (8x25GbE or 2x50GbE)

Download Intel® Quartus® Prime Design Software v19.3

1. Contact Intel® FPGA team for soft IP schedules to support Intel® FPGA + Intel® Optane™ DC persistent memory.

Intel® Stratix® 10 FPGAs DX FPGA Development Kit

- Intel® Stratix® 10 DX FPGA
  - 2.8 Million logic density
  - 4x P-Tiles, 1x E-Tile
- 4x DDR4 channels (32GB Total)
  - 2x 8GB DDR4-2400 onboard (x72 w/ ECC)
  - 2x 8GB DDR4 DIMMs or *Optane™ DCPM modules (option)*
- 2x QSFP56 cages (2x 10GbE / 25GbE / 100GbE supported)
  - FPGA Ethernet hard IP feature
- PCIe Gen4 Capability & Upgraded Virtualization features
  - 1x PCIe Gen4 x16 Endpoint via edge connector
  - 3x PCIe / UPI connectors (for cable over to selected CPU boards)*
- UPI compatible board designed to be compatible with selected Intel® Xeon® CPU platforms/CRB's
  - IP support for this feature is *restricted*. Specific customer approval & activation is required.

P/N: DK-DEV-1SDX-P-CES
Order Now

* Cables and daughter card will not be included with kit

Intel® Stratix® 10 DX FPGAs

New capabilities added to Stratix 10 accelerator portfolio

- **Up to 2x** more bandwidth via PCI-SIG compliant PCIe Gen4 x16 [2]
- Will deliver **37%** lower latency [3] via Intel Ultra Path Interconnect (UPI), the coherent interface to select Intel® Xeon® Scalable processors
- Memory controller IP to connect to select Intel® Optane™ DC persistent memory modules

Accelerates server applications:

- Coherent flexible accelerators (FPGA + Xeon® processors)
- Increase total host addressable memory via coherently attached FPGA and select Optane™ technology
- High performance, low latency for SmartNICs, Test & Measurement, etc.
- Path to Compute Express Link (CXL) and PCIe Gen5

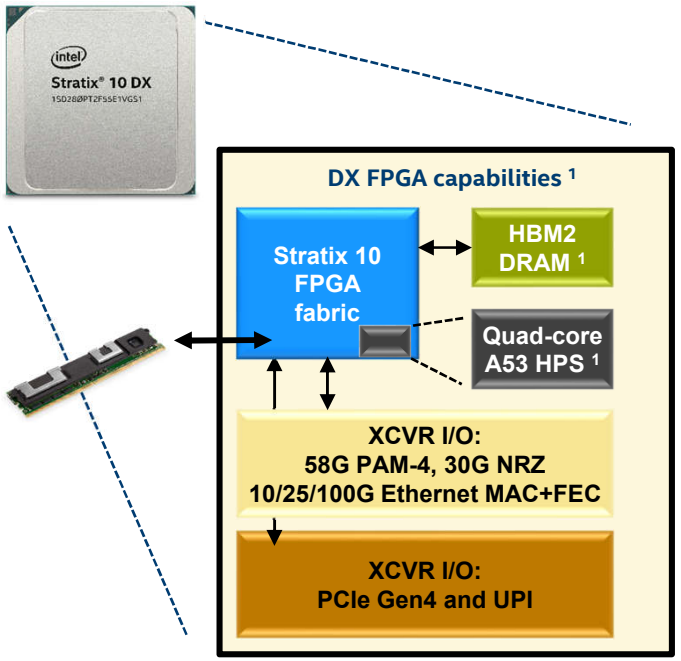

[1] Some features are not available in every member of the family. [2] Per the PCI specification for PCIe Gen.3 x16 vs. Gen4 x16. [3] UPI latency vs. PCIe roundtrip latency in Stratix 10 DX, based on Intel estimates.

Intel® FPGA PCIe Gen4 and Gen5 Solutions

Leverages Intel's expertise with PCI Express (PCIe) protocol

PCI-SIG compliant PCIe Gen4 x16 shipping now!	New PCIe virtualization and other features save cost and power!
PCIe Gen5 technology building blocks demo today!	Multi-generation PCIe technologies will ease transition to Gen.6

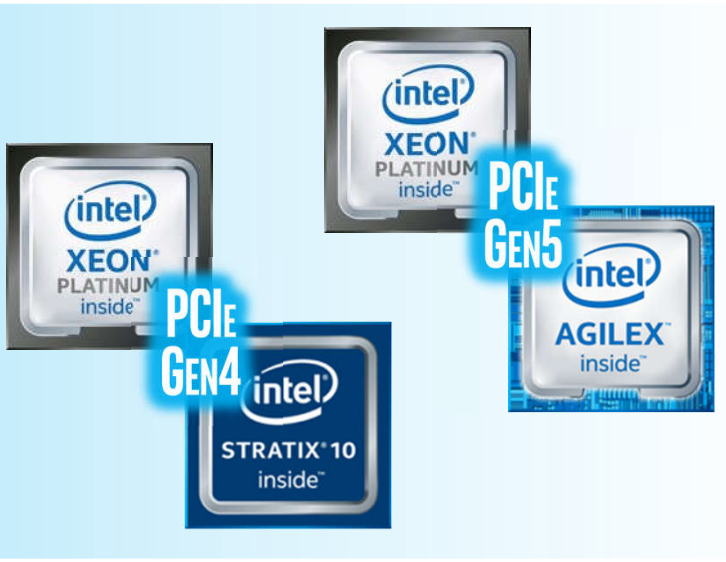

INTEL® FPGAS DESIGNED AS A PCI EXPRESS-BASED ACCELERATOR

ACCELERATING CUSTOMER INNOVATION FROM EDGE TO CLOUD WITH INTEL® PROGRAMMABLE ACCELERATION CARDS

The Data-Centric Era Is Upon Us …
… and Driving Market Change

BY 2020[1]

AVERAGE INTERNET USER	SMART HOSPITAL	AUTONOMOUS AUTOMOBILE	CONNECTED AIRPLANE	SMART FACTORY
1.5 GB DATA/DAY	3 TB DATA/DAY	4 TB DATA/DAY	40 TB DATA/DAY	1 PB DATA/DAY

[1] Amalgamation of analyst data and Intel® analysis

Application Acceleration Challenges and Needs

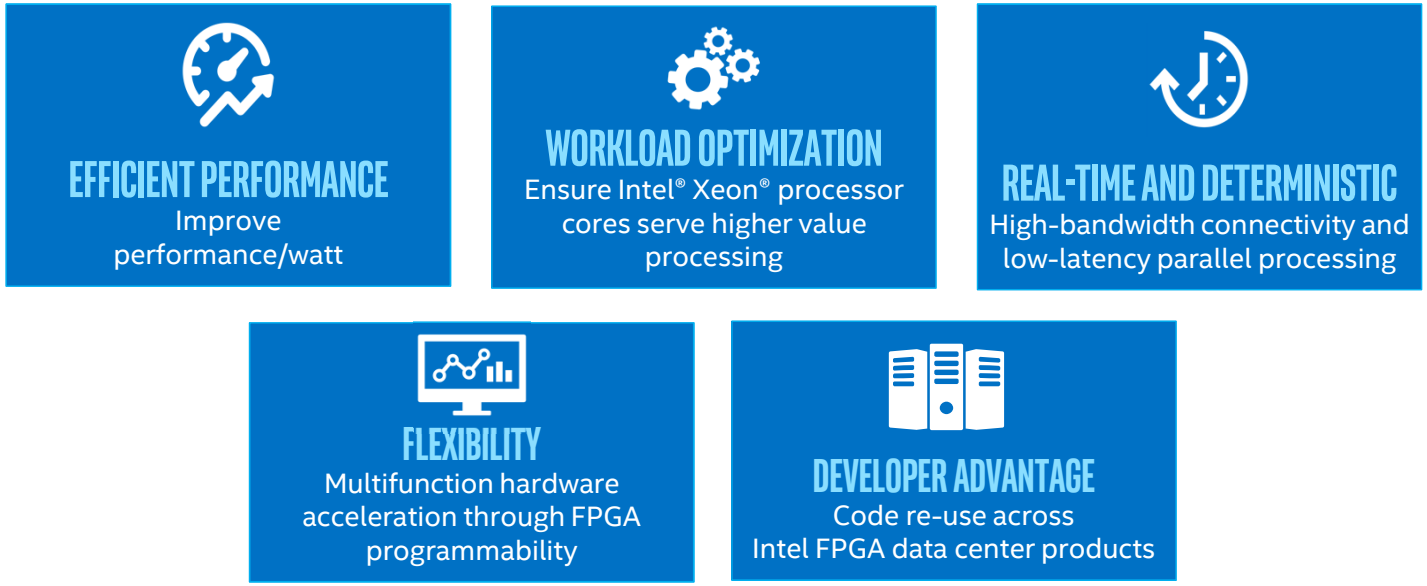

Acceleration Options

Acceleration of compute means

HETEROGENEOUS COMPUTE

DEDICATED ACCELERATORS
for maximum compute efficiency of specific, stable functions

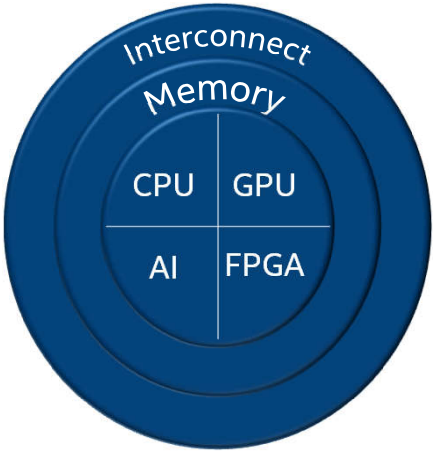

VERSATILE ACCELERATORS
for customized and changing workloads in networking, storage, and compute

Intel® FPGA PAC End-to-End Product Portfolio

EDGE ANALYTICS ACCELERATION

Intel® Vision Accelerator Design with Intel Arria® 10 FPGA

WIRELESS OR NETWORK FUNCTIONS VIRTUALIZATION (NFV) APPLICATION ACCELERATION

Intel FPGA PAC N3000 for Networking

CLOUD AND ENTERPRISE APPLICATION ACCELERATION

Intel FPGA PAC D5005 and Intel PAC with Intel Arria 10 GX FPGA for Data Center

Programmable Acceleration Solutions Ingredients

Intel® FPGA Programmable Acceleration Card (PAC) Qualified on Industry-Leading Enterprise Servers

Acceleration Stack for Intel® Xeon® CPU with FPGAs

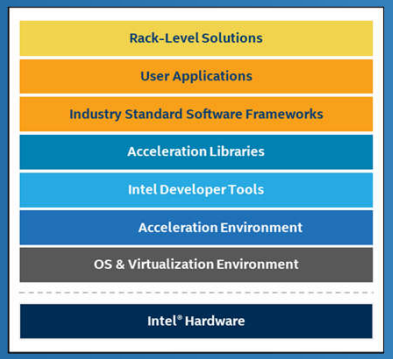

Acceleration IP from Application Experts

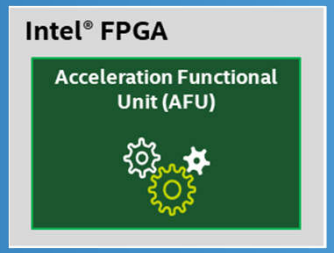

Intel® Acceleration Stack – Developer Platform for Fast Time to Market

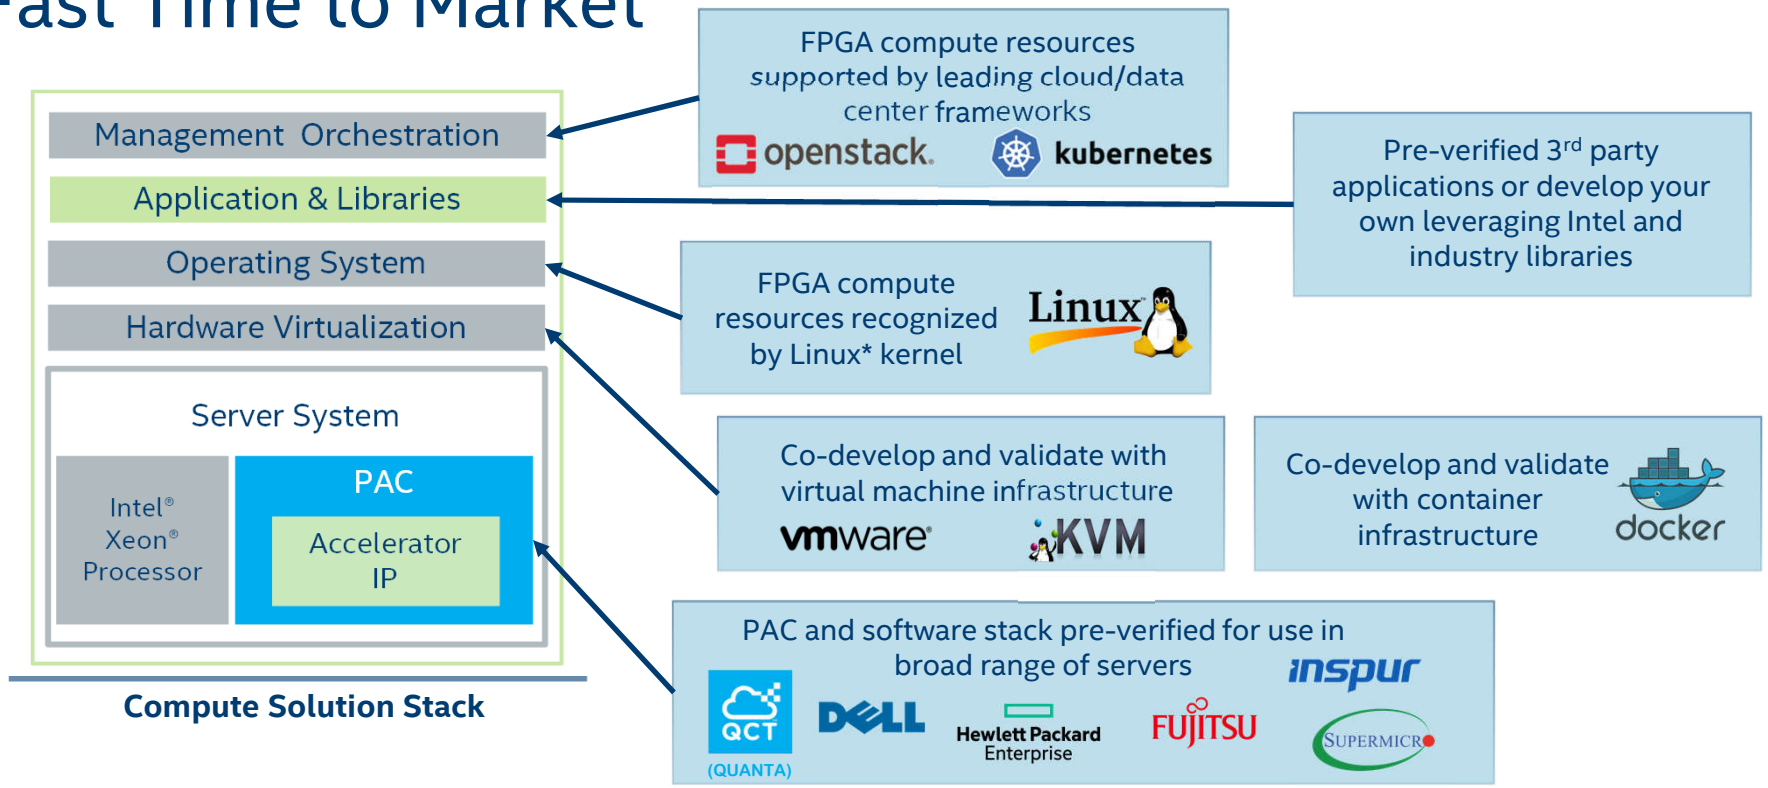

INTEL® PROGRAMMABLE ACCELERATION CARDS
ACCELERATE 5G

Cloudification of the Network

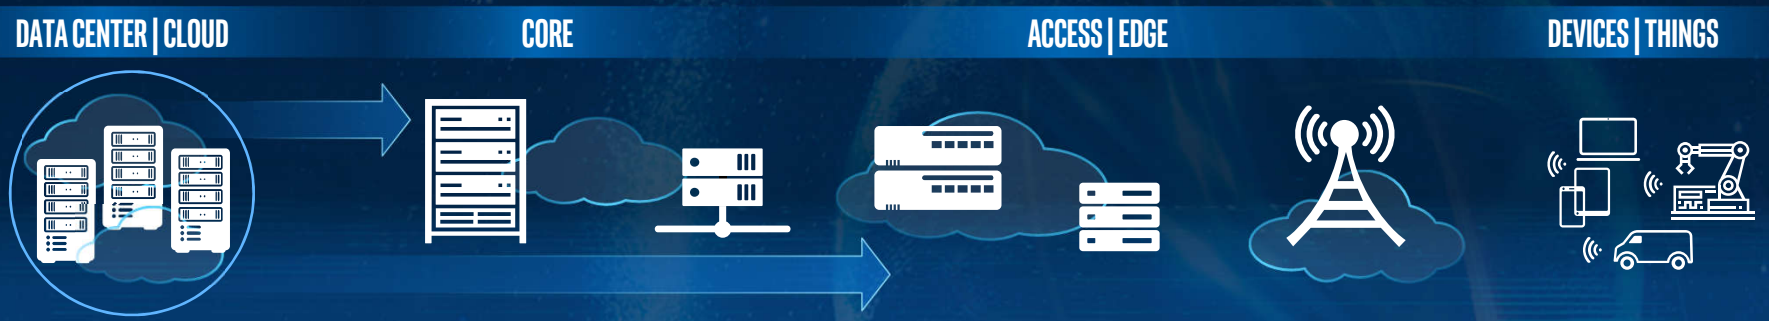

THE NETWORK MOVES TO INTEL® ARCHITECTURE (IA)

2011	2013	2015	2017	2019	2020
NFV DEFINED	1ST NFV PROOF OF CONCEPTS	20% OF COMMS SPS ADOPT NFV	INTEL DPDK MOVES TO LINUX FOUNDATION	1ST 100% CLOUD-NATIVE NETWORK	75% NETWORK WILL BE VIRTUALIZED

Acceleration Workload Solutions for NFV and vRAN

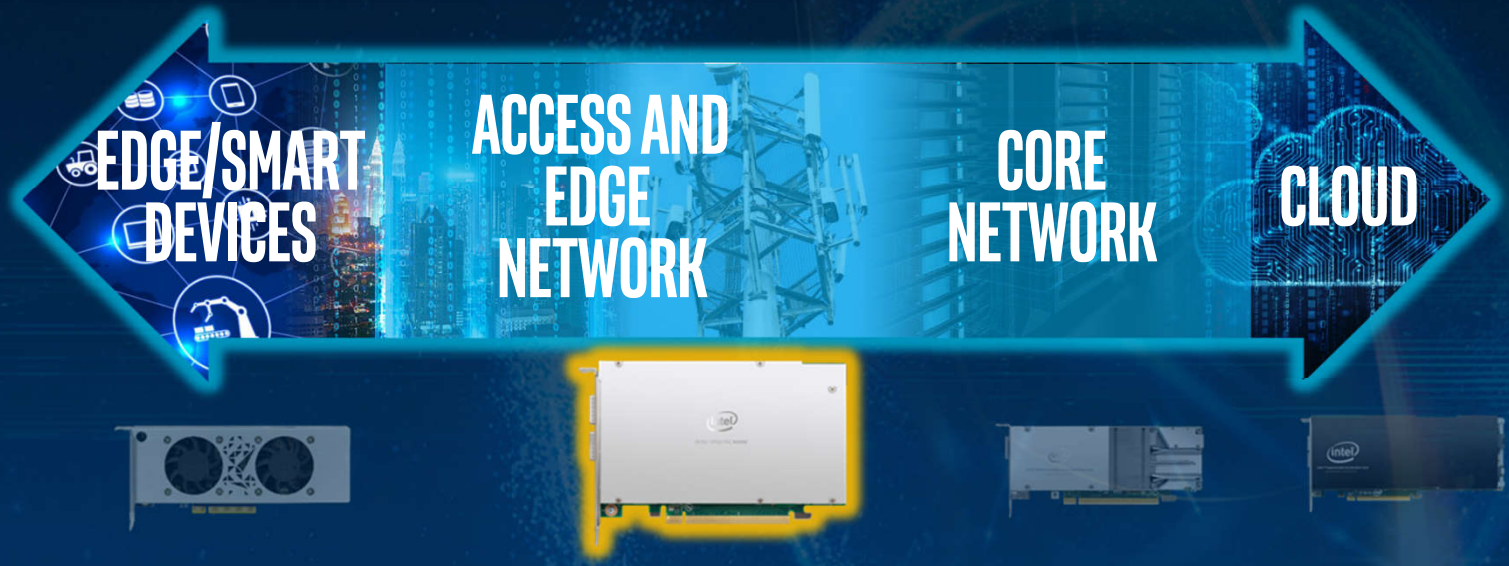

Target Workloads

Virtualized Radio Access Network (vRAN) Application

- **4G:** Turbo reference design
- **5G:** LDPC reference design
- **Fronthaul:** I/Q compression intellectual property (IP)

vCORE/NFV Application

- **NFVI:** Open vSwitch (OVS), Vector packet processing (VPP) router
- **VNF:** vBNG, vEPC, 5G NGCN
- **Security:** IPSec, vFirewall

Solution Components and Ecosystem

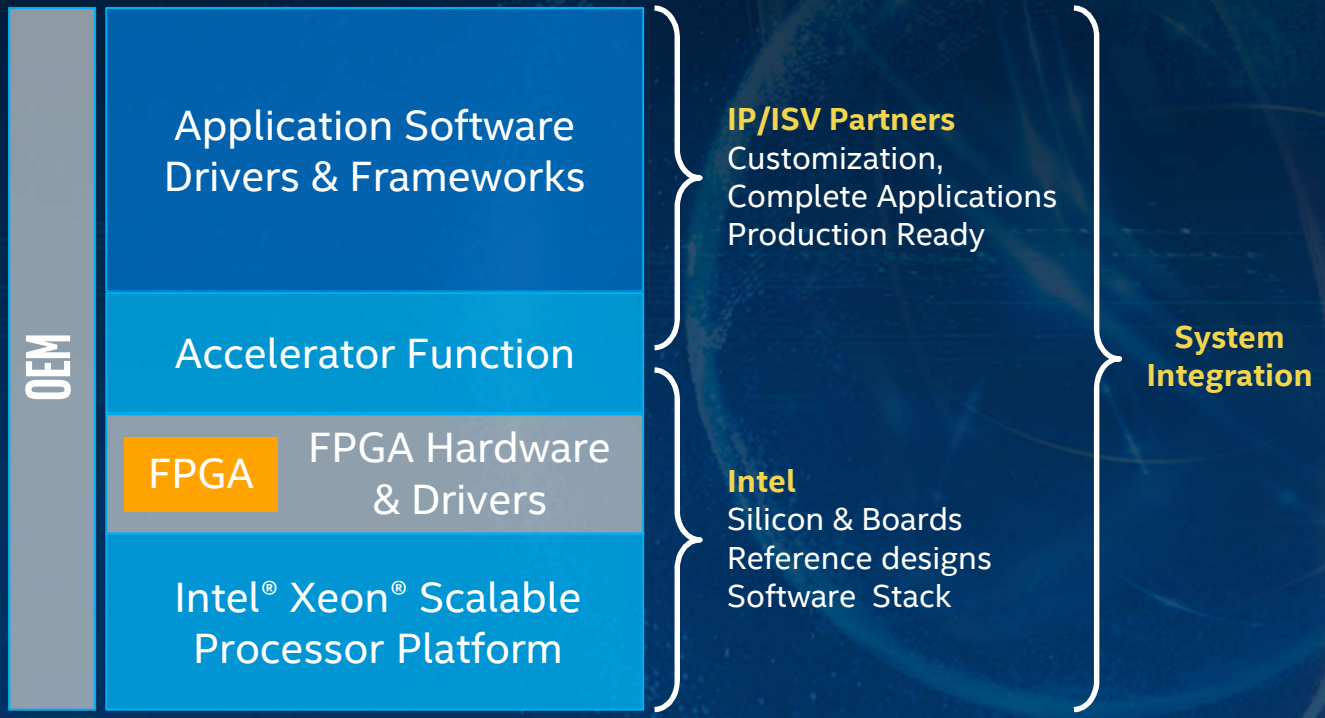

Ecosystem Enablement Model for vRAN

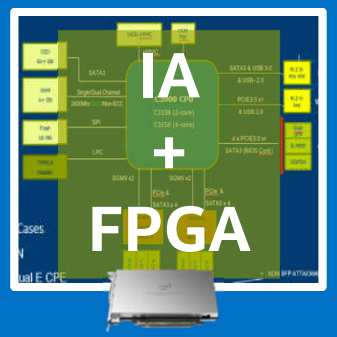

COTS Architecture, Server Platforms (Intel® Xeon® D processor, Intel® Xeon® Scalable Processor), FPGA PAC

Intel FlexRAN & vRAN Enablement Package

L1 (PHY) 4G & 5G

L2 / L3 (Stack) 4G & 5G

3rd Party SW

RRU

OEM/TEM/SI vRAN Solution (BBU on COTS Server & Software, FPGA, RF and remote radio head (RRU)

BBU SW

4G & 5G vRAN Solutions

Service Providers

Enterprises

Platforms

Reference S/W & User Workload Optimized on IA/FPGA

Commercial Product

Intel working with ecosystem to accelerate the development of vRAN solutions

Intel® FPGA PAC N3000 for Networking

Accelerates network traffic for up to 100 Gbps to support low-latency, high-bandwidth 5G applications enabling custom tailored solutions for vRAN and core network workloads with support of end-to-end industry standard and open source tools

HIGH PERFORMANCE

1.1 M logic elements

LOW LATENCY

9 GB DDR4 memory, 144 MB QDR IV

HIGH BANDWIDTH

8 X10 Gbps
4 x 25 Gbps

INTELLIGENT OFFLOAD

Two Intel® Ethernet Converged Network Adapter XL710 network interface card (NIC) for packet processing

SMALL FORM FACTOR

½ length, full height, single slot PCI Express* (PCIe*) card

Accelerating 5G vRAN with Intel® FPGA PAC N3000

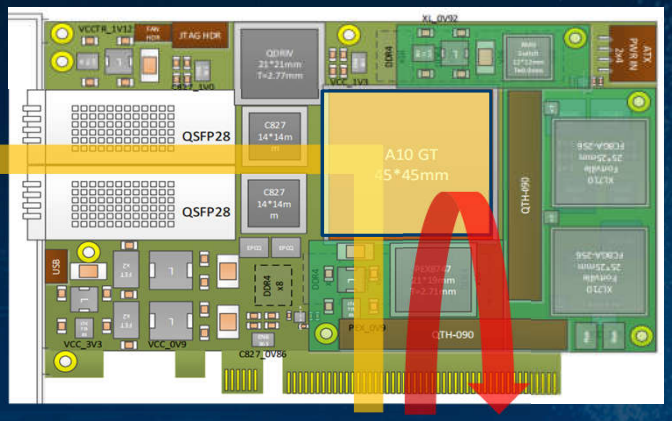

PCIe Gen3 16x

LDPC / TURBO OFFLOAD

+

FRONTHAUL

INTEL PAC SOLUTION DEPLOYED BY Rakuten

LOOK-ASIDE FORWARD ERROR CORRECTION (FEC)

LDPC/Turbo (De) Encoder, Rate (De)Matcher, (De) Interleaver, cyclic redundancy check (CRC)

BUMP-IN-WIRE FRONTHAUL

Compression or decompression for latency and bandwidth reduction

VIRTUALIZATION

Multi-channel direct memory access (DMA), virtual machine integration, load balancing functions

Accelerating 5G Core with Intel® FPGA PAC N3000

affirmed®

Affirmed Networks 5G Mobile Core

Multi-Access
4G/5G

Network Slicing

Microservices

Imperatives - Solutions

Revenue Growth via Delivery of New Services

- Industry's first virtualized 5G NGCN 200 Gbps/server solution at 50% CPU utilization
- Support higher bandwidth and lower latency 5G applications such as autonomous driving, augmented reality (AR) or virtual reality (VR) , and network slicing

OpEx and CapEx Efficiency

- Better CPU utilization with FPGA accelerated solution: Smart load balancing and CPU cache optimizations
- Lower power consumption/bit

Economies of Scale

- Build tailored solutions for different 5G use cases
- Future proof your design by seamlessly migrating 4G to 5G workloads

ENABLING COMMUNICATIONS SERVICE PROVIDERS TO MEET 5G HIGH DENSITY I/O GOALS

Load Balancing Accelerator

Overview

- Achieved 200 Gbps on 100 GbE links using Intel® Xeon® processors and Intel FPGA-based Programmable Acceleration Card
- Enabling cost-effective and scalable user plane 5G/4G core solutions for both edge and central office requirements by utilizing COTS platforms

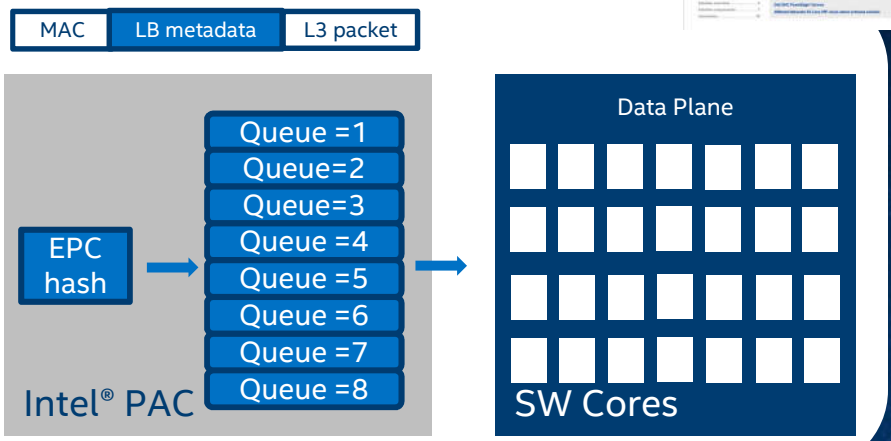

Value

- Packet processing performance improvement with optimized CPU utilization
- Frees up cores for other hosted 5G services, e.g. DPI
- Reduced CapEx/OpEx

Performance Gain

- 50% core saving on user plane forwarding
- 100 GbE termination
- 20 Gbps/ core

Source : Affirmed Networks Benchmarks

200G FPGA-Accelerated 5GCN

- First cloud-native containerized 5G core network with 100GbE interfaces (MWC 2019)
- Load balancing acceleration using Intel® PAC
- Affirmed 5GCN SA VNF
  - User plane packet forwarding
- Hardware-accelerated load balancing free up 12 physical core available for DPI and/or other applications

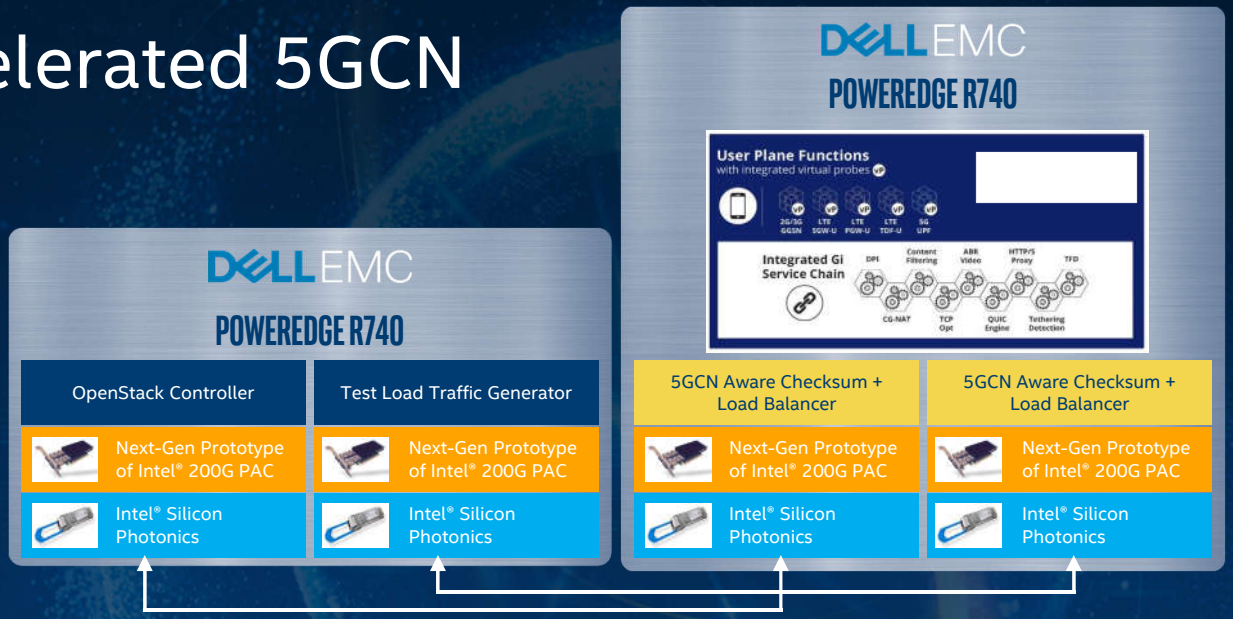

#pCPU User	Traffic Type: 50/50 Uplink/Downlink	Throughput per Server	Packet Rate per Server	Highlights
12	UDP, 640 Bytes Packet	~170 Gbps	33 Mpps	• Skylake Platinum with 200 Gbps line rate • Demo and corresponding core were based on 5GCN packet forwarding • Lots of headroom with 44 pCPUs out of 56 pCPUs in 2 socket server with 170 Gbps throughput at application level and 193 Gbps line rate.
8	UDP, 640 Bytes Packet	~160 Gbps (20 Gbps/pCore)	31 Mpps	

Source : Affirmed Networks Benchmarks

FPGA BASED DATA CENTER WORKLOAD ACCELERATORS

Intel WL Acceleration Solutions

RENIAC	swarm64	Falcon COMPUTING	CTAccel Accelerated Computing	MEGH COMPUTING	Levyx	napatech
No SQL Cassandra	**Data Warehousing**	**Genomics GATK**	**Image Processing**	**Streaming Analytics**	**Financial Black Scholes**	**Network Security/ Monitoring**
6X Performance	**4X+ Performance**	**2.5X Performance**	**3-4X Performance**	**5X Performance**	**8X Performance**	**3X Performance**
Deterministic low latency at higher Q/S	High Ingest rate with fast query/second	Speedup of Broad GATK pipeline	Transcode images faster	Real-time AI inference within Spark BigDL	Risk Analytics within Spark framework	Deep Packet Inspection at 40Gbps lossless
~ 80% TCO savings	~ 50% TCO savings	~ 60% TCO savings	~ 45% TCO savings	~ 50% TCO savings	~ 50% TCO savings	~ 75% TCO savings

More on the way.....

Problem

Data Source

- Sensors
- Web
- Operations
- Transportation
- Marketing
- Surveillance

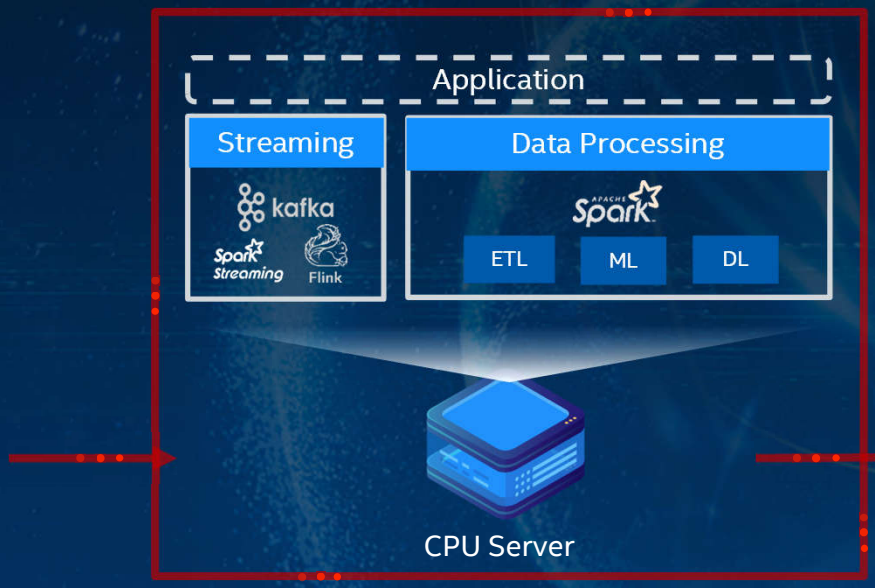

Database Processes

- Queries
- Alerts
- Analysis

Current Solutions Using Spark Streaming with ML / DL

Does not meet performance requirements

Challenges with SW only solutions:

(1) Latency is too high and does not meet real time requirements

(2) Throughput does not scale linearly with number of nodes

Acceleration Dataflow

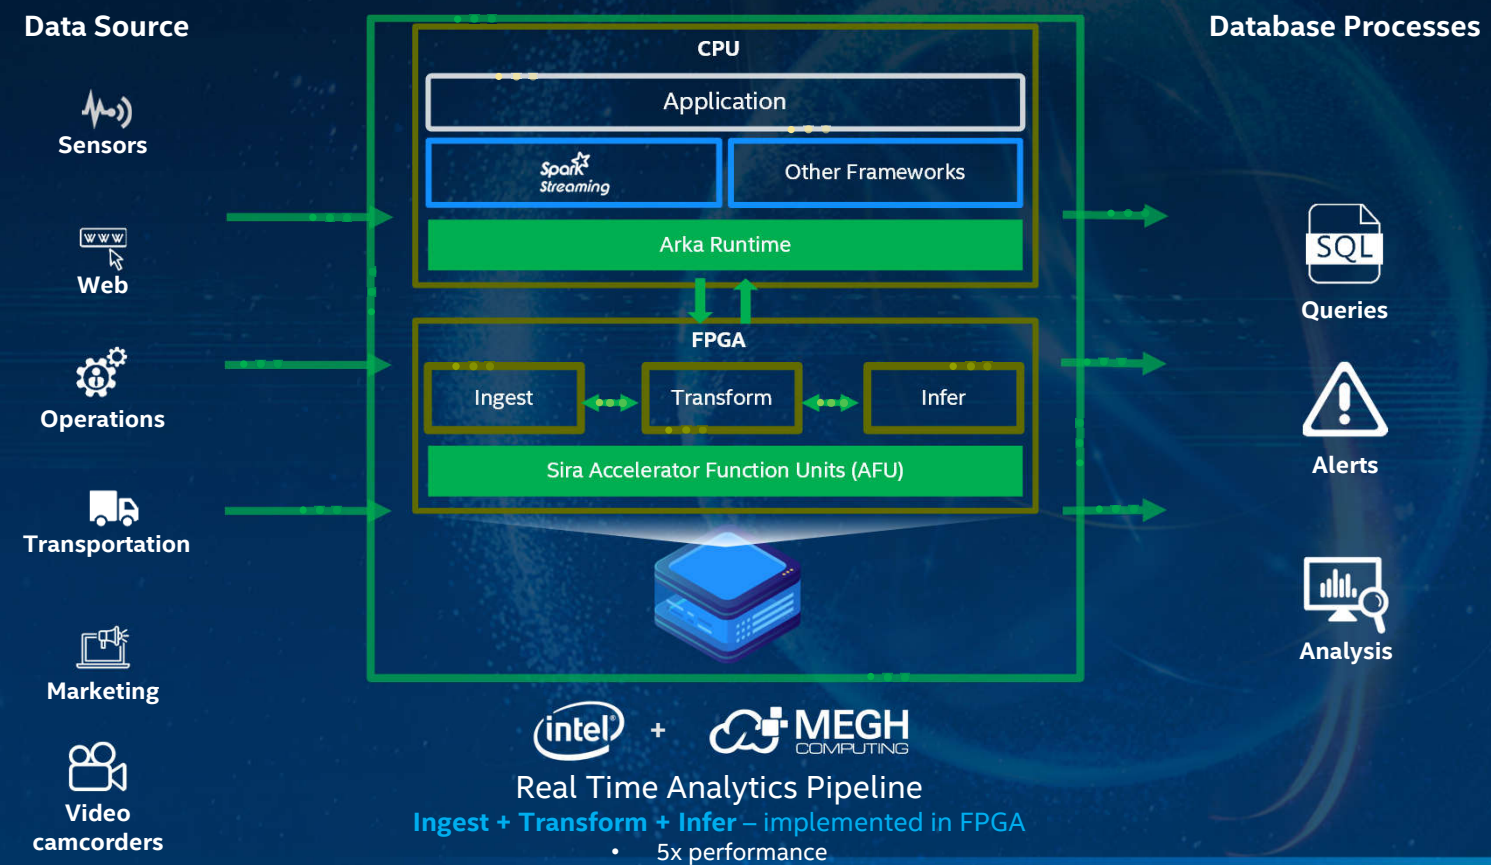

intel + MEGH COMPUTING

Real Time Analytics Pipeline

Ingest + Transform + Infer – implemented in FPGA

- 5x performance
- 2x less TCO and lower latency

use case: rt Video analytics

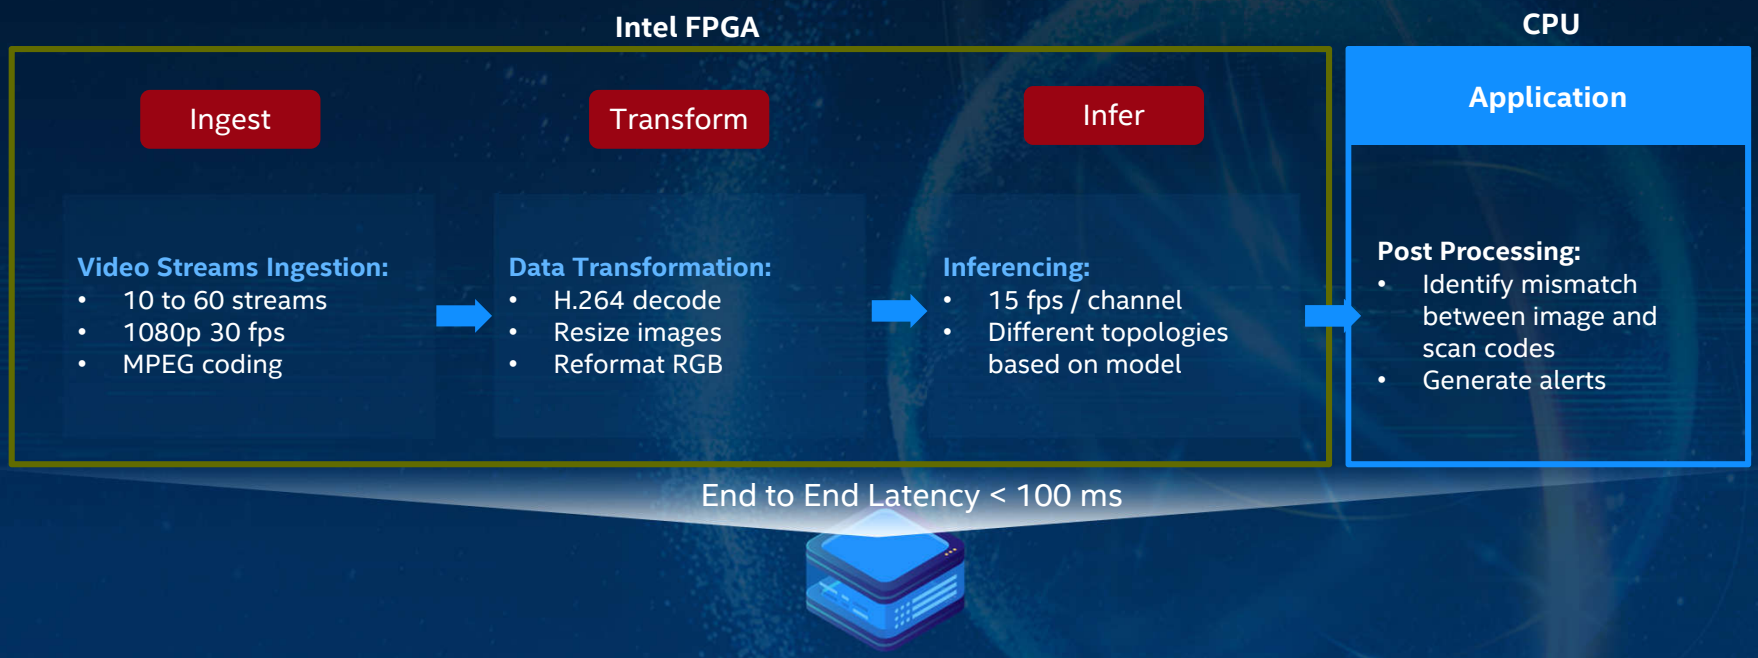

Use Case	Business Value
Fraud Prevention due to mis-scans of SKUs	Identify and prevent loss at POS
Image Classification of Fresh Produce for quality control	Identify poor quality produce and prevent packaging losses.
Better inventory management by avoiding product misplacement	Reduce inventory loss

Intel Accelerated Solutions

SOLVING REAL DATA CENTER PROBLEMS

SPARK STREAMING ACCELERATION

MEGH COMPUTING

Spark Acceleration: Real-time Streaming

Deliver >10x performance and 10x lower latency at same Total Cost of Ownership

2x	5~10x	2x
TCO Reduction	Better Performance	Latency Reduction

Running on an Intel qualified platform and Acceleration Stack

Ingest real time streams from varied sources in different formats for ETL processing

- In-line processing of network protocols, data filtering and transformations on FPGA

Process data streams using ML / DL libraries with different Data Analytics frameworks

- Initial implementation with Spark framework + BigDL with offload of inferencing to FPGA

Accelerate complete pipeline in FPGA for Low Latency and High Performance

- Delivered by differentiated technology – Arka Host SW and Sira FPGA HW libraries
- SDKs and Support Services for customization provide faster time to value.

Intel WL Acceleration Solutions

RENIAC	swarm64	Falcon COMPUTING	CTAccel Accelerated Computing	MEGH COMPUTING	Levyx	napatech
No SQL Cassandra	**Data Warehousing**	**Genomics GATK**	**Image Processing**	**Streaming Analytics**	**Financial Black Scholes**	**Network Security/ Monitoring**
6X Performance	**4X+ Performance**	**2.5X Performance**	**3-4X Performance**	**5X Performance**	**8X Performance**	**3X Performance**
Deterministic low latency at higher Q/S	High Ingest rate with fast query/second	Speedup of Broad GATK pipeline	Transcode images faster	Real-time AI inference within Spark BigDL	Risk Analytics within Spark framework	Deep Packet Inspection at 40Gbps lossless
~ 80% TCO savings	~ 50% TCO savings	~ 60% TCO savings	~ 45% TCO savings	~ 50% TCO savings	~ 50% TCO savings	~ 75% TCO savings

More on the way.....

Problem

Existing JPEG Images
Stored in Storage

Image Processing Servers
BOTTLENECK

Transcode to WebP
Many Servers needed due to resource bottleneck limitations to meet demands

Applications
- Ecommerce
- Social Networks
- News Apps

WebP
Images

View Images From
Browser or Apps

Acceleration Dataflow

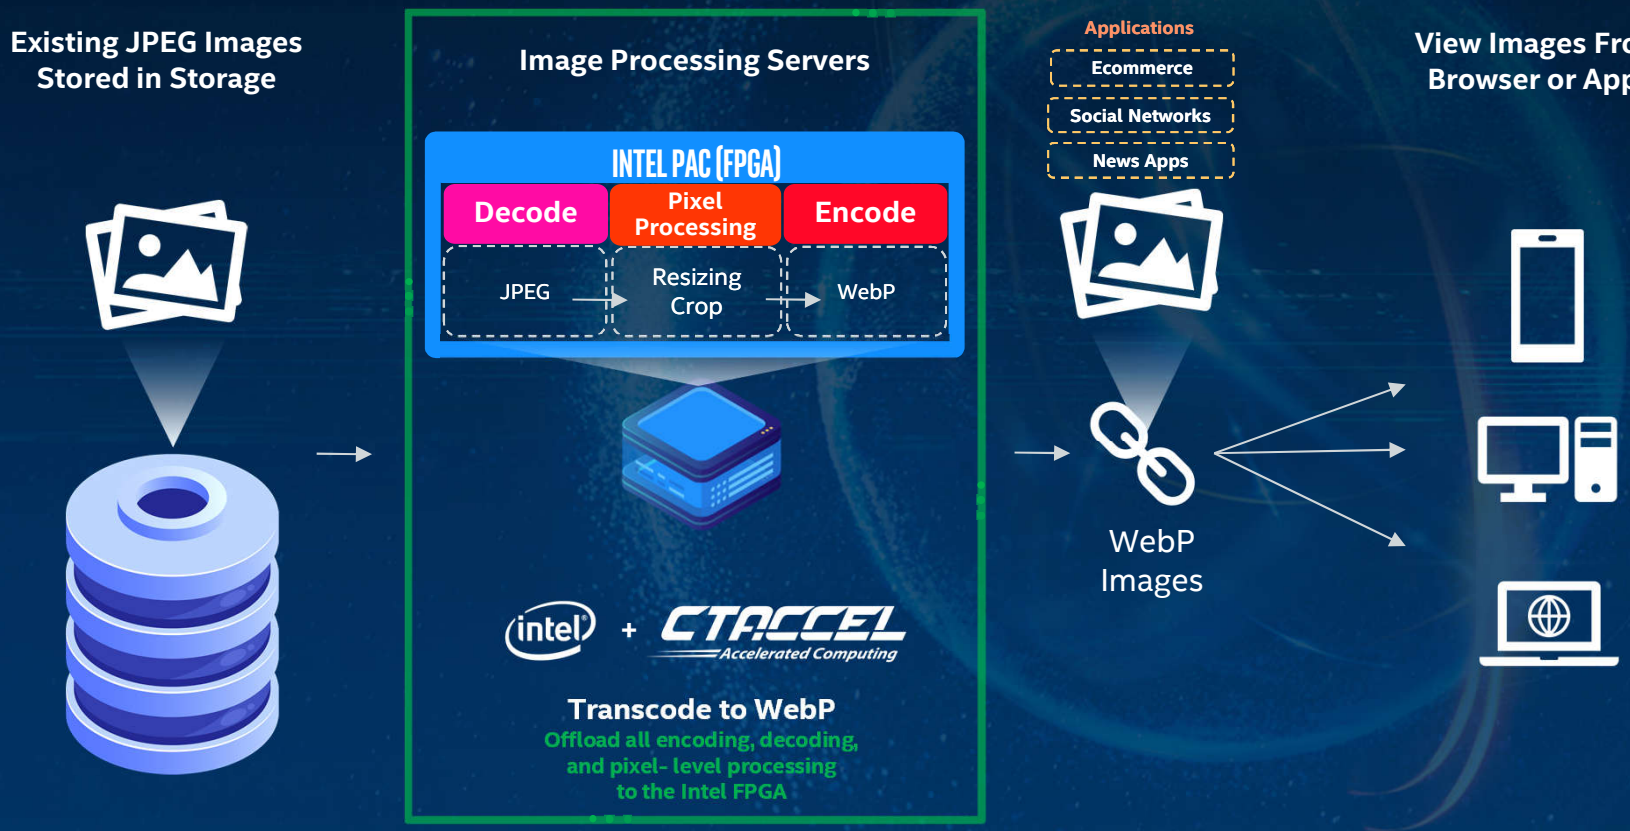

Intel Accelerated Solutions

SOLVING REAL DATA CENTER PROBLEMS

IMAGE PROCESSING ACCELERATION

CTACCEL Accelerated Computing

50%~70% TCO Reduction

5~10x Throughput Promotion

3x~5x Latency Reduction

Running on an Intel qualified platform and Acceleration Stack

Lower Power

- TDP 20W per Kernel for JPEG to Webp transcoder application (3 kernel at 60W Per RushCreek)
- TCO reduction: improve the computing density of DC, reduce racks

Industry Standard Software Compatibility for quick integration

- ImageMagick, OpenCV, GraphicsMagick, Lepton
- Allow users to migrate seamlessly from software-based implementation to CIP

Ease of Maintenance

- Can be upgraded remotely
- PR(Partial Reconfiguration) technology allows fast and easy context switch of accelerator functionality without rebooting the server

Intel WL Acceleration Solutions

RENIAC	swarm64	Falcon COMPUTING	CTAccel Accelerated Computing	MEGH COMPUTING	Levyx	napatech
No SQL Cassandra	**Data Warehousing**	**Genomics GATK**	**Image Processing**	**Streaming Analytics**	**Financial Black Scholes**	**Network Security/ Monitoring**
6X Performance	**4X+ Performance**	**2.5X Performance**	**3-4X Performance**	**5X Performance**	**8X Performance**	**3X Performance**
Deterministic low latency at higher Q/S	High Ingest rate with fast query/second	Speedup of Broad GATK pipeline	Transcode images faster	Real-time AI inference within Spark BigDL	Risk Analytics within Spark framework	Deep Packet Inspection at 40Gbps lossless
~ 80% TCO savings	~ 50% TCO savings	~ 60% TCO savings	~ 45% TCO savings	~ 50% TCO savings	~ 50% TCO savings	~ 75% TCO savings

More on the way.....

Problem

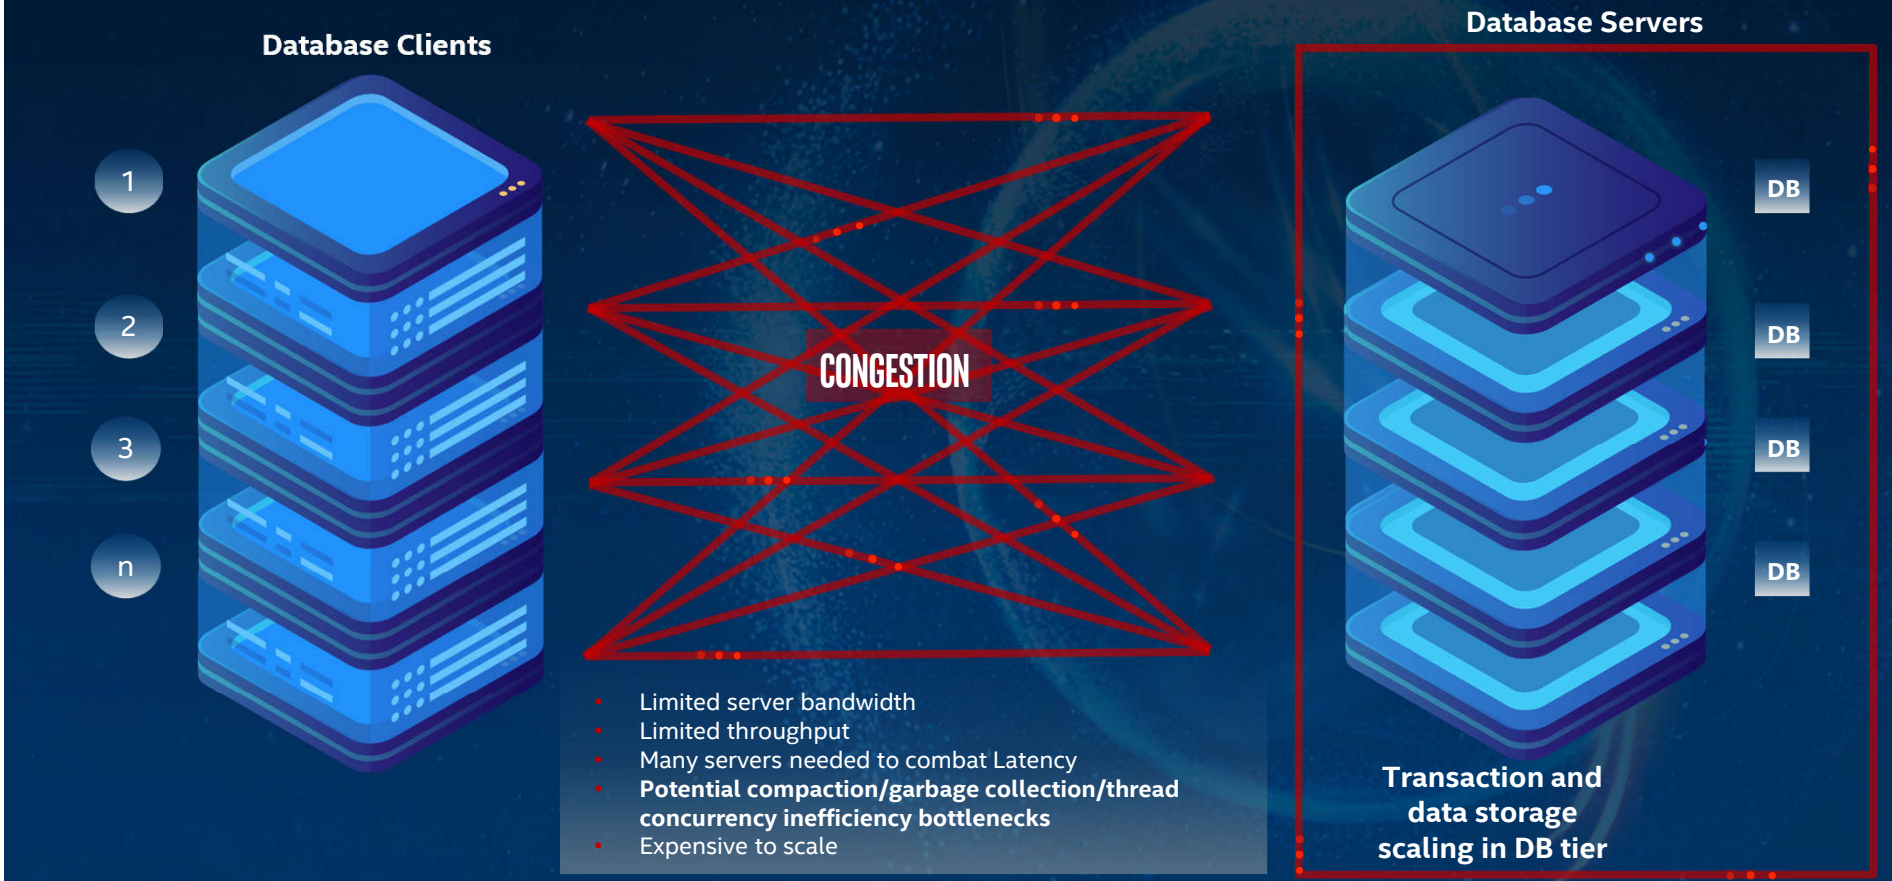

- Limited server bandwidth
- Limited throughput
- Many servers needed to combat Latency
- **Potential compaction/garbage collection/thread concurrency inefficiency bottlenecks**
- Expensive to scale

Acceleration Dataflow

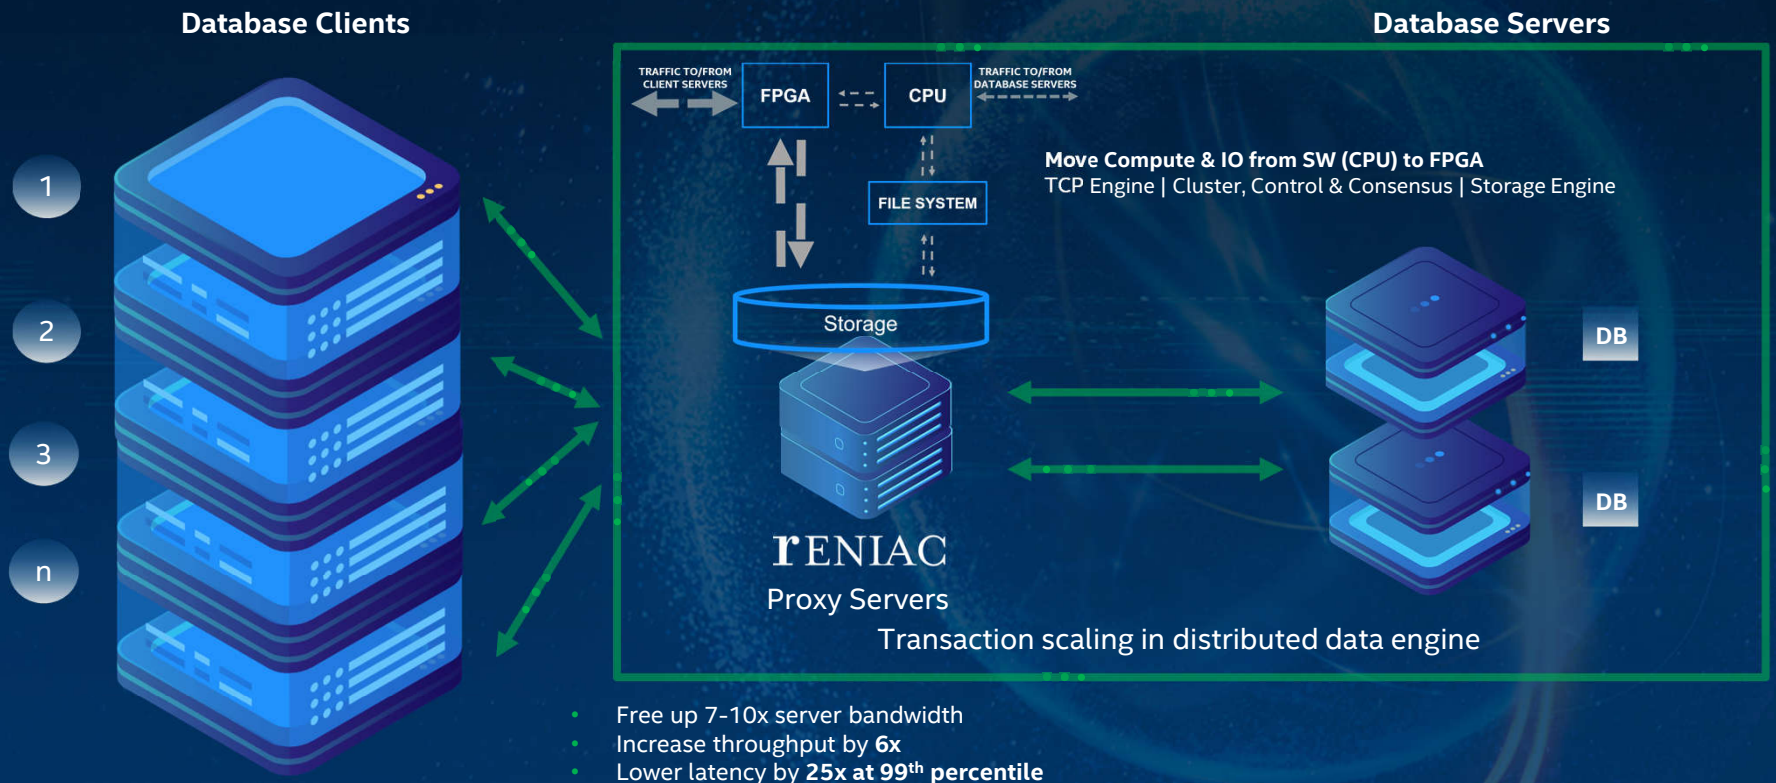

- Free up 7-10x server bandwidth
- Increase throughput by **6x**
- Lower latency by **25x at 99th percentile**
- **Reduce compaction/garbage collection/thread concurrency inefficiencies**
- Plug & Play: **No SW changes**

Intel Accelerated Solutions

SOLVING REAL DATA CENTER PROBLEMS

CASSANDRA DB ACCELERATION

rENIAC

Cassandra DB Acceleration

Up to 80%

TCO Reduction

Up to 6x

Throughput Promotion

Up to 25x

Latency Reduction

Running on an Intel qualified platform and Acceleration Stack

Decouples data and application layers, simultaneously acting as an I/O accelerator to resolve any bottlenecks
Unique ability to accelerate AI inference algorithms close to the data store and speeding up analytics
Tightly couples storage class memory to a low latency network stack
Deployed as a network service with no software change required
Free up to 75% CPU cycles for business logic, resulting in increased revenue at
Significantly lower TCO (7-10x lesser servers)

Intel WL Acceleration Solutions

RENIAC	swarm64	Falcon COMPUTING	CTAccel Accelerated Computing	MEGH COMPUTING	Levyx	napatech
No SQL Cassandra	**Data Warehousing**	**Genomics GATK**	**Image Processing**	**Streaming Analytics**	**Financial Black Scholes**	**Network Security/ Monitoring**
6X Performance	**4X+ Performance**	**2.5X Performance**	**3-4X Performance**	**5X Performance**	**8X Performance**	**3X Performance**
Deterministic low latency at higher Q/S	High Ingest rate with fast query/second	Speedup of Broad GATK pipeline	Transcode images faster	Real-time AI inference within Spark BigDL	Risk Analytics within Spark framework	Deep Packet Inspection at 40Gbps lossless
~ 80% TCO savings	~ 50% TCO savings	~ 60% TCO savings	~ 45% TCO savings	~ 50% TCO savings	~ 50% TCO savings	~ 75% TCO savings

More on the way.....

Problem

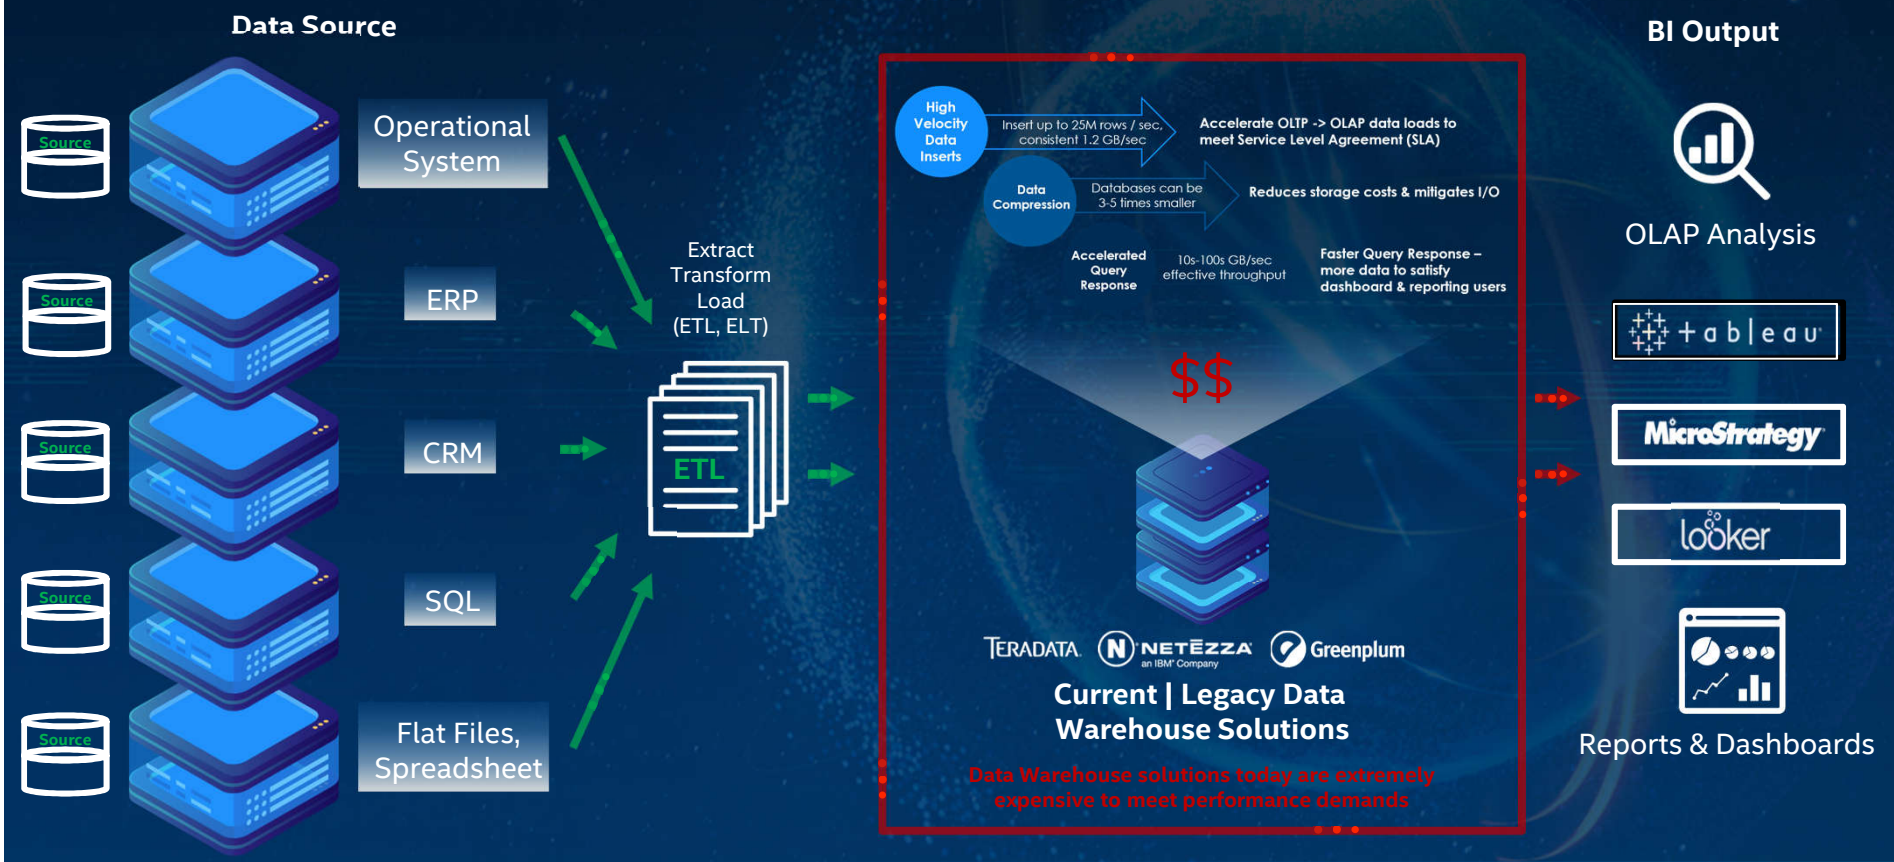
Data Source
Source
Source
Source
Source
Source
Operational System
ERP
CRM
SQL
Flat Files, Spreadsheet
Extract Transform Load (ETL, ELT)
ETL
High Velocity Data Inserts
Insert up to 25M rows / sec, consistent 1.2 GB/sec
Accelerate OLTP -> OLAP data loads to meet Service Level Agreement (SLA)
Data Compression
Databases can be 3-5 times smaller
Reduces storage costs & mitigates I/O
Accelerated Query Response
10s-100s GB/sec effective throughput
Faster Query Response – more data to satisfy dashboard & reporting users
$$
TERADATA
NETEZZA
an IBM Company
Greenplum
Current | Legacy Data Warehouse Solutions
Data Warehouse solutions today are extremely expensive to meet performance demands
BI Output
OLAP Analysis
tableau
MicroStrategy
looker
Reports & Dashboards

Acceleration Dataflow

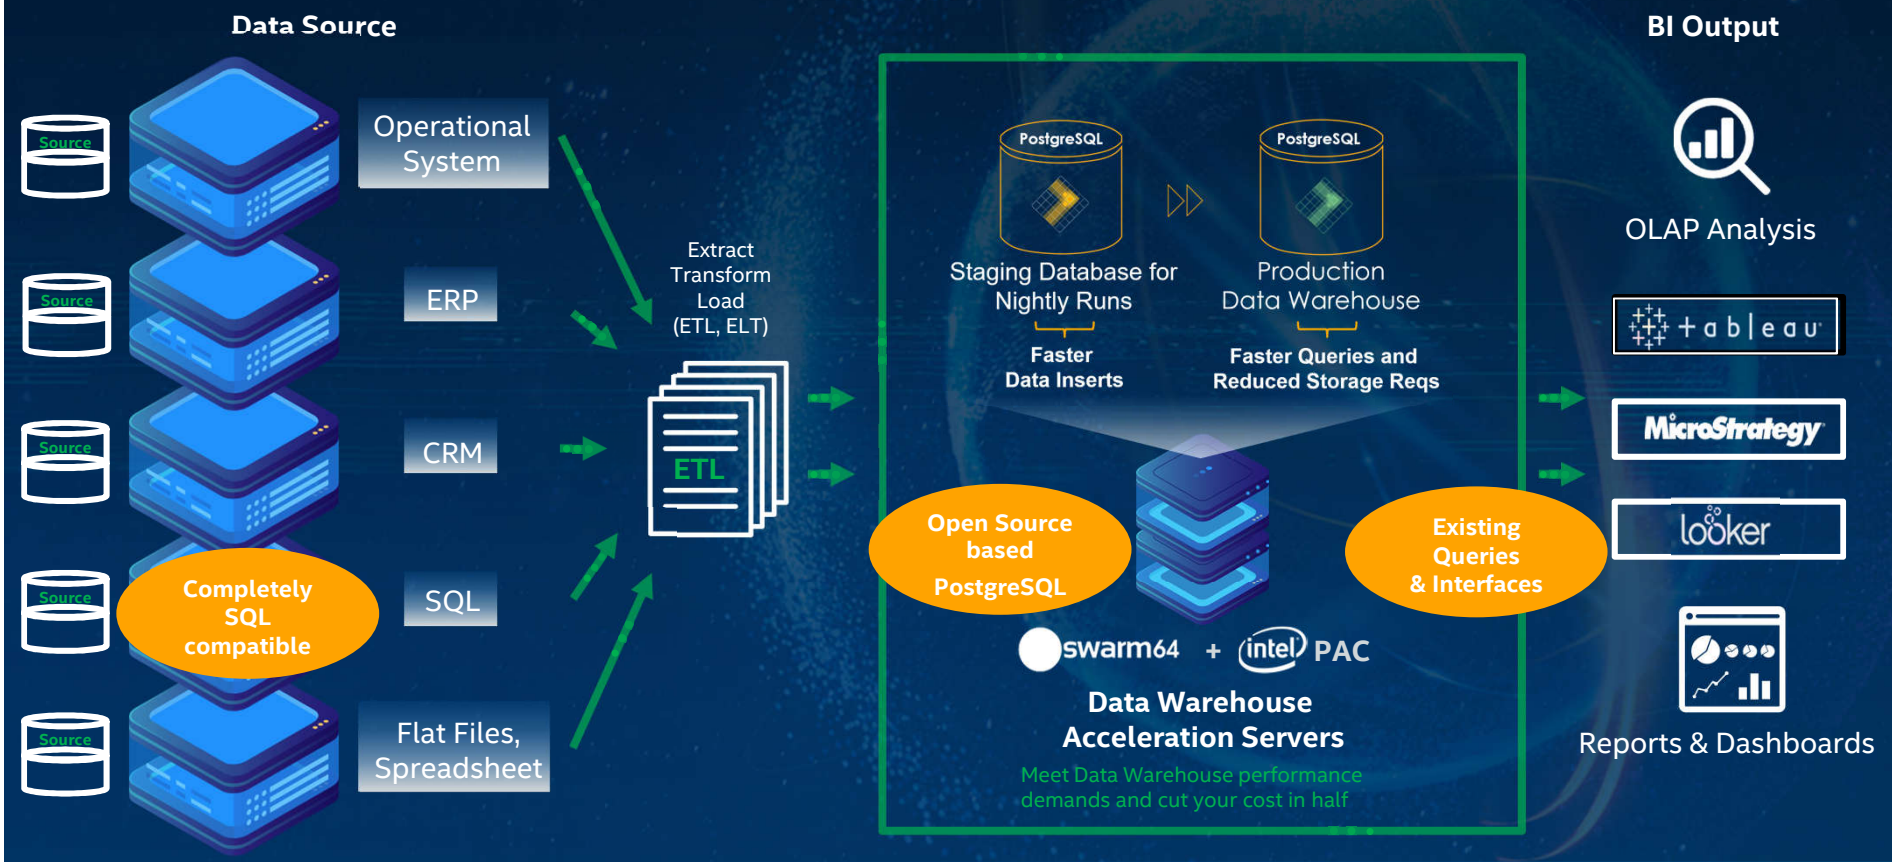

Intel Accelerated Solutions

SOLVING REAL DATA CENTER PROBLEMS

DATA WAREHOUSE ACCELERATION

swarm64

50% TCO Savings

½ TCO vs Legacy enterprise data warehouse solutions

12-70x
Faster Query Performance

1.2GB/s | 25M Rows/s
Data Insertion Rate

3-5x
Data Compression

Running on an Intel qualified platform and Acceleration Stack

- Enables customers to run data analytics using latest information more frequently
- Consolidate redundant server resources and increase capability of the IT organization
- Frees both administrative and HW data center resources for other tasks

Intel WL Acceleration Solutions

RENIAC	swarm64	Falcon COMPUTING	CTAccel Accelerated Computing	MEGH COMPUTING	Levyx	napatech
No SQL Cassandra	**Data Warehousing**	**Genomics GATK**	**Image Processing**	**Streaming Analytics**	**Financial Black Scholes**	**Network Security/ Monitoring**
6X Performance	**4X+ Performance**	**2.5X Performance**	**3-4X Performance**	**5X Performance**	**8X Performance**	**3X Performance**
Deterministic low latency at higher Q/S	High Ingest rate with fast query/second	Speedup of Broad GATK pipeline	Transcode images faster	Real-time AI inference within Spark BigDL	Risk Analytics within Spark framework	Deep Packet Inspection at 40Gbps lossless
~ 80% TCO savings	~ 50% TCO savings	~ 60% TCO savings	~ 45% TCO savings	~ 50% TCO savings	~ 50% TCO savings	~ 75% TCO savings

More on the way.....

Problem

TODAY

Cluster

Expanding Server Build-out

- Limited budgets
- More requirements
  - Compliance
  - Model Complexity
  - Security

More Compute Required

FUTURE

Cluster

Intel PACs

Densely populated servers with multiple PACs

- Denser infrastructure
- Lower costs
- More flexibility
- Higher density of compute and data

More Compute Delivered

Acceleration Dataflow

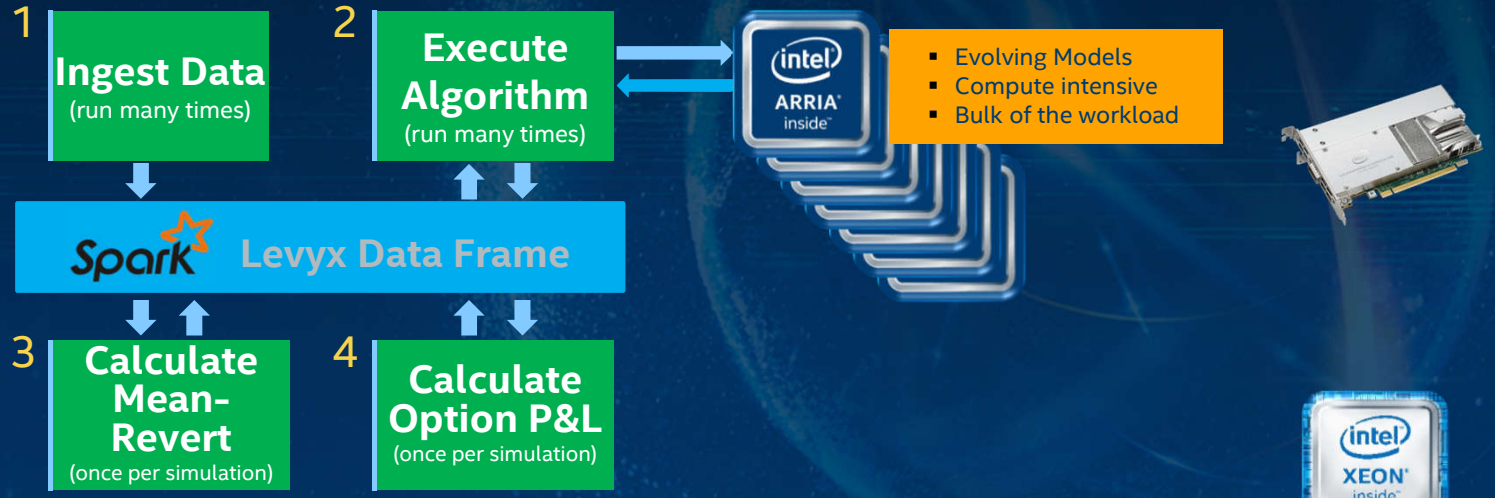

SOLVING REAL DATA CENTER PROBLEMS

financial risk managing

8X COMPUTE PERFORMANCE

vs. Server Only Solutions

Reduce TCO by 2x	All-Intel-in-One Solution	Ease-of Use	Intel-Qualified Platform and Acceleration Stack

Intel® FPGA Acceleration Hub

Intel® portal for all things related to FPGA acceleration

- **Acceleration Stack**
  for Intel® Xeon® CPU with FPGAs
- **FPGA Acceleration Platforms**
- **Acceleration Solutions and Ecosystem**
- **Knowledge Center**
- **FPGA as a Service**
- **Academia**
- **01.org ***

The Intel FPGA Acceleration Hub – www.intel.com/fpgaaccelerationhub

* 01.org is an open source community site

Legal Notices & Disclaimers

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit www.intel.com/benchmarks

Results have been estimated or simulated using internal Intel analysis or architecture simulation or modeling, and provided to you for informational purposes. Any differences in your system hardware, software or configuration may affect your actual performance.

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 www.intel.com/benchmarks

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No product or component can be absolutely secure. Check with your system manufacturer or retailer or learn more at intel.com.

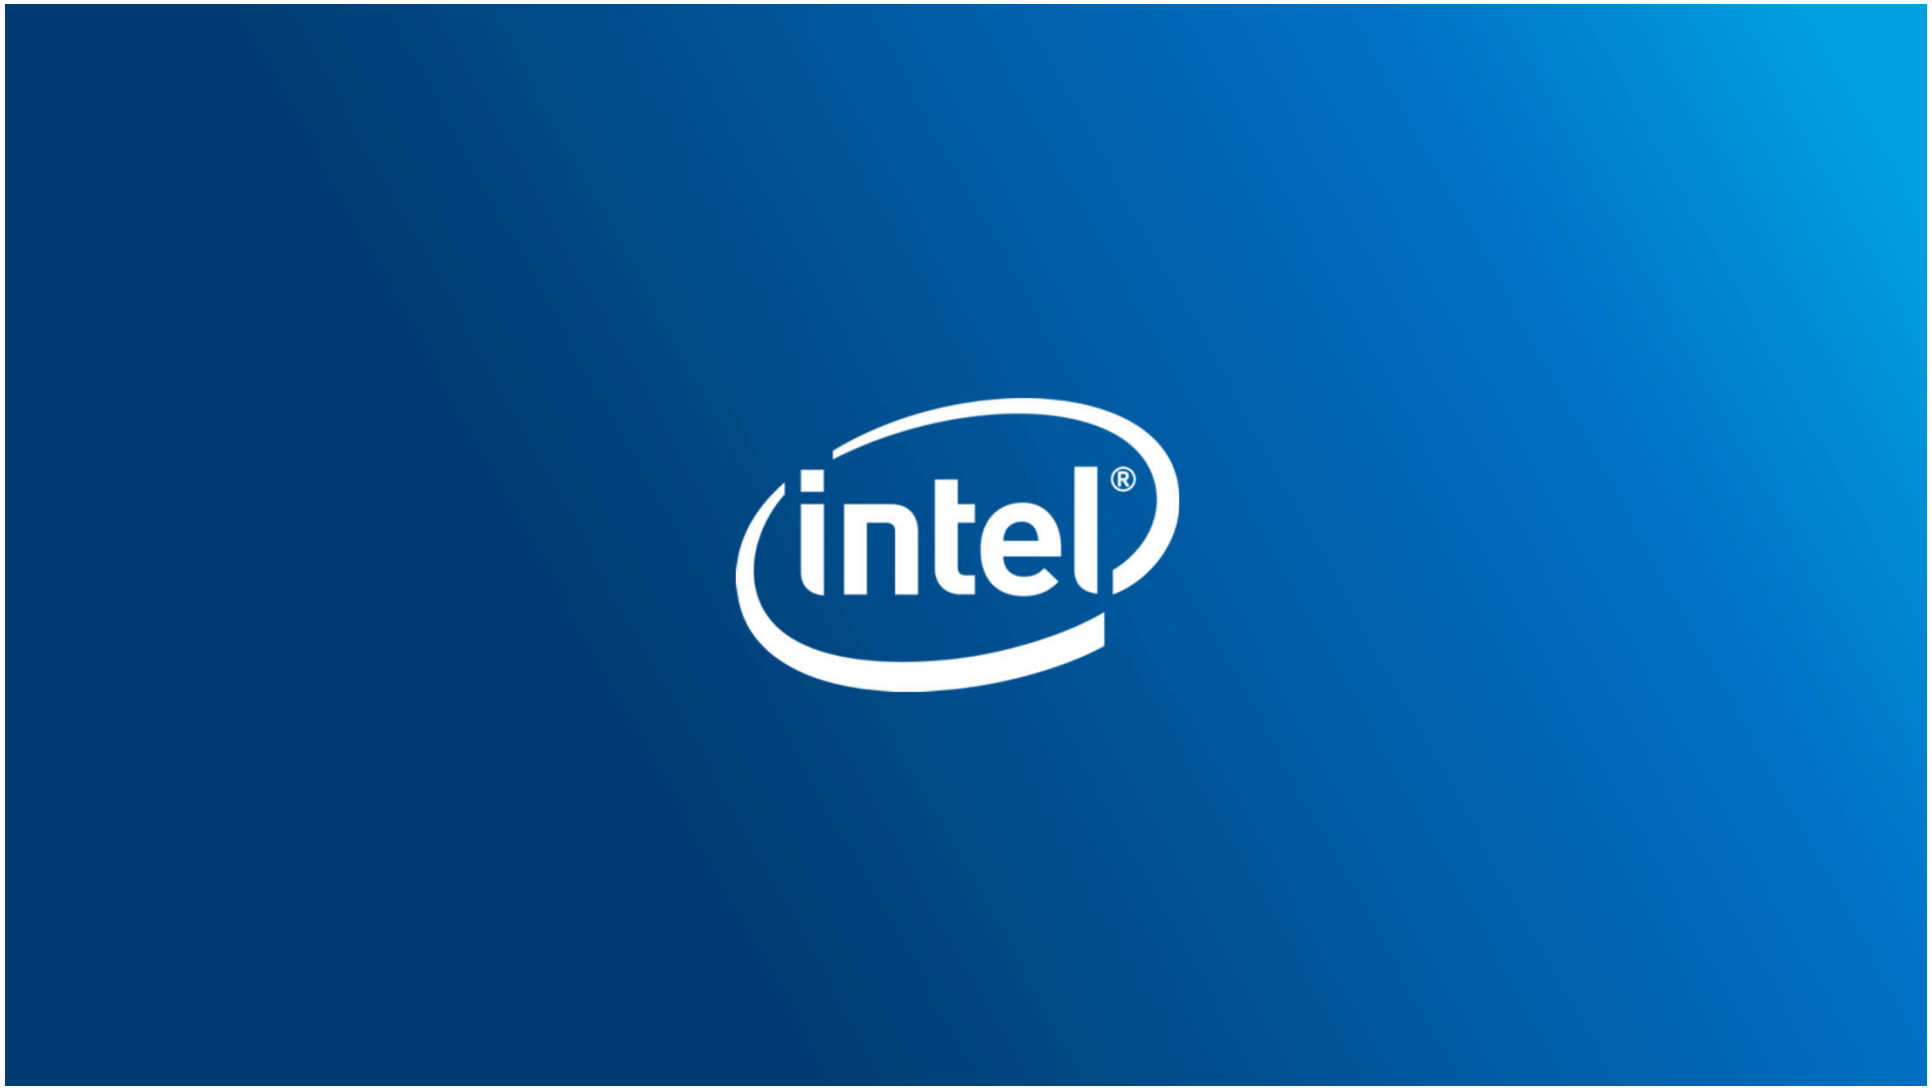
intel