



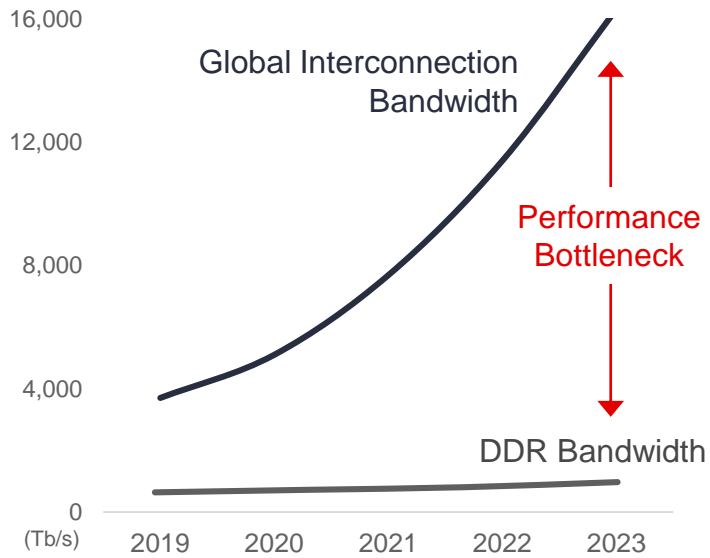
Versal HBM Series Announcement

Mike Thompson, Senior Product Line Manager,
Versal™ Premium and HBM ACAPs, Virtex® UltraScale+™ FPGAs

Bandwidth and Security Requirements Outpacing Current Processing and Memory Technologies



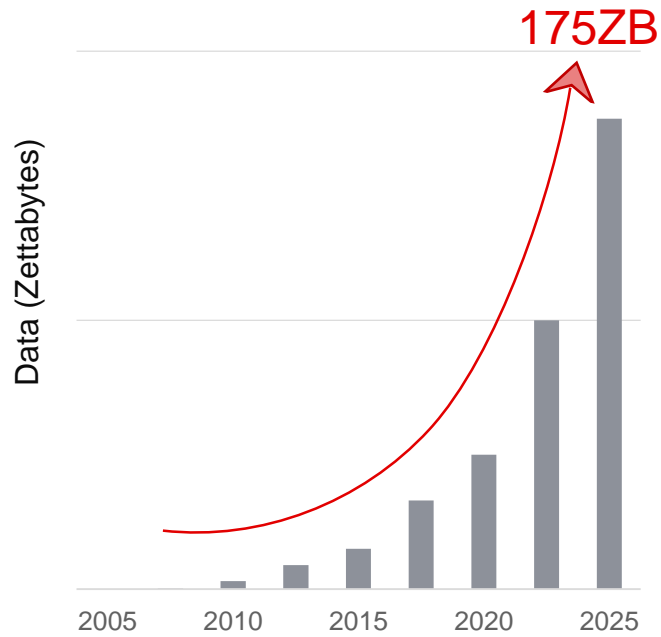
Gap Between Network Traffic and Memory Bandwidth



Source: [“Global interconnection bandwidth”](#) grow at a 45% CAGR—translating to 16,300+ Tb/s



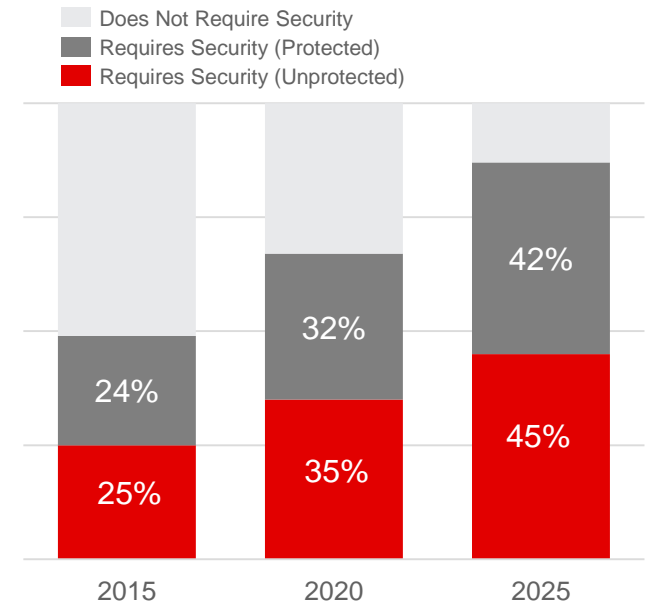
Exponential Growth of Data to be Processed



Source: [Adapted from Data Age 2025 from IDC Global DataSphere, Nov 2018](#)




Data Security Falling Short



Source: [Data Age 2025 study, April 2017, IDC](#)


Traditional Architectures Are Bottlenecked on Memory and Network Access for Real-Time Applications



Versal™ HBM Series: Solving Big Data, Big Bandwidth Problems

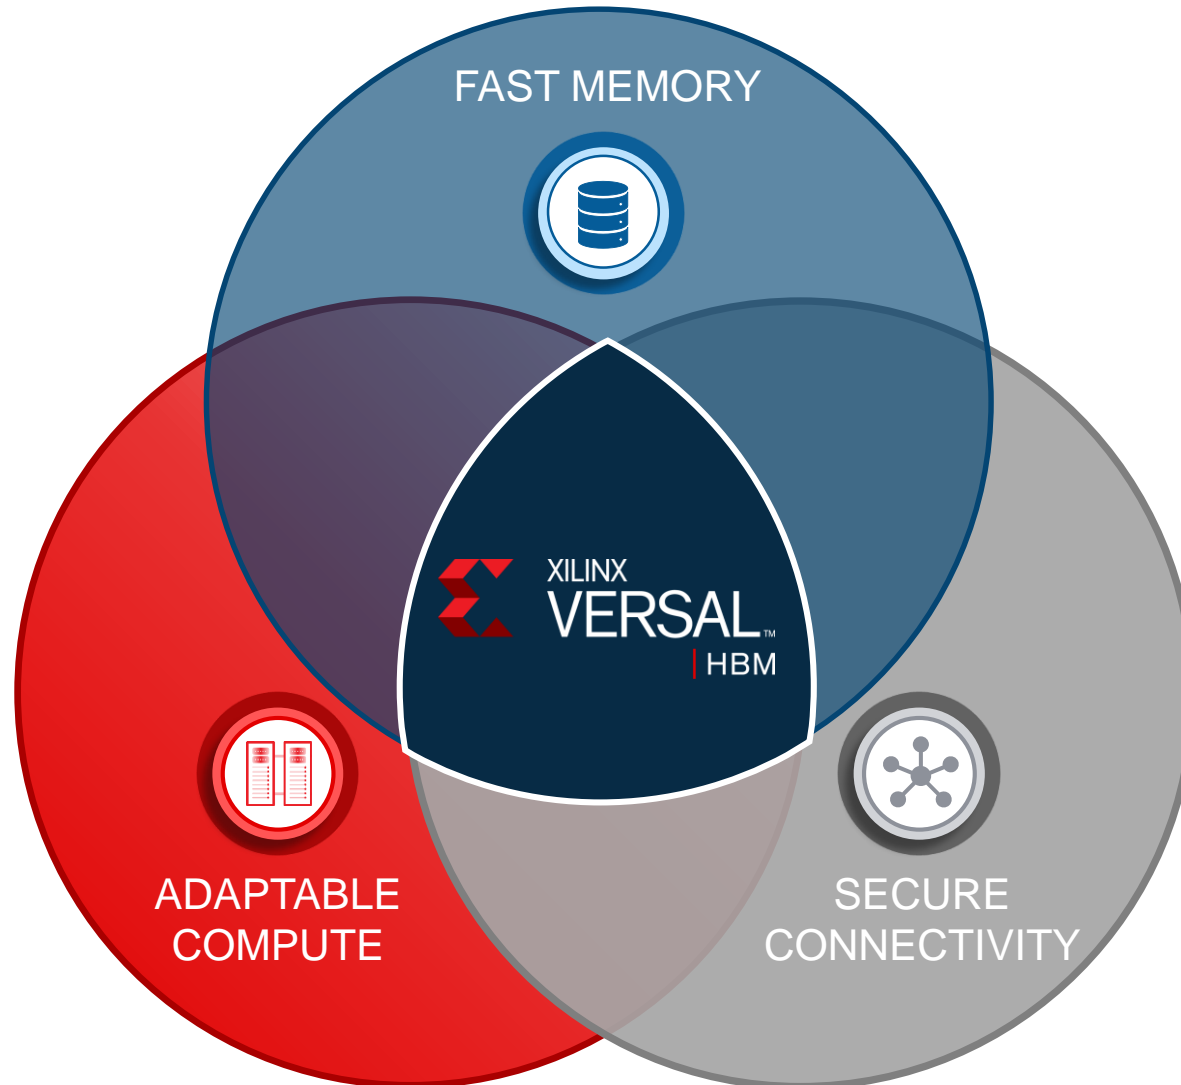


ELIMINATING MEMORY BOTTLENECKS



Maximize Performance and Minimize Power, Area, and Application Latency

Versal™ HBM Series: A Single, Converged Platform



Versal™ HBM Series

Convergence of Fast Memory, Secure Data, and Adaptable Compute



8X memory bandwidth¹ at 63% lower power alleviates network and compute bottlenecks

2X faster secure connectivity² to adapt for emerging networks

2X adaptable compute engines³ for evolving algorithms and protocols


PCIe® Gen5

112G SerDes

820GB/s HBM2e

600G Cores

400G High-Speed Crypto





1: Based on a typical system implementation of four DDR5-6400 components

2: Line rate vs. Virtex® UltraScale+™ FPGA

3: Logic density vs. Virtex UltraScale+ HBM FPGA

Target Markets for Versal™ HBM Series



Network

- ▶ Security Appliance
- ▶ Search and Look-Up
- ▶ 800G Switch / Router

Data Center

- ▶ Machine Learning Acceleration
- ▶ Compute Pre-Processing & Buffering
- ▶ Database Acceleration & Analytics

Test & Measurement

- ▶ Network Testers
- ▶ Packet Capture
- ▶ Data Capture

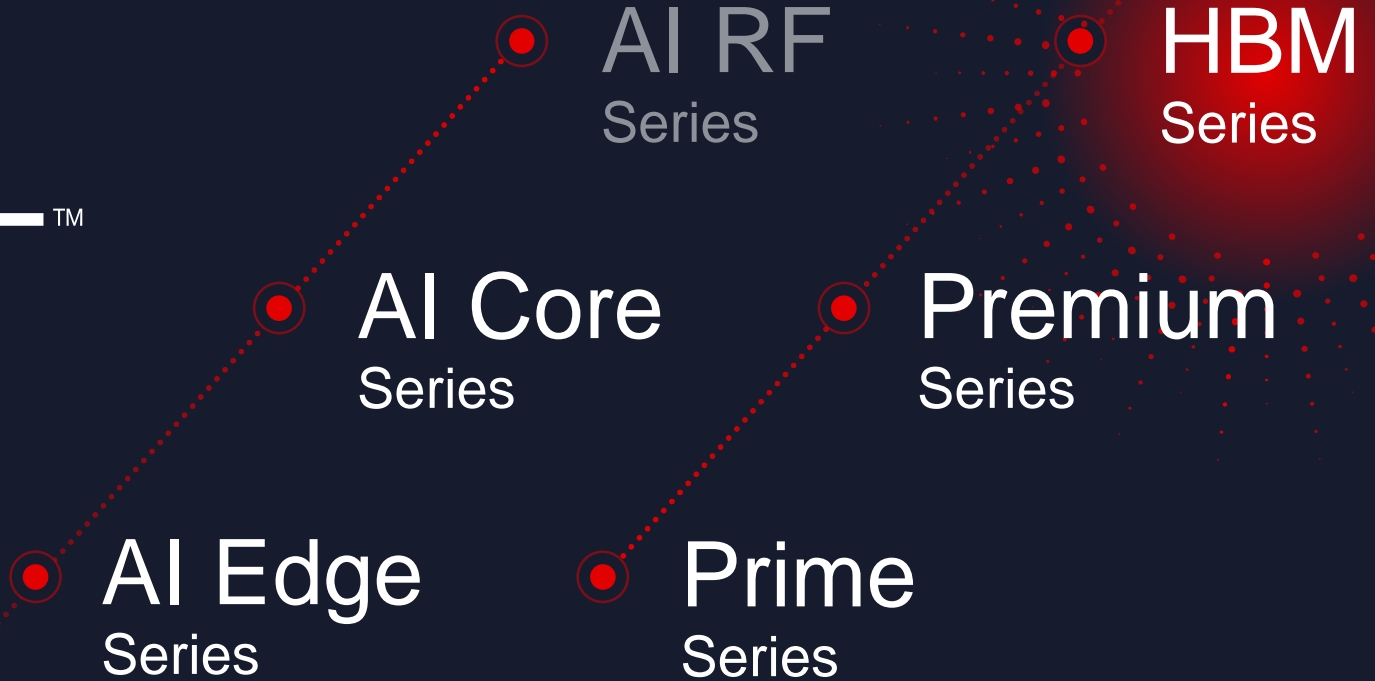
Aerospace & Defense

- ▶ Radar
- ▶ Signal Processing
- ▶ Secure Communication

For Memory Bound, Compute Intensive, High Bandwidth Applications




XILINX®
VERSAL™




Execution through Evolution

Built on a Proven Foundation




4th Gen Stacked Silicon Interconnect (SSI) Technology

- ▶ SSI technology (CoWoS) is the de facto standard for HBM integration
- ▶ Swapped out one super logic region (SLR), swapped in HBM stacks
- ▶ Modified one SLR to add integrated HBM controller



Architected for Fast Data Movement & Adaptive Processing


Scalable Transceivers



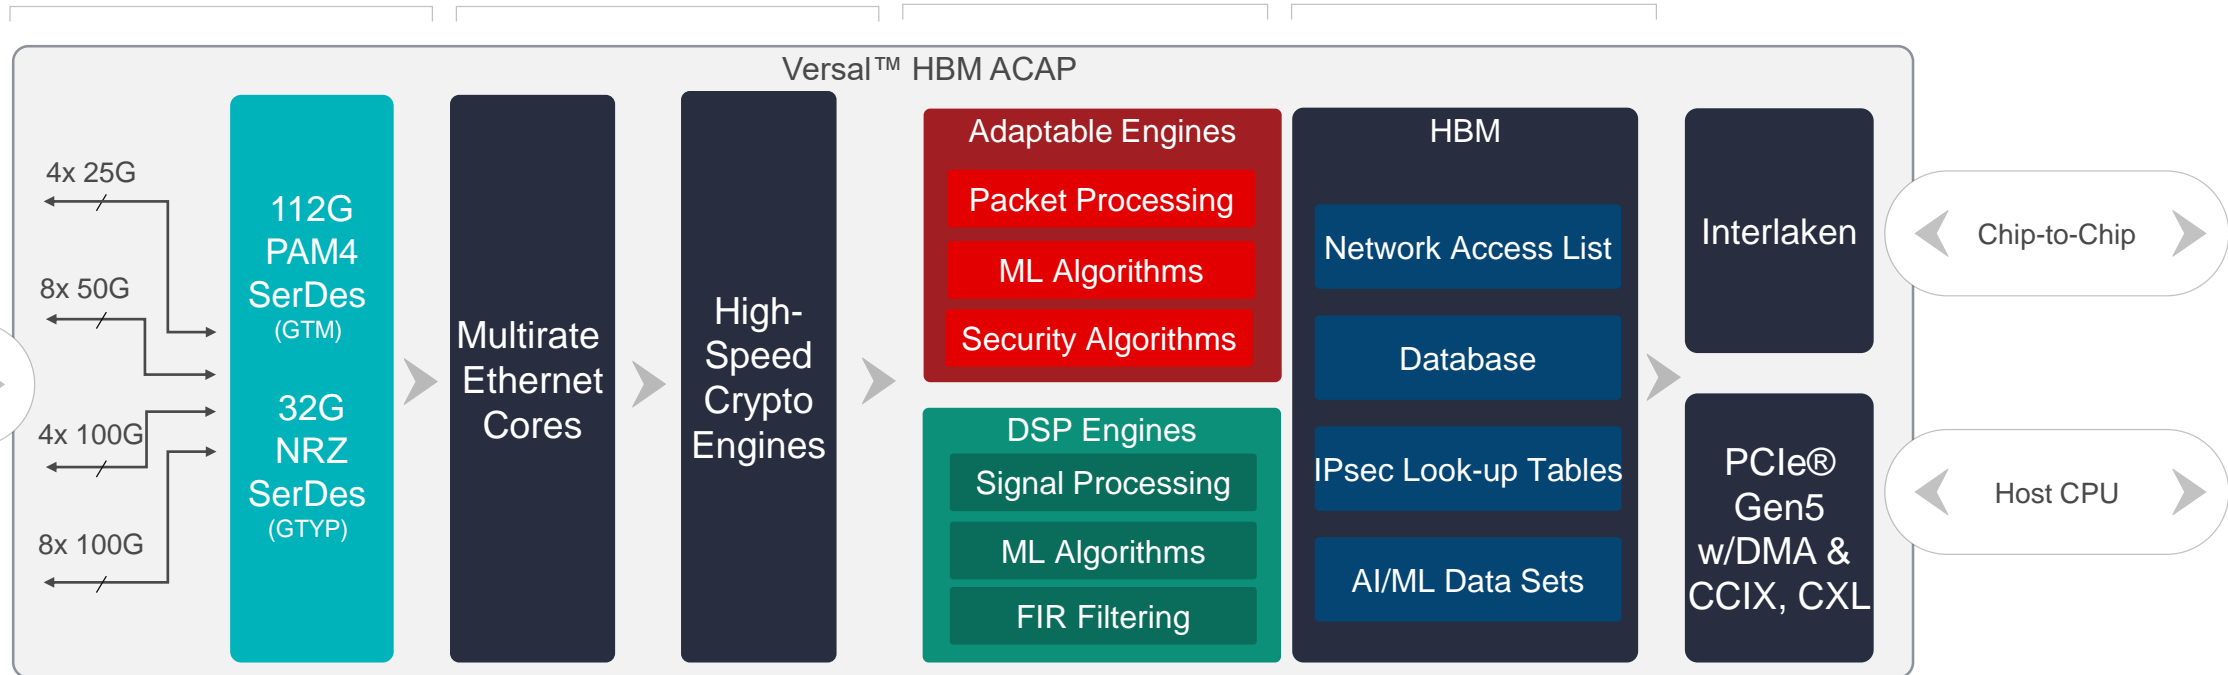
Secure Connectivity



Adaptable Hardware



Massive Memory Bandwidth




Hyper Integration of Networked IP and Memory Subsystem

14 Equivalent FPGAs
of Integrated Cores¹



 XILINX
VERSAL™
| HBM



Replaces 32 DDR5 Chips²
with Integrated HBM




1: Xilinx® Virtex® UltraScale+™ FPGAs vs. Versal™ HBM VH1782 ACAP equivalent logic density of Ethernet, Interlaken, and High-Speed Crypto Engines
2: For equivalent HBM bandwidth vs. DDR5-6400 components

Integrated HBM Eclipses Commodity Memories for Data Intensive Applications

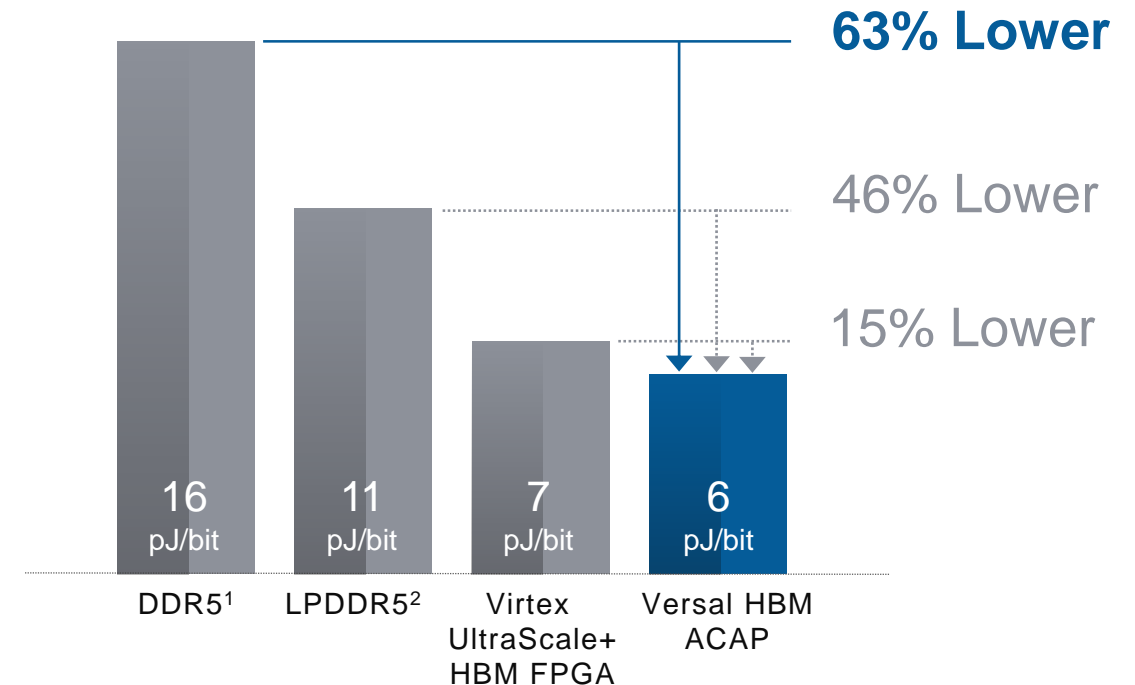
8X More Bandwidth

- ▶ Higher capacity network processing
- ▶ Higher performance AI acceleration



63% Lower Power

- ▶ Eliminates high-power I/O
- ▶ Major OpEx reduction




1: Based on a typical system implementation of four DDR5-6400 components

2: Based on a typical system implementation of four LPDDR5 components

2X HBM Capacity vs. Virtex UltraScale+ HBM FPGA

- ▶ Enables processing on bigger data sets
- ▶ Less swapping of data results in higher performance



2X Faster, Scalable Serial Bandwidth

5.6Tb/s of Total SerDes Bandwidth

Proven in
16nm/7nm Silicon

32Gb/s
NRZ




58Gb/s
PAM4



112Gb/s
PAM4


Mainstream Power-Optimized 100G Interfaces

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




Current 400G Ramp and Deployment

Enabling latest generation optics for maximum system bandwidth



Future 800G Networks on Existing Infrastructure

Industry moving towards 100G per lane optics and 800G infrastructure



Copper Cable



QSFP28

CFP8

QSFP28-DD


Optics

QSFP56-DD
4x100G, 400G

OSFP

QSFP-DD800
100G per lane

Backplane



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



2.4Tb/s of scalable Ethernet bandwidth

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



1.2Tb/s of line rate encryption throughput

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



1.5Tb/s of aggregated PCIe link bandwidth

- ▶ PCIe® Gen5 with DMA, CCIX, and CXL
- ▶ Dedicated connectivity over programmable NoC to memory



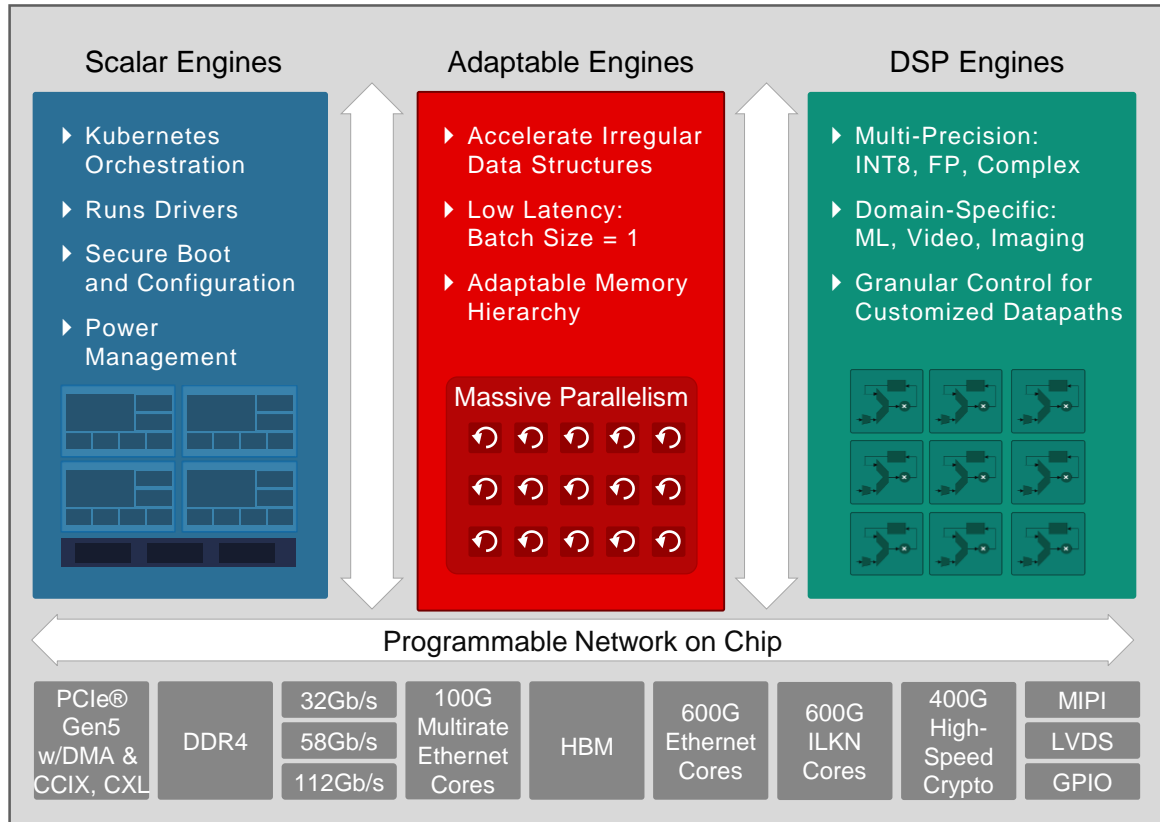
600Gb/s of off-the-shelf Interlaken connectivity

- ▶ Scalable chip-to-chip interconnect from 12.5Gb/s to 600Gb/s
- ▶ Integrated FEC for power-optimized error correction

Adaptable Acceleration for Massive Connected Data Sets

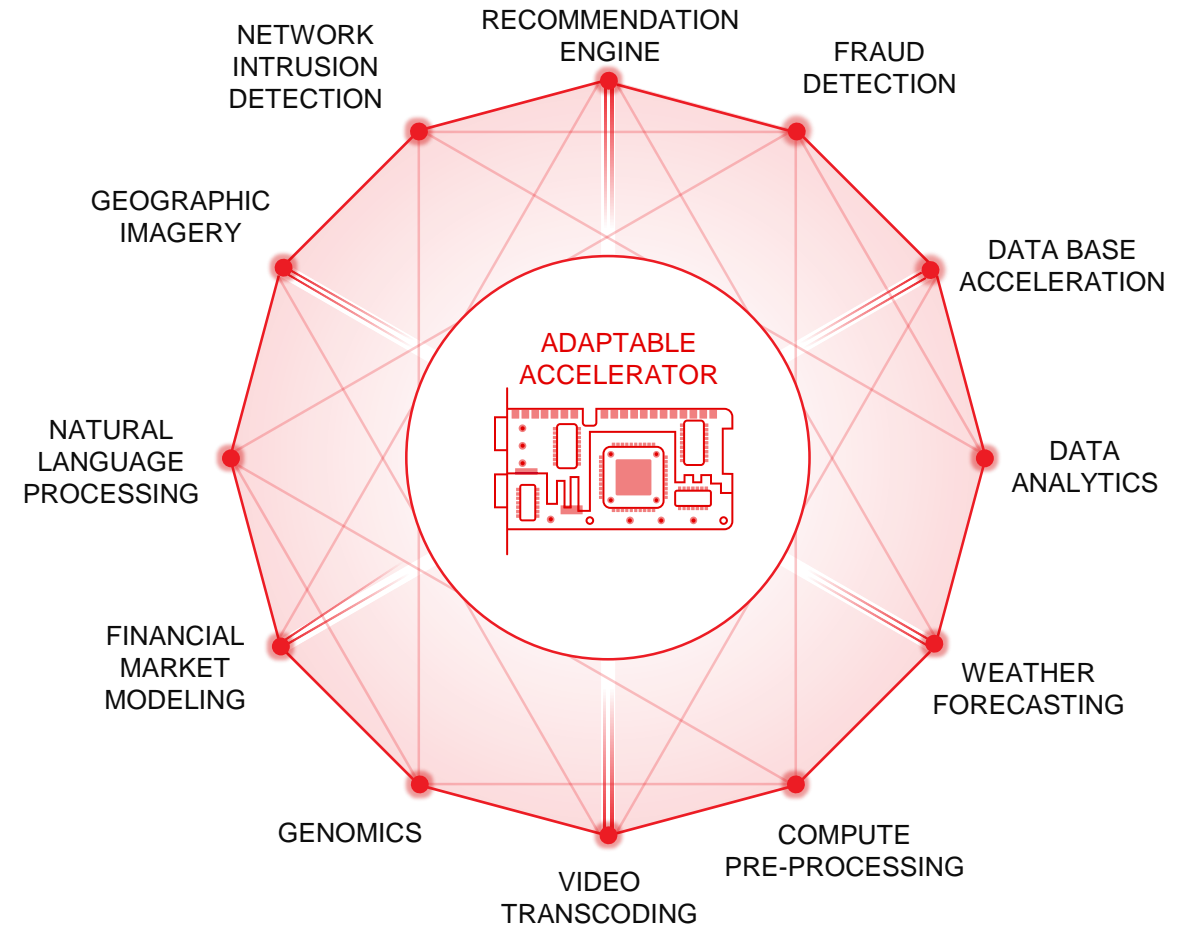
Adaptive, Heterogeneous Compute

Match the Engine to the Algorithm



Acceleration for Large Data Sets

Compute Intensive, Memory-Bound Workloads




Faster Runtimes on Bigger Data Sets

Deploy with Fewer and Lower Cost Servers

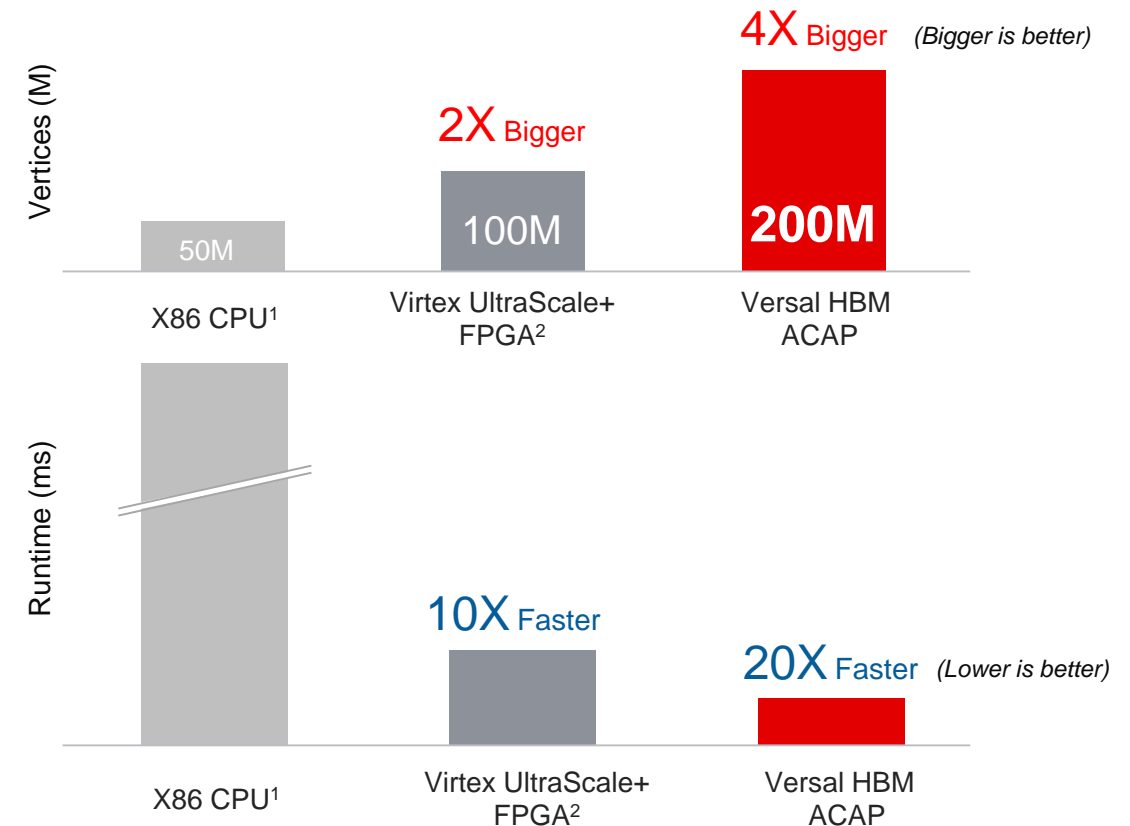
Real-Time Recommendation Engine

- ▶ Cosine similarity algorithm
- ▶ Clinical outcome predictions



Real-Time Fraud Detection

- ▶ Louvain modularity algorithm
- ▶ Detect anomalies in behavior/transactions



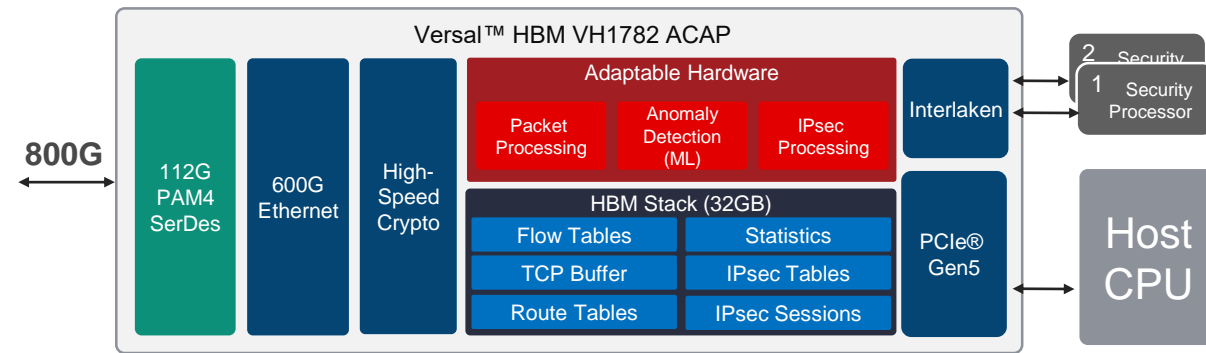
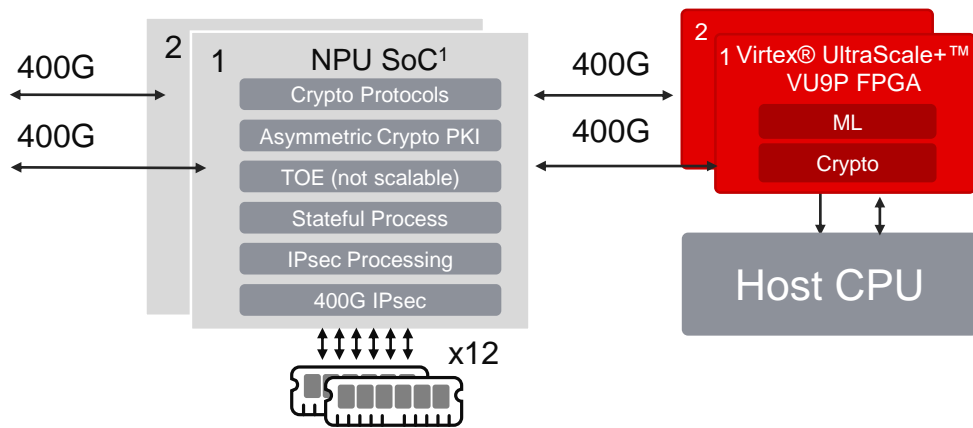
1: 3rd gen Intel Xeon gold/platinum scalable processors

2: Xilinx® Virtex UltraScale+ FPGA based Alveo™ Accelerator card

800G Next-Gen Firewall

High Performance, Low Power, ML-Enabled Network Security

Next-Gen Firewall



	NPU SoC		Versal ACAP
Session Capacity	16M	▶ 2.5X ▶	40M
Memory Throughput	250GB/s	▶ 3.3X ▶	820GB/s
Area²	16 devices = 58,569mm ²	▶ 89% Less ▶	3 devices = 6,700mm ²
Power	305W	▶ 38% Less ▶	190W
SerDes Line Rate	50G Only	▶ 2X ▶	100/50/25/10G (Greater Flexibility)

19 1: Marvell CN106XXS
2: Xilinx estimates



Users Can Get Started Now

Scalable Compute and Memory Capacity

		VH1522	VH1542	VH1582	VH1742	VH1782
Memory	HBM DRAM (GB)	8	16	32	16	32
	Total PL Memory (Mb)	509			752	
Connectivity	GTYP 32G	68			68	
	GTM 56G (112G)	20 (10)			60 (30)	
	100G Multirate Ethernet MAC	4			6	
	600G Ethernet MAC	1			3	
	400G High-Speed Crypto Engines	2			3	
Compute	System Logic Cells	3.8M			5.6M	
	Adaptable Engines (LUTs)	1.8M			2.6M	
	Intelligent Engines (DSP Slices)	7.4K			10.9K	
	Scalar Engines	Dual-Core Arm® Cortex®-A72 Application Processing Unit / Dual-Core Arm Cortex-R5F Real-Time Processing Unit				

Customers Can Get Started Now

Start Now with Versal™ Premium Series

Tools and Devices Available Now
Evaluation Boards in Early Access




Migrate to Versal HBM Series

Documentation Available **Now**

Tools Available 2nd Half of 2021

Devices Sampling 1st Half of 2022*



Evaluate Key
Architectural Blocks



Key Interfaces for
System Testing



System-Design
Methodology Guides

* Schedules subject to change

Versal™ HBM Series: Convergence of Fast Memory, Secure Data, and Adaptable Compute

8X Memory Bandwidth at 63% Lower Power¹


- ▶ HBM2e for 820GB/s of memory bandwidth
- ▶ Eliminates data movement between memory and processing
- ▶ Alleviates network and compute bottlenecks

2X Faster Secure Connectivity²

- ▶ Multi-terabit networked, power-optimized cores
- ▶ 112G PAM4 transceivers
- ▶ Adaptable to emerging network optics and protocols

2X Adaptable Compute Engines³

- ▶ Heterogeneous platform to match the engine to the workload
- ▶ Maximizes performance and adapts with evolving algorithms
- ▶ Massive CapEx/OpEx savings for cloud and network providers



Silicon Sampling in 1st Half 2022

1: Based on a typical system implementation of four DDR5-6400 components

2: Line rate vs. Virtex® UltraScale+™ FPGA

3: Logic density vs. Virtex UltraScale+ HBM FPGA



Thank You

