

Versal Premium ACAPs: Breakthrough Integration of Networked IP on a Power-Optimized, Adaptable Platform

The Versal Premium ACAP provides breakthrough heterogeneous integration, very high-performance compute, connectivity, and security on an adaptable platform with a minimized power and area footprint.

ABSTRACT

In every market across the world, continuous demand for higher bandwidth metro and core networks scales beyond what today's technologies can support. Data center-centric scientific, enterprise, and consumer applications demand more efficient, higher performance compute that scales beyond what traditional technologies can match. Discrete solutions cannot meet performance, thermal, and bandwidth requirements.

The Versal[®] Premium ACAP provides breakthrough heterogeneous integration, very high-performance compute, connectivity, and security on an adaptable platform with a minimized power and area footprint. This highly integrated platform allows users to focus on their unique core competencies and novel algorithms, rather than designing connectivity and memory infrastructure, to achieve the earliest possible time to market.

This white paper describes the unique Versal Premium ACAP features, reviews selected use cases, and outlines the benefits of the integrated hardware and software platform for hardware developers, software developers, and data scientists alike.

Introduction

Market Forces in the Always-On, AI-Everywhere, Everything-On-Demand Era

In markets across the world, continuous demand for higher bandwidth scales beyond what today's technologies and form factors can support. The demand is for more efficient, pervasive compute that scales beyond what CPU and GPU technologies can match.

This bandwidth demand is driven by explosive growth of video streaming services over general-purpose networks, wide consumer adoption of many always-connected IoT devices per person, smart infrastructure, and the continued movement towards cloud services. In addition, with the 5G rollout, endpoint bandwidth is planned to increase by an order of magnitude per connected device, and the number of connected devices is expected to increase by up to 100X the number of wireless devices in service today. This massive increase in mobile bandwidth calls for a significant increase in capacity of backhaul, metro, and core networks with limited additional physical space.

The forecasted investments for 5G-driven network core infrastructure is shown in [Figure 1](#).




Figure 1: 5G Wireless Deployment Drives Wired Infrastructure to Hit Bandwidth Constraint

Today's switch fabrics typically implement 12.8Tb/s of capacity in a 1RU box, which requires 32x400G OSFP or QSFP-DD front panel ports. Next-generation switch capacity is anticipated to grow to 25.6Tb/s in the same form factor. To support the higher capacity on already crowded front panels, the same OSFP and QSFP-DD form factors will be required to carry 800Gb/s per module. Since the electrical side of the optical modules support eight serial links, next-generation optical modules will need at least 100Gb/s per serial link. This will force system vendors to adopt a new SerDes technology.

An explosive increase in data center-centric scientific, enterprise, and consumer services drives demand for more efficient, pervasive compute that scales beyond what traditional architectures can support. Traditional processing architectures with discrete processing, connectivity, and

acceleration ICs consume too much power and do not provide sufficiently low latency to satisfy consumer expectations.

Due to multi-year development cycles of core communications equipment and compute infrastructure, system vendors need to start designing their next-generation systems now.

System Design Challenges

Systems designers face many challenges as they architect the next-generation of highest bandwidth, most secure networks, and compute infrastructure. Three key challenges include:

- An insatiable demand for higher bandwidth networks within a thermally and spatially constrained environment
- Demand for scalable, tightly integrated ICs, processing, and storage that provide very low latency to satisfy the consumers' user experience
- Requirement for the fastest time to market to win a majority of market share

To meet the earliest time to market requirements for very high bandwidth applications, system designers must plan support for emerging interface standards. Some of these emerging standards are 800GE, 400G ZR, Flex Ethernet (FlexE), and 112G Direct Attach Cable. Since standards can change as they are being ratified, it is important to develop adaptable solutions—or face significant risk of time-consuming hardware redesigns, which lead to missed opportunities for market share.

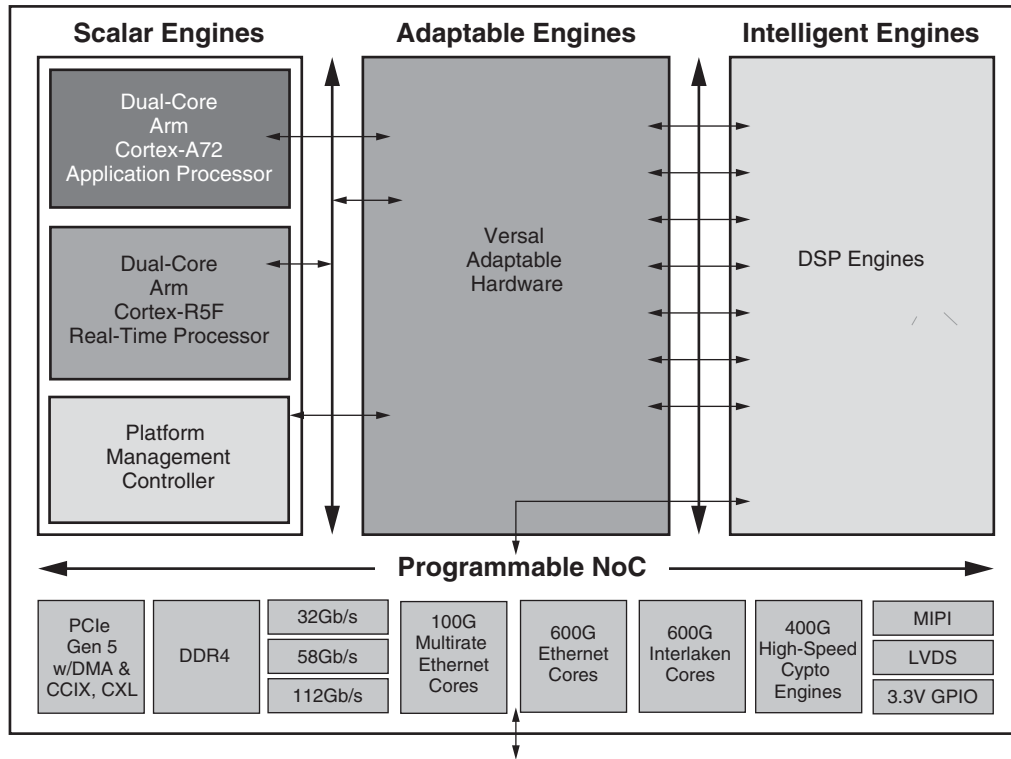
Both data center and Telco equipment already have power and thermal constraints that are difficult to satisfy with current technology. Since many functions like storage, processing, system controllers, high-speed network interfaces, and accelerators are often implemented in discrete ICs, there is a negative impact on complexity and power consumption. Designers are challenged to meet board space, power, and thermal constraints.

For data center and machine learning (ML) acceleration applications, the algorithms and tasks being performed by hardware change very frequently, especially in agile development environments. This creates the risk of having to replace hardware at an unacceptable rate. In addition, data center and ML applications have a wide variety of workloads. As a result, it is imperative that the same accelerator hardware can be used for a wide variety of workloads in data center and ML applications, or the data center is forced to develop a wide variety of accelerator modules. This inefficiency affects R&D resourcing and creates a fractured acceleration solution that is difficult to support and maintain, and increases sourcing overhead and OPEX.

Traditional discrete hardware implementations increase system complexity, power budgets, space budgets, BOM component count, and cost. This complexity increases R&D time and hence time to market for high bandwidth, high compute density, secure applications.

Versal Premium ACAP

The 7nm Versal Premium ACAP is a highly integrated hardware/software platform built on a foundation of architectural elements leveraged from the Versal AI Core and Versal Prime series. This foundation consists of next-generation Adaptable Hardware (programmable logic), Intelligent Engines (DSP), Scalar Engines (Arm® CPUs and RPs), a programmable network on chip (NoC), external memory controllers, and a variety of peripheral interfaces and I/O. Unique to the Versal Premium series are 112Gb/s PAM4 transceivers, robust Ethernet and Interlaken connectivity, High-Speed Crypto (HSC) Engines, and integrated blocks of PCIe® Gen5 with DMA, CCIX, and CXL. See [Figure 2](#).



WP519_02_030920

Figure 2: Versal Premium ACAP Block Diagram

The Versal Premium series key figures of merit are listed in [Table 1](#).

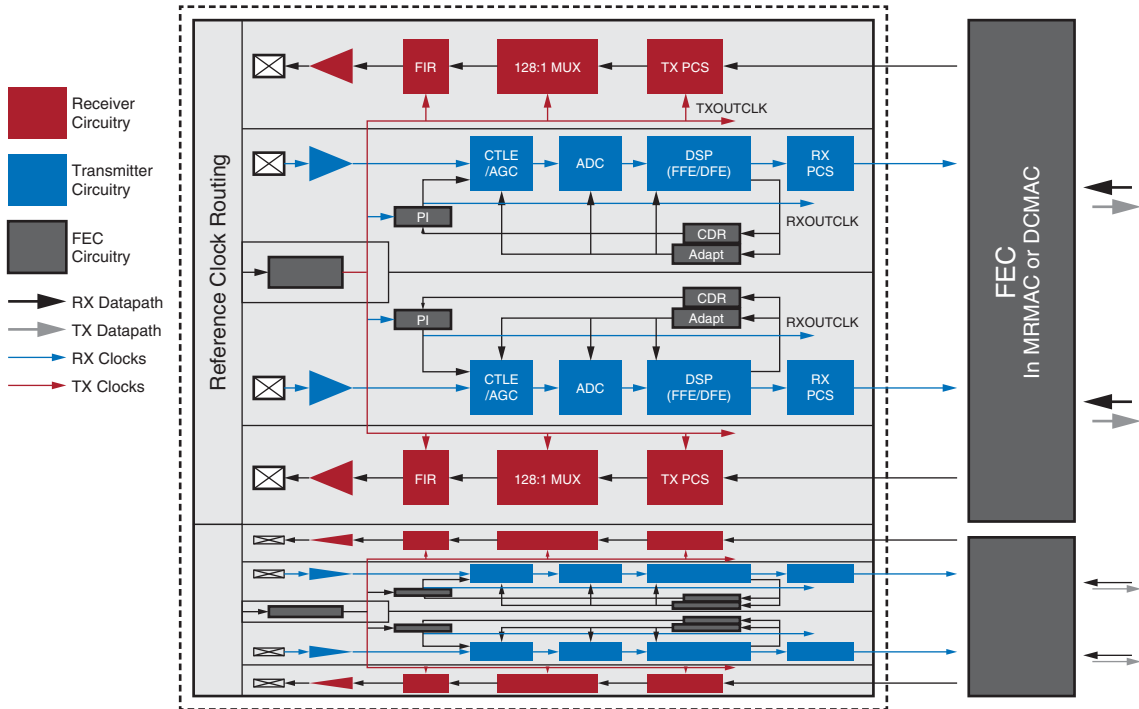
Table 1: Key Figures of Merit

Feature	Resources	Unit	Value
Adaptable Engines	System Logic Cells	# (K)	7,352
	NoC Cross-Sectional Bandwidth	Tb/s	2.5
Memory	Total On-Chip SRAM Bandwidth	Tb/s	982
	DDR4 Memory Bandwidth	Gb/s	102.4
	LPDDR4 Memory Bandwidth	Gb/s	136.5
Intelligent Engines	DSP Engine Peak Perf - INT8	TOPs	99.0
	DSP Engine Peak Perf - FP32	TFLOPs	23.1
I/O	Transceiver Bandwidth	Tb/s	9.04
Connectivity Throughput	Interlaken Throughput	Gb/s	1800
	Ethernet Throughput	Gb/s	5000
	Cryptographic (AES-256) Throughput	Gb/s	1600
Connectivity Ports	50G Ethernet Ports	#	16
	100G Ethernet Ports	#	50
	200G Ethernet Ports	#	21
	400G Ethernet Ports	#	7
	PCIe Gen5 Ports	#	10

Unique Versal Premium ACAP Features

GTM Transceivers (112Gb/s PAM4)

The Versal ACAP's GTM transceiver is an evolution of the PAM4 transceiver available in 16nm Virtex® UltraScale+™ FPGAs, with targeted improvements to reach the next line rate node of 112Gb/s. The transmitter is a 4-tap driver: 1 main driver, 1 tap of pre-emphasis, and 2 taps of post-emphasis. The receiver is an ADC-based design, meaning that after an analog AGC and CTLE stage, the data is sampled, and the remainder of the equalization takes place via the integrated DSP Engine. This DSP Engine has been specifically improved to implement 1 tap of decision feedback equalization (DFE) and 31 taps of feed forward equalization (FFE). The additional taps beyond those in the prior generation GTM transceiver help overcome the increasingly subtle channel impairments that can corrupt data at these phenomenal rates. While using PAM4 signaling to reach the highest rates, the GTM transceiver can use NRZ signaling as well at half rates. Combined with internal dividers, this means the GTM transceiver can support 10G NRZ rates, 25G NRZ and PAM4 rates, 50G NRZ and PAM4 rates, and 112G PAM4.



WP519_03_022520

Figure 3: GTM Transceiver (112G PAM4 SerDes) Block Diagram

GTYP Transceivers (32.75G NRZ)

Versal Premium ACAPs also include GTYP transceivers (32.75G NRZ), which are an evolution of 16nm Virtex UltraScale+ GTY transceivers. The Versal platform’s GTYP transceiver supports PCIe Gen5 and has been tuned to meet the challenging PCIe Gen5 electrical specifications. More than 100 other protocols are also supported. This means that the same GTYP transceiver can be used for a Gen5x8 IC pushing 256Gb/s of data, a 25G-LR Ethernet connection using all 15 taps of DFE, or a SATA Gen1 connection at 1.5Gb/s. The Fractional N feedback divider in each of the five PLLs in a GTYP transceiver enables each of these widely disparate protocols to use the same reference clock.

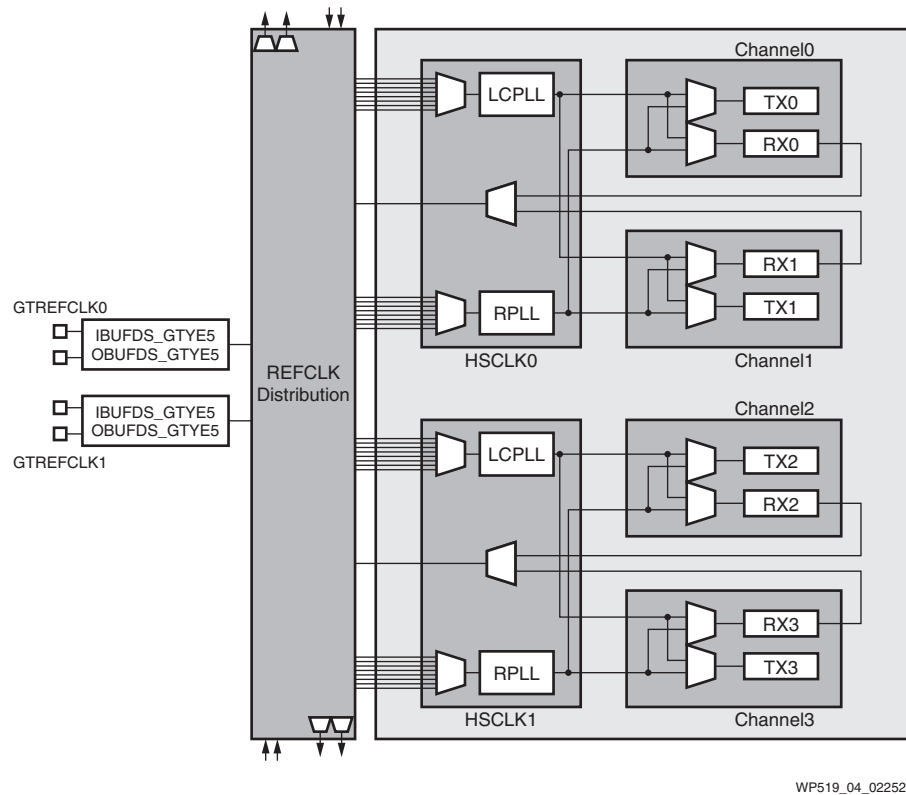


Figure 4: GTYP Transceiver (32.75G NRZ) Block Diagram

Highly Integrated, Programmable, Flexible Connectivity - 600G Ethernet, PCIe Gen5, and 600G Interlaken

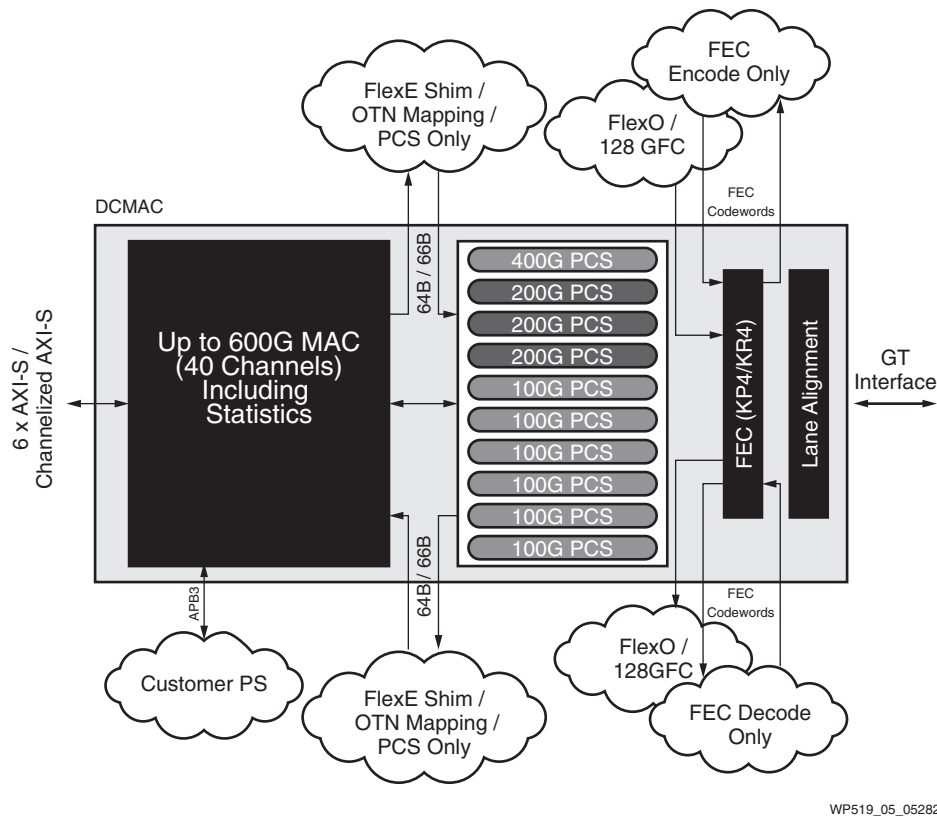
Versal Premium ACAPs integrate the most essential connectivity features for communications, data center, and test & measurement applications. Integrated connectivity features include 600Gb/s Ethernet MAC and PCS, 600G Interlaken with FEC, multirate 100Gb/s Ethernet MAC and PCS, and PCIe Gen5. These integrated connectivity features are combined in various ratios across the Versal Premium series. The inclusion of hardened connectivity cores greatly reduces the logic fabric and power consumption required to implement standardized interfaces. This gives the user the ability to implement high bandwidth connectivity functions in a smaller, lower power device.

The integrated connectivity cores are programmable for a wide variety of configurations to suit many applications. The cores, with granular flexibility, provide the ability to bypass submodules within Ethernet and Interlaken processors. Users have the ability to insert their own code at various points in the cores.

Integrated 600Gb/s Channelized, Multirate Ethernet Subsystem

To support the latest standardized and emerging high-speed networking interfaces, each of the Versal Premium ACAP's integrated 600G Ethernet MAC/PCS subsystems (DCMAC) provide up to 600G of channelized Ethernet bandwidth that can be configured for various rates as shown below. The DCMAC implements a time-sliced MAC block that supports up to 40 channels for data processing and statistics. The module also implements KP4 and KR4 forward error correction (FEC)

blocks that enable remote correction of errors encountered during transmission. Support for FlexE v2.0 standard is included along with sub-nanosecond accuracy timestamping. See [Figure 5](#).



WP519_05_052820

Figure 5: DCMAC - Integrated 600G Channelized Multirate Ethernet Subsystem Block Diagram

The DCMAC is a highly capable integrated block and includes the following key features:

- Multiple configurations
 - 1 x 400GE
 - 3 x 200GE
 - 6 x 100GE
- Time-sliced 600GE MAC block
 - Supports 40 channels for data processing and statistics
- Integrated FEC
 - Optional integrated RS-FEC (KP4/KR4)
 - FlexO FEC for OTN
 - Optional FEC only mode
- FlexE support with shim path for future proofing
- IEEE Std 1588 Timestamping: Sub-nanosecond accuracy

Integrated 600Gb/s Interlaken High-Speed Chip-to-Chip Interface

To support the latest high-speed data transmission between FPGAs, ACAPs, ASICs, and ASSPs, Versal Premium devices integrate several 600Gb/s Interlaken interfaces. These hard IP cores support channelized interfaces with built-in flow control. The Versal Premium platform's GTM (112G PAM4) and GTYP (32.75G NRZ) transceivers allow PCB, backplane, and cable PHY layers. See [Figure 6](#).

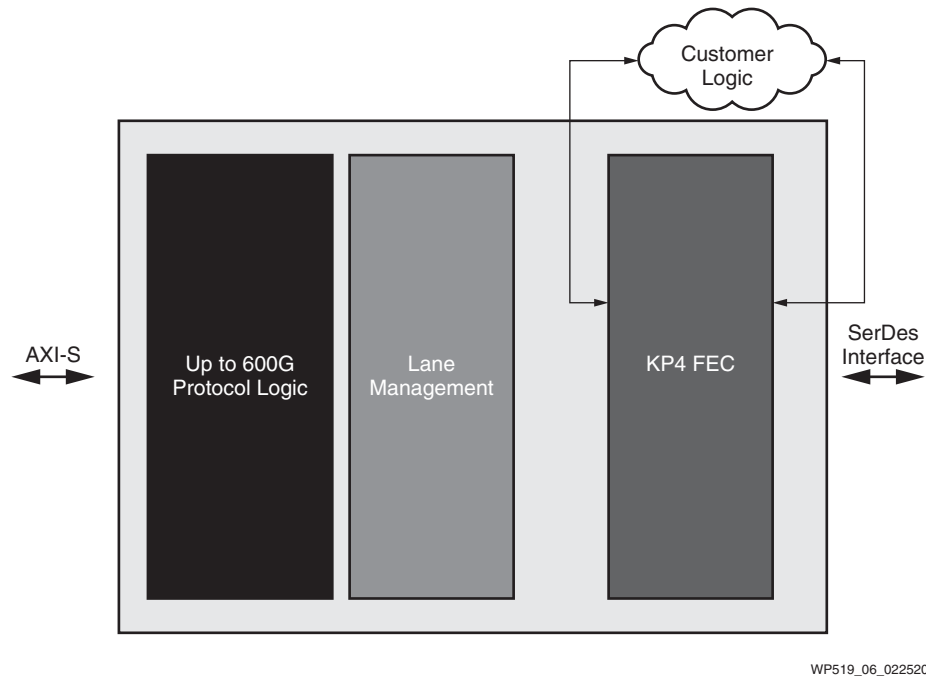


Figure 6: Integrated 600G Interlaken with FEC Block Diagram

Following is a feature summary.

- The transceivers have up to 600G of channelized bandwidth and adaptable rates, including:
 - 12 x 56.42G
 - 24 x 28.21G
 - 24 x 12.5G
 - Other rates and widths are programmable
- Flexible AXI-S User Interface: Configurable from 2048b to 512b data width
- Optional, integrated RS-FEC
 - 100G/50G RS-FEC (KP4)
 - FEC-Only mode
- Lane decommissioning for all configurations

Integrated Blocks with PCIe Gen5

Versal Premium ACAPs include two types of integrated blocks for PCIe:

- PCIe Gen5 (PL PCIE5) with Compute Express Link (CXL)⁽¹⁾
- PCIe Gen5 with DMA and CCIX (CPM5)

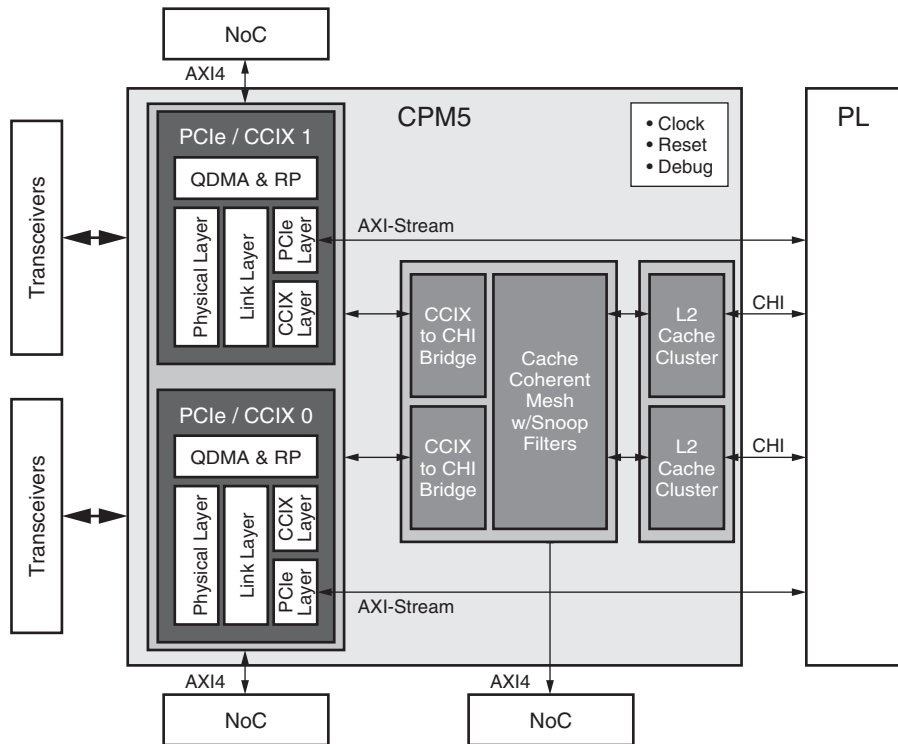
These integrated blocks, in conjunction with the new GTYP transceiver, support data rates from 32GT/s per lane (commonly called Gen5) through 16GT/s per lane (Gen4), 8GT/s per lane (Gen3), 5GT/s per lane (Gen2), and 2.5GT/s per lane (Gen1).

The CPM5 block contains two sub-blocks of controllers for PCIe designs, each with an integrated queue-based direct memory access (QDMA) controller. Each controller is independently customizable and enables direct connection to the NoC as well as the Adaptable Hardware's programmable logic fabric. The CPM5 also includes sub-blocks to support Cache Coherent Interconnect for Accelerators (CCIX) designs, which enable cache-coherent acceleration with any CCIX-enabled processor for compute acceleration applications. The CPM5 block has access to 16 GTYP transceivers.

Maximum link width configurations supported for CPM5 are 2*Gen5x8 or Gen4x16. The links can interoperate with devices that are compliant with any revision of the PCI Express Base Specification. The Versal architecture has the addition of switching and bridging application support in the CPM,

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

along with Endpoint and Root Port support from previous architectures. The CPM5 block diagram is in [Figure 7](#).



WP519_07_030420

Figure 7: CPM5 PCIe with DMA & CCIX Block Diagram

The PL PCIE5 is an enhancement to the PL PCIE4 deployed in Versal AI Core ACAPs and Versal Prime ACAPs, upgrading the protocol revision. Like the CPM5, PL PCIE5 blocks work with the GTYP transceivers to support the full range of data rates. Maximum link width configurations supported for PL PCIE5 are Gen5x4, Gen4x8, and Gen3x16. The links can interoperate with devices that are compliant with any revision of the PCI Express Base Specification. New in the Versal architecture is the addition of support for switching and bridging applications in PL PCIE blocks, in addition to the Endpoint and Root Port support from previous architectures. The PL PCIE5 block diagram is in [Figure 8](#).

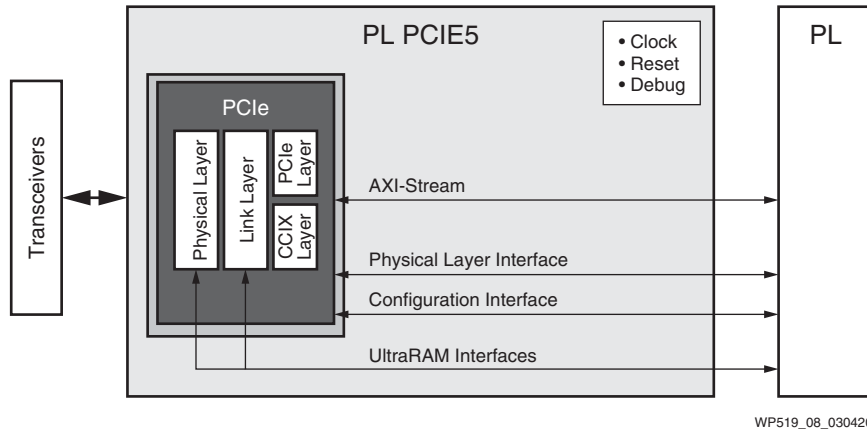


Figure 8: PL PCIe5 Block Diagram

High-Speed Crypto (HSC) Engine

Each of the High-Speed Crypto (HSC) Engines in the Versal Premium ACAP implements an AES-GCM-256/128 engine that provides up to 400Gb/s of bulk encryption capability on up to 40 channels, which can be connected to the 600G Ethernet subsystem. Up to four HSC engines are available to provide up to 1.6Tb/s of line-rate encryption.

Each HSC core supports both MACsec and IPsec, and the bandwidth can be channelized into 1x400G, 2x200G, or 4x100G channels with up to 128 security associations (SA) per 100G. Additional SAs can be implemented with soft logic in the Adaptable Hardware. See [Figure 9](#).

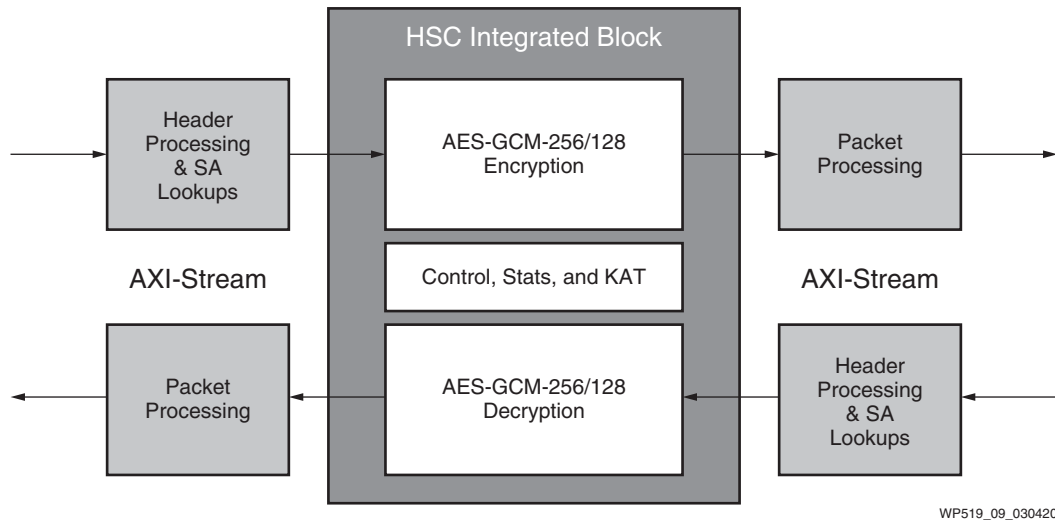


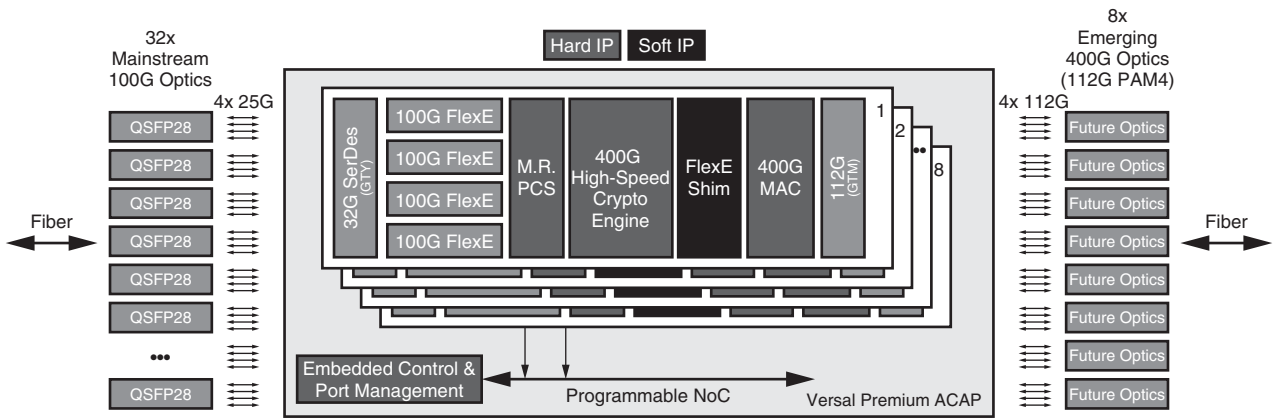
Figure 9: Integrated 400G High-Speed Crypto Engine

Selected Use Cases

Data Center Networking: 3.2Tb/s Data Center Interconnect Platform

Much of today's exploding bandwidth demand ultimately has an impact on the data centers that are providing content and services, and process data and transactions for a wide variety of users and applications. Not only is there tremendous bandwidth and services growth on the links in and out of data centers, but also on the data center interconnect (DCI) links that connect geographically distributed data center sites. DCI equipment must be able to accommodate a variety of server-side and transport-side optics and protocols with the ability to adapt to emerging and evolving standards in a secure, cost efficient platform.

Versal Premium ACAPs enable implementation of a 1RU system or a single card that provides 3.2Tb/s capacity with 1.6Tb/s line rate encryption that can support a wide variety of standardized and emerging protocols and optics. With the breakthrough integration of connectivity and cryptographic cores, this is possible in a single Versal Premium ACAP. The implementation, which bridges mainstream 100G optics and emerging 400G optics, is shown in the block diagram. See [Figure 10](#).



WP519_10_030520

Figure 10: 3.2Tb/s Data Center Interconnect Block Diagram

In this example, 32 channels of 100G FlexE are implemented server-side with 4x25G NRZ connections to QSFP28 optics. On the line side, 8 channels of 400G Ethernet are implemented with 4x112G PAM4 connections to future optics, which are anticipated to be in QSFP-DD or OSFP form factors. The hardened Ethernet blocks are used to implement 100G FlexE with RS-FEC and multirate PCS on the server side as well as 400GE with KP4 FEC on the line side. 1.6Tb/s of AES256 line rate encryption is implemented in 4x400G HSC Engines. Dual-core Arm Cortex® -A72 processors are used to implement control and port management functions for efficient and readily programmable system and network management. All of these features are implemented in the Versal Premium ACAP's integrated hard IP blocks, which enable ASIC-class performance per watt, and minimize power and the silicon area dedicated to these functions.

The remaining feature that requires logic fabric is the FlexE shim, which enables bonding, channelization, and sub-rating on one or more Ethernet PHYs. For example, the shim allows a 2x100GE interface to carry independent 10GE, 25GE, and 50GE channels.

Wired Networking: 2.4Tb/s Client Interface Card

For transport applications that service common packet-based client interfaces, Versal Premium ACAPs provide an efficient method of bridging and encapsulating a wide variety of clients' digital traffic and services into industry-standard OTN wrappers. In this example, a 2.4Tb/s capacity client interface card is implemented in the Versal Premium ACAP by leveraging integrated channelized Ethernet, Interlaken, and FEC as well as the 112G and 58G PAM4 GTM transceivers and 32.75G GTYP transceivers. The Adaptable Hardware logic fabric is used to implement ODUk mapping, segmentation and reassembly (SAR), and ODUk overhead functions. See [Figure 11](#).

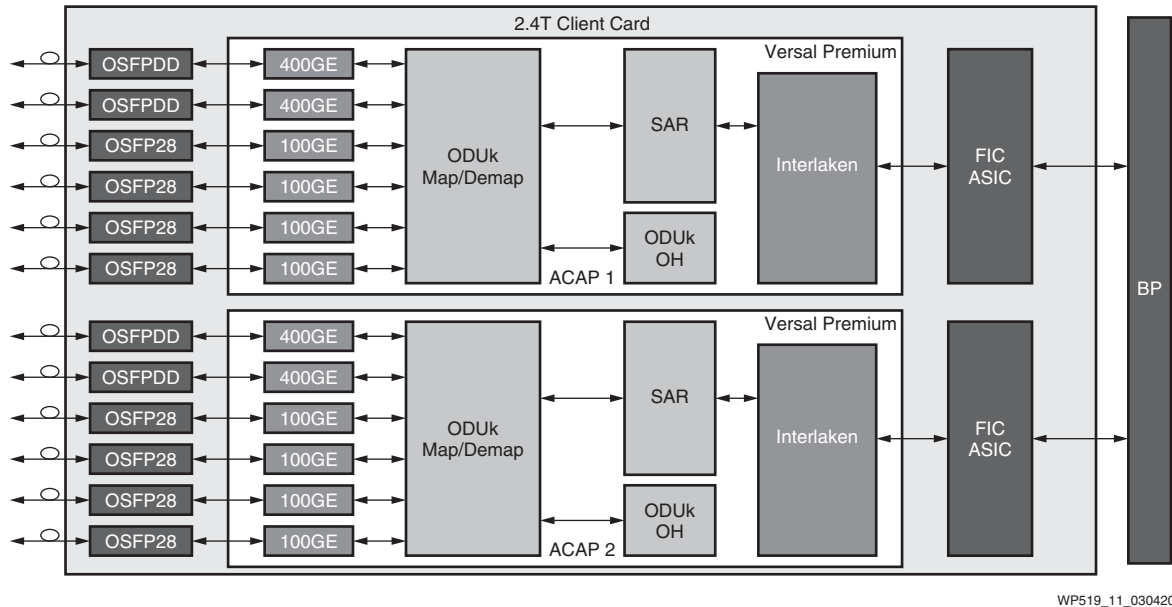


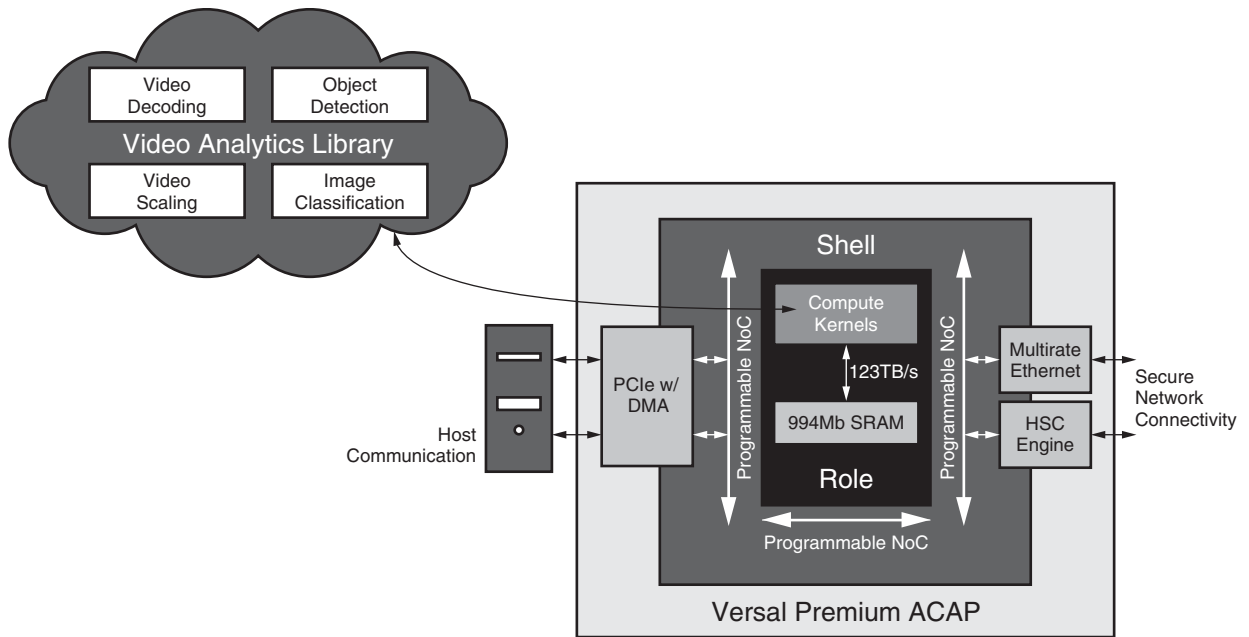
Figure 11: 2.4Tb/s Client Interface Card Block Diagram

Benefits of Versal Premium ACAP Implementations of Transport Client Cards

Client cards can readily adapt to dynamic client interface requirements. New client-facing optics standards can be adopted, such as 800GE or 400GE implemented with 112G PAM4 channels, by upgrading optical modules while leveraging the same hardware platform. Implementation of Ethernet, FECs, and Interlaken in the Versal Premium ACAP's dedicated hard IP enables ASIC-class power efficiency while freeing up logic fabric resources for mapping, overhead, and SAR functions as well as for implementing differentiating functions.

Retail Analytics: Adaptable Accelerator for Video Content Analytics

Retail businesses are facing a multi-faceted set of challenges for loss protection (LP), cost effective real-time inventory management, and maximizing revenue through marketing personalization. A data-driven video content analytics system can be used to efficiently identify LP risks, offer automated, real-time and actionable insight into inventory, and provide the ability to tailor the customer experience in a way that maximizes sales. The Versal Premium ACAP offers the ability to build a video analytics solution on a single platform for identification, extraction, and classification of video metadata. See [Figure 12](#).



WP519_12_030520

Figure 12: Retail Video Content Analytics Accelerator

In this use case, the Versal Premium ACAP's shell and role capabilities allow designers to focus on application and ML algorithm implementation in a role while utilizing the ACAP's off-the-shelf shell connectivity and security features. The role is implemented in the ACAP's Adaptable Hardware and DSP Engines as software programmable compute kernels that are loaded with functions sourced from a video analytics library. These libraries can accelerate core smart retail functions like object detection, image classification functions, and video encoding, decoding, and scaling. A wide variety of CNNs can be implemented in Xilinx's deep learning processing units (DPUs). The DPUs are also in the Versal platform's Adaptable Hardware, which allows the DPU to be continuously optimized and updated, keeping up with the fast changing world of AI. Most importantly, up to 1Gb of on-chip SRAM is available immediately adjacent to the compute kernels that provide up to 123TB/s of memory bandwidth, which is critical for AI acceleration of CNNs, RNNs, MLPs, and others. In addition, ACAPs readily accommodate a batch size of 1. This eliminates the memory bottleneck and batch size throughput limitations that are endemic to GPU and CPU-based architectures and allows processing up to 13,000 images/sec for Resnet50.

The Versal Premium ACAP's shell provides the platform's connectivity and cryptography. Host communication is provided by the integrated blocks for PCIe that are capable of Gen5 rates and are paired with integrated DMA engines. Network connectivity is handled by the Ethernet subsystem, which provides up to 5Tb/s of channelized bandwidth. The HSC Engines enable secure communication with AES-256/128 encryption, MACsec, and IPsec. All shell and role functions are interconnected via the programmable NoC, which is capable of up to 2.2Tb/s of cross-sectional bandwidth.

A video analytics system built on the Versal Premium ACAP is a robust, adaptable, very high-performance platform that provides actionable intelligence for retailers to minimize loss and maximize revenue.

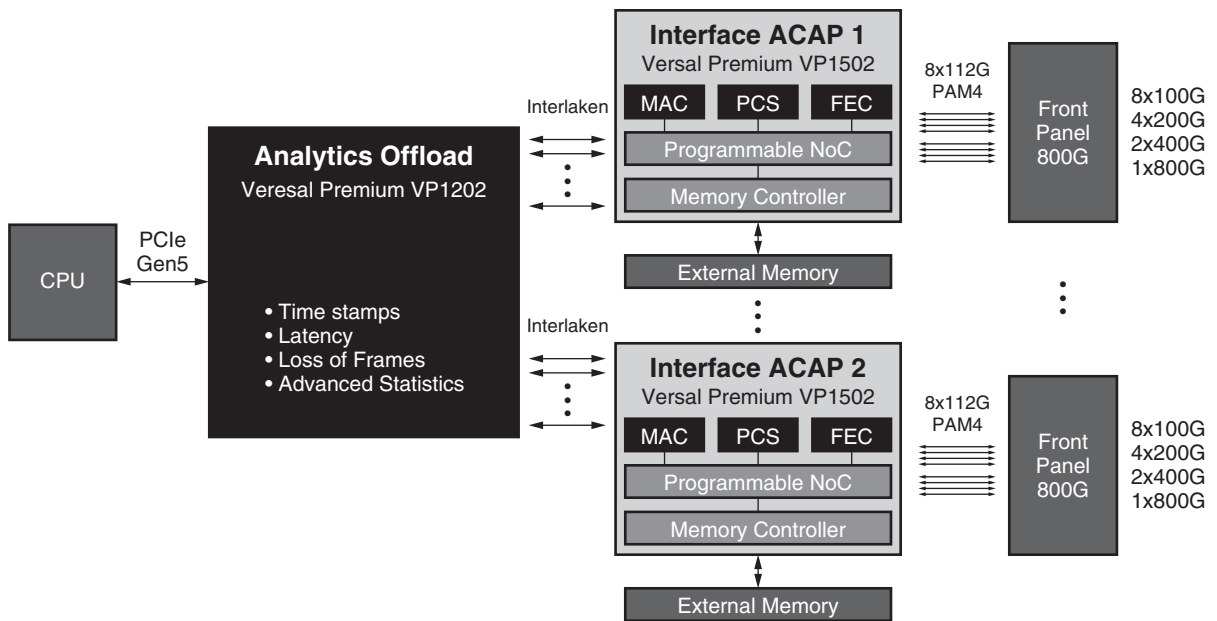
Network Test: 3.2Tb/s Capacity 800G L2-L3 Network Tester

Cloud and enterprise data centers use servers, switches, routers, appliances, and other equipment from multiple vendors to build their infrastructure. To ensure the equipment interoperates and can appropriately handle the intended traffic, the equipment needs to be tested at specific OSI layers. There are typically two levels of testing: Layer 2-3 (L2-L3) testing for equipment like Ethernet switches, bridges, and routers; and layer 4-7 (L4-L7) testing for application-level testing of equipment like routers, video servers, media gateways, and firewalls. While Versal Premium ACAPs are well suited for tester implementation at all layers, here the focus is on 800Gb/s L2-L3 test equipment.

L2-L3 test equipment requires highly customized protocol logic to implement non-standard functions that enable fault injection (often called "jamming") and analysis of all features in a given protocol. These testers are space and power constrained, so there is a drive for integration of BOM components in power efficient ICs that allow the ability to accommodate highly customized protocol logic.

Versal Premium ACAP's Adaptable Hardware is very well suited for implementing this highly customized protocol logic because it is highly programmable and has the logic capacity to implement very large protocol logic blocks.

Figure 13 shows a block diagram of an 800G L2-L3 network tester, which supports channels that are 800G with up to 4 channels for a capacity of 3.2Tb/s.



WP519_13_030420

Figure 13: 3.2Tb/s Capacity 800G L2-L3 Network Tester

In [Figure 13](#), the Interface ACAPs implement PCS, MAC, and FEC functions. Each Interface ACAP can implement either one stream of 800G or multiples of lower speed streams such as 8x100G or 2x400G. On the back end, an Analytics Offload ACAP implements the following functions at line rate:

- Frame Loss %
- Round-trip Latency
- Packet Jitter
- Inter-Arrival Time
- Sequence Errors
- Time Stamps

The Interface ACAPs can be implemented in a mid-density Versal Premium device while the Analytics Offload ACAP can be implemented in the smallest Versal Premium device. This partitioning represents a typical tester architecture, and some vendors can choose to implement all the functionality in one large Versal Premium ACAP.

Benefits of Versal Premium ACAPs for L2-L3 Network Testers

Integration of 112G PAM4 eliminates the need for external gearboxes (sometimes called bitmuxes) to bridge the front panel interface to mainstream rates and protocols within the tester. This enables vendors to support the latest protocols and fastest rates. It also reduces board complexity, device count, and BOM cost. Using Versal Premium ACAPs in place of gearboxes enables full visibility into the serial signal's eye opening and margin, and enables analysis of various types of jitter that gearbox-based implementation cannot provide.

Versal Premium ACAP's integrated KP4 FEC blocks enable support for 400G+ protocol FEC. Because the FEC is implemented in the main die, it is readily accessible by user logic and built-in Ethernet MAC and PCS blocks via standard routing resources. This architecture provides two distinct advantages when compared to a chiplet implementation, where the MAC/PCS and FEC are tied directly to the SerDes. First, custom protocols have direct access to the FEC without borrowing the connectivity to the chiplet that would otherwise be used for the high-speed SerDes datapath. Second, the power consumption and associated thermal dissipation is distributed throughout the fabric, eliminating hot spots and enabling maximum performance throughout the device.

Integrated (hardened) blocks for PCIe Gen5 with DMA enable power-efficient PCIe interfaces to be implemented with no impact on the logic resources. This maximizes logic fabric availability for implementing custom protocol logic.

Versal Premium Series Benefits

Versal Premium ACAP integration enables massively reduced cost, complexity, and time to market while enabling customers to more acutely focus on their unique core competencies.

Versal Premium ACAPs integrate a wide variety of critical features in networked, ASIC-like integrated blocks that free up logic resources for implementing customers' key, differentiating, custom accelerators and logic in Adaptable Hardware. Integration enables a significantly reduced footprint, power budget, and component count by minimizing the number of ICs required to be implemented in hardware for many applications.

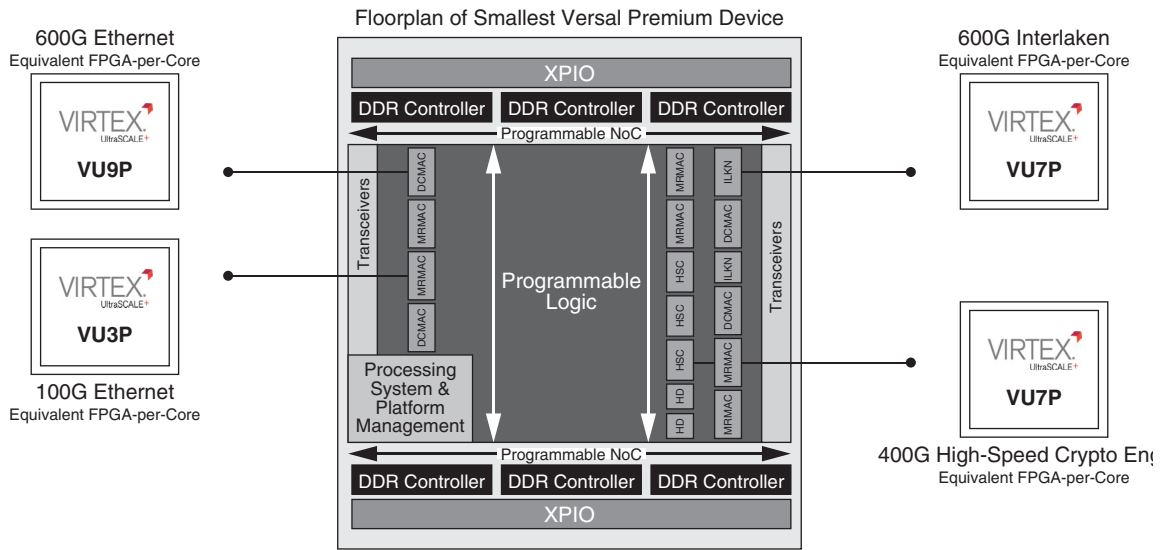
More importantly, Versal Premium ACAPs are pre-engineered for efficient data movement in, out, and through the device, many high-speed communication interfaces, custom-hierarchy and off-the-shelf memory subsystems, security processing, application and real-time processing, and software-controlled platform management. All of these subsystems can be offloaded to hard IP via the programmable NoC.

Versal Premium ACAP's massive, networked hard IP integration enables a greatly simplified system design that frees architects and R&D teams to focus on optimally implementing their own key algorithms and protocols—not the infrastructure—as they strive to create their next big thing. This in turn enables the most rapid commercialization of the user's core competencies in an adaptable accelerator with massive data movement and processing capacity within a best-in-class performance and bandwidth-per-watt envelope.

22X Equivalent Logic Capacity

The scale of hard IP integration in Versal Premium devices enables a 22X increase in compute density versus Xilinx's previous generation of high-end FPGA, the 16nm Virtex UltraScale+ device. In this example, the equivalent logic resources of Virtex UltraScale+ FPGAs are compared to the Versal Premium ACAP VP1802's hardened IP cores. To implement the 600G channelized Ethernet MAC, 100G multirate Ethernet MAC, 400G HSC Engine, and 600G Interlaken with FEC, a design would require 22 Virtex UltraScale+ FPGAs. The Versal Premium ACAP VP1802 implements all of these functions in hard IP while providing nearly 2X the logic capacity of the Virtex UltraScale+ VU13P with a dramatically reduced power profile. See [Figure 14](#).

Multi-FPGA Functionality & Integration in Smallest Device



Note: Equivalent logic density represents all available configurations of integrated Ethernet, Interlaken, and Cryptographic cores.

WP519_14_030520

Figure 14: Integrating Equivalent Logic of 22 Virtex UltraScale+ High-End FPGAs into a Single Versal Premium VP1802 ACAP

112G PAM4 Transceivers Doubles Bandwidth Density and Reduces Latency by Half

Core, metro, and DCI networking applications that require 100Gb/s of bandwidth or more are constantly driven for more efficient use of front panel rack space. Business, commercial, and consumer applications demand the lowest possible latency for the best user experiences. Using 112G PAM4 transceivers for core, metro, and DCI infrastructure allows a doubling of bandwidth density per port and a 50% reduction in latency for transmitting a given payload of data versus 16/14nm's 58Gb/s PAM4 technology. Higher port density reduces OPEX via power savings and rack space savings and allows a doubling of bandwidth per unit volume in Telco and data center applications while preserving legacy infrastructure. Minimized latency through higher bandwidth allows applications to be more responsive, which drives customer satisfaction and an improved user experience. This helps mitigate latency impacts when interconnecting geographically distributed data centers. See [Figure 15](#).

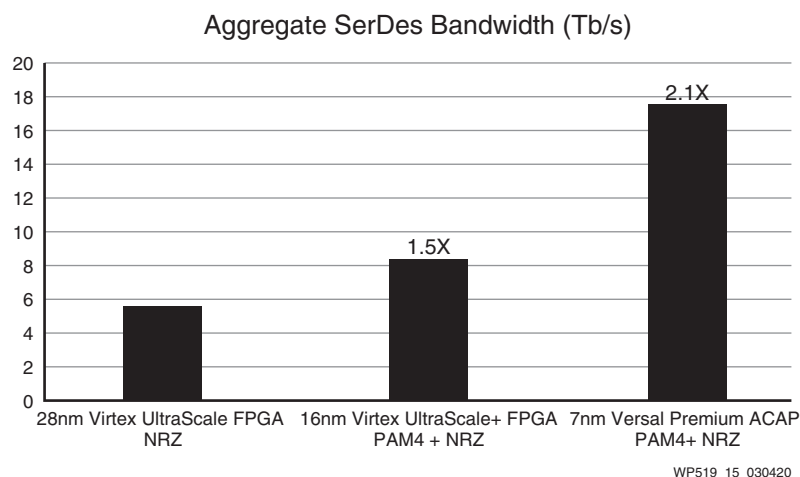


Figure 15: Aggregate Xilinx SerDes Bandwidth per Node

Scalability + Adaptability: One Platform for a Diverse and Dynamic Optics Market

When planning system and card designs that integrate optical interfaces, designers are faced with trade-offs between supporting the latest highest bandwidth and density optics standards versus support for a variety of standard optics already in production and widely adopted. In addition, designers must accommodate access rates from 1G to 10G in systems that are engineered for multi-Terabit capacity. Versal Premium ACAPs allow system software to scale services from the lowest rate interfaces, such as 1GE, to the highest rate interfaces such as the upcoming 800GE on the same hardware/software programmable platform with a unified, consistent user interface. R&D teams can execute more efficiently due to the common look and feel of the ACAP hardware/software platform. The Versal Premium ACAP's inherent adaptability enables integration of the latest protocols and optical standards into existing infrastructure even if the systems are deployed prior to the standards being finalized.

DSP Flexibility and Performance Boost

DSP-related applications are very prevalent, touching almost all market segments from 5G communications and test, avionics, and hyperscale data center AI inference. One thing these applications require in common is the type of underlying functions and compute, predominantly multiply and accumulates functions (MACs) for FFTs and matrix multiply operations. Versal Premium ACAPs support MACs using many different data types, e.g., integer, single, or mixed precision floating point, and can be dynamically tuned for precision, accuracy, and power consumption. The Versal Premium ACAP's DSP resources can be reprogrammed to adapt to workloads that change over time, or as algorithmic implementations evolve.

The DSP Engines in the Versal architecture are based on DSP58 blocks. The architecture builds on previous generations of DSP48 slices to deliver increased performance versus Virtex UltraScale+ FPGAs for the most commonly used operands, e.g., INT8, 32-bit floating point, 18-bit complex math, and more. These data types are all now natively supported in the Versal platform's DSP Engines. Versal Premium devices have up to 14,000 low-power DSP blocks that combine high speed with small size. The DSP resources enhance the speed and efficiency of many applications beyond

digital signal processing, such as wide dynamic bus shifters, memory address generators, wide bus multiplexers, and memory-mapped I/O registers.

To demonstrate the increased DSP performance, the following is a comparison between the previous generation high-end FPGA with the most DSP resources, Virtex UltraScale+ VU13P FPGA, versus the Versal Premium ACAP with most DSP resources (VP1802). See [Figure 16](#).




Figure 16: DSP Performance of Versal Premium ACAP (VP1802) versus Virtex UltraScale+ FPGA (VU13P)

Summary

For metro and core network and leading-edge compute acceleration applications, Versal Premium ACAPs offer breakthrough integration of networked, power-optimized IP cores on a platform that adapts to new standards and algorithms. The Versal Premium ACAP platform brings very high-performance compute fabric immediately adjacent to up to 1Gb of on-chip SRAM accessible at rates up to 123TB/s for next-generation compute applications. Integrated connectivity provides up to 5Tb/s of Ethernet connectivity for metro and core network applications, as well as for fastest possible data movement for data-hungry compute applications. The Versal Premium platform's Adaptable Hardware provides over 7 million system logic cells, allowing users to implement highly differentiated logic and the latest algorithms. This heterogeneous platform provides users with pre-engineered connectivity, processors, custom memory hierarchies, security, and customizable compute fabric that enable the fastest time to market to capture market share in the most challenging compute and networking applications.

For detailed descriptions of the full description of features that are common to the Versal architecture, refer to [DS950: Versal Architecture and Product Data Sheet: Overview](#).

Revision History

The following table shows the revision history for this document:

Date	Version	Description of Revisions
10/12/2022	1.0.1	Updated Table 1 and Figure 5 .
03/10/2020	1.0	Initial Xilinx release.

Disclaimer

The information disclosed to you hereunder (the “Materials”) is provided solely for the selection and use of Xilinx products. To the maximum extent permitted by applicable law: (1) Materials are made available “AS IS” and with all faults, Xilinx hereby DISCLAIMS ALL WARRANTIES AND CONDITIONS, EXPRESS, IMPLIED, OR STATUTORY, INCLUDING BUT NOT LIMITED TO WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, OR FITNESS FOR ANY PARTICULAR PURPOSE; and (2) Xilinx shall not be liable (whether in contract or tort, including negligence, or under any other theory of liability) for any loss or damage of any kind or nature related to, arising under, or in connection with, the Materials (including your use of the Materials), including for any direct, indirect, special, incidental, or consequential loss or damage (including loss of data, profits, goodwill, or any type of loss or damage suffered as a result of any action brought by a third party) even if such damage or loss was reasonably foreseeable or Xilinx had been advised of the possibility of the same. Xilinx assumes no obligation to correct any errors contained in the Materials or to notify you of updates to the Materials or to product specifications. You may not reproduce, modify, distribute, or publicly display the Materials without prior written consent. Certain products are subject to the terms and conditions of Xilinx’s limited warranty, please refer to Xilinx’s Terms of Sale which can be viewed at <http://www.xilinx.com/legal.htm#tos>; IP cores may be subject to warranty and support terms contained in a license issued to you by Xilinx. Xilinx products are not designed or intended to be fail-safe or for use in any application requiring fail-safe performance; you assume sole risk and liability for use of Xilinx products in such critical applications, please refer to Xilinx’s Terms of Sale which can be viewed at <http://www.xilinx.com/legal.htm#tos>.

Automotive Applications Disclaimer

XILINX PRODUCTS ARE NOT DESIGNED OR INTENDED TO BE FAIL-SAFE, OR FOR USE IN ANY APPLICATION REQUIRING FAIL-SAFE PERFORMANCE, SUCH AS APPLICATIONS RELATED TO: (I) THE DEPLOYMENT OF AIRBAGS, (II) CONTROL OF A VEHICLE, UNLESS THERE IS A FAIL-SAFE OR REDUNDANCY FEATURE (WHICH DOES NOT INCLUDE USE OF SOFTWARE IN THE XILINX DEVICE TO IMPLEMENT THE REDUNDANCY) AND A WARNING SIGNAL UPON FAILURE TO THE OPERATOR, OR (III) USES THAT COULD LEAD TO DEATH OR PERSONAL INJURY. CUSTOMER ASSUMES THE SOLE RISK AND LIABILITY OF ANY USE OF XILINX PRODUCTS IN SUCH APPLICATIONS.

© Copyright 2020–2022 Xilinx, Inc. Xilinx, the Xilinx logo, Alveo, Artix, Kintex, Spartan, UltraScale, Versal, Virtex, Vivado, Zynq, and other designated brands included herein are trademarks of Xilinx in the United States and other countries. AMBA, AMBA Designer, Arm, Arm1176JZ-S, CoreSight, Cortex, and PrimeCell are trademarks of Arm in the EU and other countries. PCI, PCIe, and PCI Express are trademarks of PCI-SIG and used under license. All other trademarks are the property of their respective owners.