




Xilinx Versal™ Premium

HotChips32, Aug 18, 2020

Martin Voogel, Yohan Frans, Matt Ouellette


Jason Coppens, Sagheer Ahmad, Jaideep Dastidar, Ehab Mohsen,
Faisal Dada, Mike Thompson, Ralph Wittig, Trevor Bauer, Gaurav Singh

Agenda



- ▶ Versal Overview
 - What is Versal
 - Versal Premium overview
- ▶ Key blocks & Features
 - SerDes, Ethernet, High Speed Crypto, Interlaken
 - NOC, Memory Interfaces, PCIe-Gen5, CCIX, CXL
- ▶ Device Architecture
 - SLR, Programmable Logic
 - SSIT, Chiplets
- ▶ Use Cases

Versal Overview



Adaptive Compute Acceleration Platform



ADAPTIVE

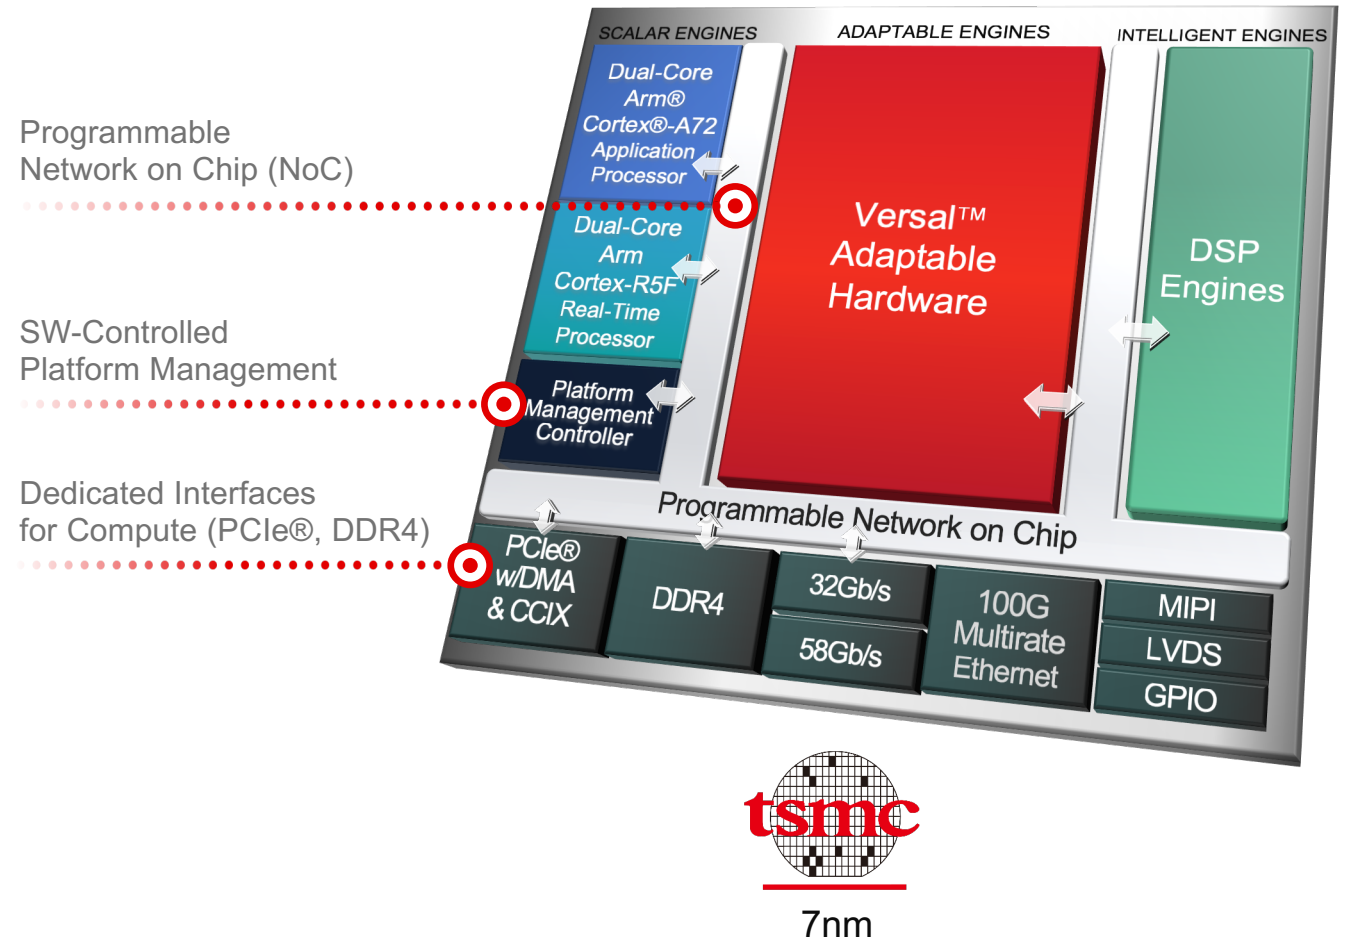
- ▶ Adaptable to diverse workloads
- ▶ Future-proof algorithms

COMPUTE ACCELERATION

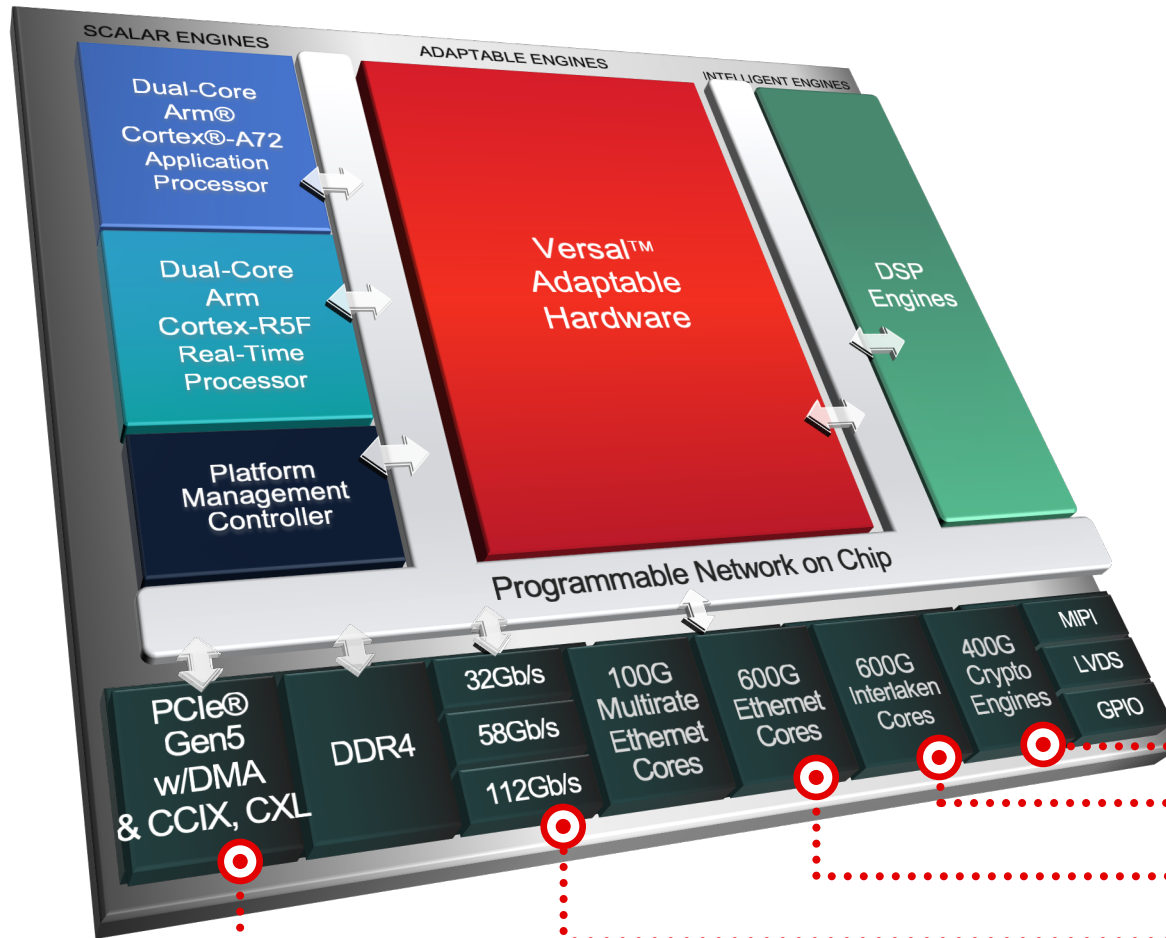
- ▶ Scalar Engines
- ▶ Adaptable Engines
- ▶ Intelligent Engines

PLATFORM

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



Versal Premium: Integration of Protocol Engines



XILINX
VERSAL™
PREMIUM

- 400G High-Speed Crypto Engines
- 600G Interlaken Cores
- 600G Ethernet Cores
- 112G PAM4 Transceivers
- PCIe® Gen5 w/DMA & CCIX, CXL

Integrated Shell with Dedicated Connectivity

'Shell': Pre-Built Infrastructure for Cloud Connectivity


- ▶ Hardens all connectivity to data center infrastructure
- ▶ Reduces design & compile time and timing closure cycles
- ▶ 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



Integrated Programmable Network-on-Chip

High bandwidth terabit Programmable NoC

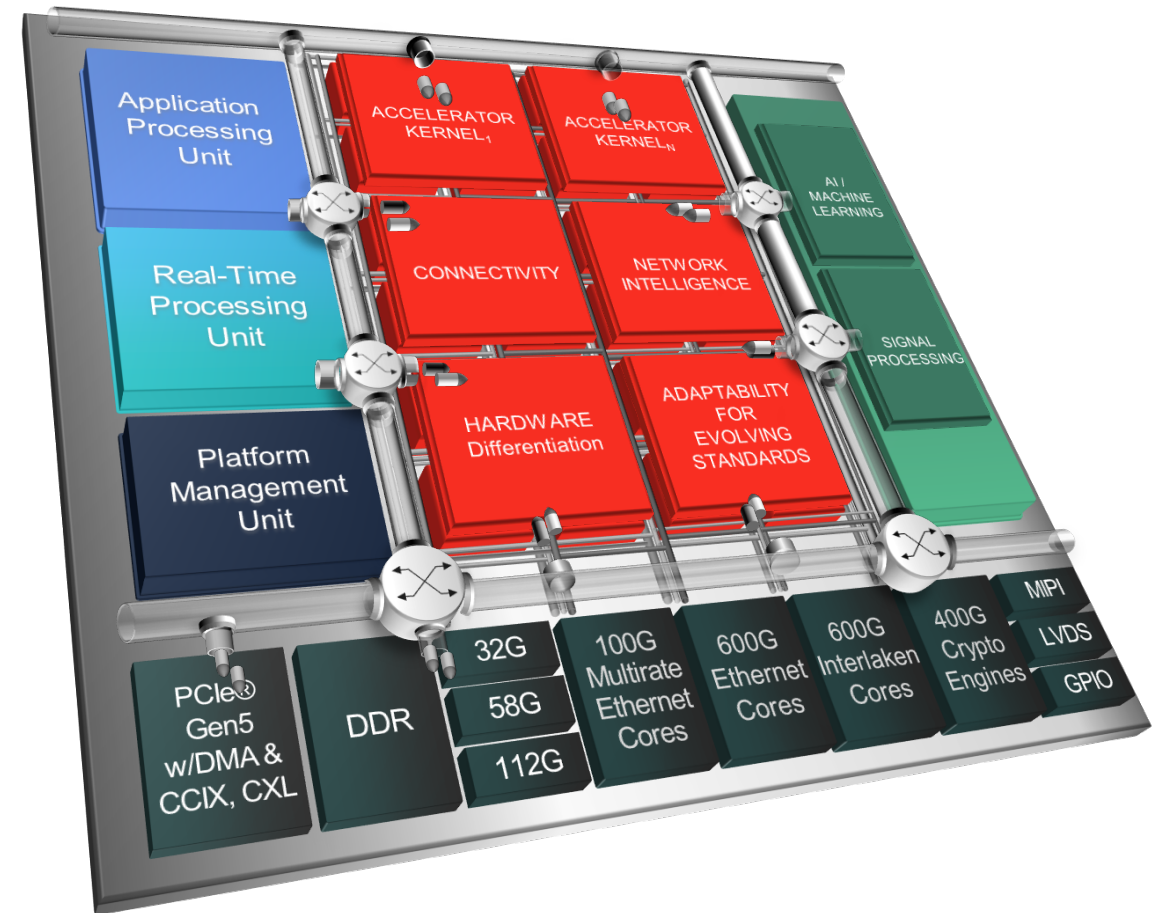
- ▶ Meet timing of critical paths
- ▶ Guaranteed QoS (bandwidth vs. latency)

Eases IP and Kernel Placement

- ▶ Simplifies connectivity of IP and peripherals
- ▶ Easily swap kernels at NoC port boundaries

Programming Framework for

- ▶ Memory mapped access to all resources
- ▶ Built-in arbitration between engines and memory





“22 Equivalent FPGAs” on Largest Versal Premium Device

600G Ethernet
Equivalent Logic Density per Core




100G Ethernet
Equivalent Logic Density per Core



600G Interlaken
Equivalent Logic Density per Core



400G Crypto
Equivalent Logic Density per Core



IP	kLUTs	Quantity	Total KLuts	Equivalent Device
MRMAC	394	8	3,153	VU3P
DCMAC	1,182	7	8,276	VU9P
ILKN	788	3	2,364	VU7P
Crypto	788	4	3,153	VU7P
Total		22	16,945	

Hardening Protocol Engine IP saves 17M LUTs

Key Blocks & Features

Adaptable Use of Building Blocks

Start with traditional PL

In Versal, NoC & PS
Backbone are always present

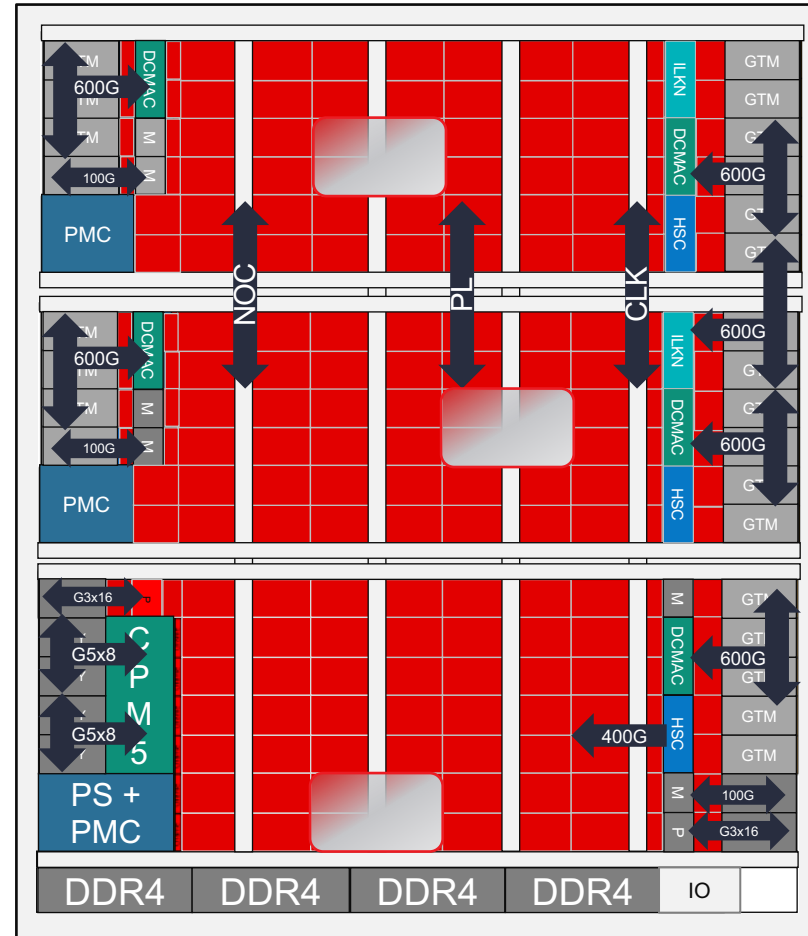
Devices have:

Memory Controller slots
(DDR4,...)

Protocol Engine slots (CPM,
PCIe5, MRMAC, DCMAC,
HSC, ILKN)

SerDes slots (GT, GTM)
Decided at tapeout time

Memory Controllers and NoC
together form **Configurable**
Memory Subsystem



Engines and GTs together
with PL form Adaptable HS
interfaces
Programmability within
blocks and between blocks

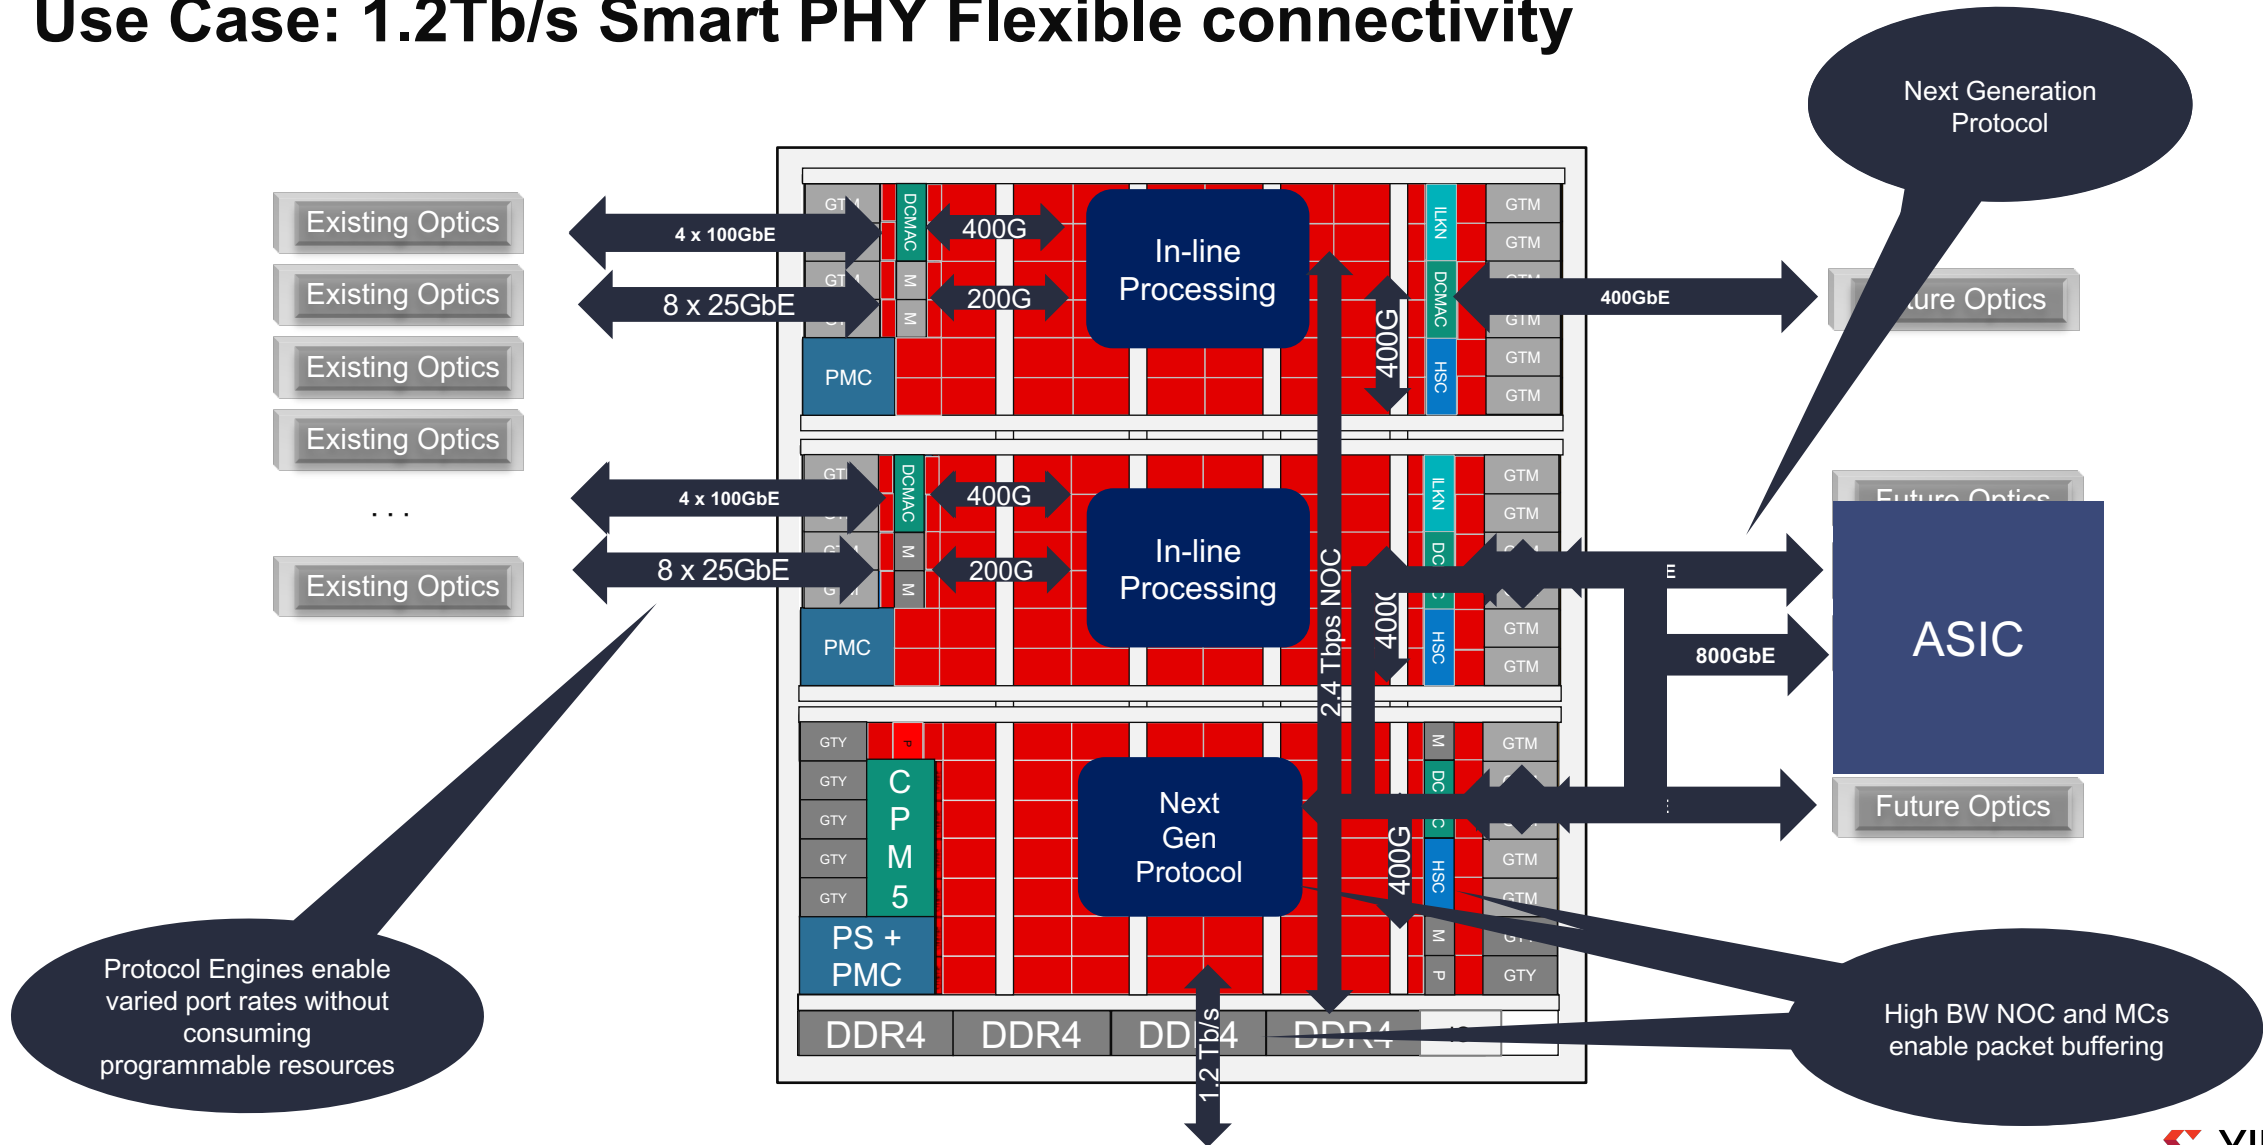
Flexible Interfaces may span
across chiplet boundaries

Swap in & out Kernels which
are adaptable

Result is an **Adaptable SoC** with enormous **Bandwidth**

Protocol Engines Enable Flexible High BW applications

Use Case: 1.2Tb/s Smart PHY Flexible connectivity

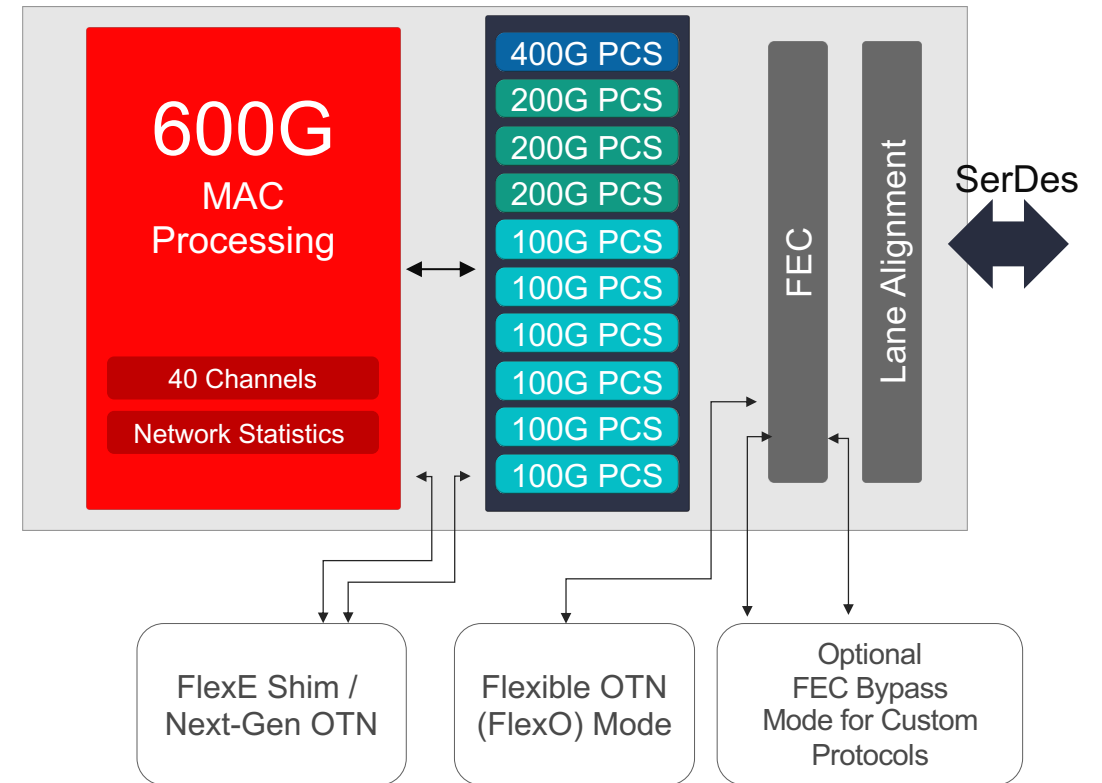


600G Multirate Ethernet (DCMAC)

Flexible, Dynamic Connectivity

- ▶ Standard Ethernet 1x400GE, 3x200GE, 6x100GE
- ▶ Channelized MAC support flexible MAC Rates
- ▶ Sub-Functions available for wide range of applications

Data Rates	Data Path Functions	Integrated PCS Options
1 x 400GE	MAC+PCS+FEC, PCS Only, FEC Only	IEEE 802.3 CL119 RS(544,514) FEC
3 x 200GE		IEEE 802.3 CL119 RS(544,514) FEC
6 x 100GE		No FEC IEEE 802.3 CL91 RS(528,514) "KR4" IEEE 802.3 CL91 RS(544,514) "KP4" FEC IEEE P802.3ck CL161 RS(544,514) "Interleaved" FEC
100G, 200G, 400G FlexO	FEC Only	RS(544,514) FEC with optional 256/257 Transcode Interface
128G FC		RS(544,514) FEC 112.2 Gbps



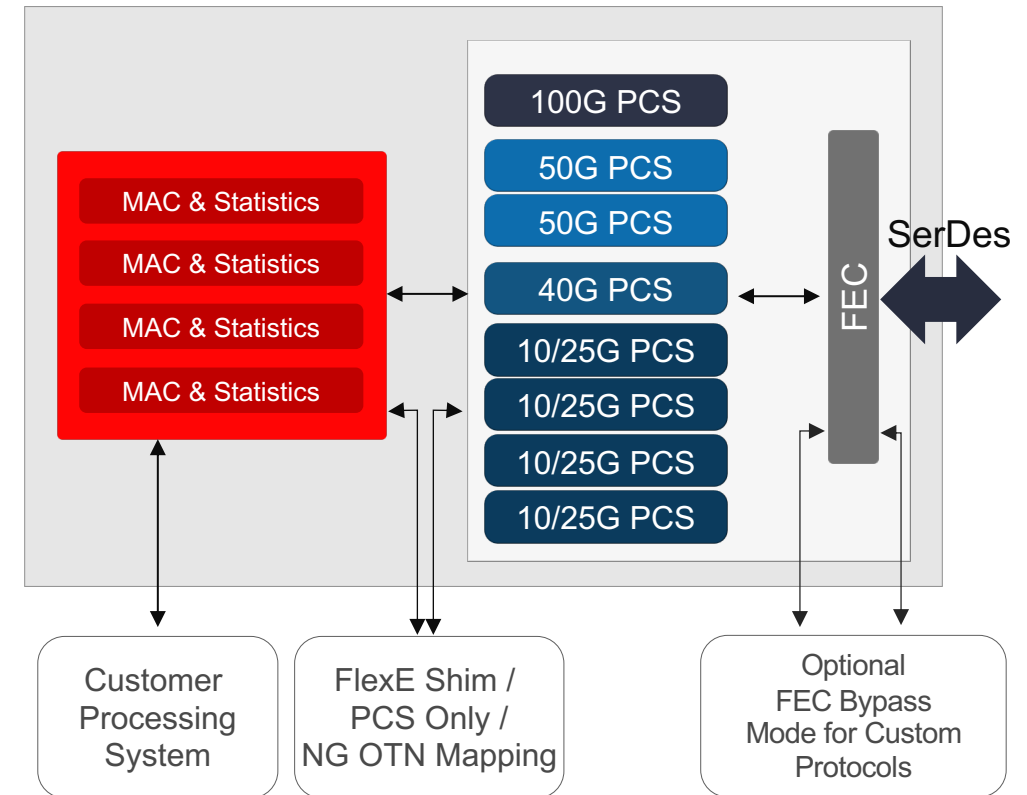
Leverage Sub-Blocks for OTN & Custom Protocols

100G Multirate Ethernet (MRMAC)

Flexible Connectivity for Mainstream Infrastructure

- ▶ Ethernet rates 1x100GE, 2x50GE, 1x40GE, 4x25GE, 4x10GE
- ▶ Sub-Functions available for custom protocols

Data Rates	Data Path Functions	Integrated PCS Options
1 x 100GE	MAC+PCS+FEC, PCS Only, FEC Only	No FEC IEEE 802.3 CL91 RS(528,514) "KR4" IEEE 802.3 CL91 RS(544,514) "KP4" FEC
2 x 50GE		IEEE 802.3 CL134 RS(544,514) FEC, Consortium RS(528,514), CL74 FEC
1 x 40GE		IEEE 802.3 CL74
4 x 25GE		IEEE 802.3 CL108 RS(528,514) FEC, Consortium RS(528,514), CL74 FEC
4 x 10GE		IEEE 802.3 CL74
100G FlexO	FEC Only	RS(544,514) FEC with optional 256/257 Transcode Interface
128G FC		RS(544,514) FEC 112.2 Gbps



Leverage Sub-Blocks for
T&M & Proprietary Backplane

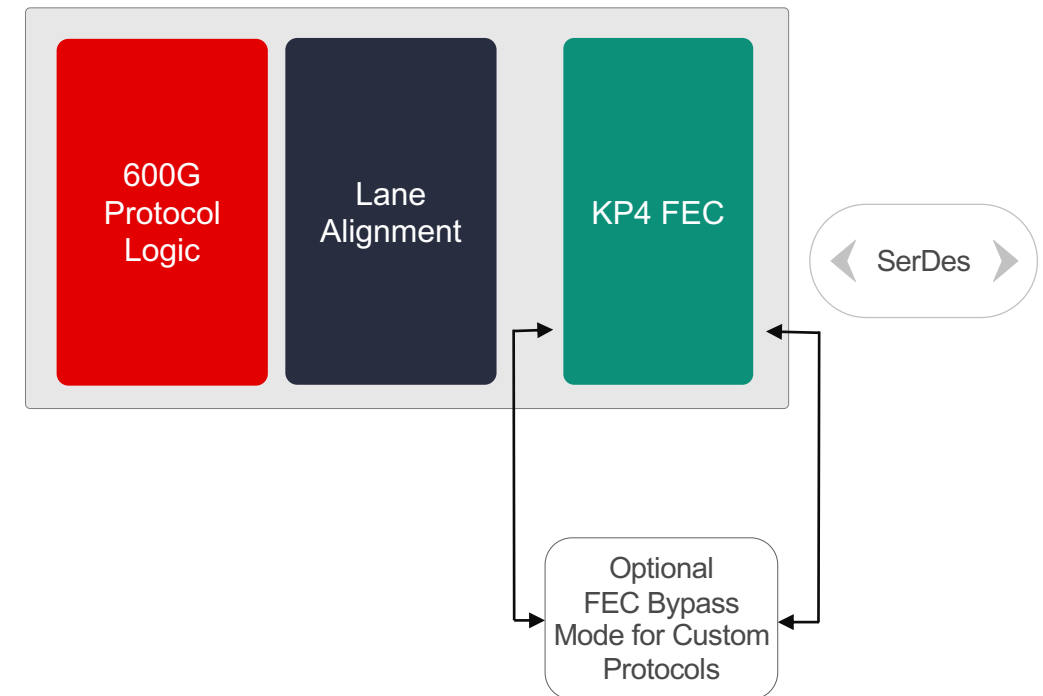
600G Interlaken with FEC for Scalable Chip-to-Chip Interconnect

Flexible, Configurable, and Scalable Interconnect

- ▶ Multiple combination of number of lanes and lane rates (12x56.42G, 24x28.21G, 24x12.5G)
- ▶ Aggregate data width to scale from 12.5Gb/s to 600Gb/s
- ▶ Flexible AXI-S User Interface w/ 2048b to 512b data width

Error Correction and Integrity Checking

- ▶ Integrated RS-FEC for power-optimized error correction
- ▶ Pairs of lanes share 100G KP4 FEC
- ▶ Support optional FEC-Only mode and bypass FEC mode
- ▶ Built-in flow control for reliable data transfer

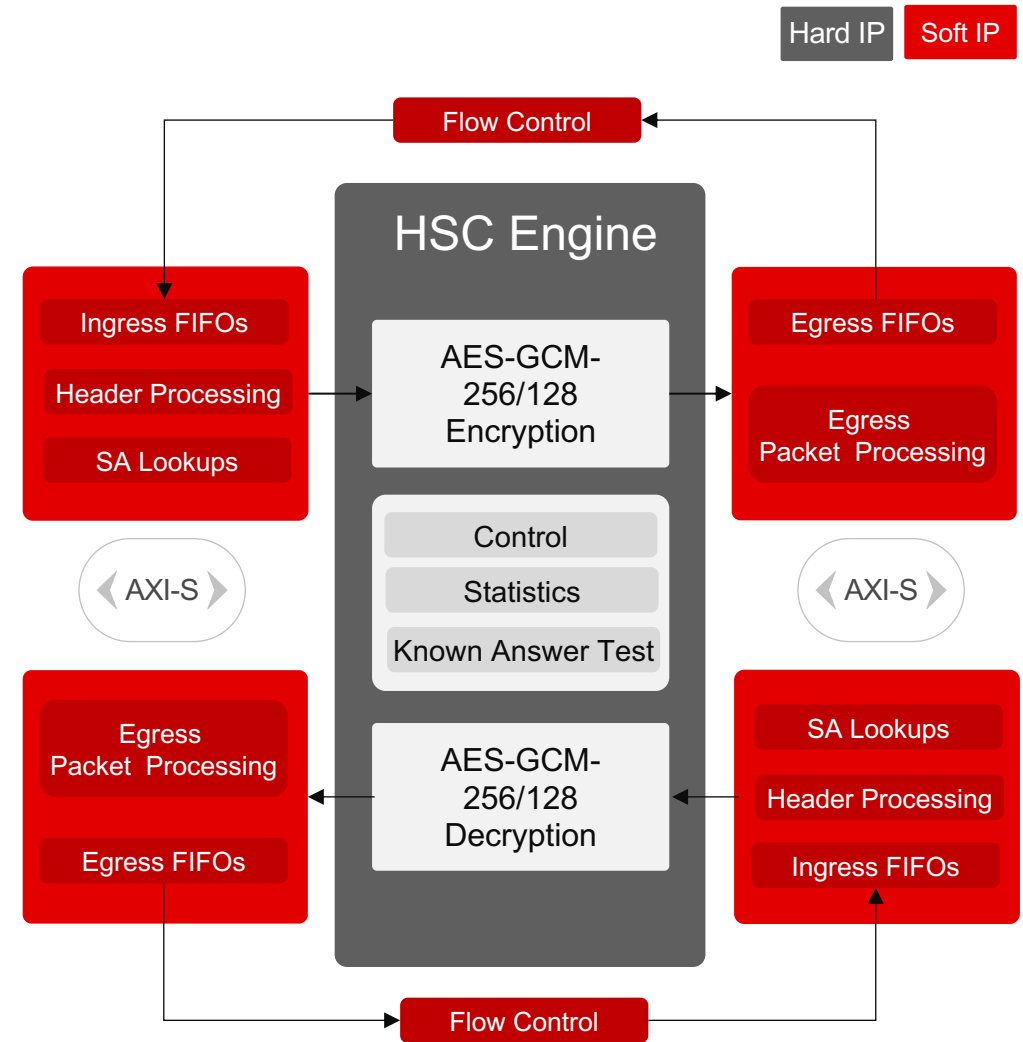


400G High-Speed Crypto Engines (HSC)

Up to 1.6Tb/s of encrypted line rate throughput

- ▶ World's only hardened 400G Crypto Engines in an adaptable platform
- ▶ AES-GCM-256/128 engine for encryption/decryption
- ▶ Integrity and Confidentiality operating modes
- ▶ 1024 SAs internal storage, expandable via external storage

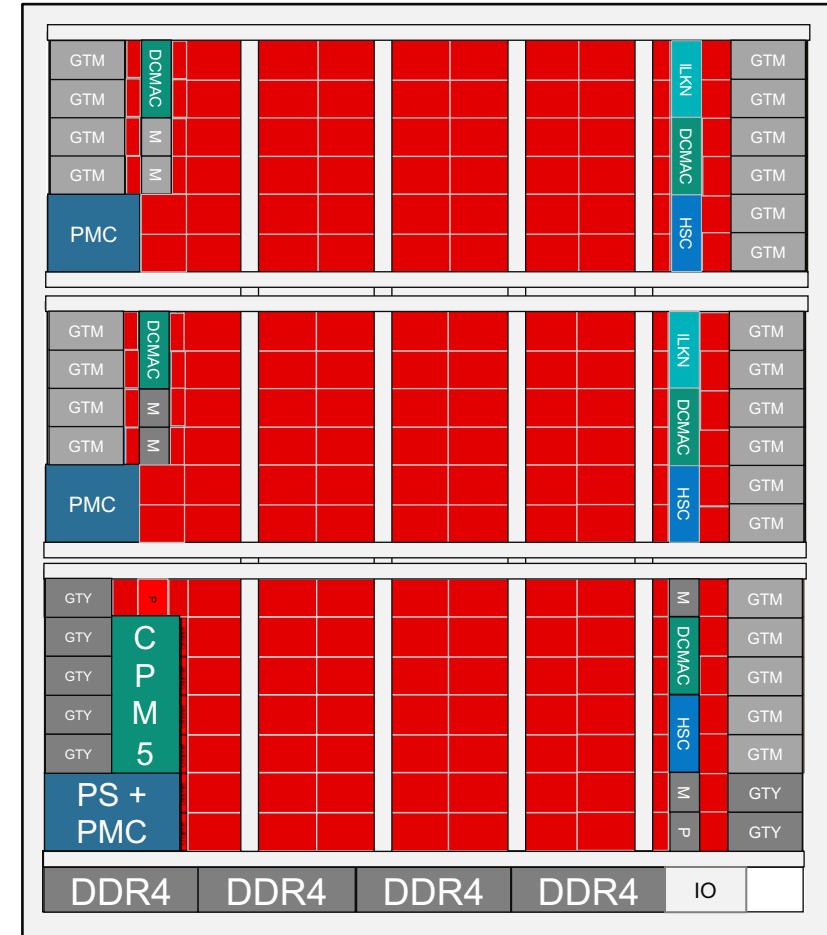

Data Rate	Crypto Functions
1 x 400GE	MACsec, IPsec, BulkCrypto, BulkECB
2 x 200GE	MACsec, IPsec, BulkCrypto, BulkECB (configurable per port)
4 x 100GE	MACsec, IPsec, BulkCrypto, BulkECB (configurable per port)
Combinations of 100, 200G (totalling 400G)	MACsec, IPsec, BulkCrypto, BulkECB (configurable per port)
40 Time-Slice Channels (totalling 400G)	MACsec, IPsec, BulkCrypto, BulkECB (configurable per Channel)



MACsec Use Model

High-Speed SerDes Interface

- ▶ Programming part of SoC shell (NPI), ready before Programmable Logic configuration
- ▶ Support many protocols with as few blocks as possible
- ▶ Two flavors of Serdes



GTYP vs GTM SerDes

- ▶ Long-Reach 56-112Gbps PAM4 requires ADC/DSP-based architecture
 - Area is larger than conventional serdes – need smaller serdes for lower-end products
- ▶ Many protocols require high-BW CDR
 - Supporting this on ADC/DSP-based SerDes is a challenge – currently in R&D for future products

GTYP

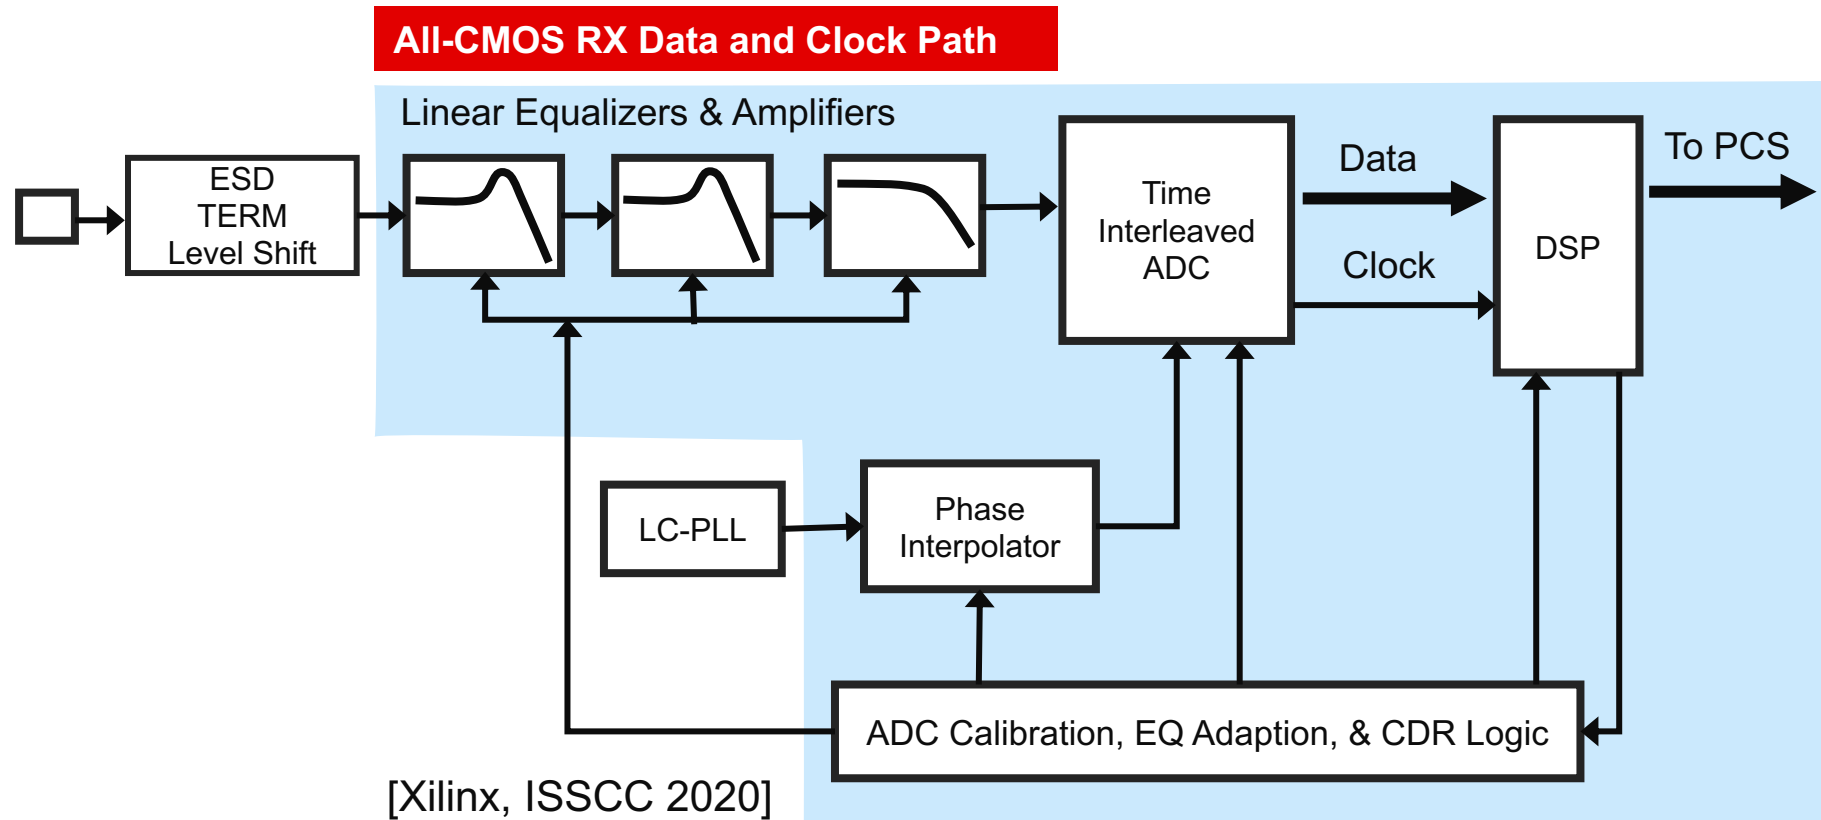
- ▶ 100+ Protocols up to 32.75Gbps NRZ
 - PCIE Gen-5
 - 10/25/40/100G Ethernet
- ▶ High Bandwidth CDR to track SSC and reference clock frequency wander
 - SRIS for Storage-over-PCIE
 - SATA/DisplayPort
- ▶ Analog/Mixed-Signal DFE-based Receiver

GTM

- ▶ Optimized for high-performance communication protocols
 - 400GE
- ▶ Up to 112Gbps PAM4
 - +/- 100ppm
- ▶ Configurable ADC/DSP-based Receiver

GTM Transceiver Architecture


- ▶ Hybrid Analog/Digital Equalization
- ▶ 7nm FinFET motivates All-CMOS RX circuit architecture
 - CMOS Inverters as building blocks for analog blocks (e.g. equalizers, amplifiers, phase interpolators)




GTM Configurable ADC & DSP

Allows power/performance optimizations over various use cases

56-112Gbps PAM4 Long Reach
Backplane, Direct-Attach Cable




56-112Gbps PAM4 VSR
Chip to Optical Module



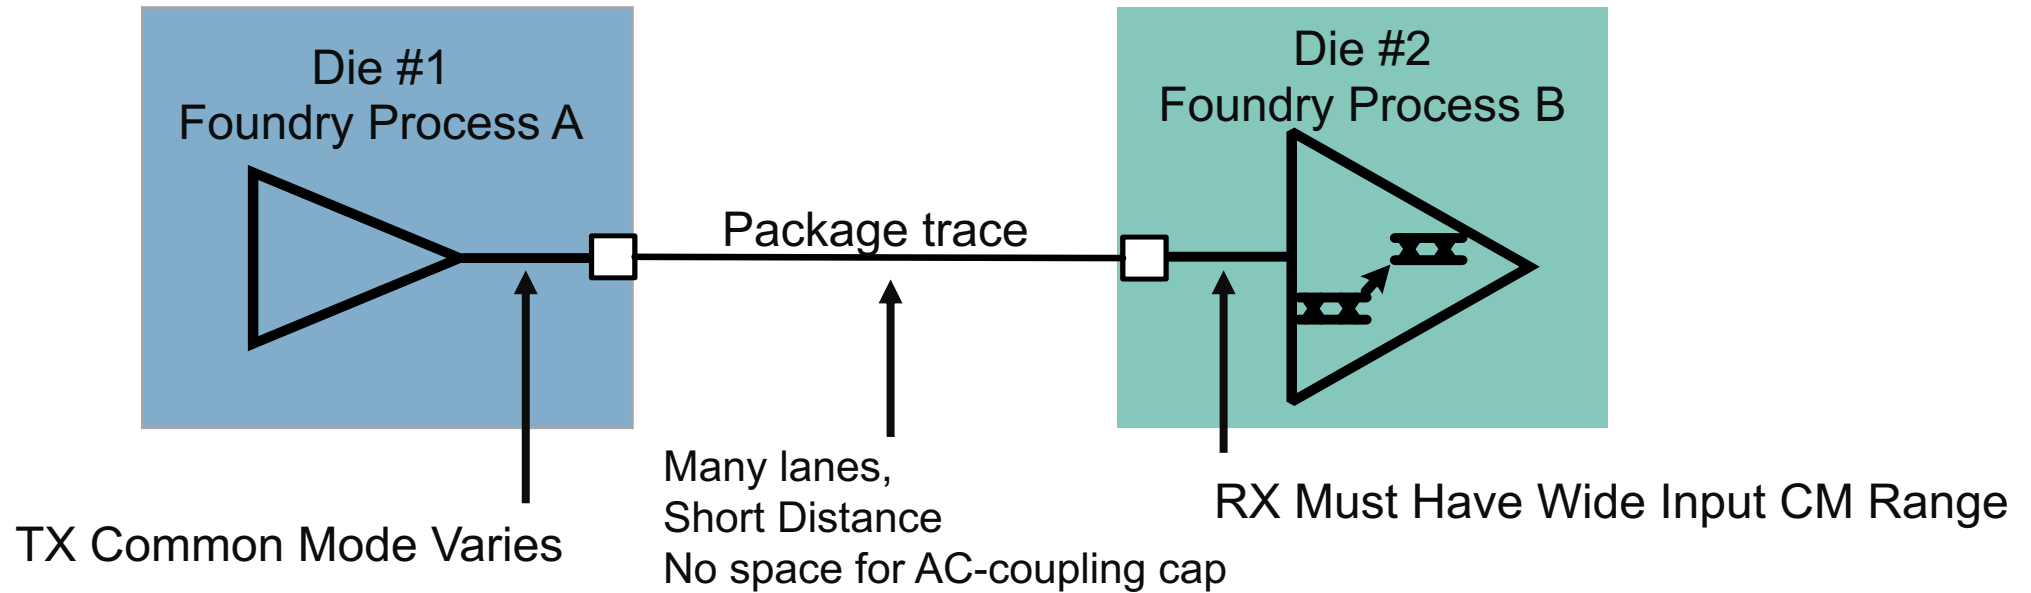
Legacy lower-rate NRZ
(10Gbps NRZ, 25Gbps)



112Gbps PAM4 XSR
Die-to-die interface for MCM



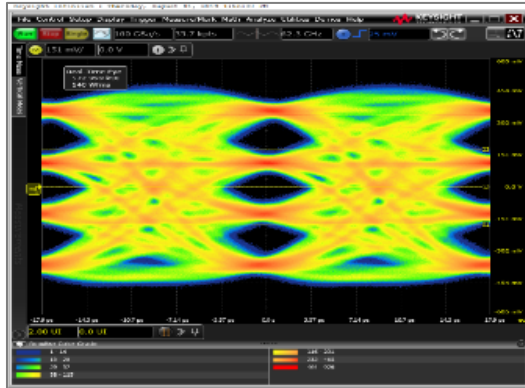
112Gbps XSR Die-to-die Interface Design Challenges



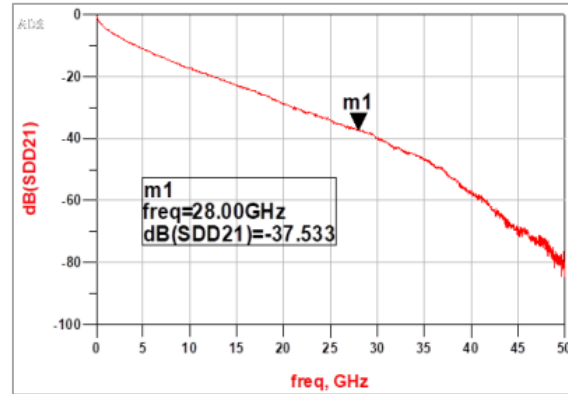
- ▶ 112Gbps RX must support wide input common-mode range
- ▶ Low Power, Low Latency
 - ADC set to lowest resolution
 - DSP bypassed or doing minimal work
 - Need to achieve low BER without FEC

GTM Long Reach Performance: Silicon Results

TX Eye Diagram



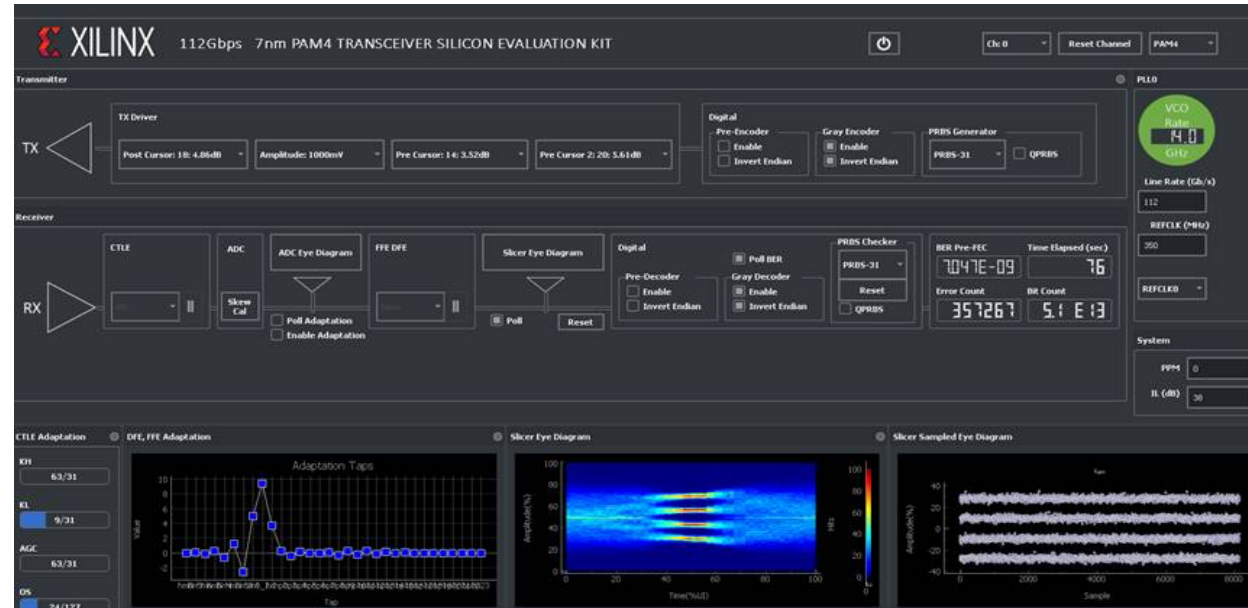

37.5dB Channel



Test Setup



RX eye (PRBS-31)

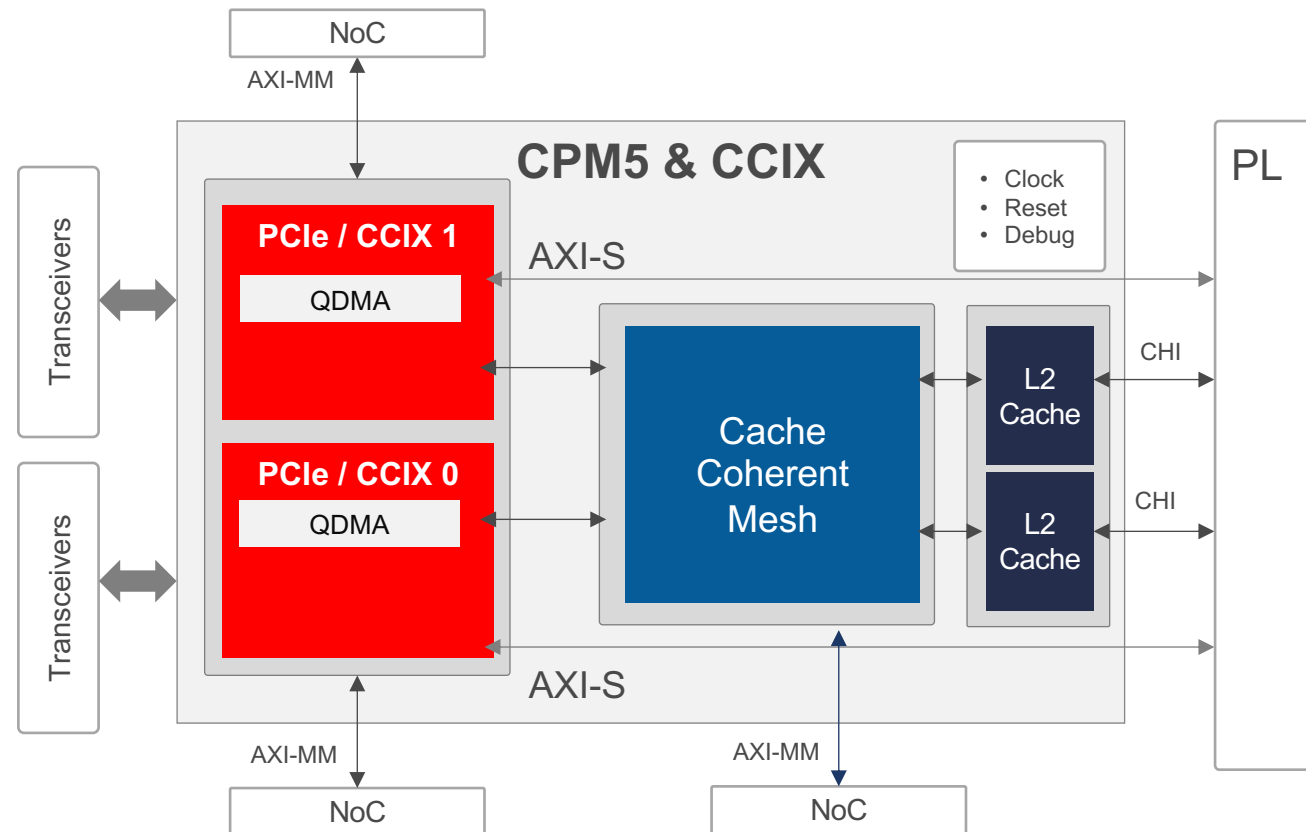


Protocol Engines & SerDes Summary

- ▶ Protocol Engines and Serdes Enable Flexible High BW Applications
 - User can focus on implementing in-line processing in the Programmable Logic
- ▶ Protocol Engines
 - DCMAC : 600G multi-rate, channelized ethernet
 - MRMAC : 100G multi-rate, channelized ethernet
 - ILKN : 600G multi-channel chip-to-chip interface with built-in flow control
 - HSC : 400G Crypto Engine
- ▶ Serdes Blocks supporting 100+ protocols
 - GTYP : Optimized for protocols \leq 32Gbps NRZ
 - GTM : Optimized for 56-112Gbps Long Reach PAM4

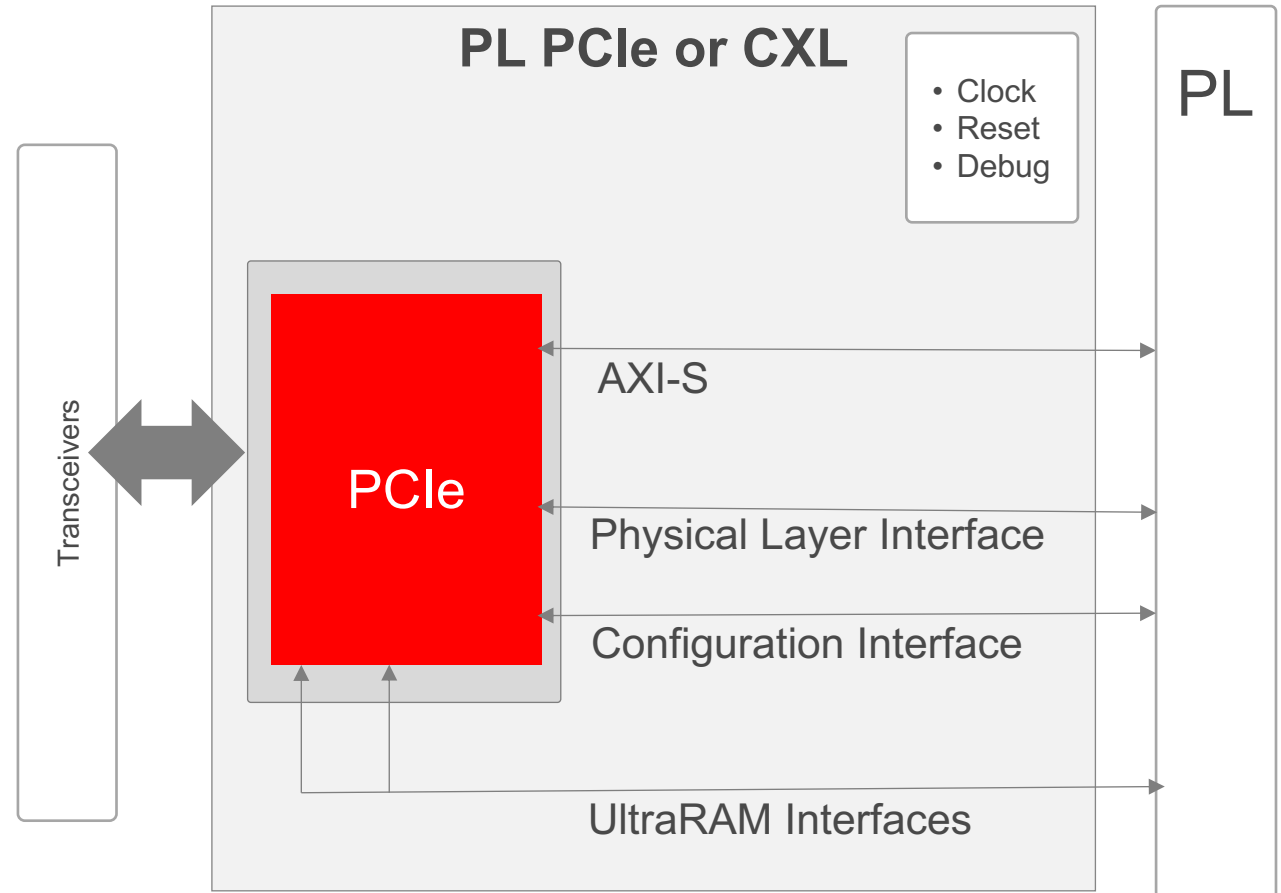
Integrated PCIe Gen5, DMA, CCIX

- ▶ CPM5 in Hardened SoC Shell
- ▶ 2x8 Gen5, CCIX ESM 25G, Gen4 x16
- ▶ PCIe and CCIX Features:
 - Integrated QDMA with SR-IOV support
 - Integrated Coherent CCIX Subsystem:
 - 4.3Tbps Coherent Mesh Interconnect, with integrated 1.2Tbps L2 Cache
 - 1Tbps Coherent Accelerator-attached Memory
 - 800Gbps Coherent Ports for PL Accelerators
 - 2 instances each of CCIX RA and HA or SA – available in any combination
 - Available at boot
- ▶ Dedicated connectivity over NoC to DDR



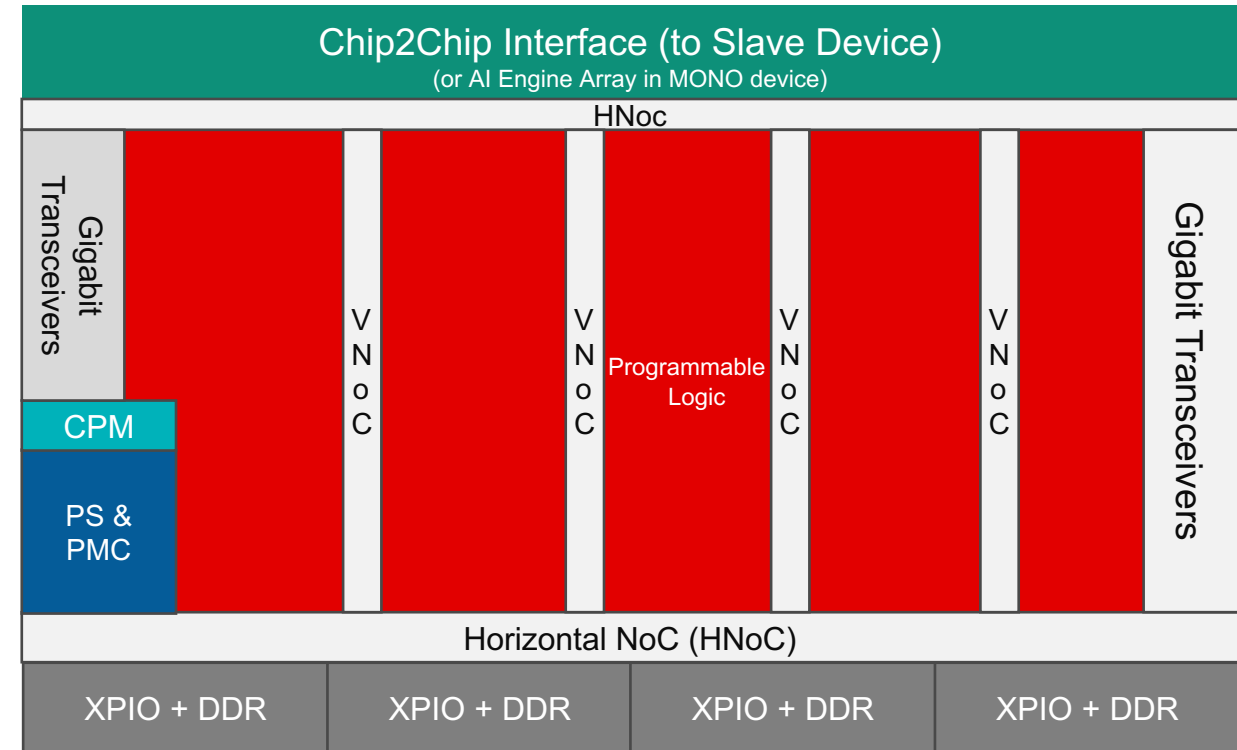
PL-Distributed PCIe Gen5 and CXL

- ▶ Hardened PCIe/CXL Controllers in PL
- ▶ Gen5 x4, Gen4 x8, Gen3 x16
- ▶ Backward pin-compatible with PCIe soft cores
- ▶ CXL: Combination of integrated IP and PL
- ▶ Root & End Point
- ▶ Upstream Port & Downstream Port
- ▶ SoC with up to 8 PCIe-G5 controllers
 - Drive up to 8 NVMe end points
 - 256 GByte/s aggregate BW



Versal NoC

- ▶ Versal device Topology & Routes are not fixed
 - Versal is a family of devices with varying NoC topologies
 - Dictated by use case
 - Routes are established at runtime by programming
- ▶ NoC topology is scalable and not regular
 - Driven by device constraints
 - Same reusable components
- ▶ HNoC & VNoC Channels
 - Each physical channel with 8 Virtual Channels
 - >1Tbps bidirectional bandwidth per row for 4 channels
 - >0.5Tbps bidirectional bandwidth per column for 2 channels
- ▶ Programmable Logic requires programmability
 - Data width, QoS, address mapping and routes are programmable



NoC Protocol, Components, Channels

▶ New NPP Packet Protocol

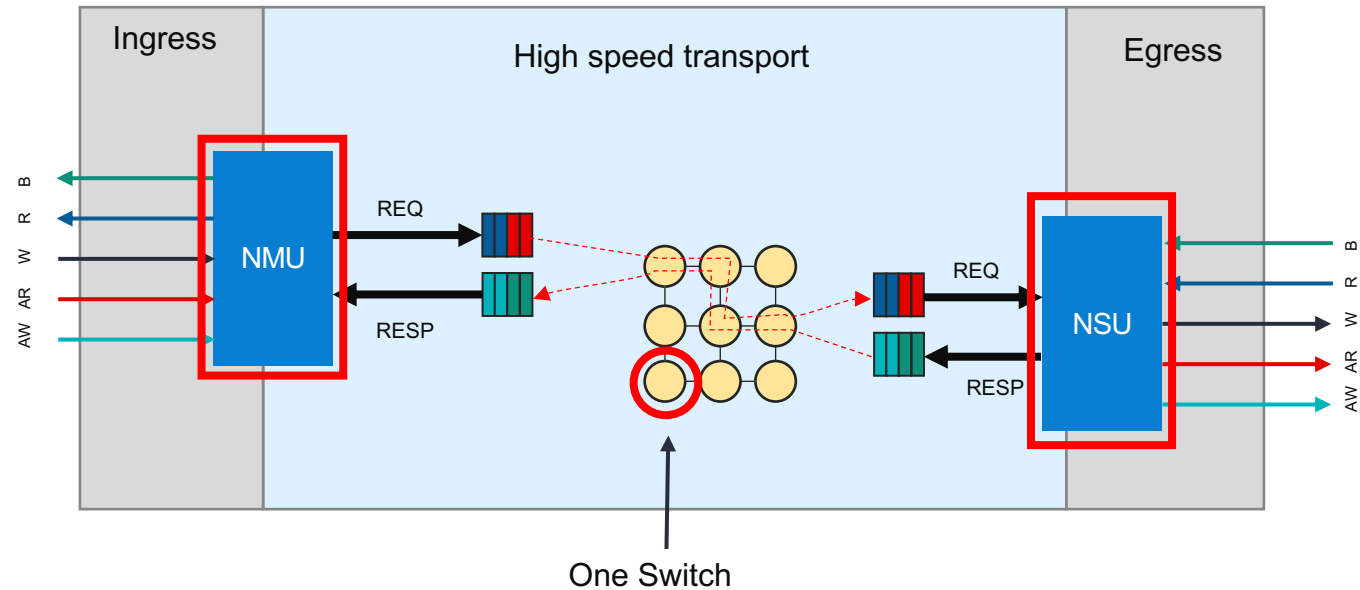
- Maps AXI Memory-mapped and Stream
- Can be used to map other protocols
- End-to-end ECC protection

▶ NoC consists of 3 key components

- NMU: Ingress from user protocol (AXI), maps to NPP
- NSU: Egress to user protocol (AXI), maps from NPP
- Core Switch: has routing table, does QoS & switching, defines which NMU talks to which NSU
- Rows & columns implemented using these components

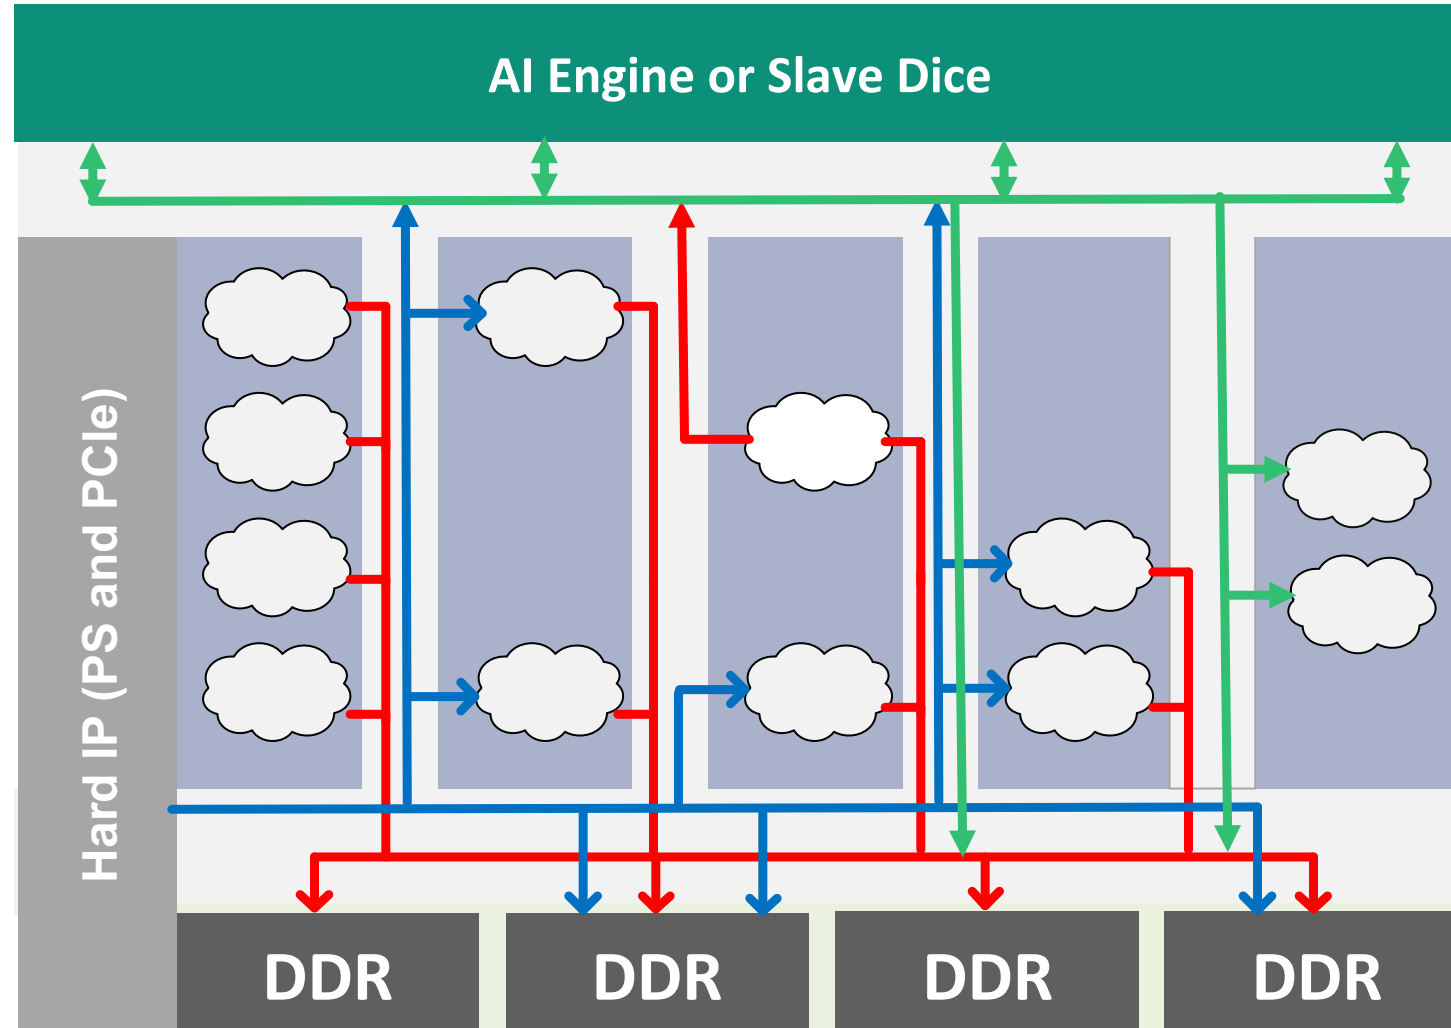
▶ Physical and Virtual channels

- Each row and column have multiple Physical Channels
- VCs (8) provide independent non-blocking traffic flow
- Traffic is packetized



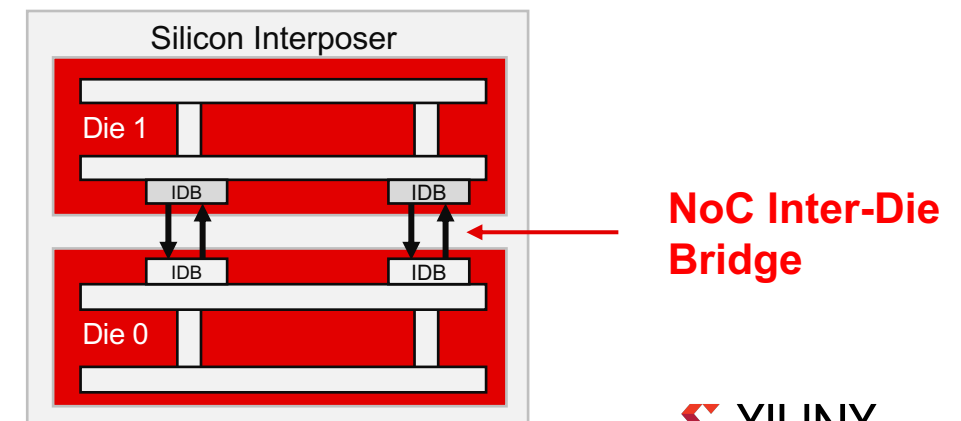

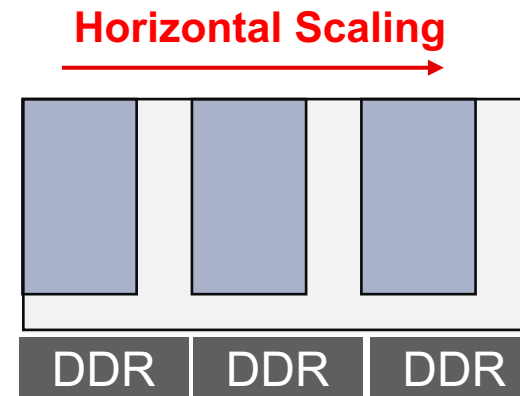
Versal NoC Use Case Example

- ▶ AI/ML Data Flow
 - Host (PCIe) reads/writes to DDR or PL
 - PS reads/writes DDR, and controls PL & AIE
 - PL accesses DDR and AIE
 - AIE accesses DDR and PL
- ▶ Traffic Sources & Destinations:
 - Masters: PS, PCIe, 10x PL, AIE Array
 - Slaves: 4x DDR, 7x PL, AIE Array
- ▶ NoC provides high-speed interconnect
 - No PL resources consumed for Routing, Switching, and QoS
- ▶ NoC compiler runs in seconds
- ▶ Easier PL design and timing closure



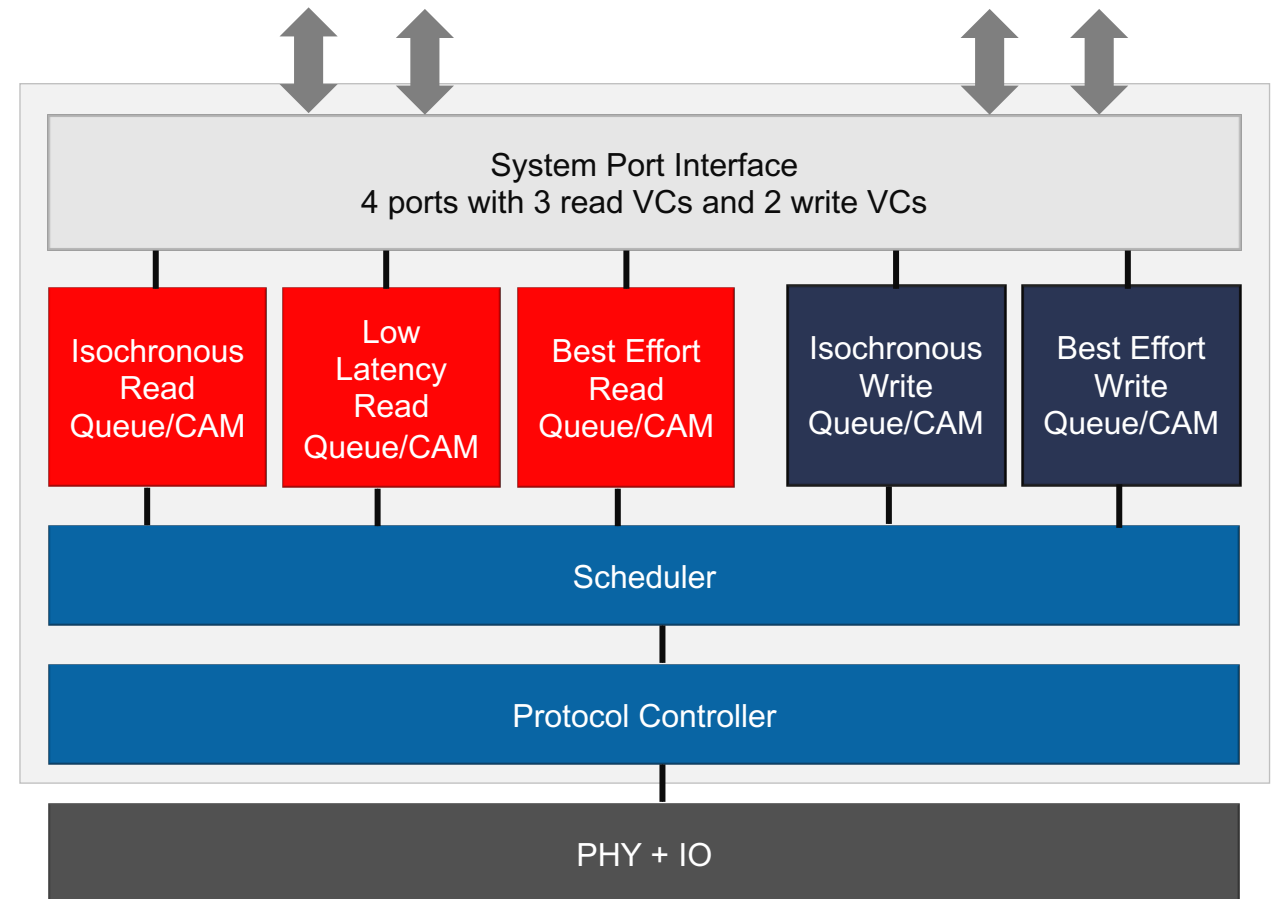
Versal NoC Scaling

- ▶ Bigger devices have more NoC resources
- ▶ Wider devices have more VNoC bandwidth
 - Wide devices have more DDR at bottom and Accelerators on top
 - More Fabric -> NoC connections as device width grows
 - One VNoC column per 64b DDR channel
- ▶ Taller devices have more Programmable Logic
 - More PL-NoC connections as device height grows
 - Each Fabric-row in VNoC has one ingress and one egress NoC port
- ▶ Die-to-die connectivity in SSIT devices
 - NoC extends vertically across multiple dice in SSIT devices
 - Source synchronous, single data rate links across interposer




Memory Subsystem and IO

- ▶ Configurable Bandwidth and QoS Allocation
 - 4 Ports, Per Port: 3 read VCs, 2 write VCs
 - Enables Real-time guarantees (Isochronous), Throughput guarantees (Low Latency), Best Effort
 - Transaction Reordering, Buffering
- ▶ 256b DDR w/ 4x 64b or 8x 32b channels
 - Optimized for 64b or 32b memory channels
 - 32b granularity more efficient for some use-cases
 - DDR4 up to 3200 and LPDDR4 up to 4266 Mbps
- ▶ Parallel IOs
 - Up to 702x High Performance XPIOs for DDR, MIPI, ...
 - Up to 122x High Density multiprotocol IOs for up to 3.3v



Device Architecture



Device Architecture

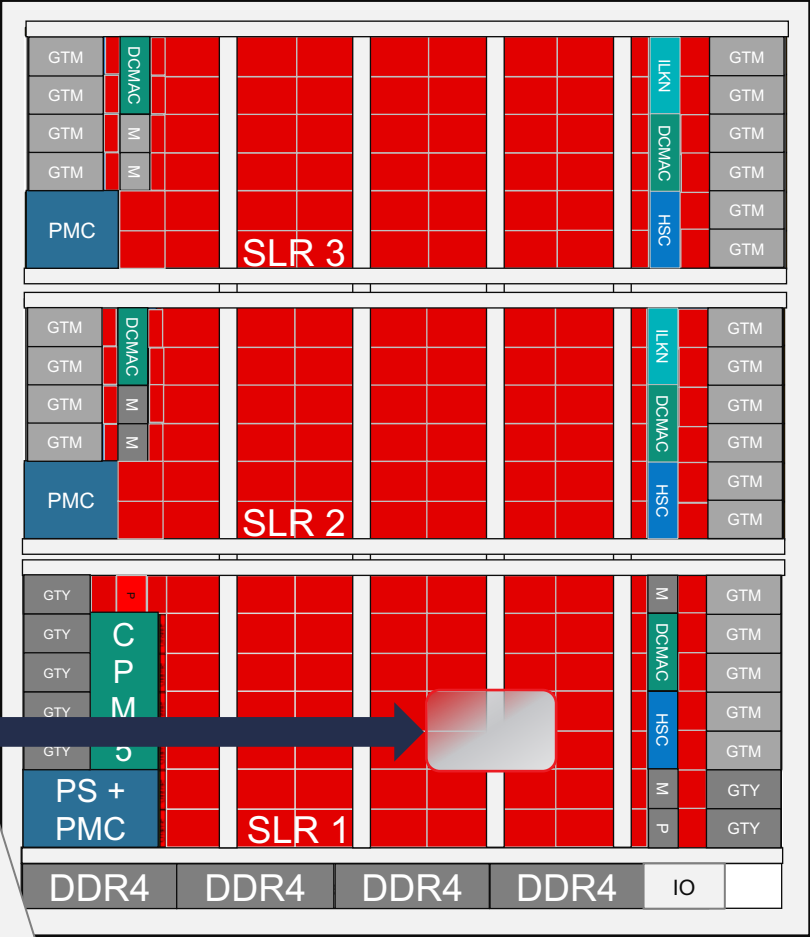

- ▶ SLR (Super Logic Region): Mono ACAP device
- ▶ ACAP: PL surrounded by hardened SoC Shell
 - PS, CPM/PCIe/CCIX/CXL, DDR4, NOC
 - SoC Shell up & running before Programmable Logic programming
- ▶ SSIT combines multiple SLRs in single device

- Kernels may be relocated between identical PL regions within device
- Relocation may happen w/o recompile, identical timing & bitstream-segment

High bandwidth terabit Programmable NoC

- Meet timing of critical paths
- Guaranteed QoS (BW vs. latency)



- Eases IP & Kernel Placement
- Simplifies connectivity of IP and peripherals
- Easily swap kernels at NoC port boundaries



Kernels may be shared between identical PL regions between (selected) devices

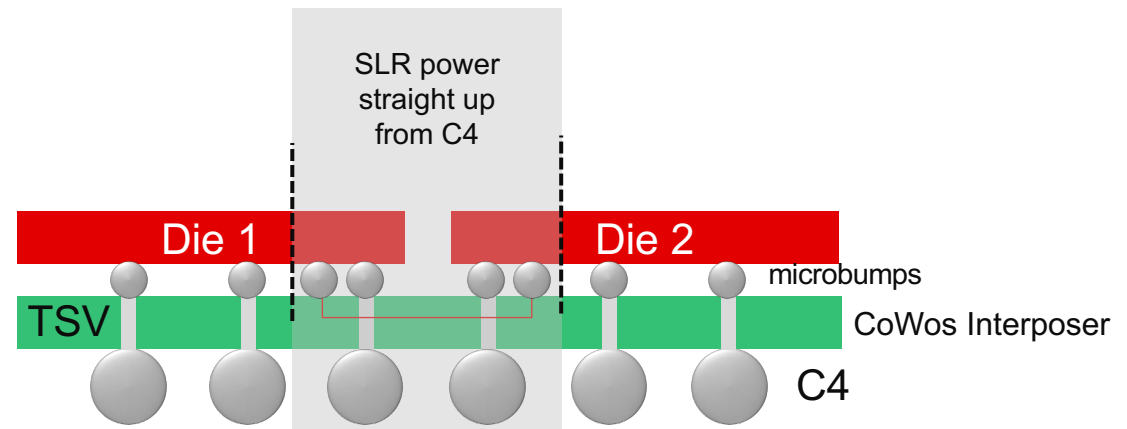
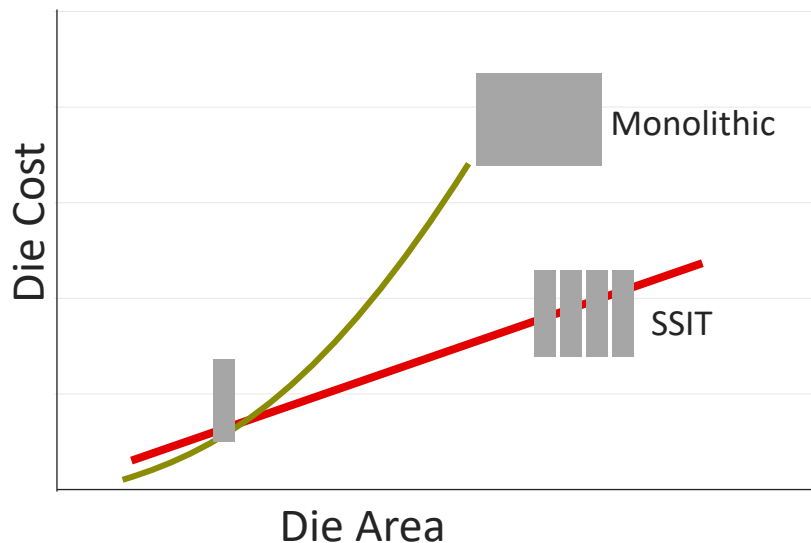
4th Gen Stacked Silicon Interposer Technology (SSIT)

- ▶ TSMC CoWoS (Chip-on-Wafer-on-Substrate)
 - Passive Interposer, for adequate wire-densities & -performance (65nm, mature)
- ▶ Large logic capacity with larger-than-reticle devices; more-than-Moore
 - Enable devices size up to 2x reticle limit; reticle size constraint from fab Litho steppers, has not changed over process nodes
 - Super-Interposer requires 2 exposures and interposer wire “stitching”
- ▶ Reduce #tapeouts for MAX #devices with various SLR combos/counts
- ▶ Heterogeneous SLR chiplets
 - “Bottom” with IO/DDR4 at bottom, “Mid” without IO/DDR4, future “Top” may have IO or HBM on top



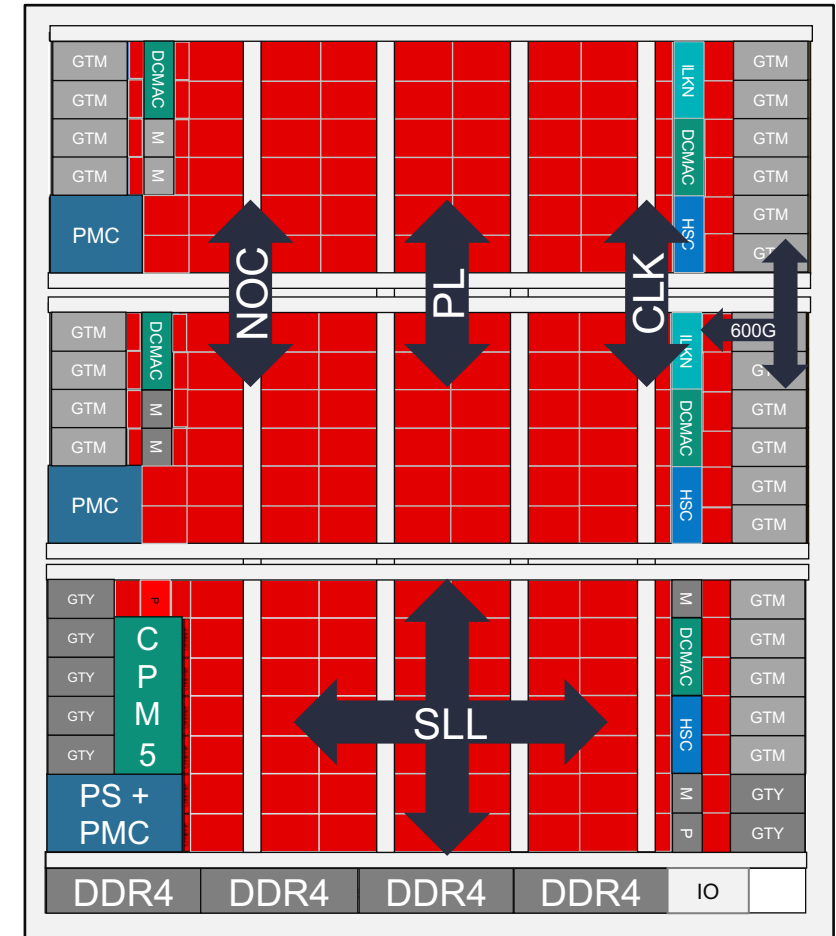
4th Gen Stacked Silicon Interposer Technology (SSIT) (cont'd)

- ▶ Favorable economics
 - Known Good Dice on Interposer; Cost scales linearly with total device-area (exponentially for monolithic die)
 - Enables us to keep SLR sizes far smaller than reticle size
- ▶ Chiplet beachfront area covered with C2C wires & u-bumps still have power supplied through TSVs



SSIT Connectivity

- ▶ NoC, PL-interconnect, CLKs extend across chiplet boundaries
 - PL BW: 15.3Tbps over ~20K wires
 - NoC BW: 2.2Tbps over 4 VNoCs
 - CLKs: 120 global clks
- ▶ SerDes interfaces can extend across chiplet boundaries
 - 3 GTM Quads for 600G interfaces (56G per channel)
 - REFCLKs forwarded
- ▶ Additional intra-chiplet interconnect from interposer metals (Super Long Lines)
 - HOR BW: 5.4Tbps
 - VER BW: 3.1Tbps



SSIT Connectivity Improvements

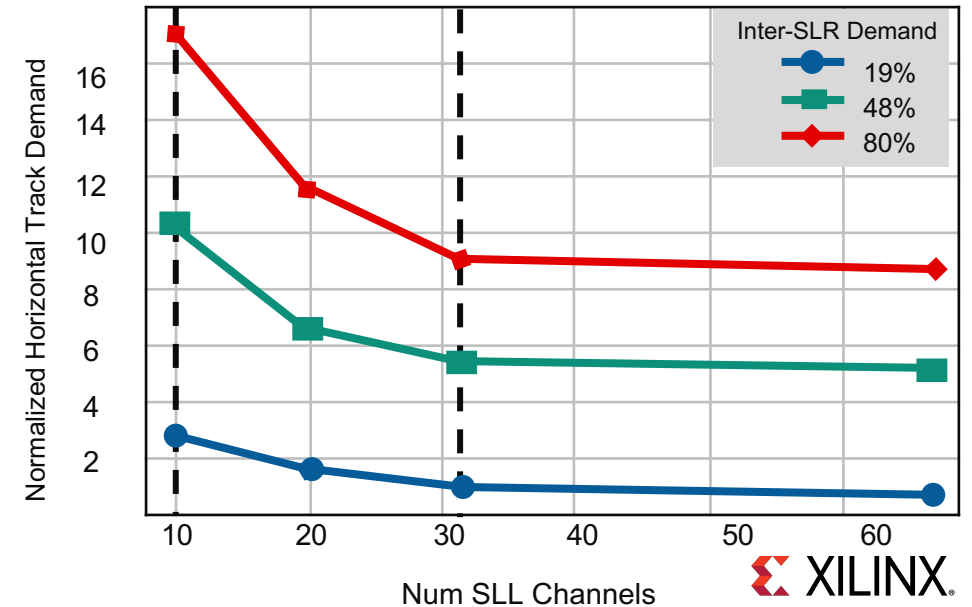
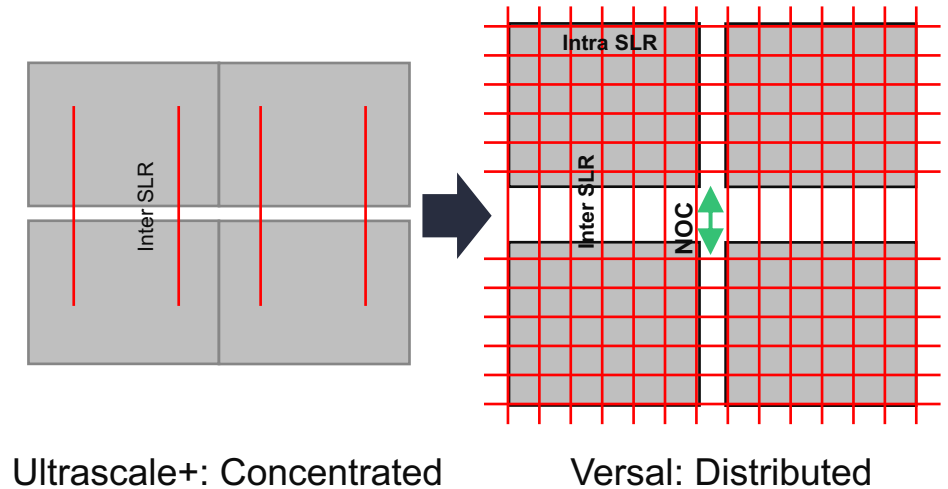
▶ Inter-SLR SLL distributed vs concentrated

- For increased #SLL-channels, HOR track demand reduces
- HOR congestion reduces 50% (at 80% SLL use)
- 5% better routability
- 40% redux in HOR wire length


▶ Added intra-SLR connectivity

- Interposer wires
- Long wires on Interposer 30% faster than regular interconnect for similar distances
- Alleviates routing congestion; 8% HOR wire length redux, 6% VBR wire length redux

▶ Added VNOC-VNOC connectivity



Use Cases



Versal Premium Delivering “More than Moore”



2.6X

System Protocol Bandwidth
for Fastest and Most Secure Networks

2.6X

INT8 Compute Density
for Adaptable Acceleration


3.3X

Logic Capacity
for Data Center Acceleration

Highly Integrated

HW/SW Platform

For Productivity



Architecture	Device	SLR	Available Logic Capacity ¹	OCM Capacity MB	OCM BW TB/s	INT8 TOPs	FP32 TFLOPs	SerDes BW Tbps	System Protocol BW Tbps ²
Virtex Ultrascale+	VU13P	4	1030	63	55	38	0.7	4.2	2.9
Versal Premium	VP1802	4	3360	124	123	99	23.0	9.0	7.6

¹ Available Logic Capacity minus “shell” resources – memory controllers, memory interconnects, PCIe and DMA

² System Protocol BW represents device-level Ethernet, PCIe and Interlaken hard IP bandwidth

For the Most Demanding Compute & Bandwidth Applications

Wired



- ▶ Metro/Core Transport
- ▶ Data Center Interconnect
- ▶ Security Appliances

Data Center



- ▶ Fine-Grain Acceleration
- ▶ Latency Sensitive Workloads
- ▶ Sort/Search/Hash

Test & Measurement



- ▶ Network Testers
- ▶ Mobile Testers
- ▶ PCIe Protocol Analyzers

Aerospace & Defense



- ▶ I/Q Complex TMACs
- ▶ Radar
- ▶ Avionics

Use Case: 3.2Tb/s Encrypted Data Center Interconnect

Scalable SerDes for Diverse Optics

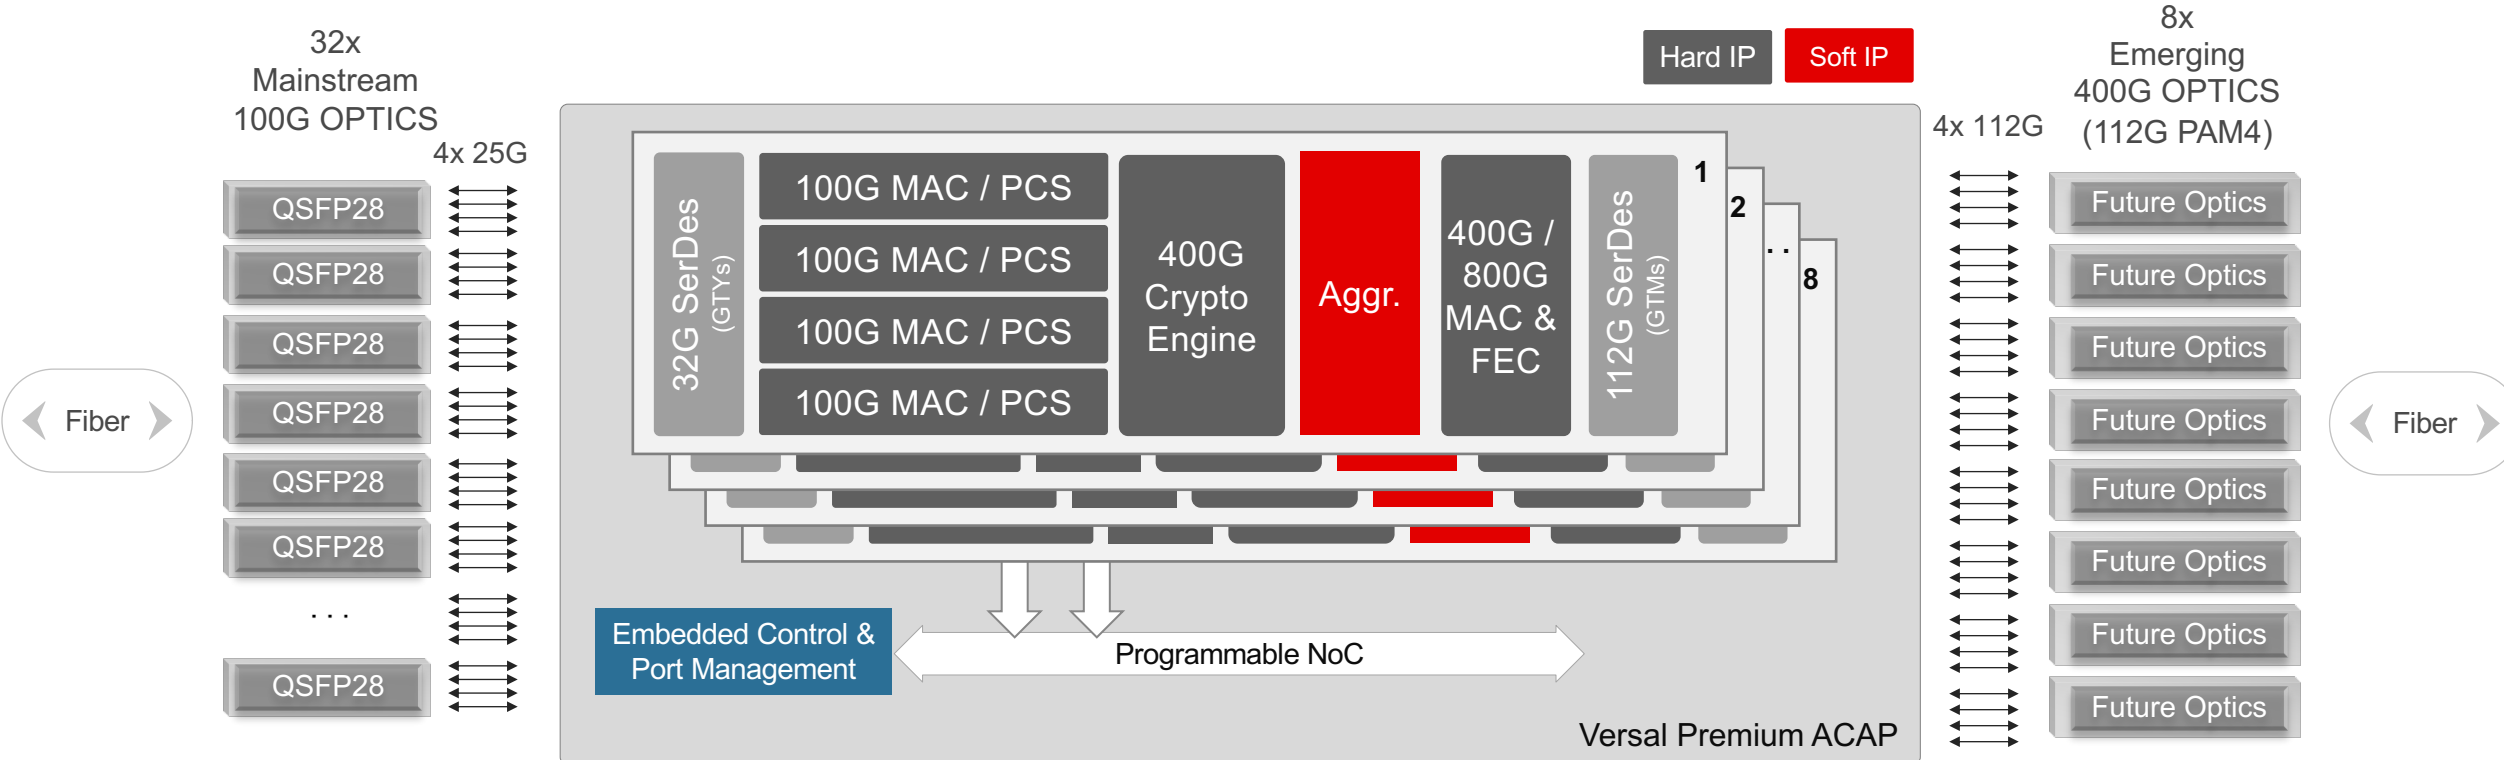
Broad mix of 32G/58G/112G SerDes to bridge between 100G and 400G Optics

Terabit Processing

Hardened IP for ASIC-class performance/watt connectivity & encryption

HW Adaptability

Adapt to Emerging and Evolving Standards (400GE/ 800GE)



Use Case : Scale-Out Fabrics & In-Line Data Processing

Serial Interfaces

Up to 112G PAM4 for fastest possible data movement

ScaleOut Protocols and RDMA

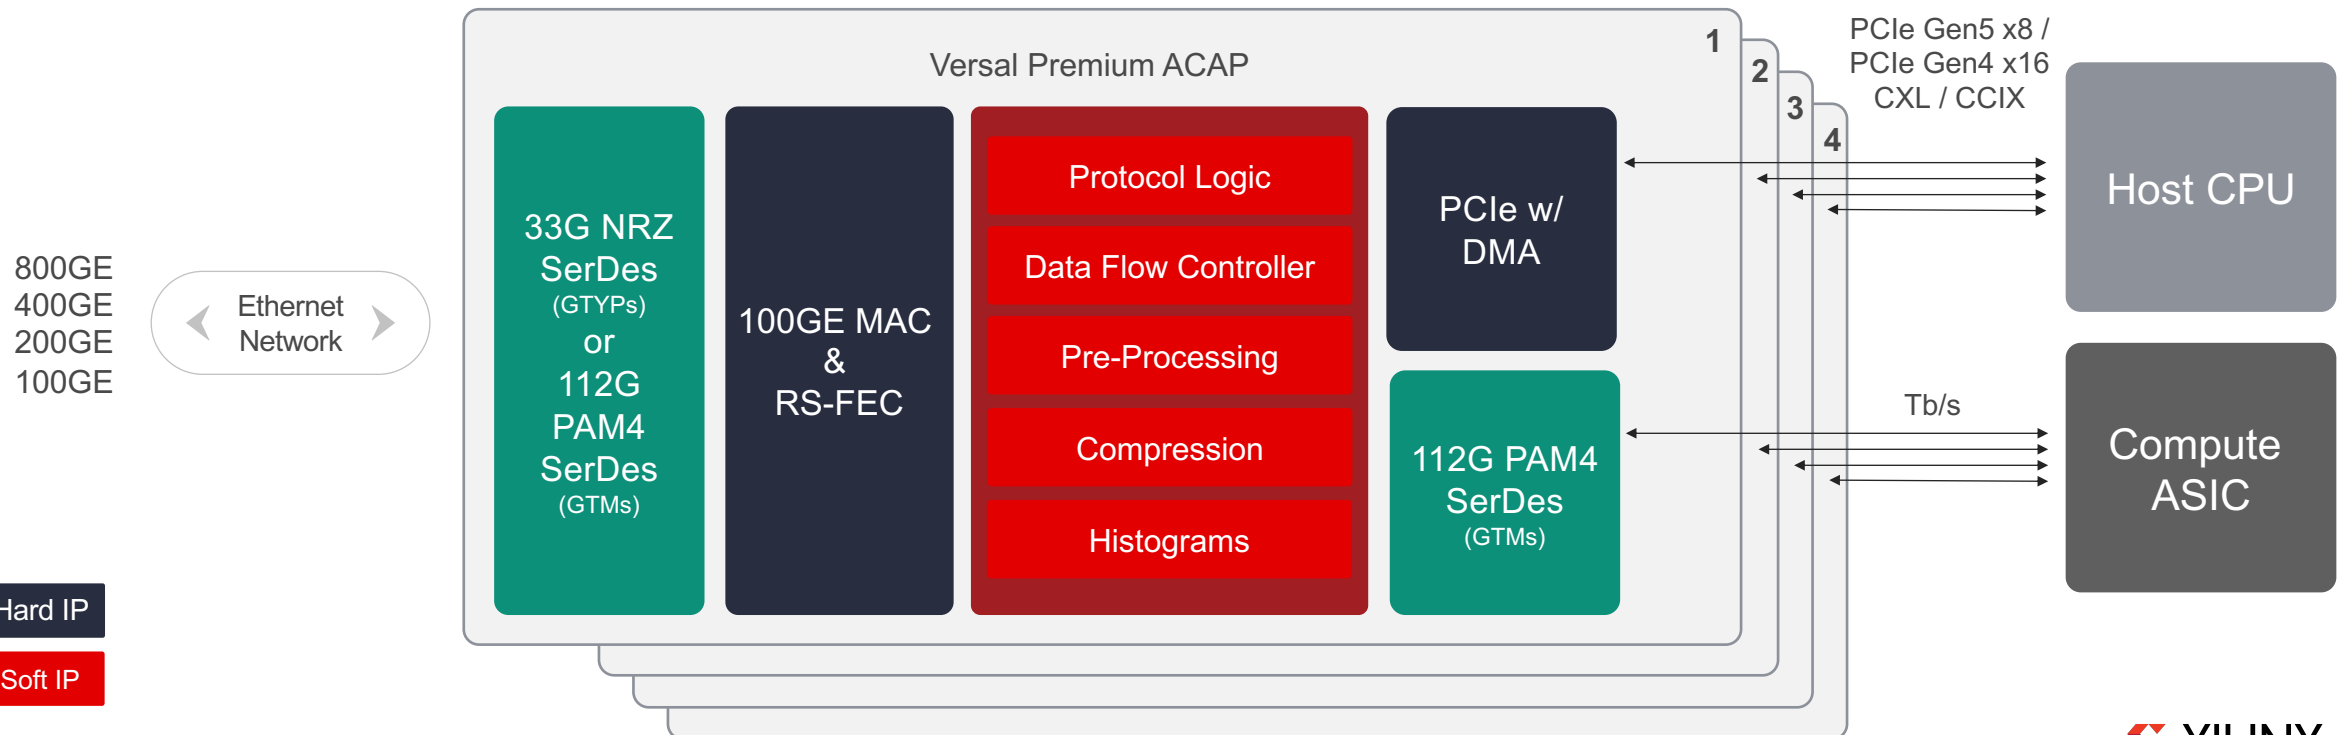
Adaptable Hardware

Adapts with changing pre-processing algorithms

In-Line Data Processing at hardware speed

Integrated PCIe® w/ CXL / CCIX

Dedicated interfaces for lower power and faster time to market



Hard IP

Soft IP

Use Case: 3.2Tb/s Capacity 800G L2-L3 Network Tester

Integrated 112G PAM4 SerDes

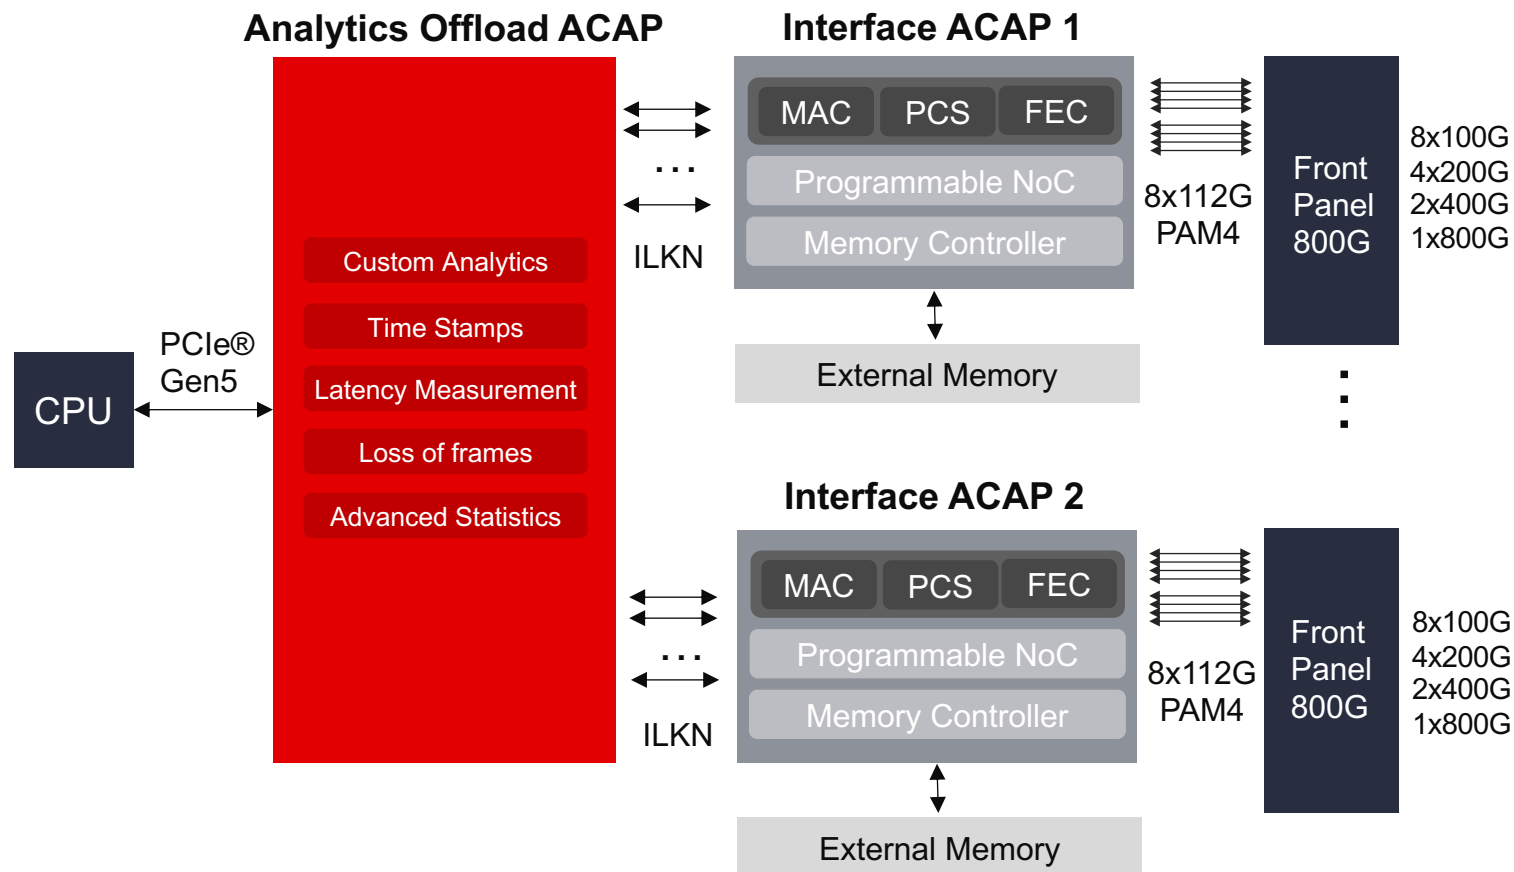
Supports emerging optics and standards
Eliminates external gearboxes

Integrated FEC


Accessible to custom logic for power
efficient error correction

Adaptable Hardware


Enables highly customized protocol logic



Summary



Versal Premium Device Family and Features



Manufacturing Technology	
Process Technology	TSMC 7nm FF
Interposer Technology	CoWoS
# Transistors	92 B

Max Performance	
Adaptable Engine Peak Perf	107 INT8 TOPS
DSP Engine Peak Perf	99 INT8 TOPS
DSP Engine Peak Perf	23 FP32 TFLOPS
Regfile BW	315 TB/s
OCM BW	123 TB/s
SerDes BW	9.7 Tb/s
LPDDR4 BW	136 Gb/s

Architectural Resources	
LUT6	3.4M
Logic Cells	7.4M
BRAM [Mb]	174
URAM [Mb]	717
LUTRAM [Mb]	103
Total On-die Memory [Mb]	993
DSP	14.4K
GTY (28G)	68
GTM (PAM4 56G, up to 112G)	140
XPIO	702
HDIO	122
CPM5 (PCIe5 w DMA/CCIX/CXL)	yes
PCIE G5 cores	8
Interlaken cores	3
MRMAC cores	8
DCMAC cores	8
High Speed Crypto cores	5
Hard DDR4 MC controllers 64/72b	4

Summary

- ▶ Versal Premium adds High Speed Networking Engines
 - 600G Ethernet, 600G Interlaken, 400G HS Crypto, 112G PAM4 SerDes
- ▶ Hardened SoC shell up and running before logic programming
 - High BW NoC with hardened DDR4, Processor, PCIeG5 with DMA/CCIX/CXL, 112G SerDes
- ▶ Hardened NoC enables dynamic linking & loading of kernels
- ▶ 4th Gen SSIT die stacking enables very large devices economically
- ▶ Result is an adaptable SoC with large bandwidth & compute density

Acknowledgements

- ▶ Wired Silicon Engineering
- ▶ Silicon Architecture
- ▶ Product Planning
- ▶ Product Marketing Teams
- ▶ Office of CTO



Thank You

