XILINX

Versal™ Premium Series Announcement

Mike Thompson
Senior Product Line Manager, High-End ACAPs & FPGAs
Xilinx

© Copyright 2020 Xilinx

Explosion of Data from Diverse Applications & Workloads Puts Tremendous Pressure on the Core

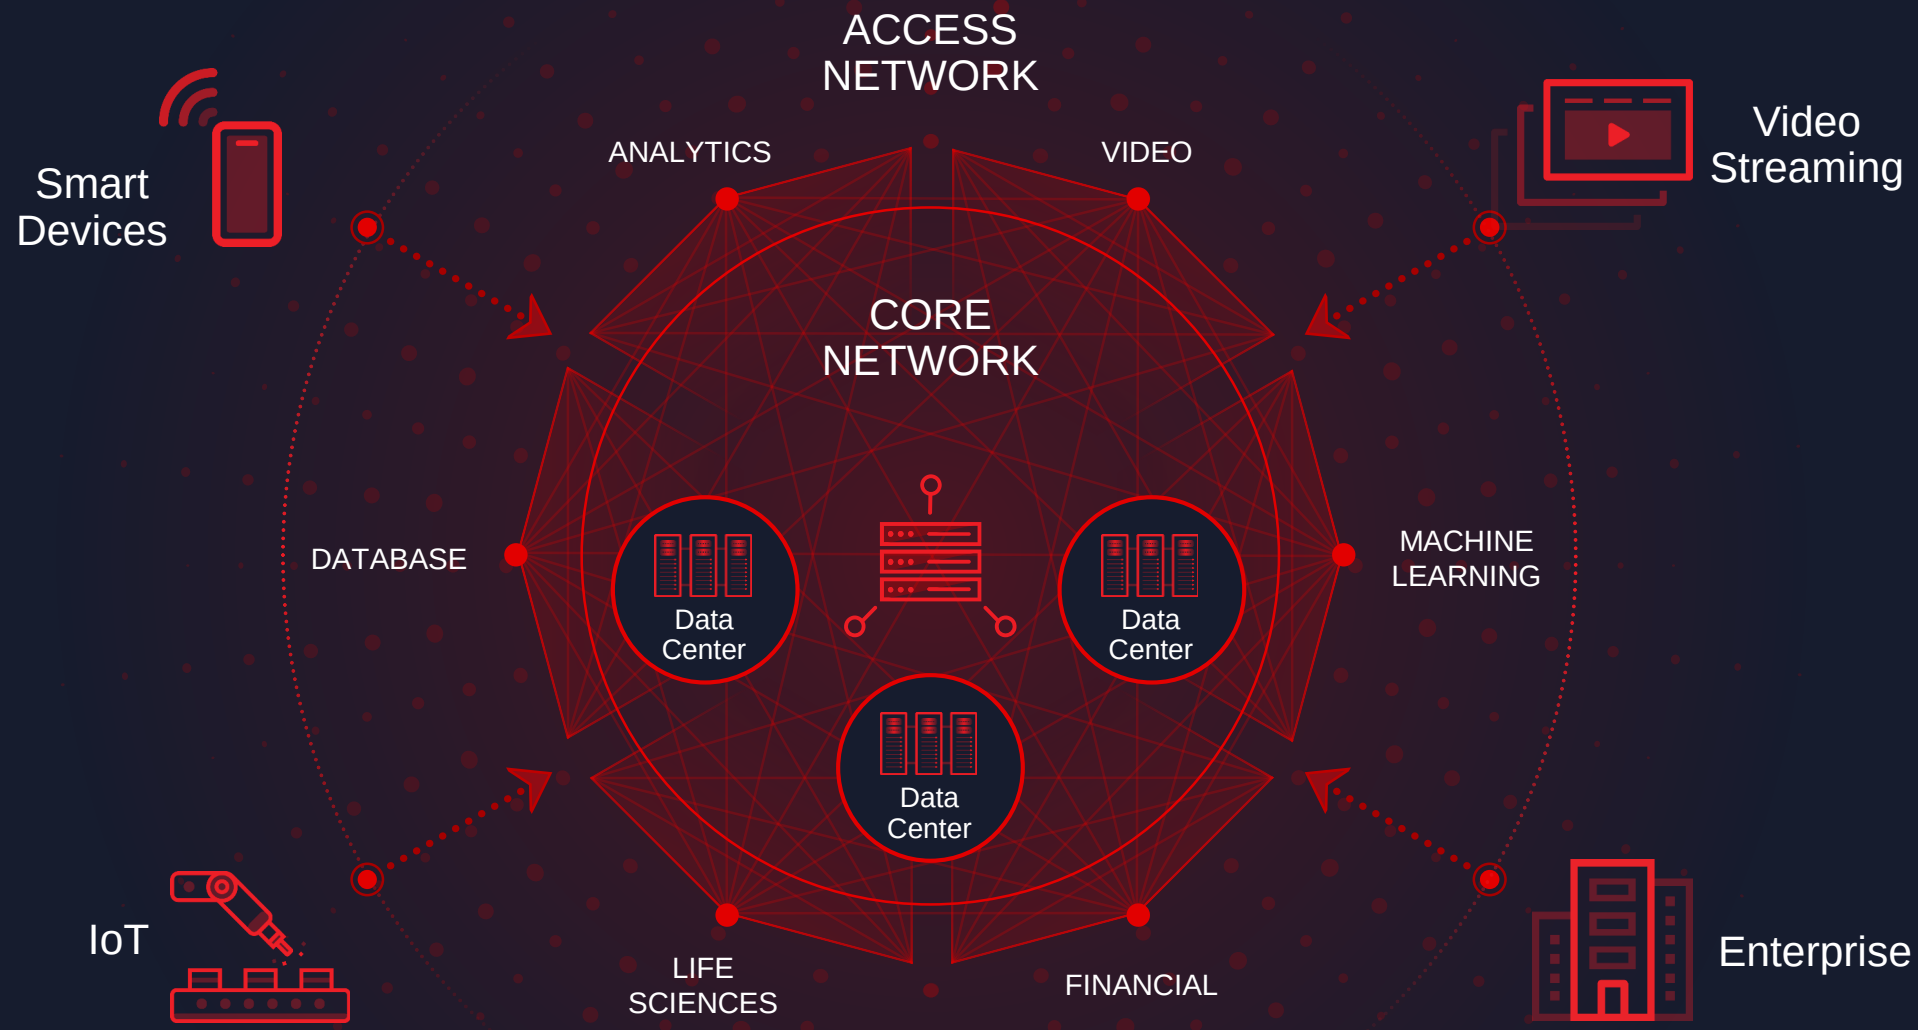

© Copyright 2020 Xilinx

XILINX

Explosion of Data from Diverse Applications & Workloads Puts Tremendous Pressure on the Core

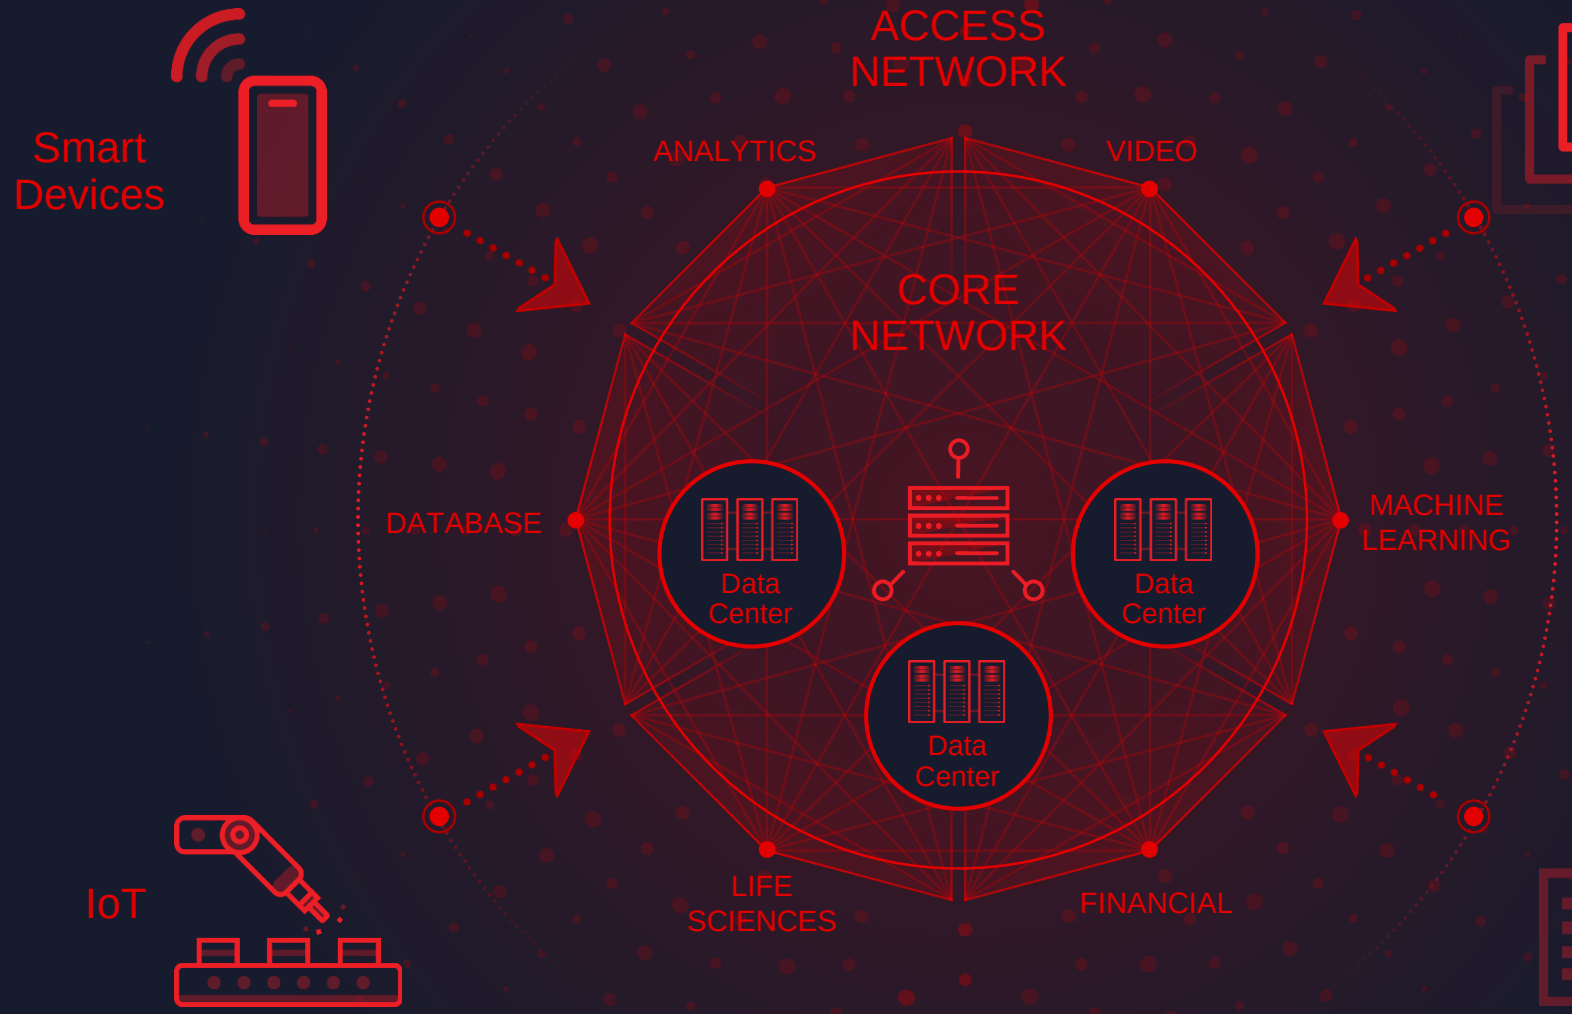

51%

Bandwidth Growth CAGR[1]

1: Equinix, Global Interconnection Index, 2019

© Copyright 2020 Xilinx

XILINX

Explosion of Data from Diverse Applications & Workloads Puts Tremendous Pressure on the Core

Smart Devices

100X

Area Traffic Capacity[1]

IoT

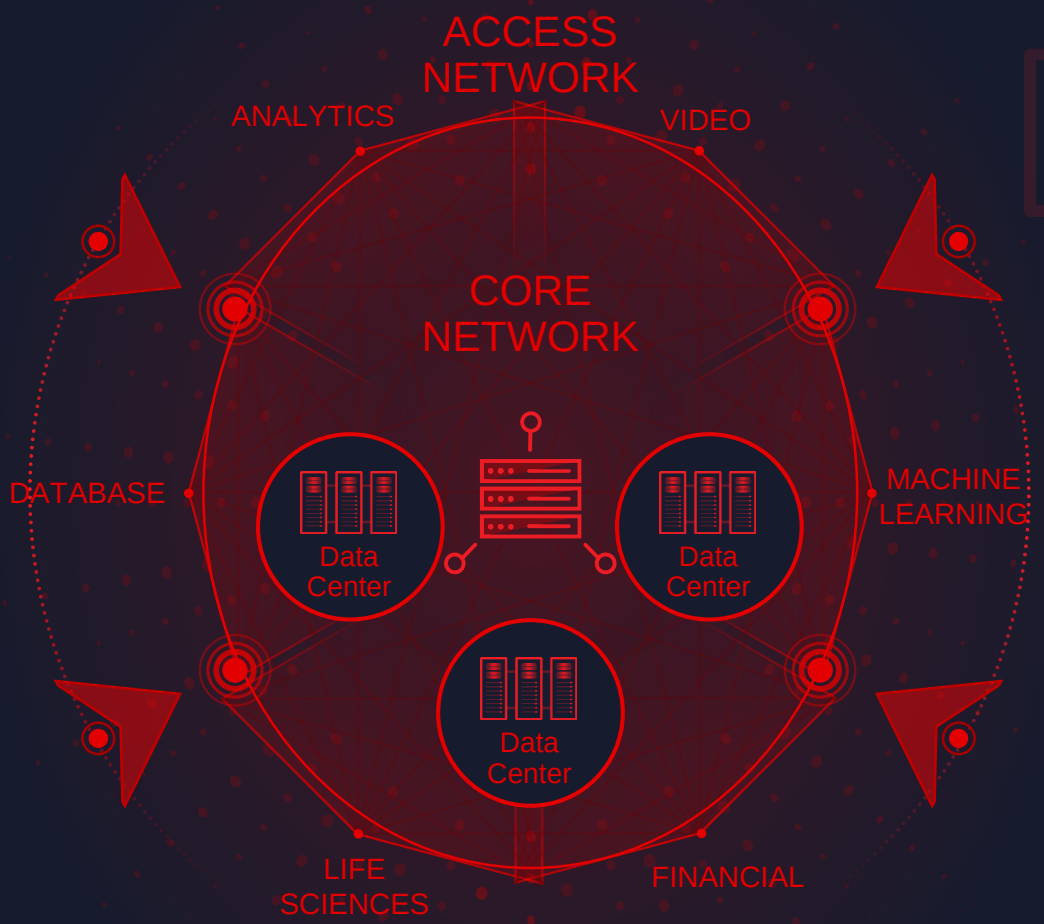

Video Streaming

51%

Bandwidth Growth CAGR[1]

Enterprise

1: Equinix, Global Interconnection Index, 2019

© Copyright 2020 Xilinx

XILINX

Data Explosion Driving Network Transformation

Core Network[1] Growth

313% CAGR Forecast for 5G Core

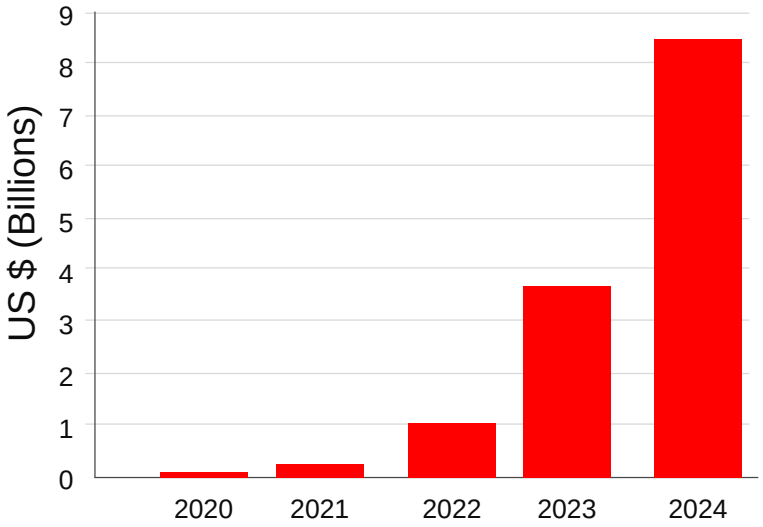

Security & Analytics[2]

Highest Priority

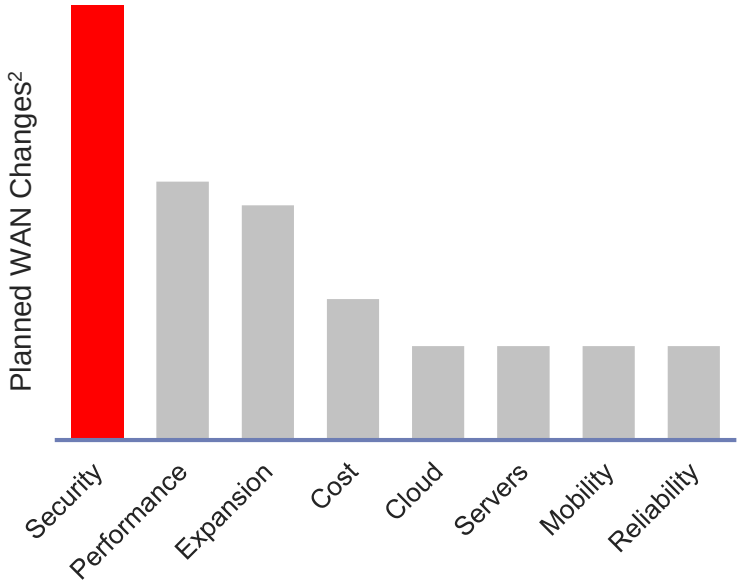

Compute vs. Bandwidth[3]

Port Speeds Surpassing Moore's Law

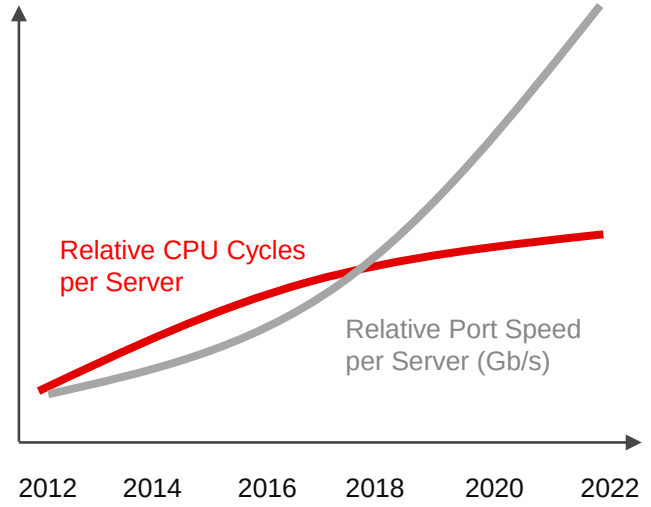

1: ABI Research, "5G Next-Generation Core and Service-Based Architecture" 2: IHS Markit, Top Changes Planned Among Network Operators 3: Xilinx Estimates

© Copyright 2020 Xilinx

XILINX.

Introducing Versal Premium
Adaptive Compute Acceleration Platform

3X
Bandwidth
for Fastest and Most Secure Networks

2X
Compute Density
for Adaptable Acceleration

Highly Integrated
HW/SW Platform
For Productivity

Bandwidth & compute density comparisons based on 14nm/16nm FPGAs

© Copyright 2020 Xilinx

Versal™ Premium is the Newest ACAP

© Copyright 2020 Xilinx

AI RF Series

HBM Series

AI Core Series

Premium Series

AI Edge Series

Prime Series

© Copyright 2020 Xilinx

XILINX

Adaptive Compute Acceleration Platform
A New Device Category

XILINX VERSAL™ | PRIME

Foundational Device Series

ADAPTIVE
- Adaptable to diverse workloads
- Future-proof algorithms

COMPUTE ACCELERATION
- Scalar Engines
- Adaptable Engines
- Intelligent Engines

PLATFORM
- SW programmable silicon infrastructure
- Pre-engineered connectivity
- Platform available at boot

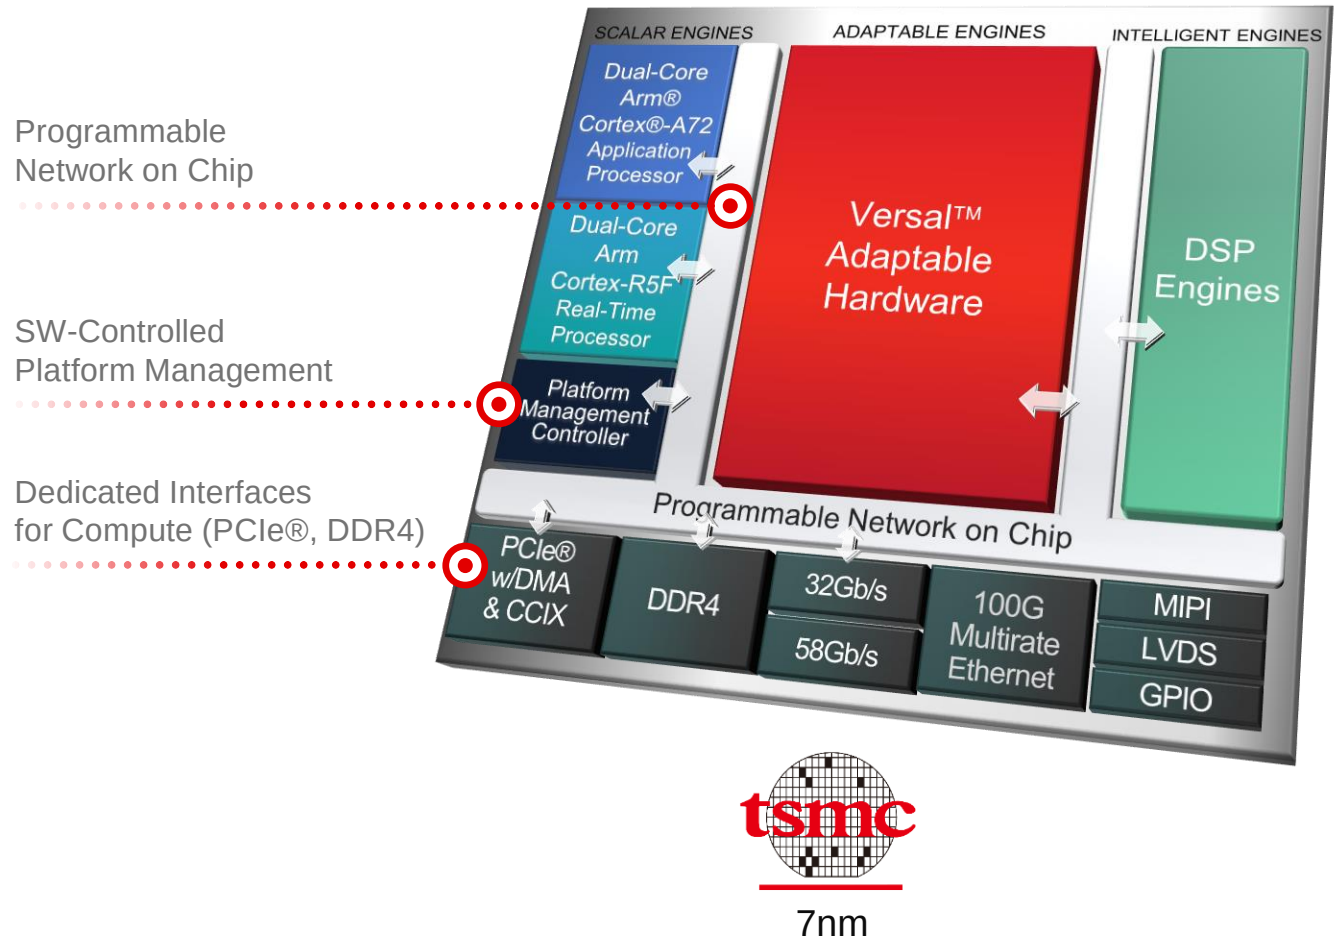

© Copyright 2020 Xilinx

XILINX

Breakthrough Integration of Networked, Power-Optimized Cores on an Adaptable Platform

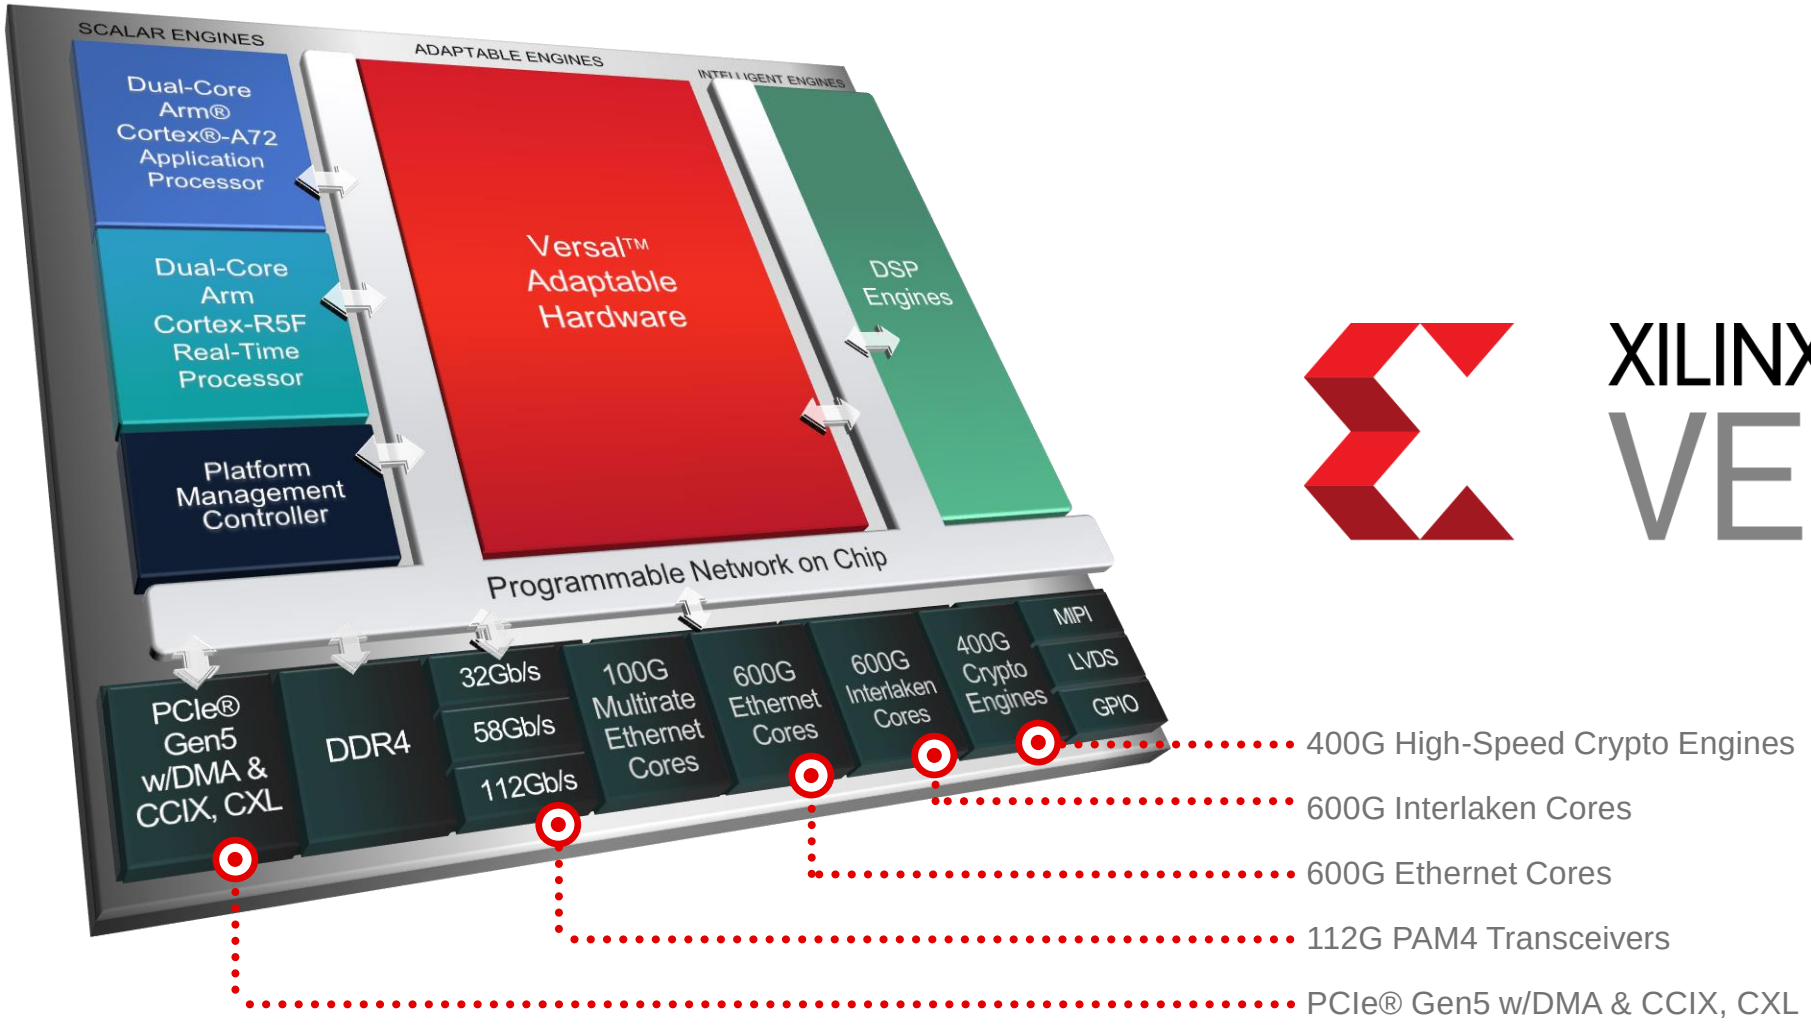

XILINX VERSAL™ | PREMIUM

© Copyright 2020 Xilinx

XILINX

Integration of Networked IP Delivers Equivalent Logic Density of 22 FPGAs

- Developers can focus on differentiation (vs. design infrastructure & connectivity)
- Surpasses limitations of Moore's Law for next-generation bandwidth and processing
- Enables greatly reduced CAPEX and OPEX

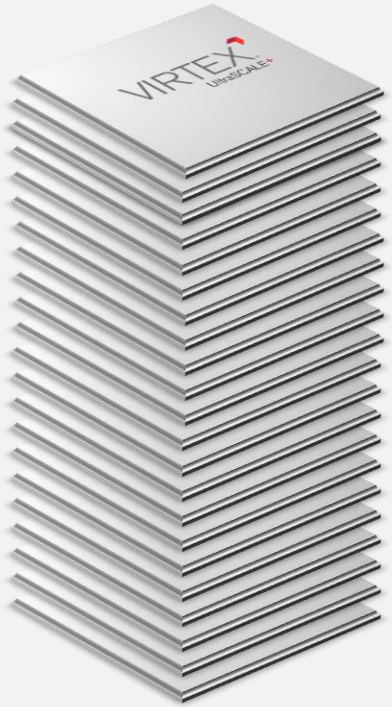

22 Equivalent 16nm FPGAs

22 Equivalent FPGAs in Versal™ Premium Integrated Cores[1]

One Versal Premium ACAP (VP1802)

1: Equivalent logic density of Ethernet, Interlaken, and Crypto cores

© Copyright 2020 Xilinx

XILINX

Integrated HW/SW Platform for All Developers

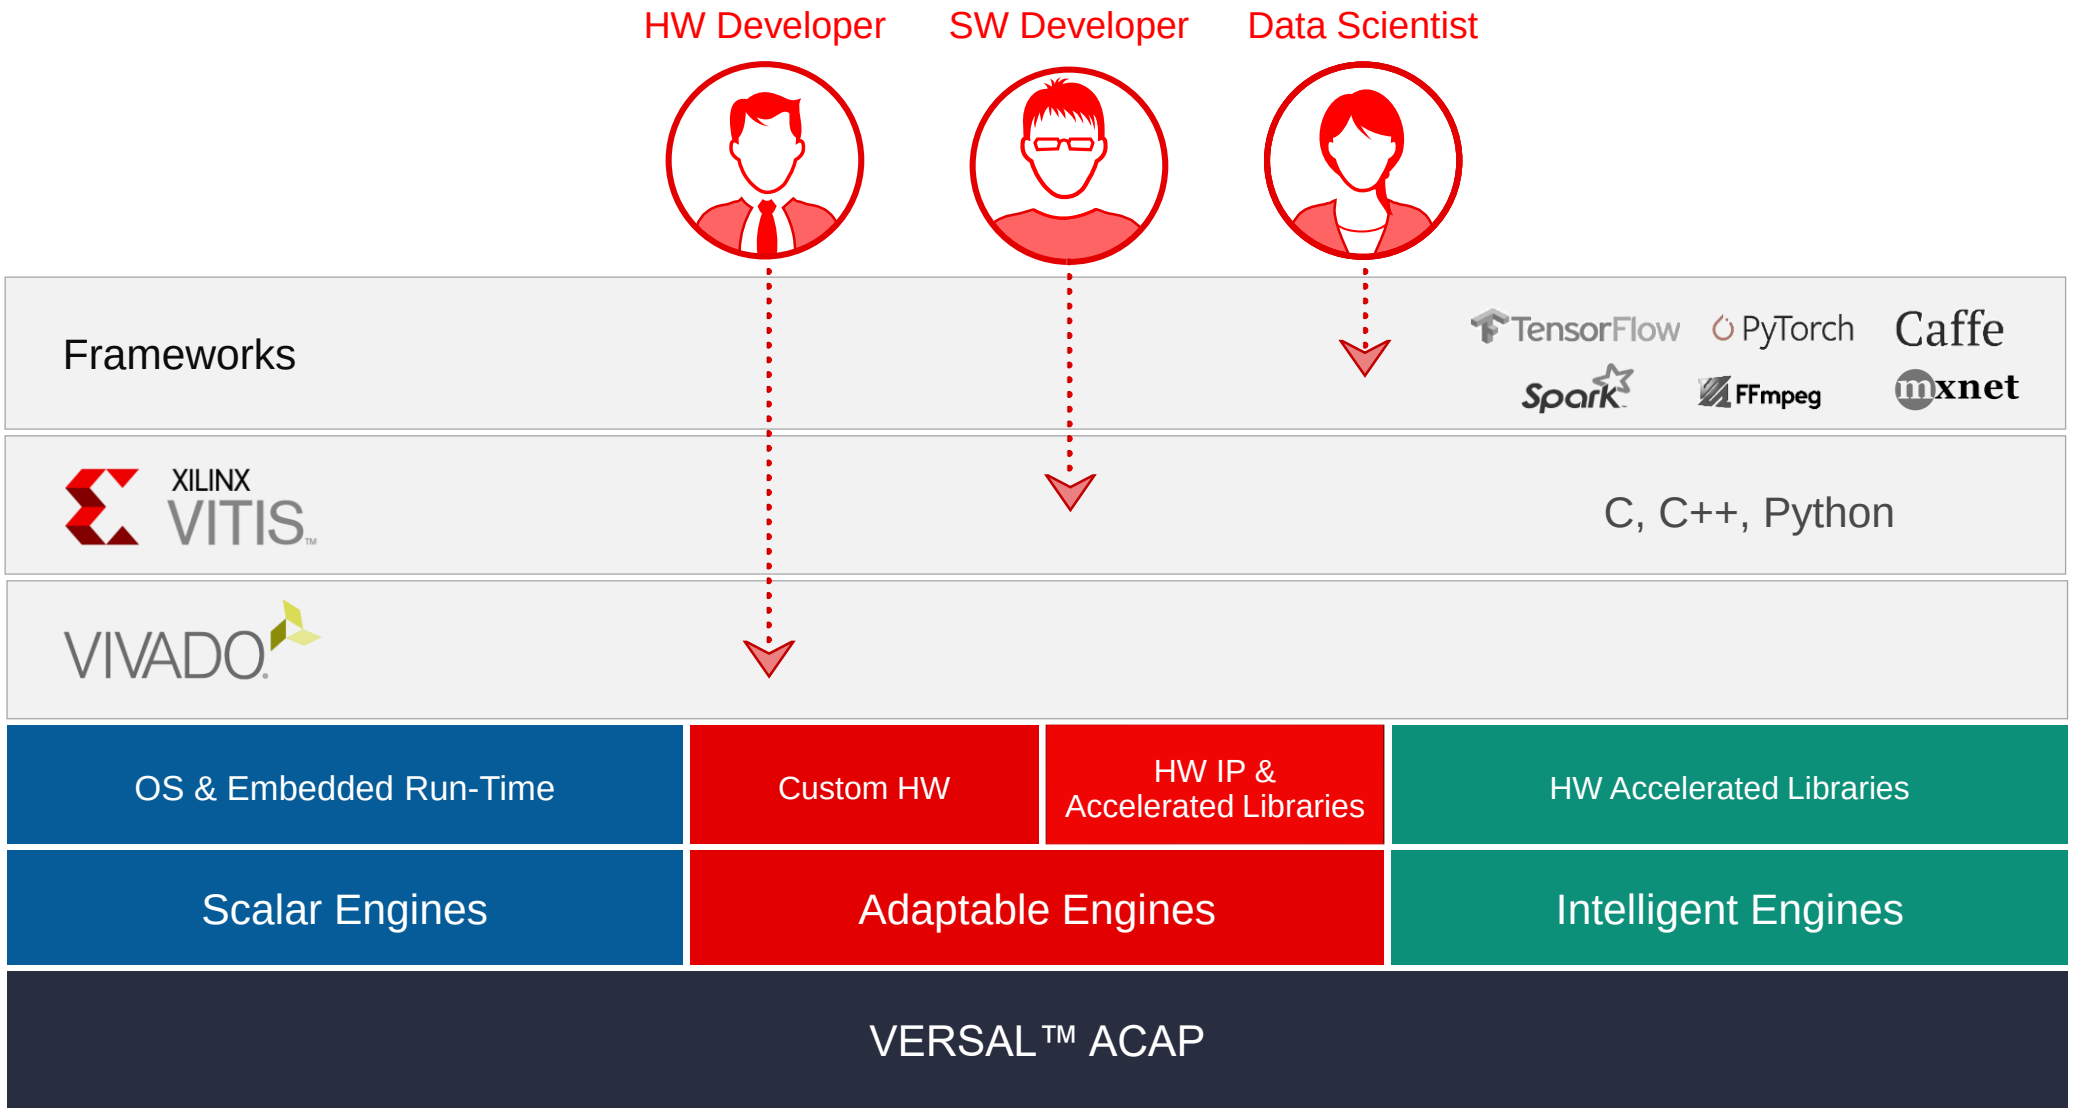

© Copyright 2020 Xilinx

XILINX

Enabling the Fastest, Most Secure Networks

© Copyright 2020 Xilinx

Next-Generation Infrastructure Demands Power-Optimized Throughput and Compute

Demands for Greater Bandwidth Density

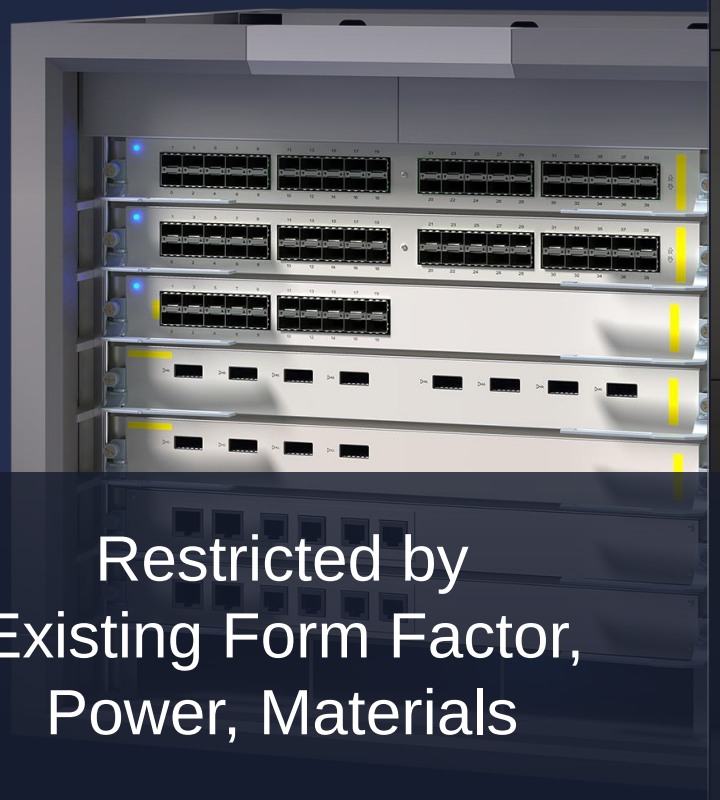

Restricted by Existing Form Factor, Power, Materials

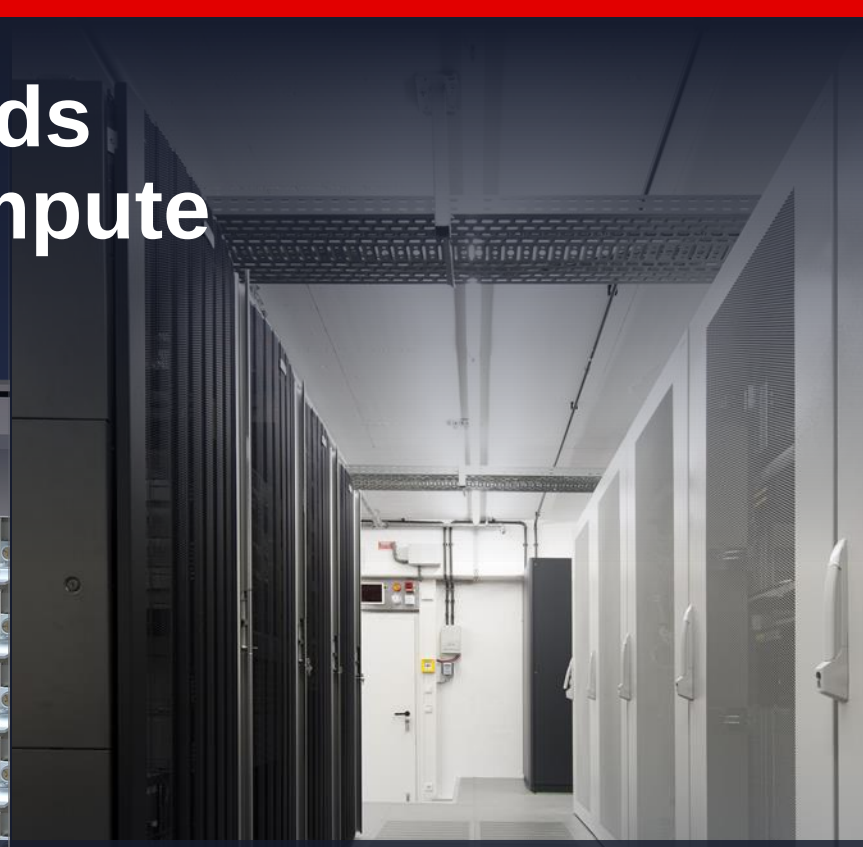

Limited Floor Space

© Copyright 2020 Xilinx

XILINX

Delivering Power-Optimized Bandwidth Density

Higher Bandwidth Density per Line Card

Scalable Transceivers

Dedicated Connectivity

Adaptable Hardware

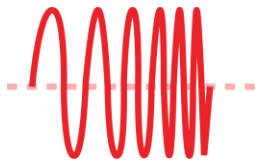

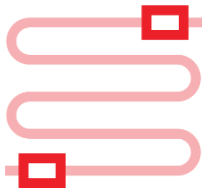

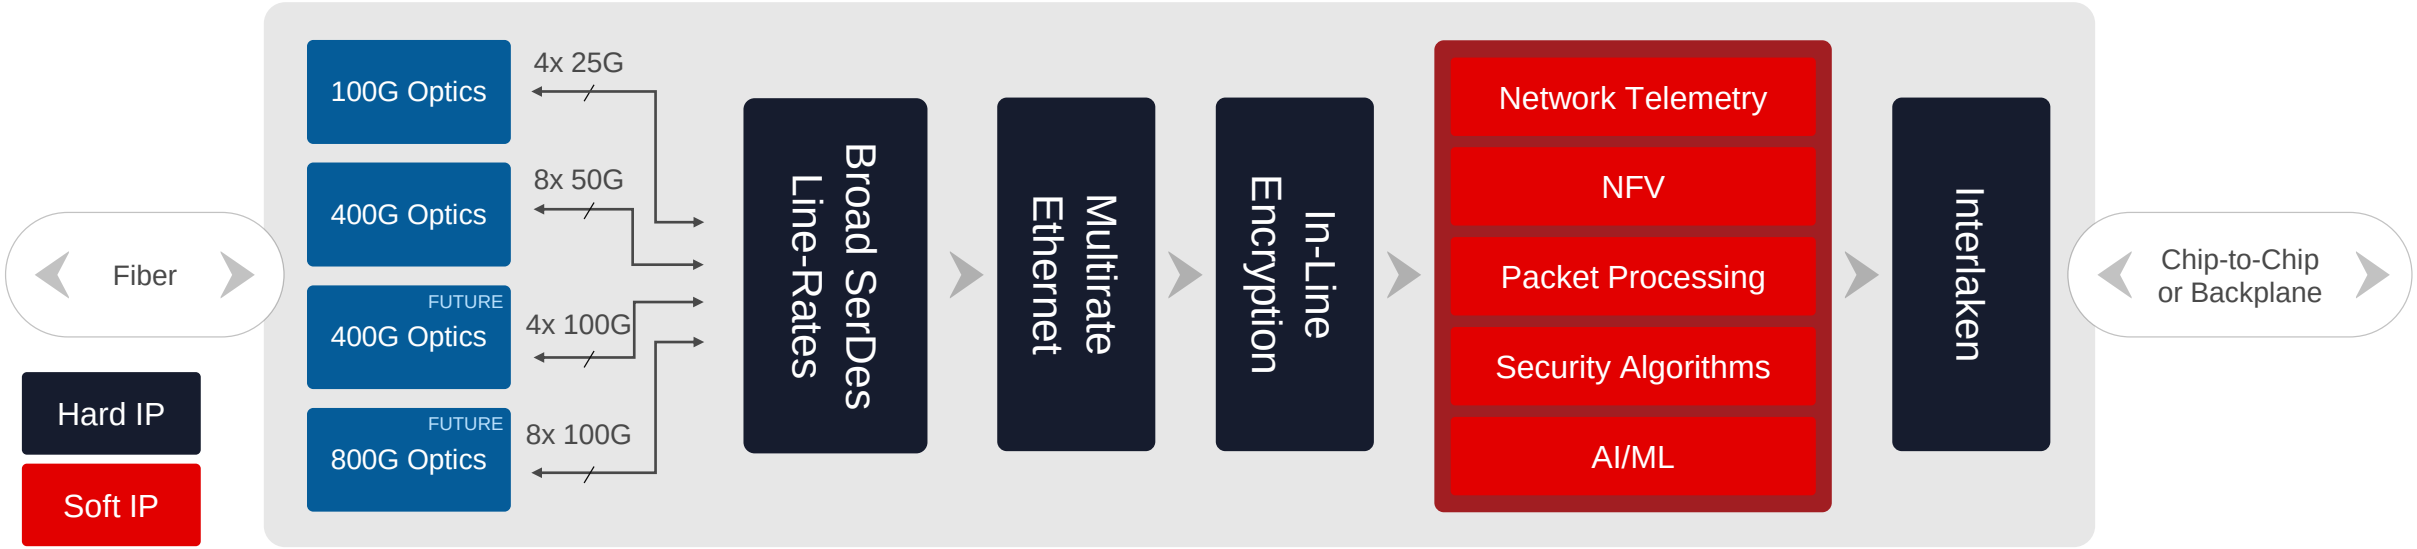

© Copyright 2020 Xilinx

XILINX

9Tb/s of Scalable, Adaptable Serial Bandwidth

Proven in
16nm/7nm Silicon

32Gb/s
NRZ

Mainstream Power-Optimized 100G Interfaces

Cost-effective 10/25/40/50/100G Ethernet with backward compatibility

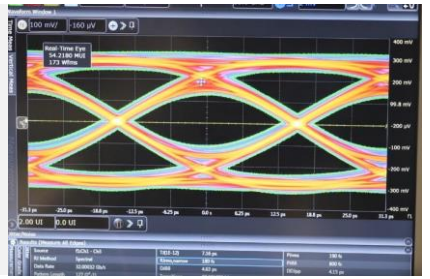

58Gb/s
PAM4

Current 400G Ramp and Deployment

Enabling latest generation optics for maximum system bandwidth

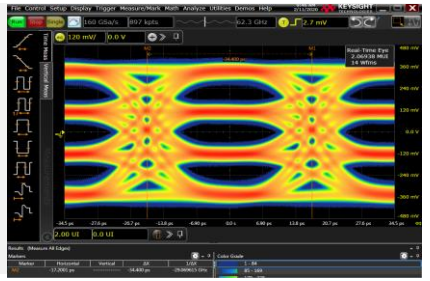

112Gb/s
PAM4

Future 800G Networks on Existing Infrastructure

Industry moving towards single-lane 100G optics and
800G infrastructure

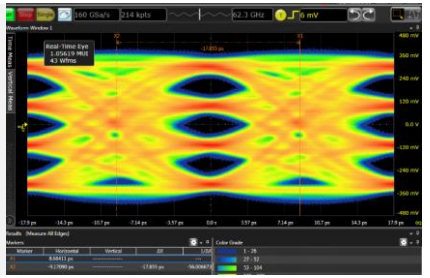

Copper Cable

Optics

Backplane

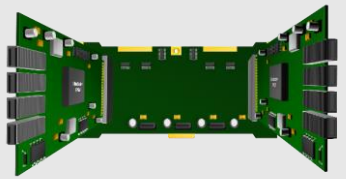

Dedicated Connectivity IP for Secure Networking

5Tb/s of scalable Ethernet throughput

- For next-gen 400G and 800G infrastructure in the core network
- Multirate: 400/200/100/50/40/25/10G with FEC
- Multi-standard: FlexE, Flex-O, eCPRI, FCoE, OTN

1.8Tb/s of off-the-shelf Interlaken connectivity

- Scalable chip-to-chip interconnect from 10Gb/s to 600Gb/s
- Integrated RS-FEC for power-optimized error correction

1.6Tb/s of encrypted line rate throughput

- World's only hardened 400G Crypto Engine on an adaptable platform
- AES-GCM-256/128, MACsec, IPsec

One Platform

From Access → Core

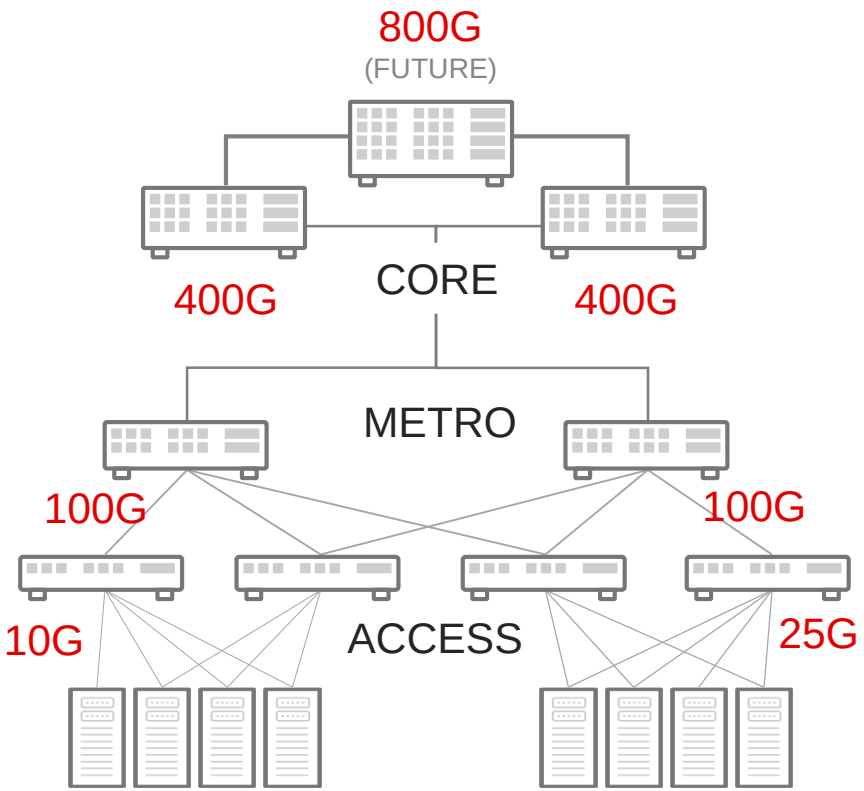

Pre-Built Connectivity for Fastest Time to Market and ASIC-Class Power/Performance

© Copyright 2020 Xilinx

XILINX

Programmable Logic for HW Differentiation, Evolving Standards, and AI/ML

For Differentiation and Future-Proofing

- World's highest logic density 7nm platform
- Differentiate, e.g., in-band network telemetry, vRAN
- Adapt to standards and protocols

AI for Network Anomaly Detection

- Intrusion detection and malware identification
- Adaptable AI algorithms for emerging threats
- Xilinx Random Forest IP now available

AI for Provisioning and Network Performance

- Auto-detection and correction of performance bottlenecks
- Self-provisioning for maximizing uptime (MLP)

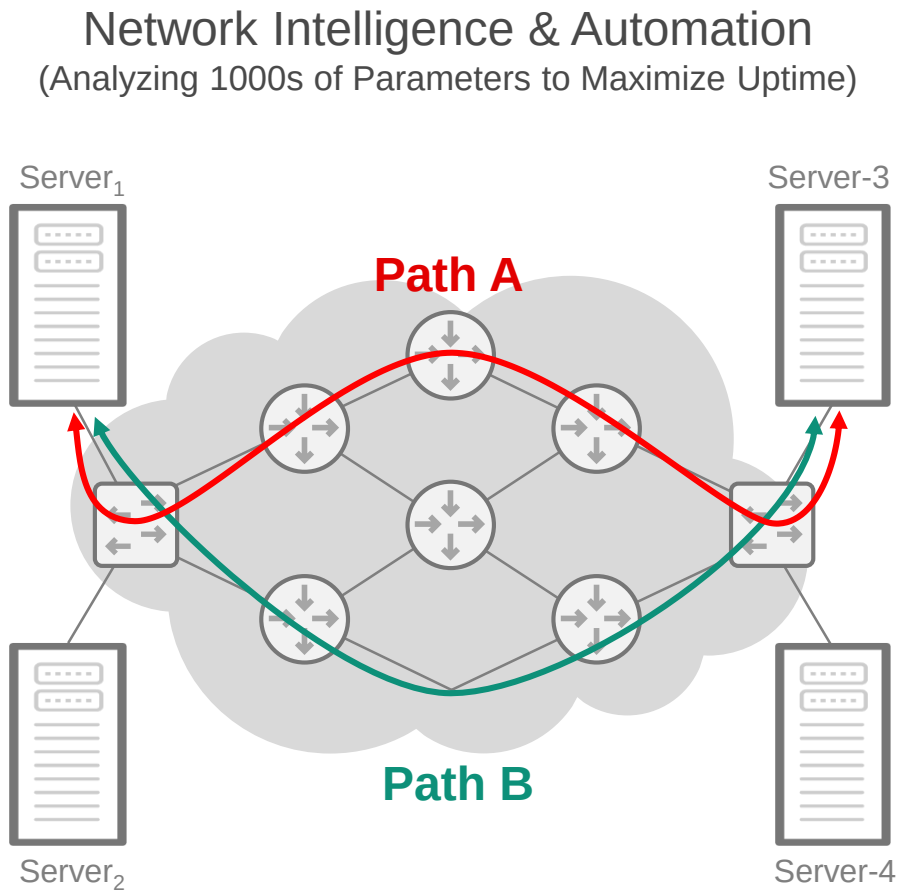

© Copyright 2020 Xilinx

XILINX.

Industry-Leading Multi-Terabit Throughput vs. ASSPs

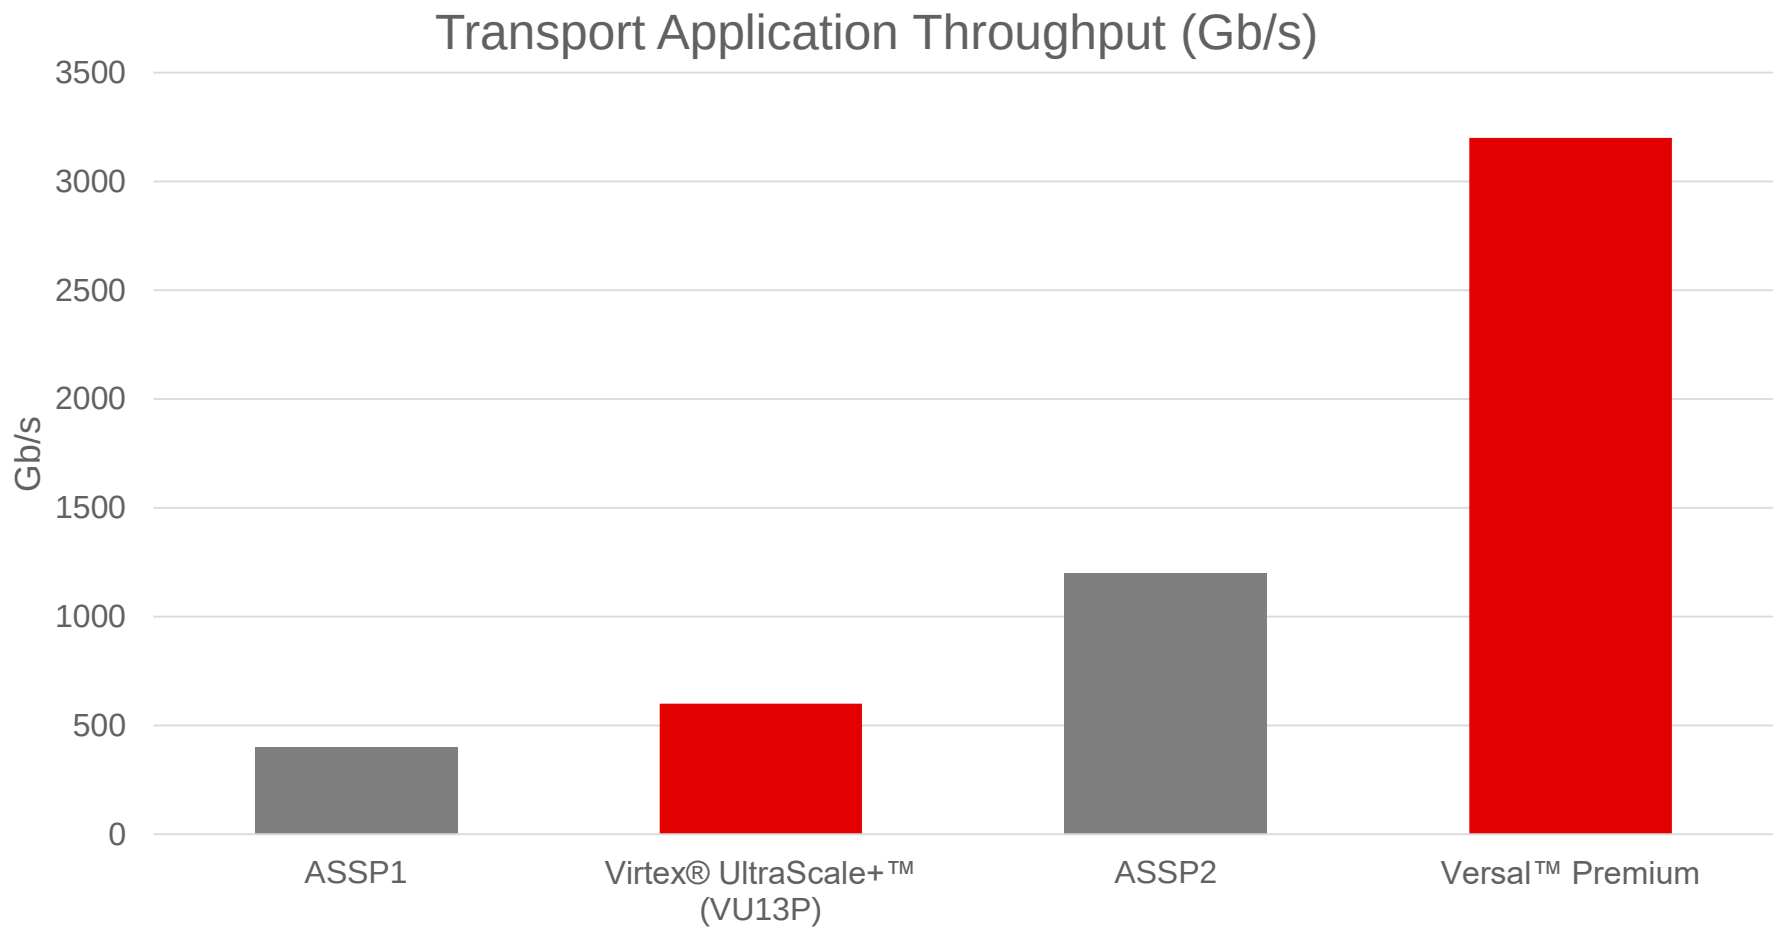

ASSP$_1$: https://www.microsemi.com/product-directory/multi-service-otn-processors/4227-pm5990-digi-g4.
ASSP$_2$: https://www.microsemi.com/product-directory/multi-service-otn-processors/5056-pm6010-digi-g5-otn-processor

© Copyright 2020 Xilinx

Single-Chip 800G DCI Throughput Under 100 Watts

< Half the Power, Half the Footprint

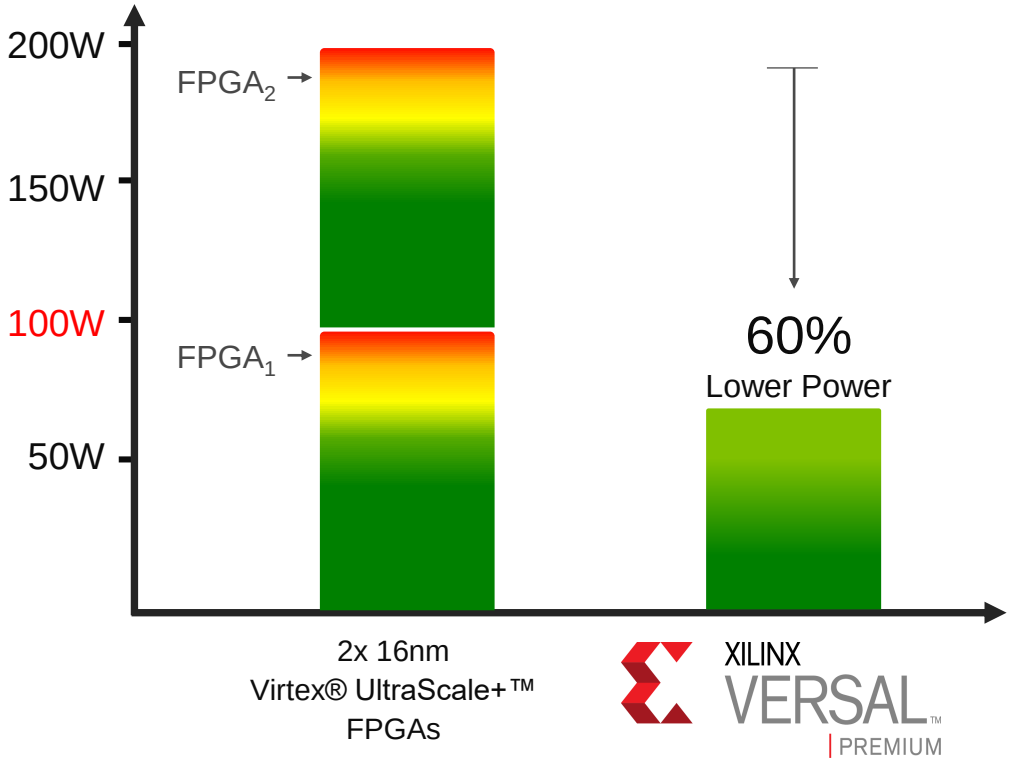

2X Bandwidth Density

Previous Gen FPGAs

Versal Premium

Same Bandwidth
at Half the Rack Space

© Copyright 2020 Xilinx

XILINX

Vivado Unlocks the Integration of Versal Premium

Modular IP Integration through Vivado and NoC

- Graphically connect hard/soft IP using Vivado® IP Integrator
- Streamlined, push-button flow with NoC Compiler
- NoC guarantees timing for critical interconnect paths

System integration in days vs. months

- Easily scale number of 100G, 400G, 600G cores
- Scale Design for 400G and Beyond

High-Speed, Unified Debug Environment

- High-bandwidth, SerDes-based debug and trace
- 1000X faster readback vs. traditional FPGAs
- Cohesive debug across heterogeneous engines

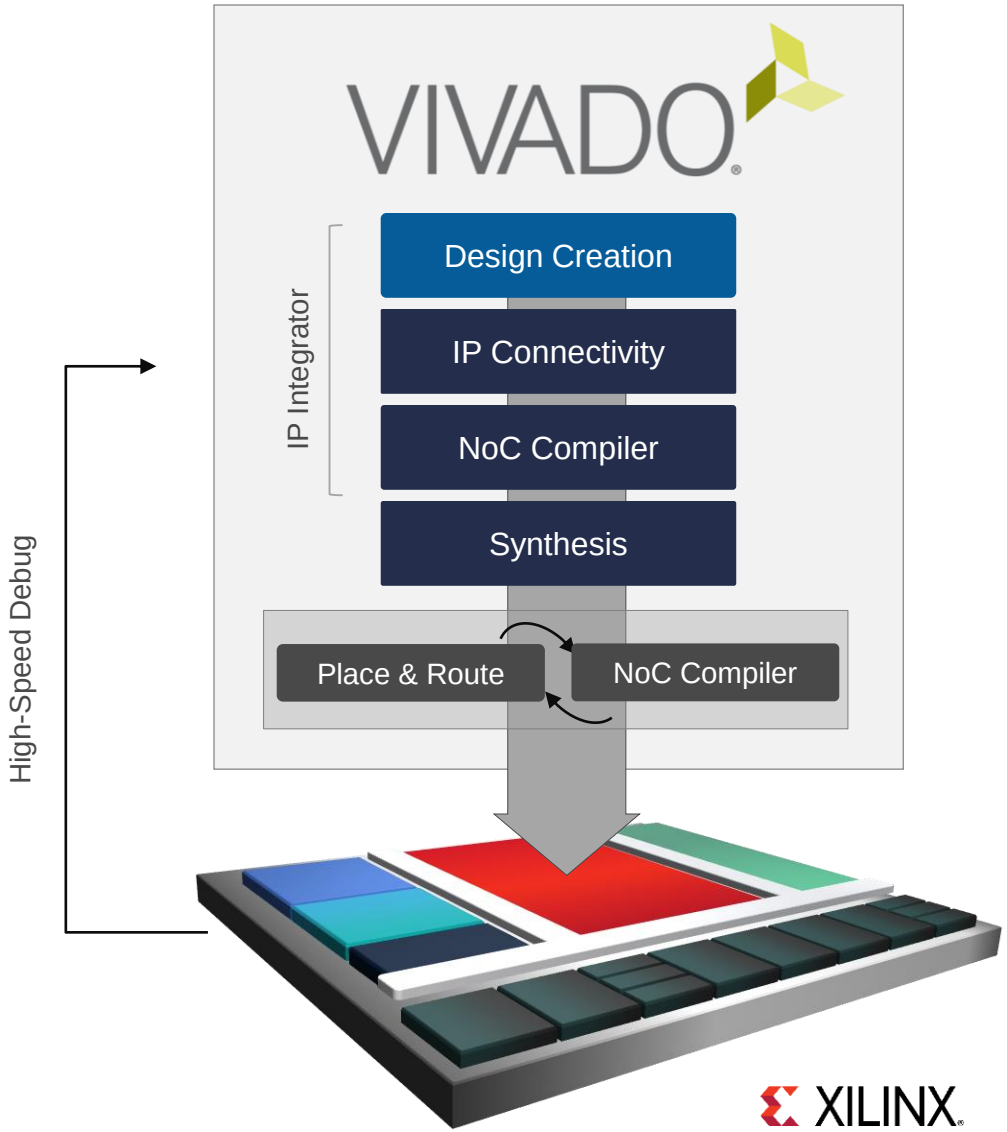

© Copyright 2020 Xilinx

XILINX

Highest Compute Density with Adaptable Acceleration

© Copyright 2020 Xilinx

XILINX

Acceleration Challenges for Hyperscale Cloud Providers

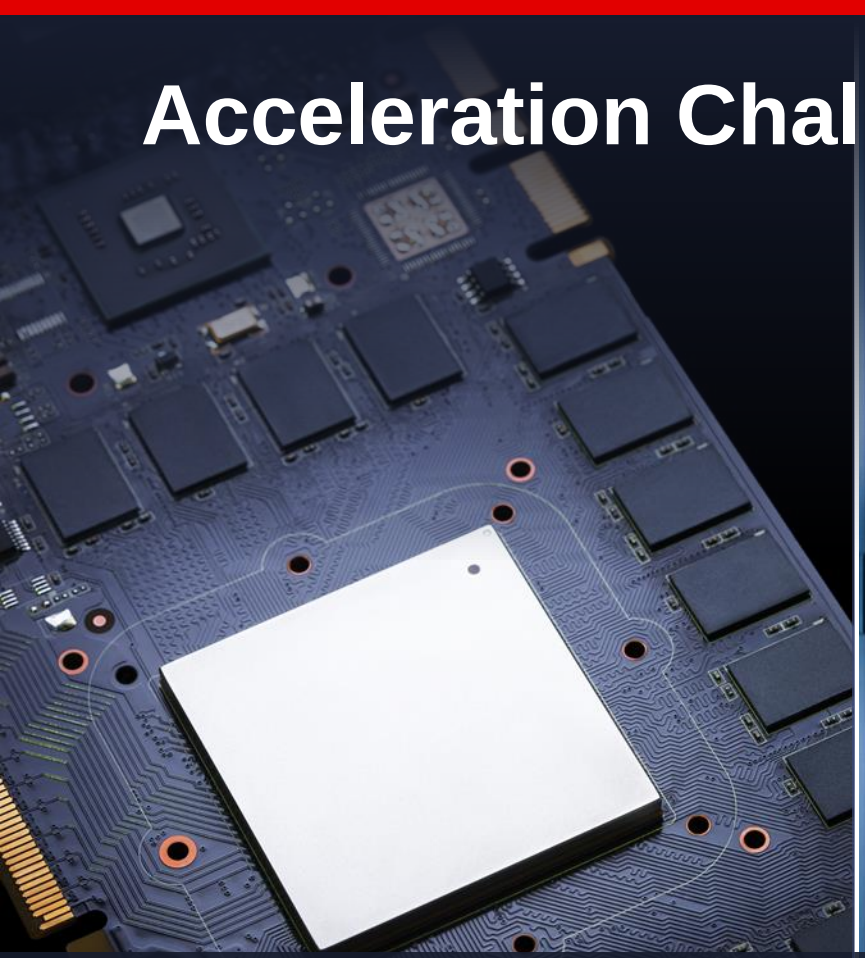

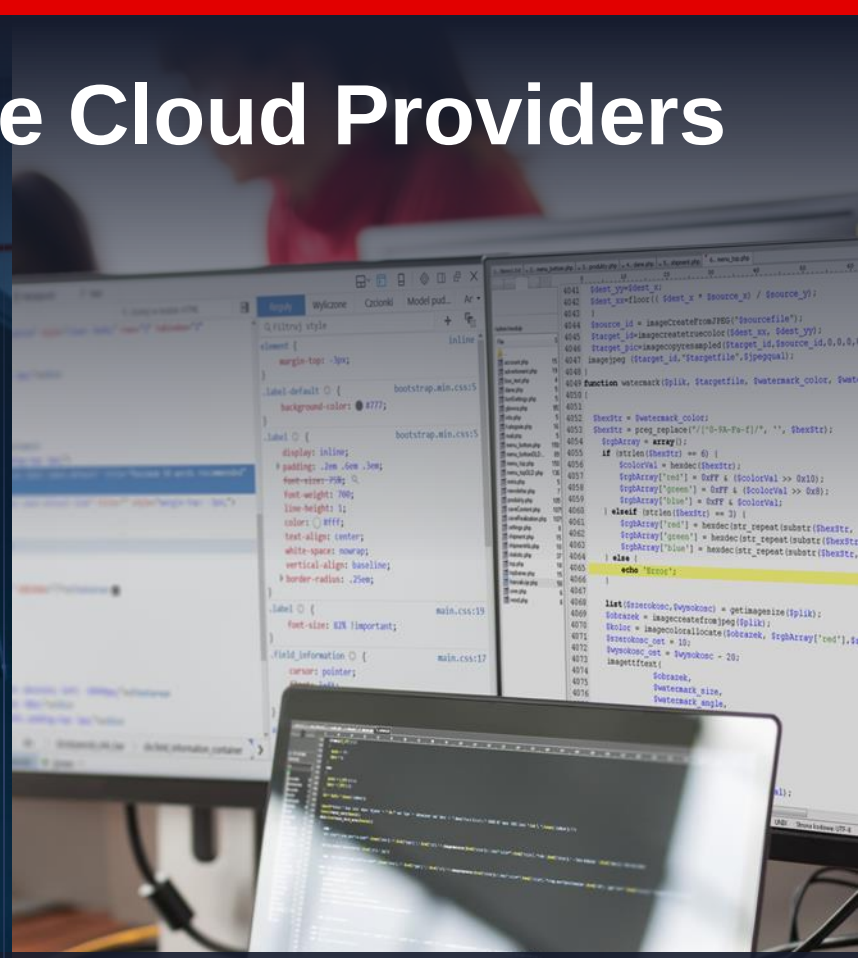

Accelerator Performance and Throughput

Integration and Connectivity to Cloud Infrastructure

Software Solution Stack

Cloud Providers Consider Overall TCO of an Accelerator Solution

Workload Provisioning with Dynamic Function eXchange

Dynamic Function eXchange (DFX)
Swaps Kernels in Milliseconds
8X faster than 16nm FPGAs

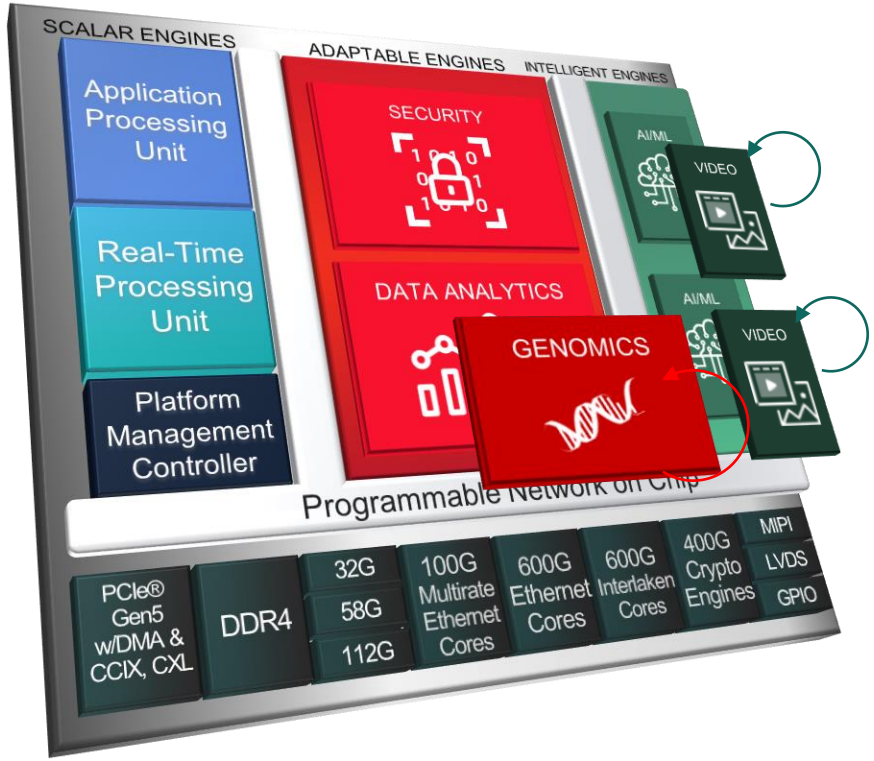

Dynamically **Provision** Accelerators
for Most Efficient Use of Infrastructure

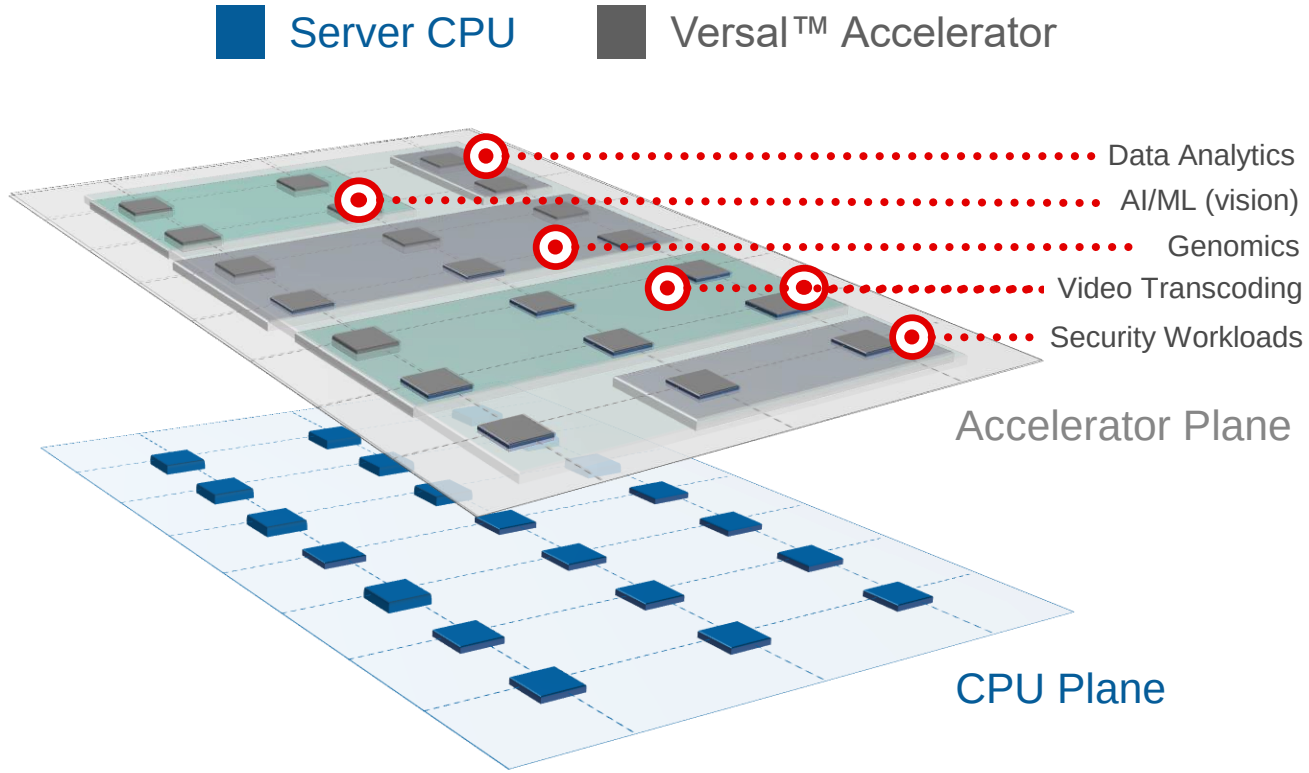

Reduced TCO and Latency for Superior Orchestration and User Experience

© Copyright 2020 Xilinx
XILINX.

Key to Acceleration: On-Chip Memory Bandwidth and Capacity

Up to 1Gb of Tightly Coupled Memory for Performance, Power, Latency

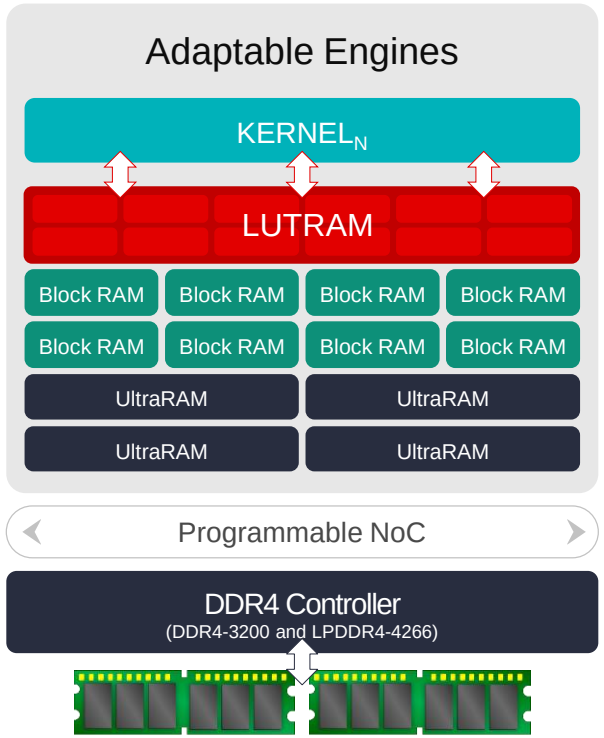

On-Chip Memory Bandwidth (TB/s)

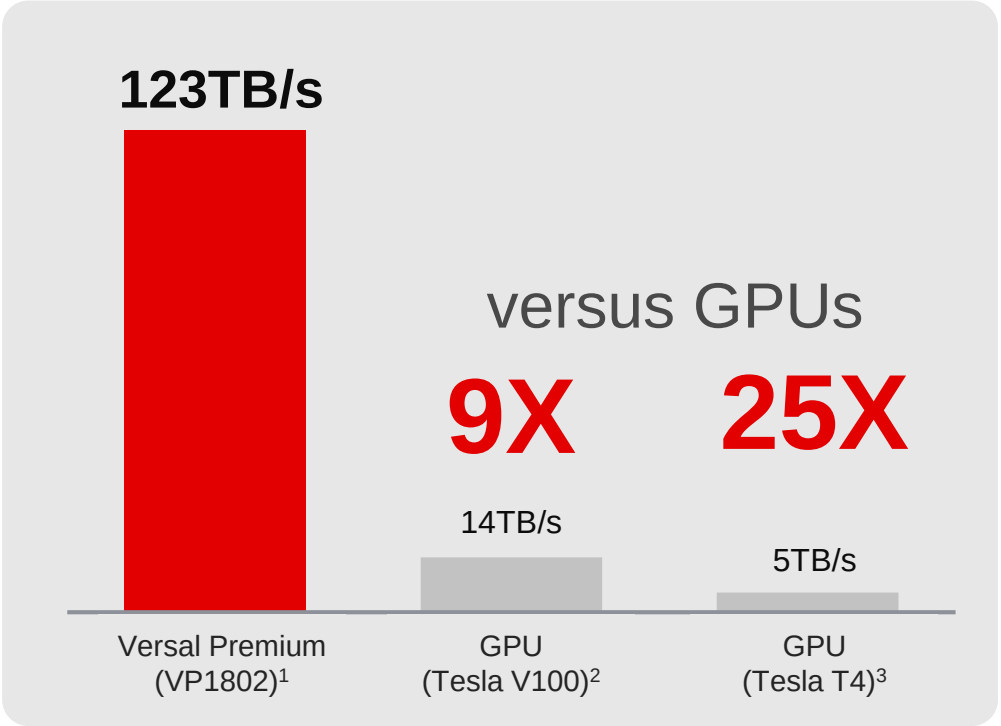

Versal™ Premium ACAP unlocks performance that GPUs can't achieve

1: Memory bandwidth assumes largest Versal Premium device, all available block RAM and UltraRAM at their maximum rates, 72-bit dual-port configuration
2: "Dissecting the NVidia Volta GPU Architecture via Microbenchmarking"- https://arxiv.org/pdf/1804.06826.pdf
3: "Dissecting the NVidia Turing T4 GPU via Microbenchmarking" - https://arxiv.org/pdf/1903.07486.pdf

© Copyright 2020 Xilinx
XILINX.

Heterogeneous Engines + Memory Bandwidth Deliver Breakthrough Performance for Diverse Workloads

Image Classification

ResNet50 (224x224)

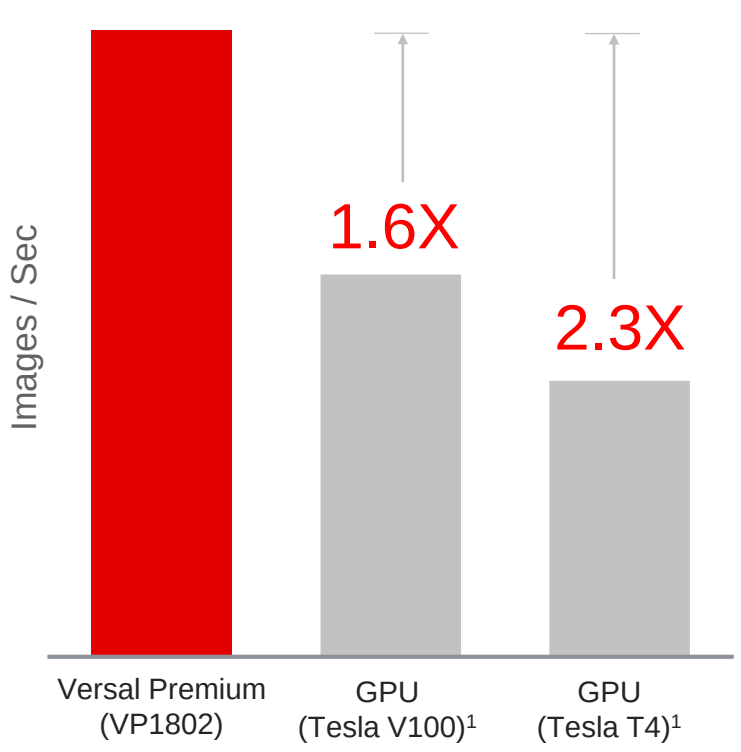

Object Detection

Yolov2 (608x608)

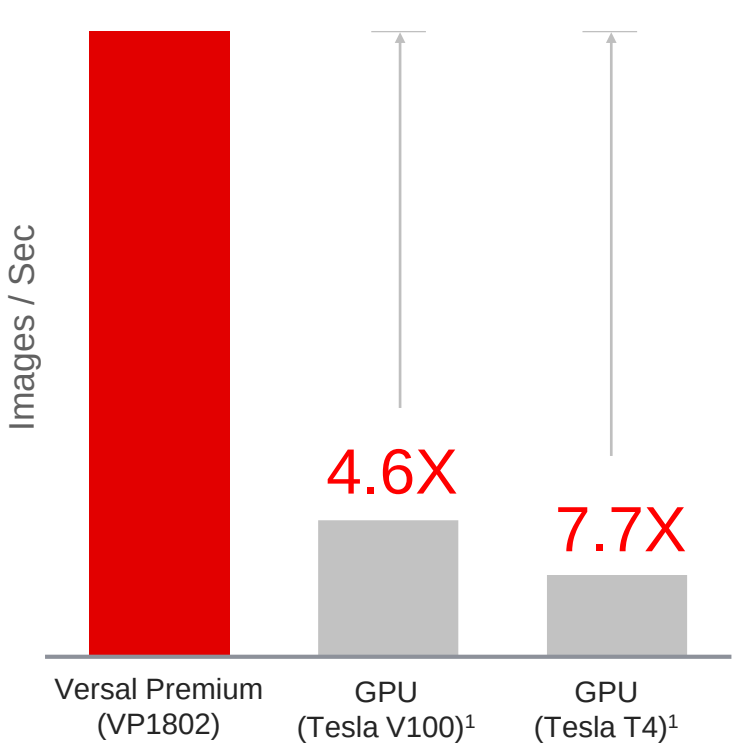

Anomaly Detection (AI)

Random Forest

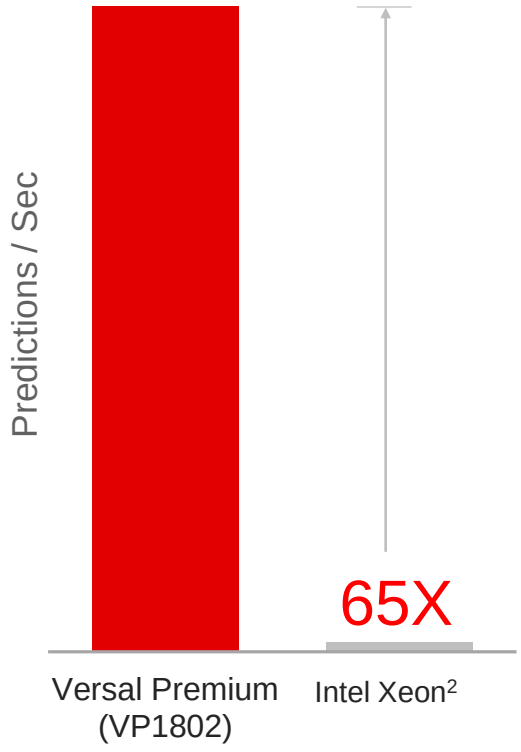

1: NVidia Data Center Deep Learning Product Performance, https://developer.nvidia.com/deep-learning-performance-training-inference

2: Xilinx Estimates, 2nd Generation Intel Xeon Scalable Processors (“Cascade Lake”)

© Copyright 2020 Xilinx

XILINX.

Integrated Shell for Dedicated Connectivity & Cloud Deployment

‘Shell’: Pre-Built Infrastructure for Cloud Connectivity

- Hardens all connectivity to data center infrastructure
- CPU-host and system memory communication available at boot
- Features PCIe® Gen5 for next-gen host communication

‘Role’ for HW Kernels and Compute Acceleration

- Simplified kernel placement and timing closure
- Easily swap kernels for “Virtualized Accelerators”

Streamlined HW Development and Deployment

- Versal™ built from the ground-up to simplify accelerator development
- HW designers spend less time on connectivity-to-cloud infrastructure

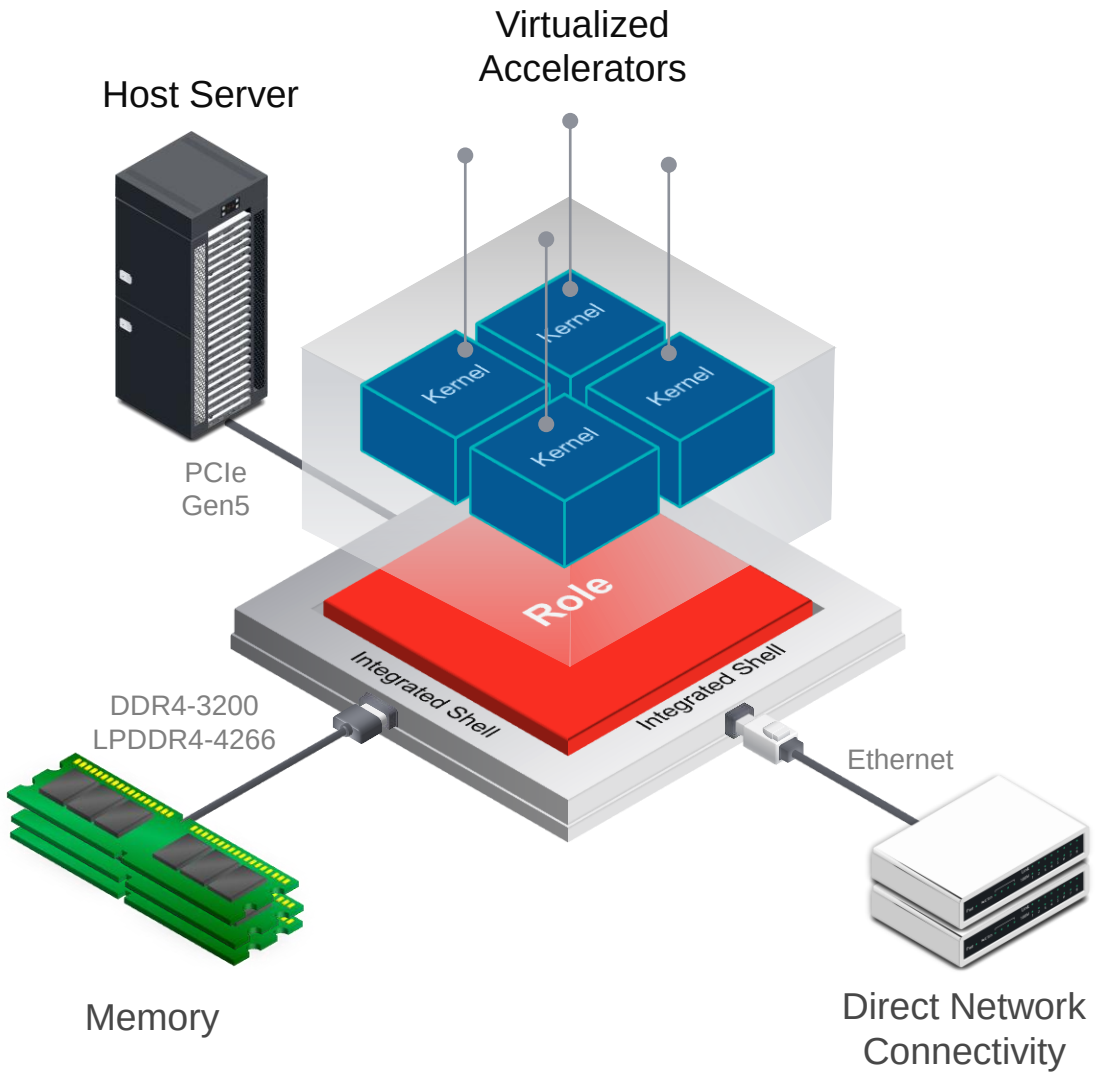

© Copyright 2020 Xilinx

XILINX.

Integrated Shell Frees More Logic for Customization

Virtex® UltraScale+™ VU9P

200K LUTs Used for Infrastructure

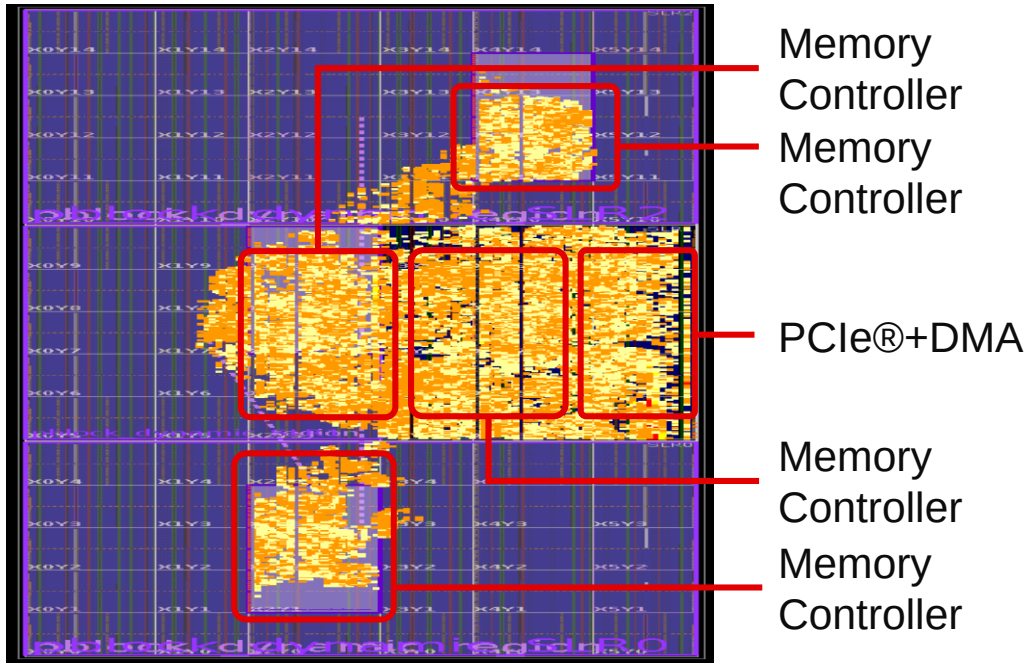

Versal™ Device

Zero LUTs Used for Infrastructure

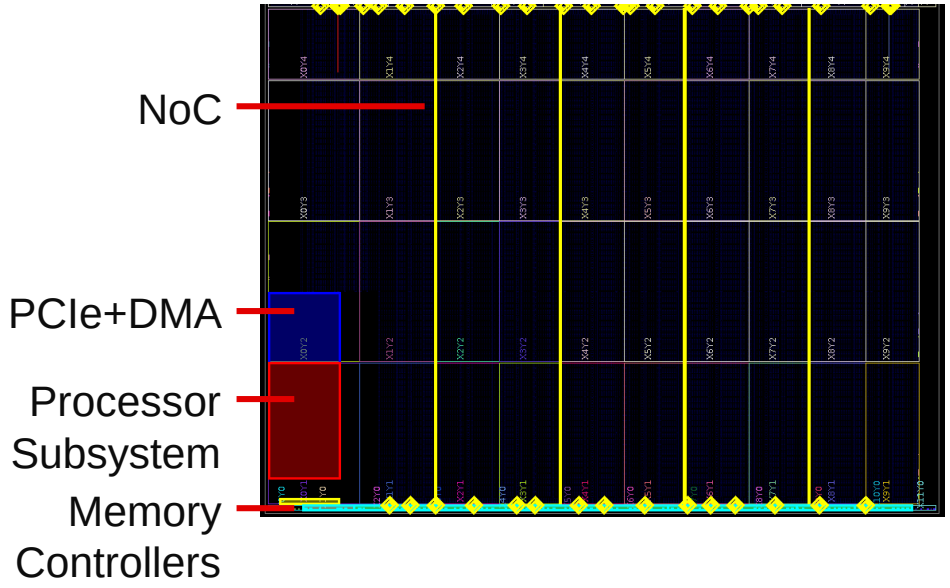

Heterogeneous Integration Reduces TCO and Enables Greater Differentiation

© Copyright 2020 Xilinx

XILINX.

Vitis Unified Software Platform

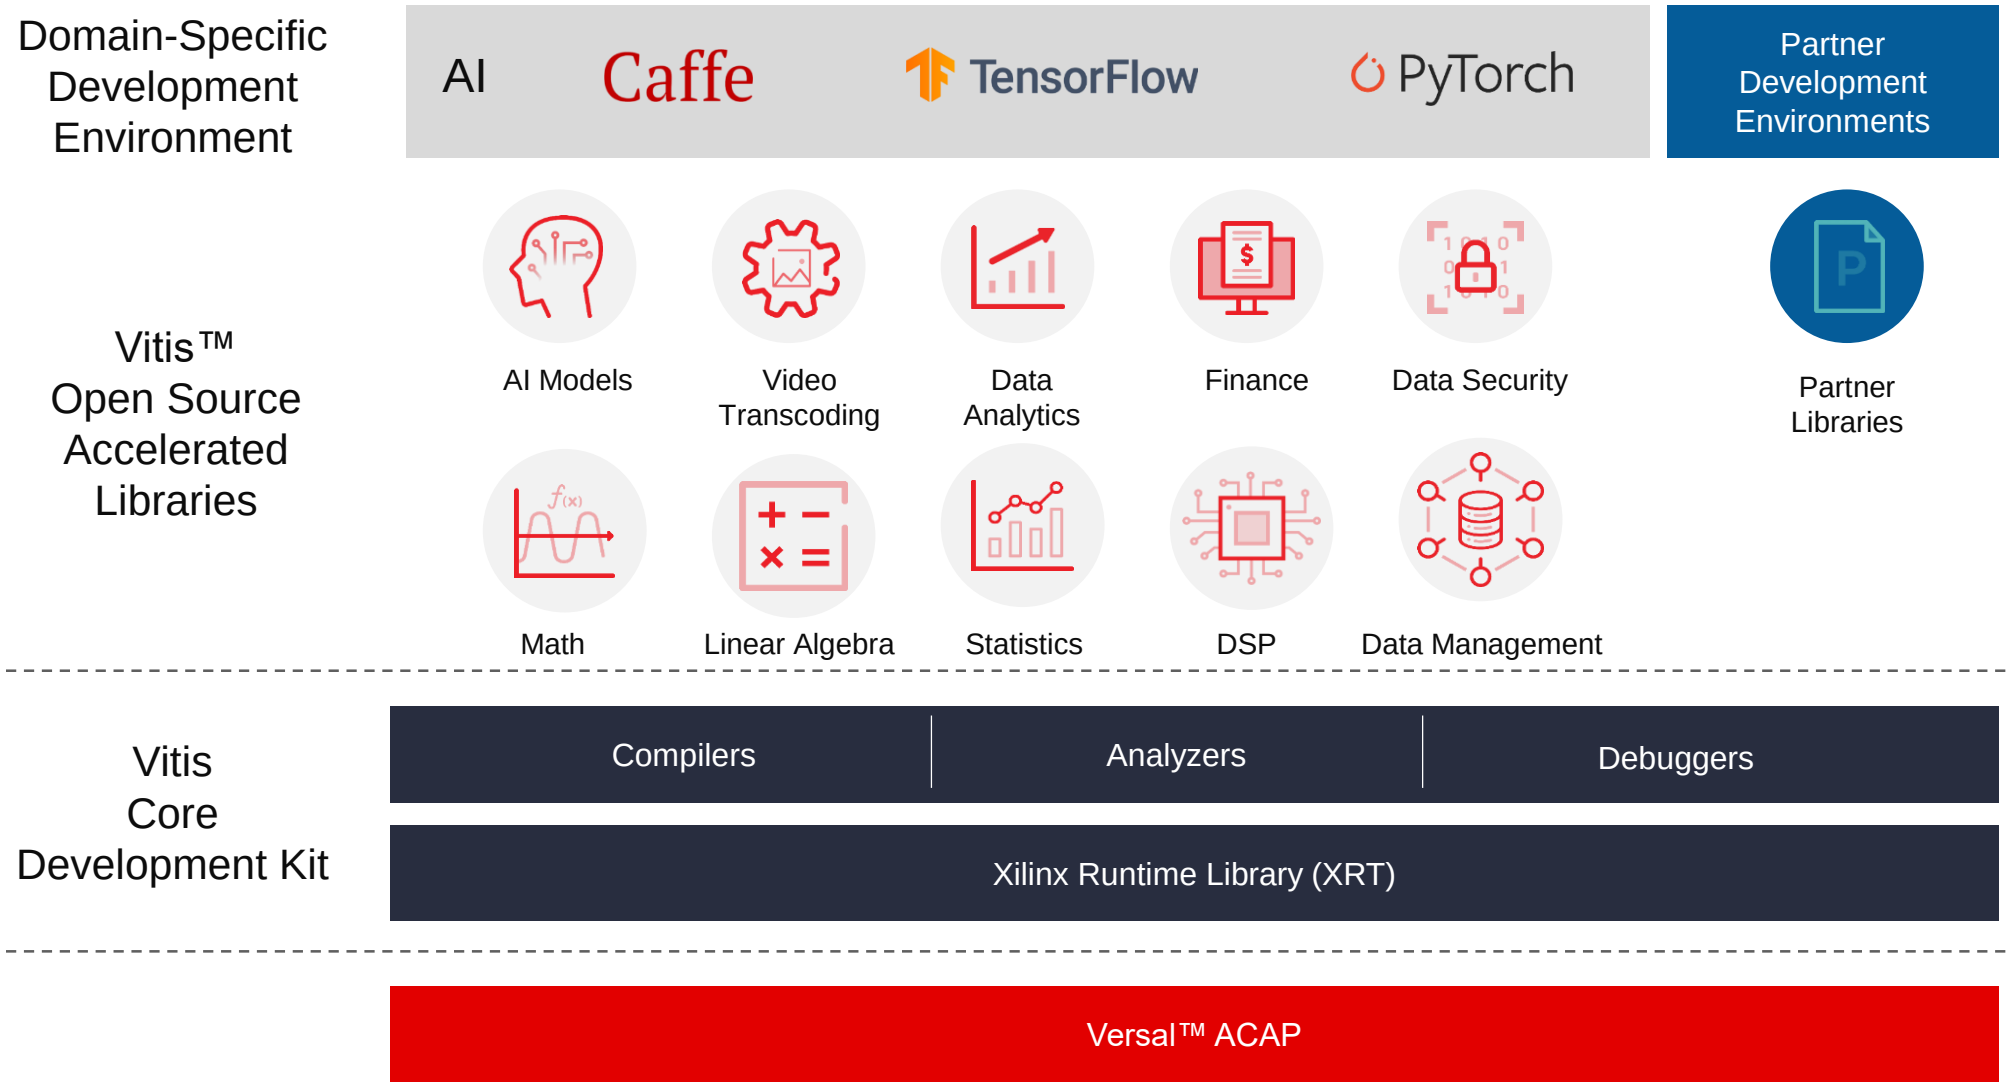

© Copyright 2020 Xilinx

XILINX.

Product Portfolio and Getting Started

© Copyright 2020 Xilinx

Versal™ Premium Portfolio: Scalable for Network & Cloud

Network 

Access/Metro

Regional/Core

Cloud 

Scalable Cloud Acceleration

		VP1102	VP1202	VP1402	VP1502	VP1552	VP1702	VP1802
Engines	System Logic Cells	1.6M	2.0M	2.2M	3.8M	3.8M	5.6M	7.4M
	Adaptable Engines (LUTs)	720K	900K	1M	1.7M	1.8M	2.5M	3.4M
	Intelligent Engines (DSP Slices)	1.9K	4K	2.7K	7.4K	7.4K	11K	14K
	Scalar Engines	Dual-Core Arm® Cortex®-A72 Application Processing Unit / Dual-Core Arm Cortex-R5F Real-Time Processing Unit						
IP Cores	PCIe® Gen5x8 w/DMA & CCIX	-	2	-	2	2	2	2
	PCIe Gen5x4 w/CXL[1]	2	2	2	2	8	2	2
	100G Multirate Ethernet MAC	6	2	8	4	4	6	8
	600G Ethernet MAC	4	1	6	3	1	5	7
	600G Interlaken	2	0	3	1	0	2	3
	400G High-Speed Crypto Engines	3	1	4	2	2	3	4

1: CXL implemented via a combination of hard and soft IP

© Copyright 2020 Xilinx

XILINX.

Customers Can Get Started Now

Versal Premium

- Documentation Available Now
- Tools Available 2nd Half 2020
- Silicon Shipping 1st Half 2021

Start Prototyping Now
With Versal Prime Eval Kits

Pin Migration to Versal Premium

Evaluate Key Architectural Blocks

Key Interfaces for System Testing

System-Design Methodology Guides

© Copyright 2020 Xilinx

XILINX.

Breakthrough Integration of Networked, Power-Optimized Cores on an Adaptable Platform

3X Bandwidth for Fastest, Most Secure Networks

- 9Tb/s serial bandwidth with 112G PAM4 transceivers
- 5Tb/s of integrated Ethernet throughput
- 1.6Tb/s of line-rate encryption

2X Compute Density for Adaptable Acceleration

- Highest logic capacity & DSP density for hardware acceleration
- Massive memory capacity & bandwidth eliminates acceleration bottlenecks

Highly Integrated HW/SW Platform for Productivity

- Integrated shell for dedicated cloud connectivity
- Unlocked by Vitis™ Unified SW Platform & Vivado® Design Suite
- Complete solution stack for HW & SW developers

XILINX VERSAL™ PREMIUM

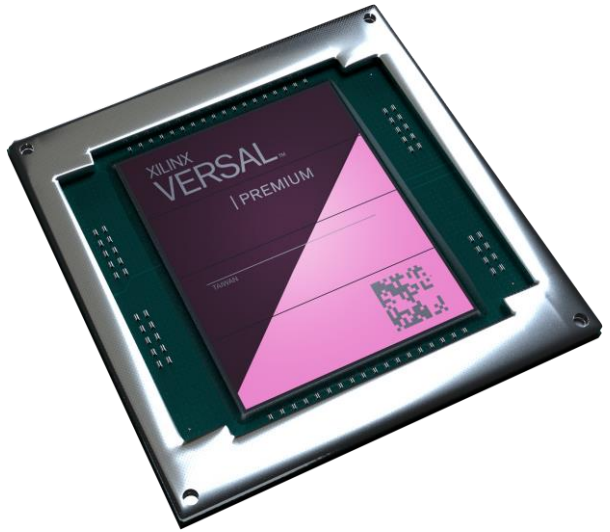

Silicon Sampling in 1st Half 2021

Customers Can Get Started Now

Bandwidth & compute density comparisons based on 14nm/16nm FPGAs

© Copyright 2020 Xilinx

Thank You

© Copyright 2020 Xilinx