Disruptive Innovation

Kirk Saban

Vice President
Product & Technical Marketing

Mountains of Unstructured Data

One Architecture Can't Do It Alone

This is the Era of Heterogeneous Compute

Today's Developer Needs

Software programmability

Performance for a diverse range of applications

Adaptability to keep pace with rapid innovation

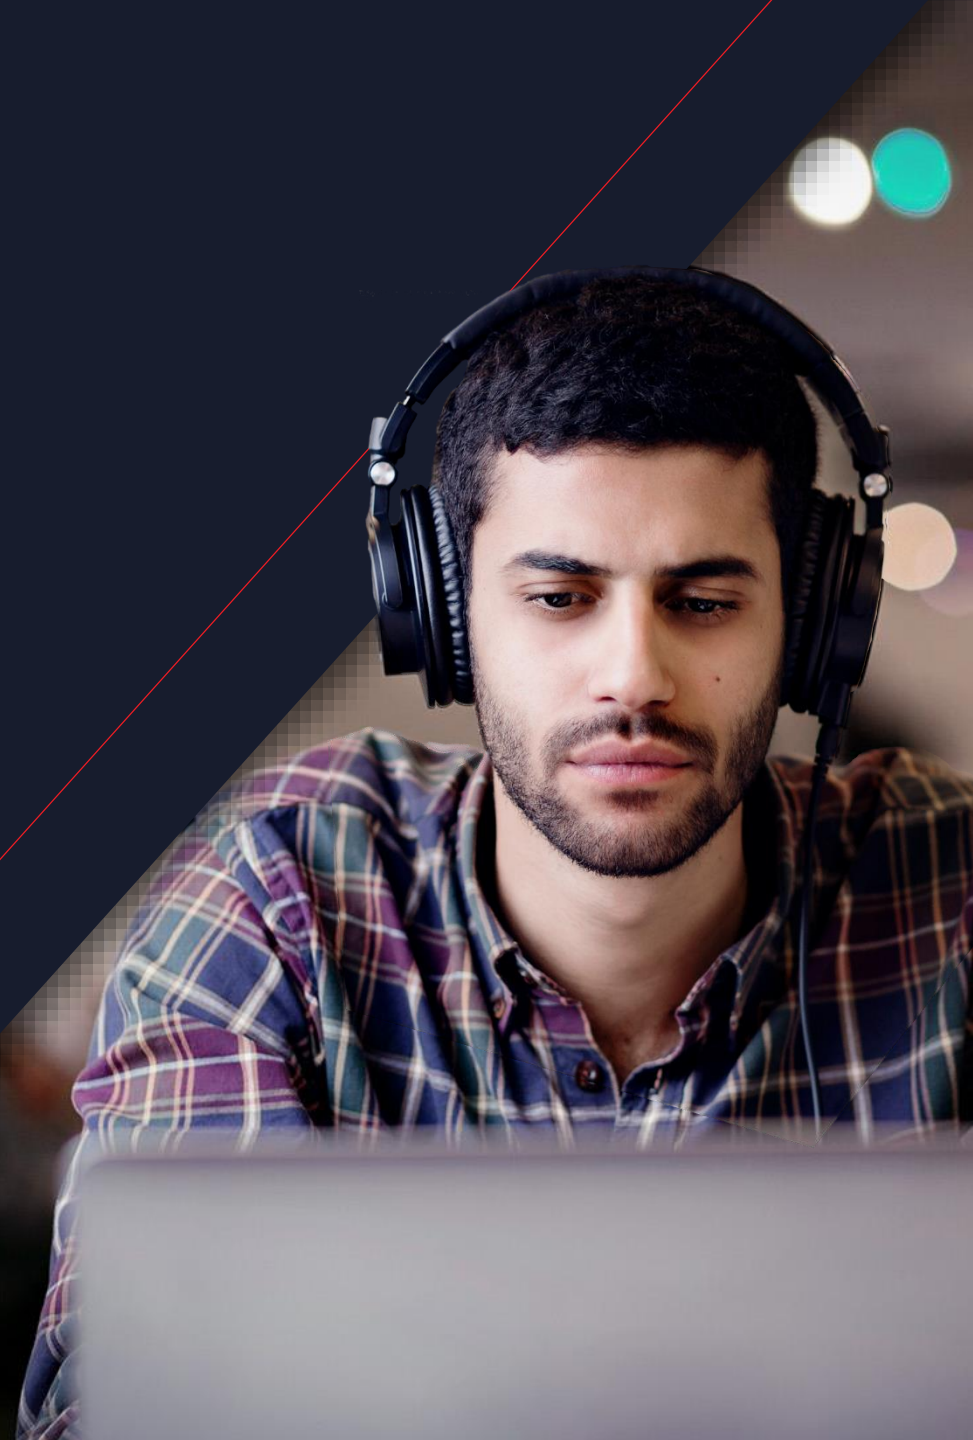

Today’s Solutions

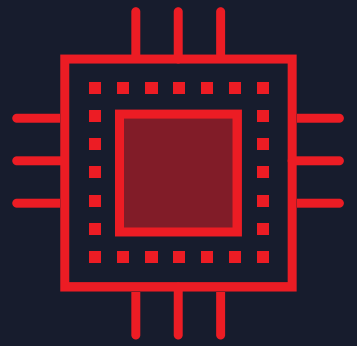

CPUs

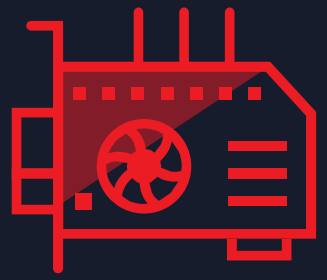

Fixed Function Accelerators

ASICs/ASSPs/GPUs

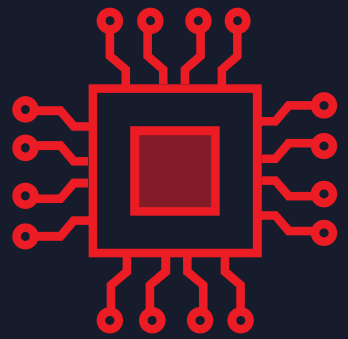

FPGAs

Disruptive Innovation Needed: Enter ACAP

A new class of devices for today's challenges

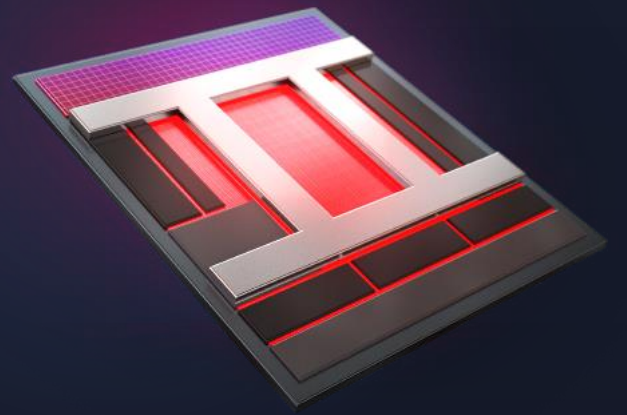

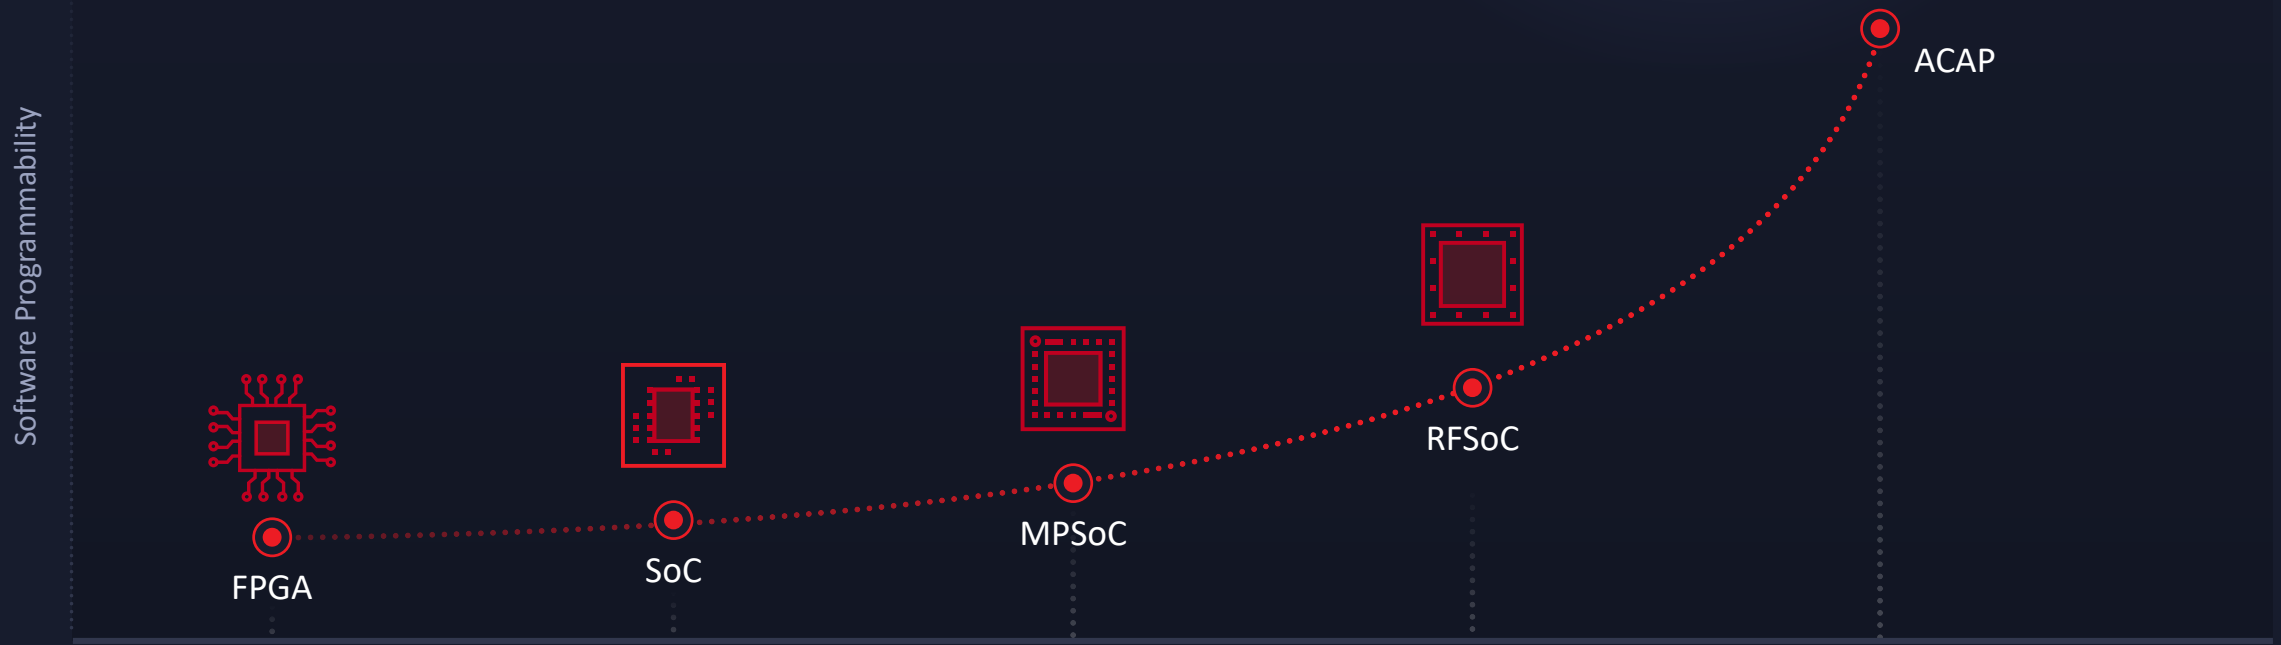

XILINX®

VERSAL™

The Industry's First ACAP

Heterogeneous Acceleration

For Any Application

For Any Developer

tsmc

7nm
FinFET

Versal ACAP Technology Tour

- Scalar Processing Engines
- Adaptable Hardware Engines
- Intelligent Engines
  SW Programmable, HW Adaptable
- Breakout Integration of Advanced Protocol Engines

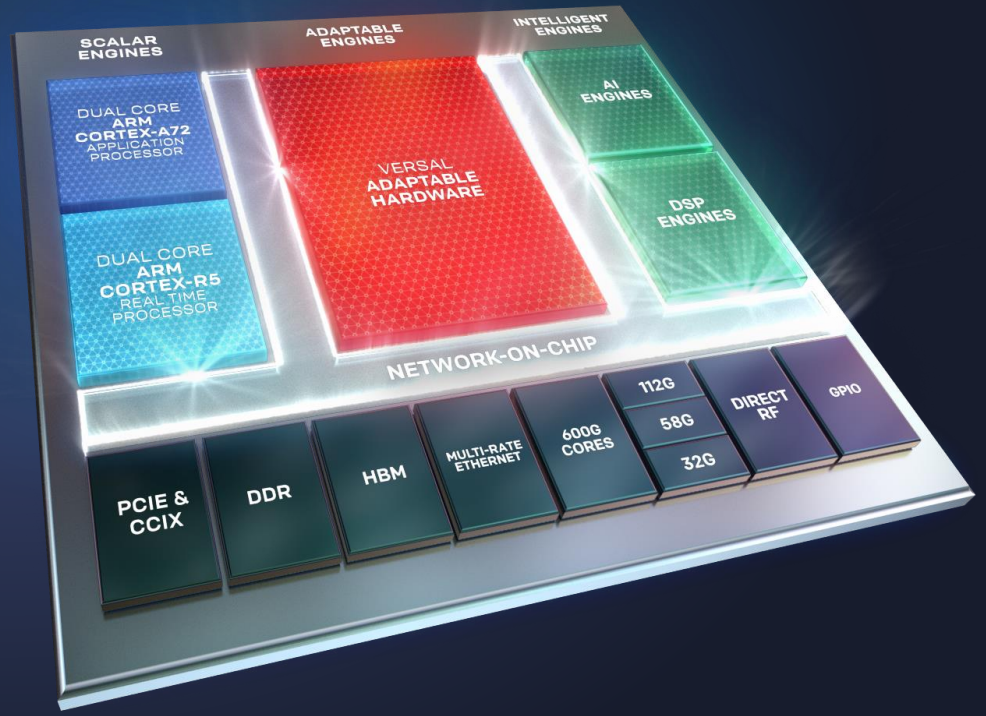

Scalar Processing Engines

Arm Cortex-A72
Application Processor

Arm Cortex-R5
Real-Time Processor

Platform Management Controller

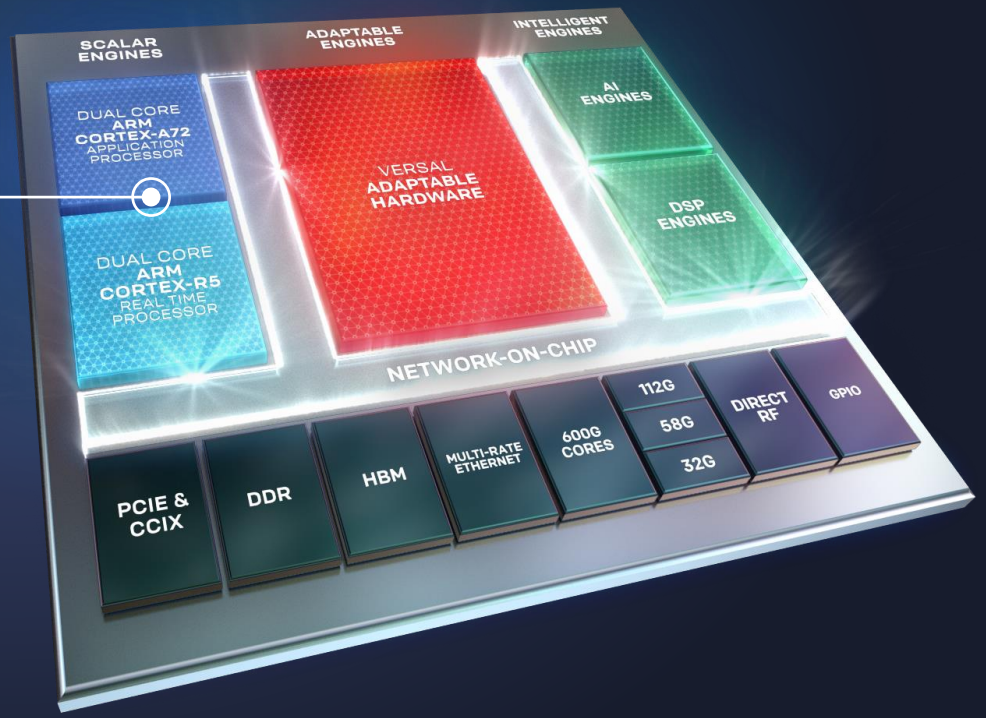

Adaptable Hardware Engines

Re-architected foundational HW fabric for greater compute density

Enables custom memory hierarchy

8X Faster Dynamic Reconfiguration (“on-the-fly”)

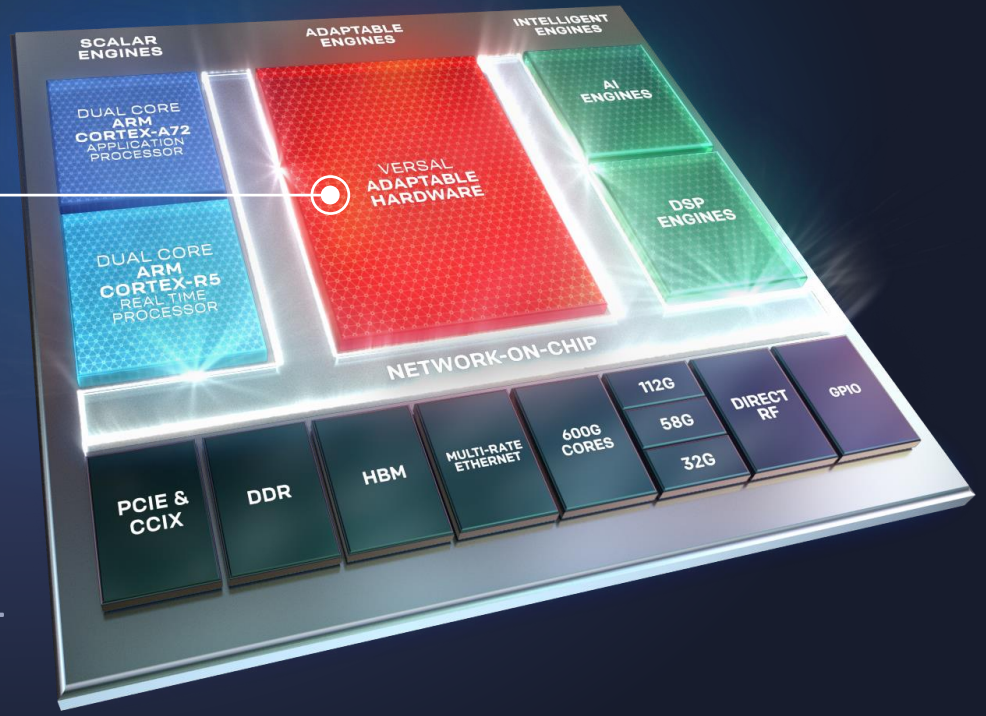

Intelligent Engines

DSP Engines

High-precision floating point & low latency

Granular control for customized datapaths

AI Engines

High throughput, low latency, and power efficient

Ideal for AI inference and advanced signal processing

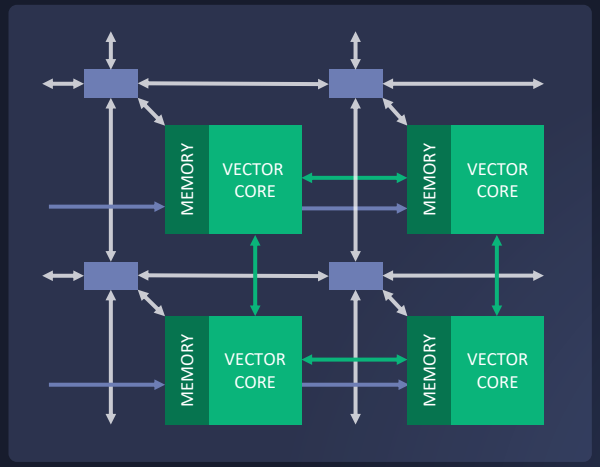

AI Engines

Optimized for AI Inference and Advanced Signal Processing Workloads

- Vector processor array w/ tightly coupled memory
- Direct access to adaptable hardware enables custom memory hierarchy
- Software programmable

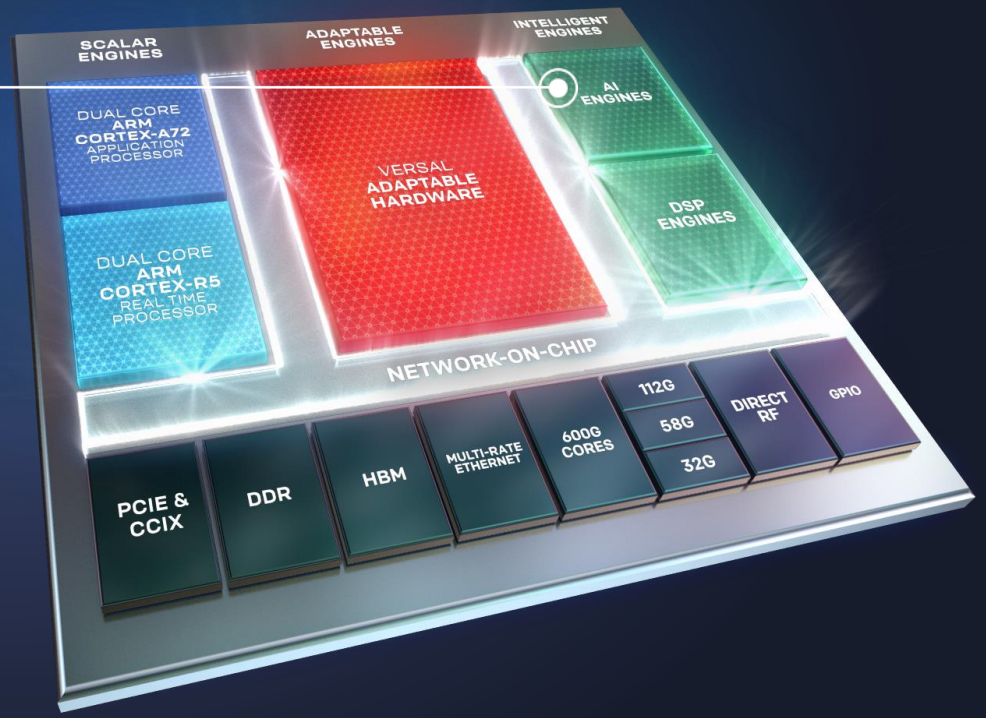

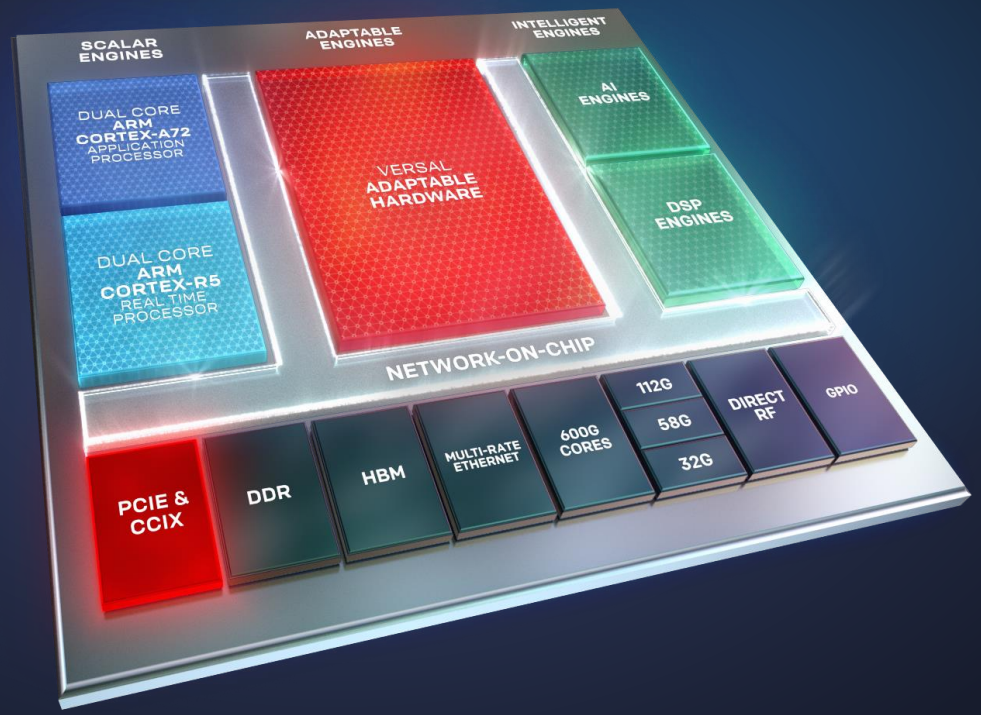

Integrated Host Interfaces

- PCIe Gen4x16
- Integrated AXI-DMA
- CCIX for seamless acceleration to server-class CPUs

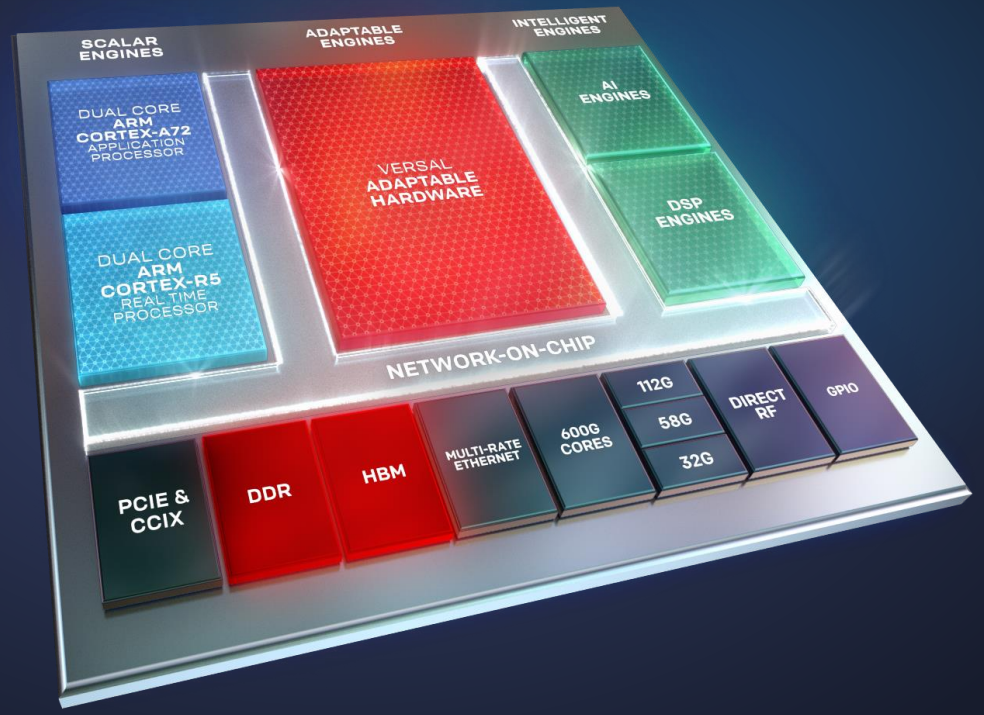

Scalable, Integrated Memory Controllers

- DDR4-3200
- LPDDR4-4266
- High Bandwidth Memory (HBM)

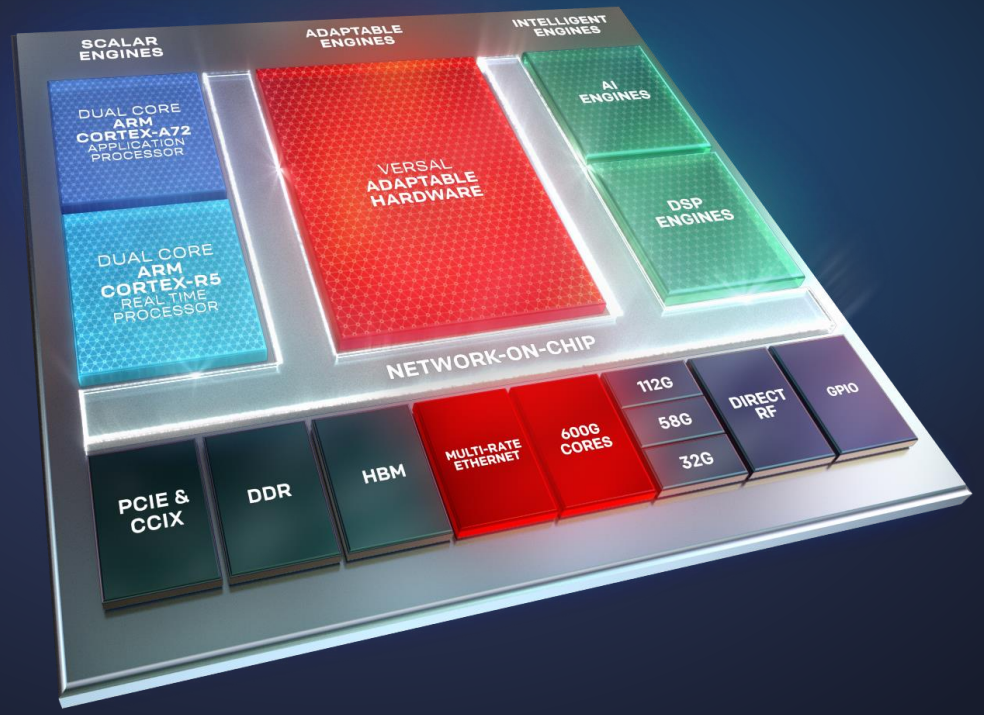

Integrated Protocol Engines

- 100G Multirate Ethernet
- 600G Ethernet and Interlaken
- 600G Cryptographic Engines (AES/IPSEC/MACSEC)

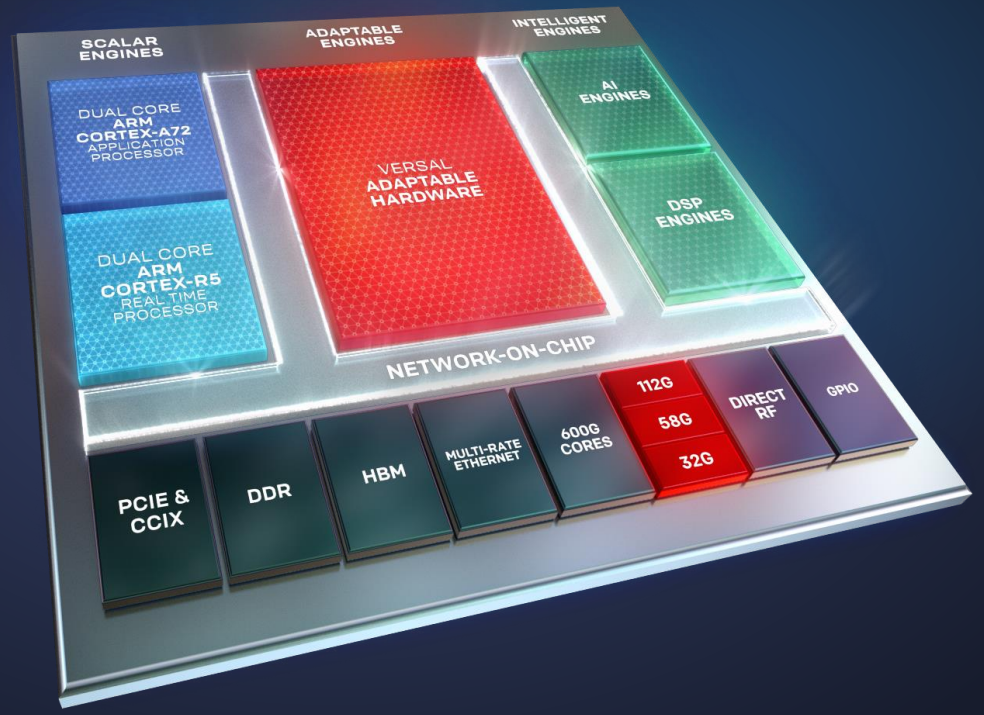

Broadest Range of Transceivers

- 32G power optimized for edge applications
- 58G PAM4—Now in mainstream devices
- 112G PAM4—Industry's highest performance

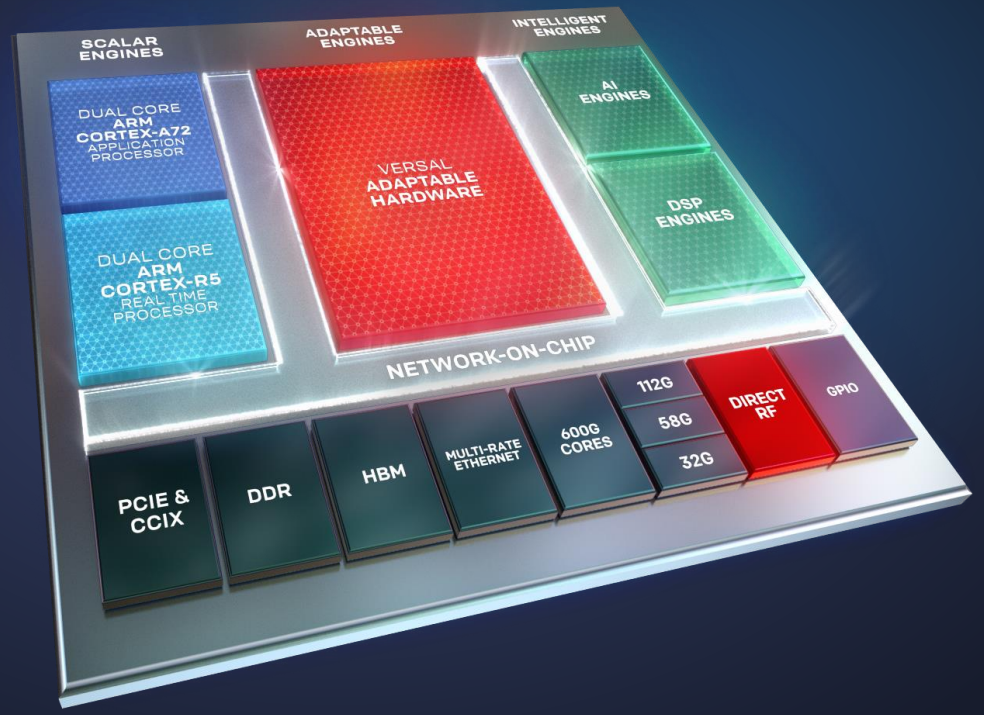

Integrated RF Signal Chain

- Next-generation multi-GSPS direct RF-ADC/DAC
- Integrated DDC/DUC
- SD-FEC for 5G and DOCSIS

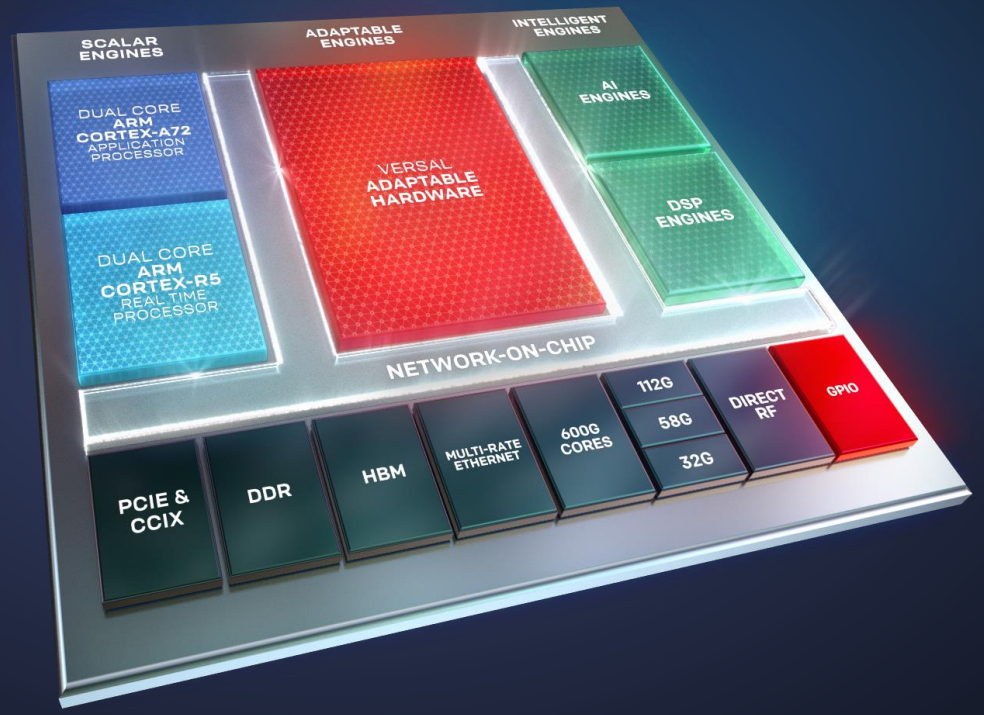

Programmable I/O Interfaces

- MIPI D-PHY >3Gb/s for sensors
- NAND and storage-class memory
- LVDS and general-purpose I/O

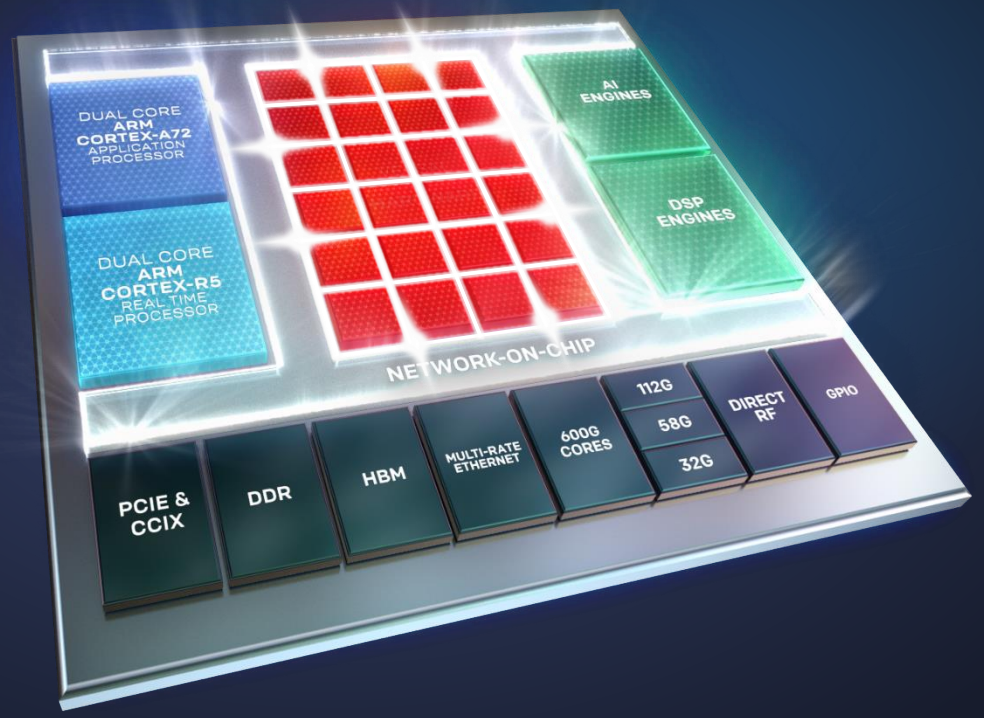

Network-on-Chip (NoC)

Ease of Use

Inherently software programmable
Available at boot, no place-and-route required

High Bandwidth and Low Latency

Multi-terabit/sec throughput
Guaranteed QoS

Power Efficiency

8X power efficiency vs. soft implementations
Arbitration across heterogeneous engines

NoC Enables Software Programmability

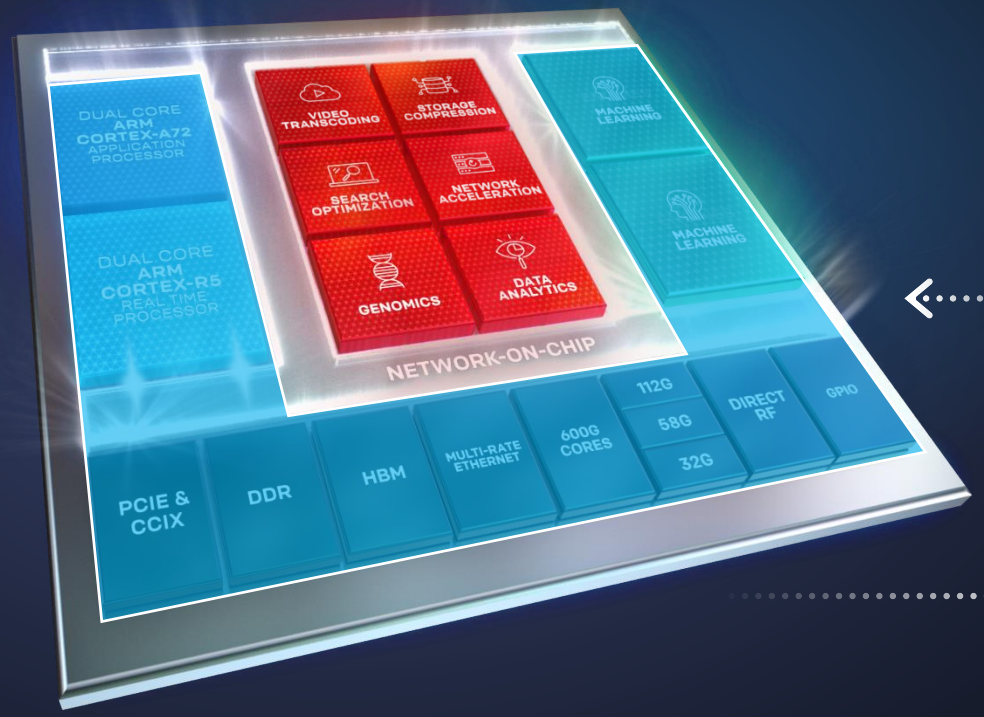

Data Transfer between Engines and Memory

Seamless Integration

Versal for Multi-Market Applications

CLOUD

NETWORK

EDGE

Data Center

Wired

Wireless

Endpoints

AI ADOPTION ACROSS MARKETS

XILINX®

VERSAL™

AI RF Series

HBM Series

AI Core Series

Premium Series

AI Edge Series

Prime Series

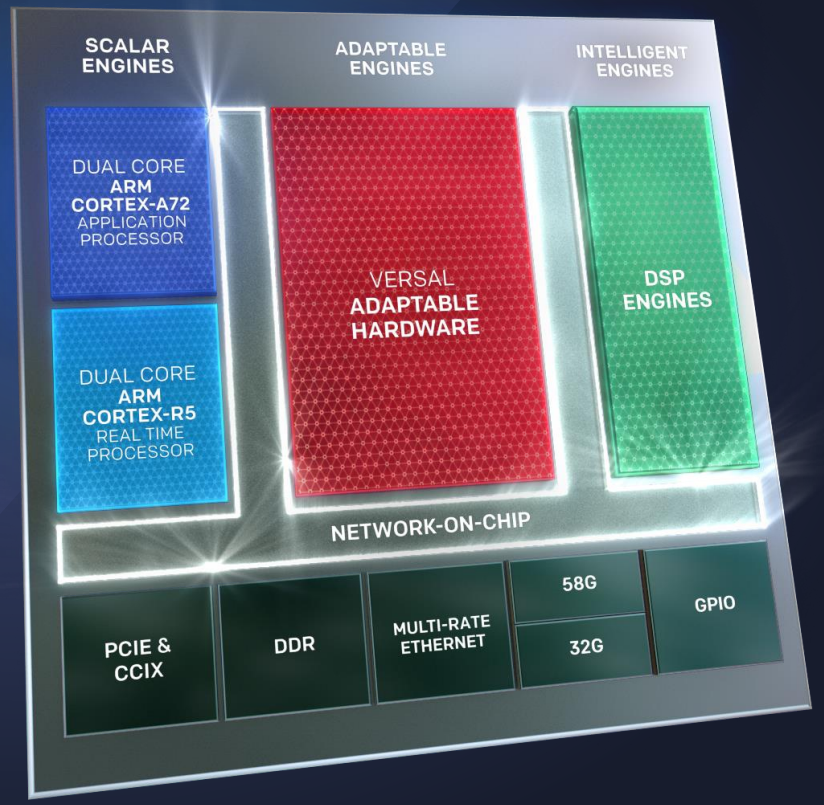

VERSAL

Prime Series

Broad Applicability Across Multiple Markets

Mid-range series in the Versal portfolio

Optimized for connectivity

For inline acceleration and diverse workloads

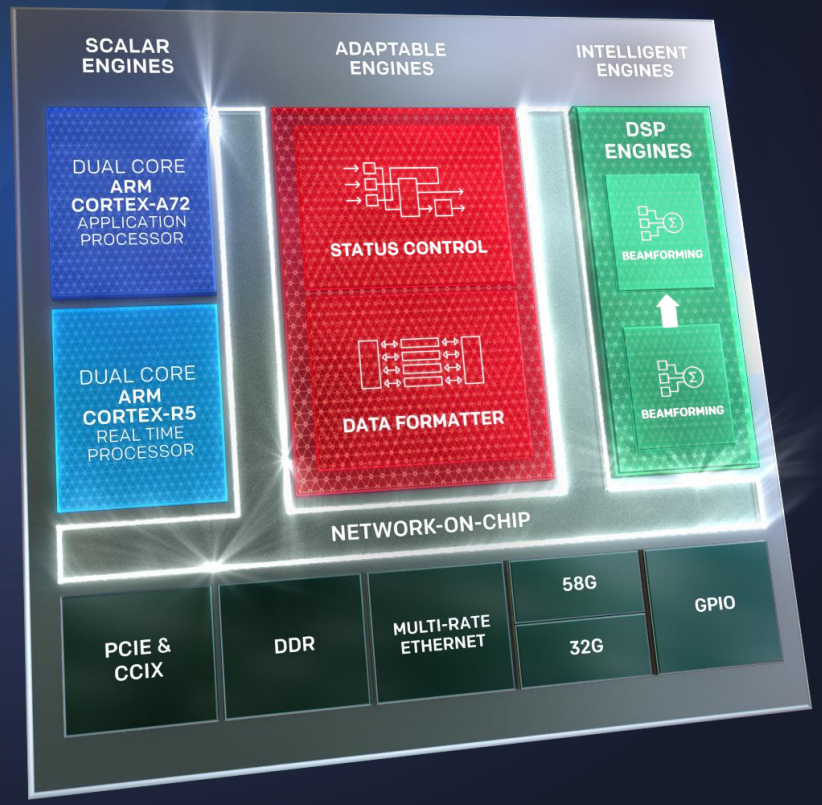

Versal Prime Series

Intelligent Engines in Radar Beamforming

DSP Engines for diverse, fixed & floating point signal processing workloads

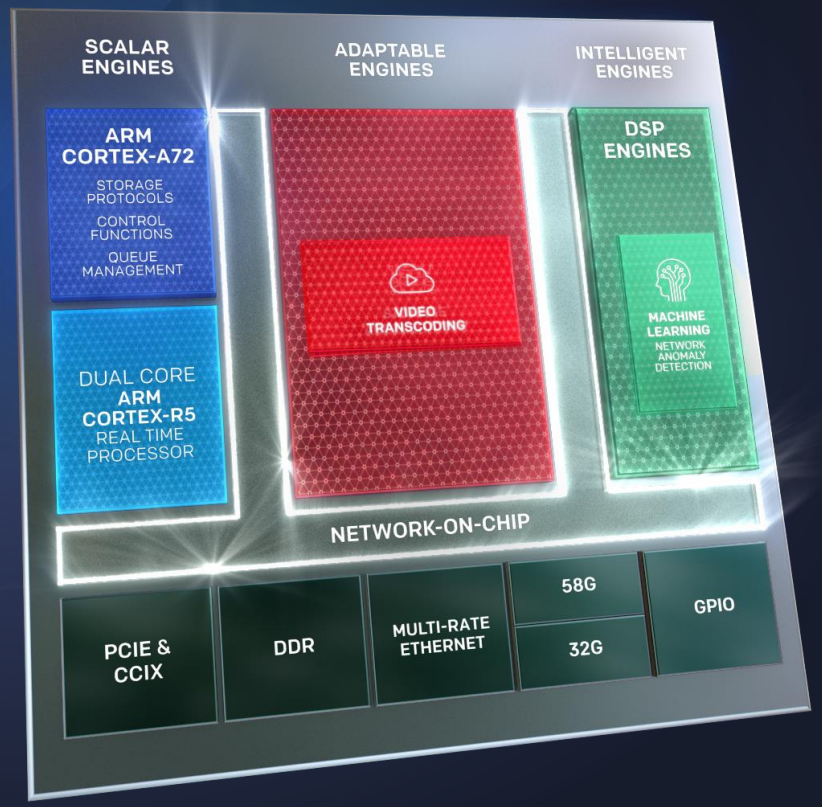

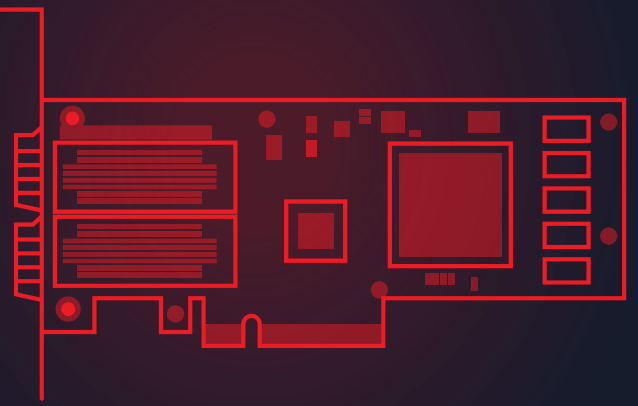

Network Attached Acceleration

Support for multiple network-attached workloads

Ability to combine workloads with AI inference

Network Attached Accelerator Workloads

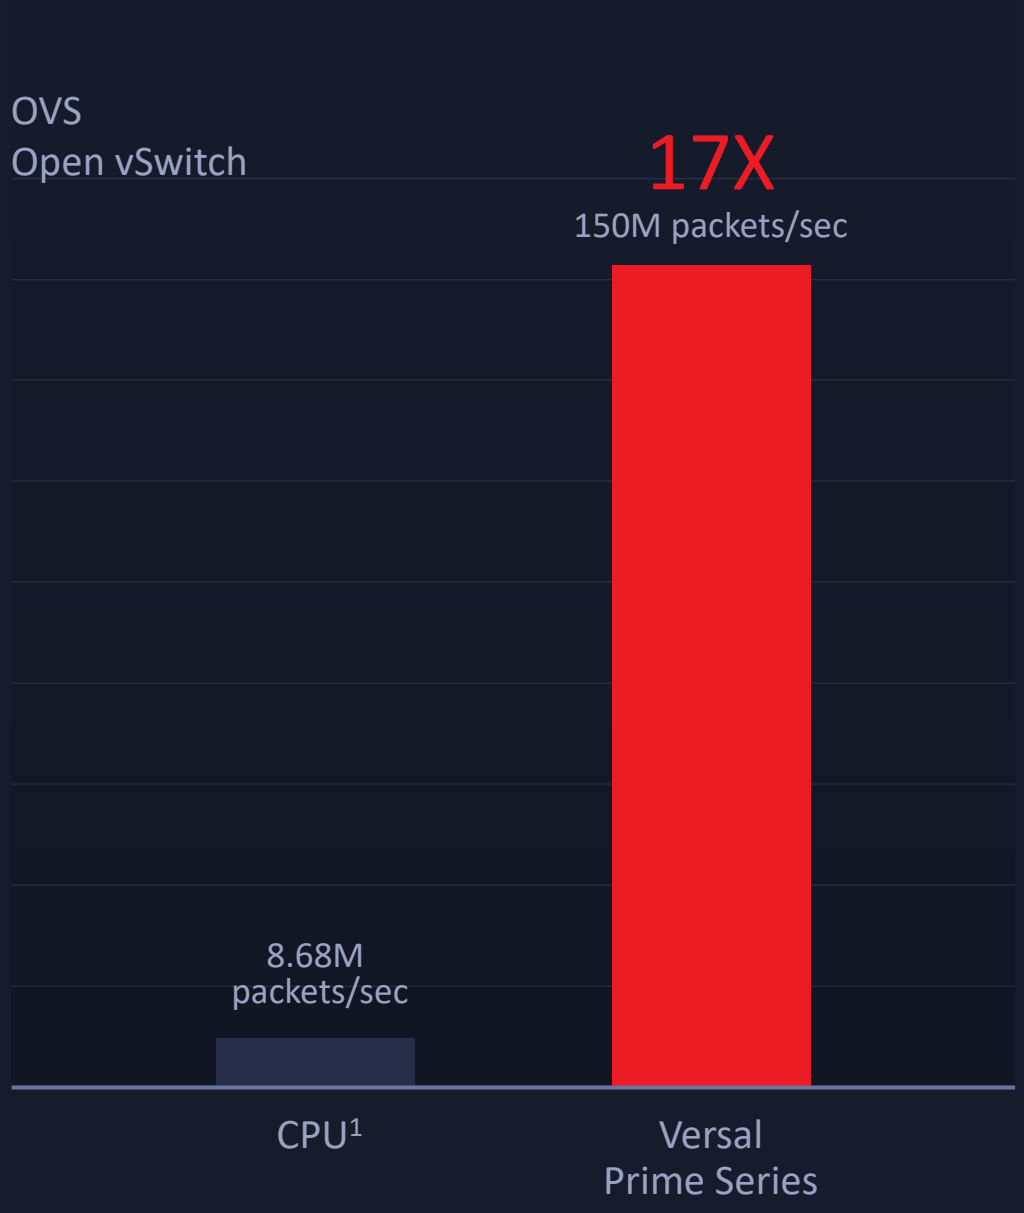

1: Assuming 4 Xeon cores at 2.17/Mp/s per core for zero packet loss; Source: "Red Hat's Perspective on OVS HW Offload Status ", Open vSwitch Fall Conference 2017

More Applications with Versal Prime Series

- Communications Test Equipment
- Data Center Network and Storage Acceleration
- Nx100G Ethernet and OTN Networking
- Broadcast Switches
- Medical Imaging
- Avionics Control

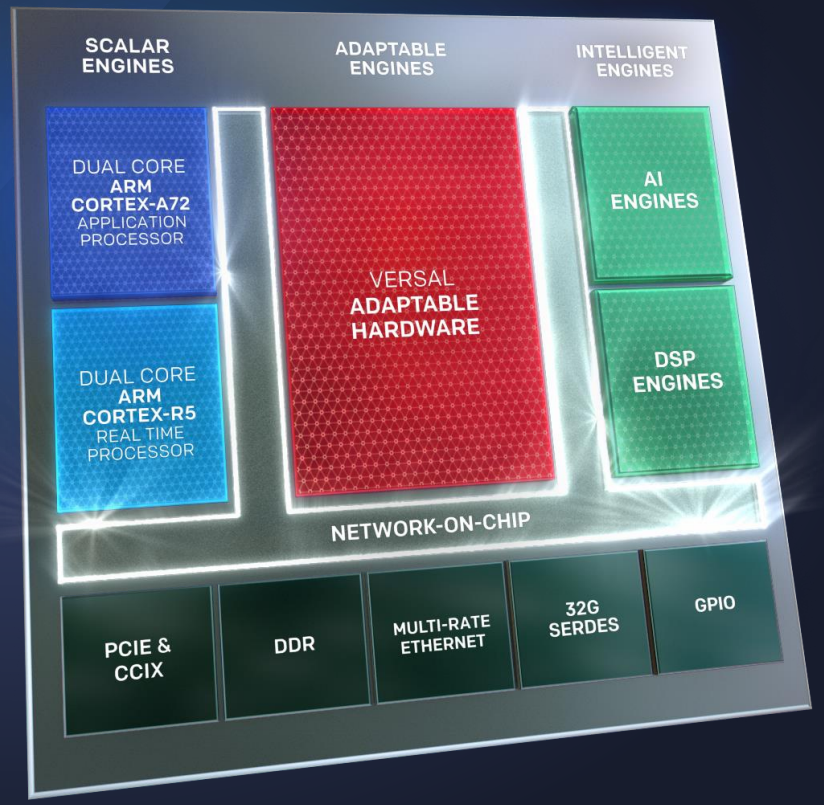

VERSAL

AI Core Series

Breakthrough AI Inference Throughput

Portfolio's highest throughput for low latency inference

Optimized for cloud, networking, and autonomous applications

For highest dynamic range of AI and workload acceleration

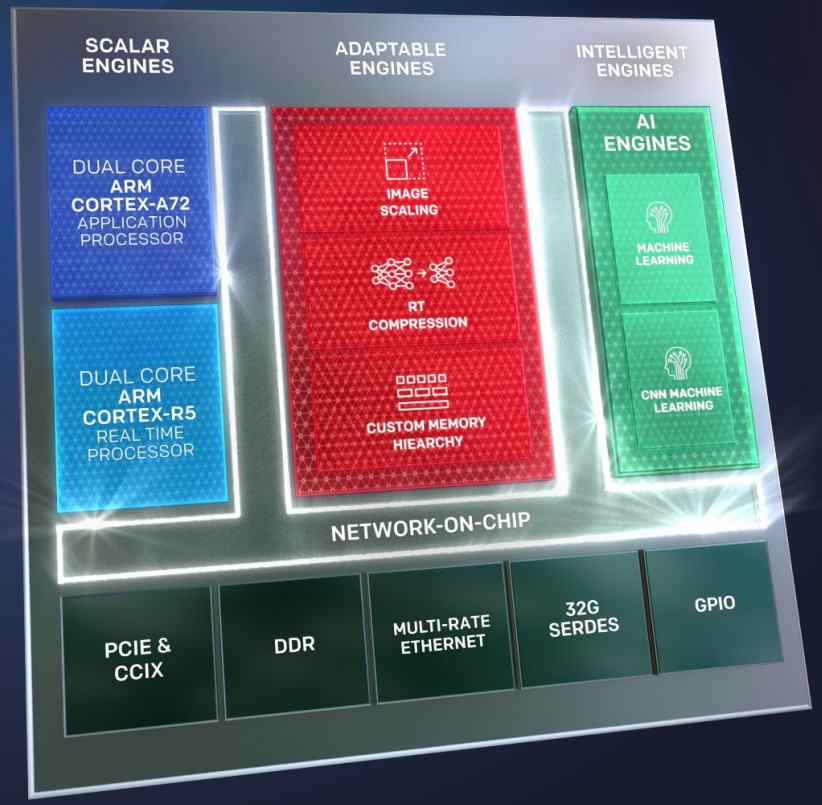

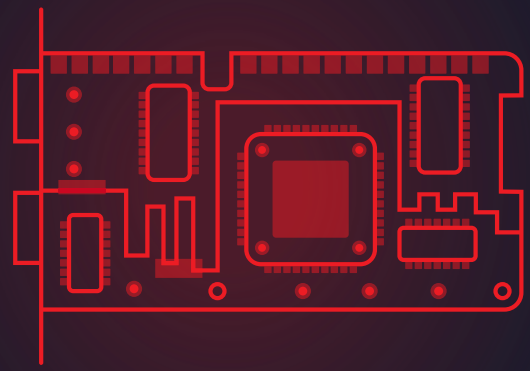

VERSAL AI Core Series

AI Engines and Adaptable Hardware Maximize AI Inference

Massive bandwidth across heterogeneous engines for optimal performance

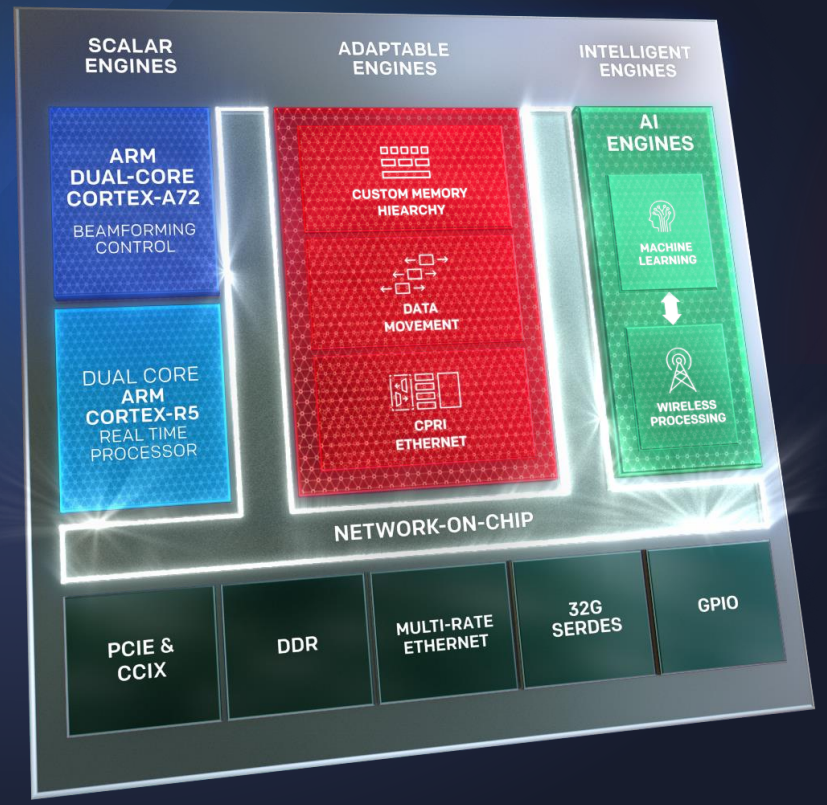

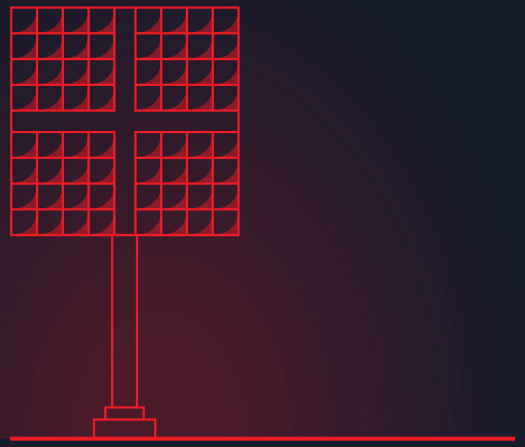

VERSAL AI Core Series

For 5G Wireless Compute with AI Inference

AI Engines have ability to combine inference with wireless compute

Mixed Workloads on AI Engine

UNDER DEVELOPER CONTROL

Combine Workloads on AI Engine

5X

MAX Wireless Compute[1]

- RRH DPD
- DDC/DUC
- Baseband Beamforming

8X

MAX AI Inference[1]

- Self-Organizing Networks
- Anomaly Detection
- Scheduling

1: Compared to Zynq UltraScale+ RFSoCs

IN SUMMARY

Versal ACAP

Heterogeneous Acceleration

For Any Application

For Any Developer

Delivers

Disruptive Innovation

Software Programmability

Hardware Adaptability

Whole Application Acceleration

XILINX®
VERSAL™

Getting Started

- Visit www.xilinx.com/versal
  - Watch ACAP Intro Video
  - Subscribe to the mailing list for the latest news
- View documentation and resources
  - Data Sheet Overview
  - Product Tables
  - Versal Architecture and AI Engine White Papers
- Contact your local Xilinx sales representative

Xilinx.com/Versal