XILINX.
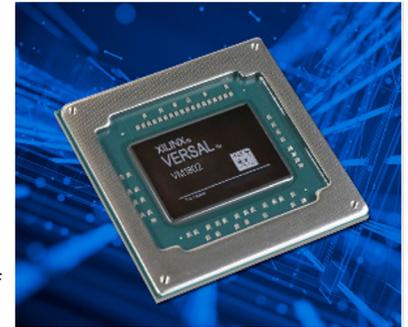
Storage Acceleration with Versal Prime Series

> Highest compute density storage acceleration device in its class

> The only 7nm/10nm adaptable hardware package deployable for EDSFF

> Scalable acceleration and custom datapath features

CHALLENGE

The latest smart storage devices for enterprise and data center deployments seek to increase the inline compute capacity per storage drives and scale the number of drives per storage rack—offloading increased functionality from CPUs and accelerating system-level compute. The industry is now moving beyond traditional 2.5-inch flash and hard disk drives in server and storage systems to new enterprise & data center small form factor (EDSFF) cards for deployment in 1U racks to meet these demands.

The area and power requirements for next-generation storage form factors limit what can be deployed (118.75mm x 33.75mm, 25W total system power with asymmetric enclosure). There is a direct trade-off between the area and power consumed by a controller / accelerator device and the number of media components that can be integrated within EDSFF specifications. To create viable enterprise storage accelerators that run at PCIe® line rates, adaptable hardware must meet strict footprint and power envelope requirements while delivering enough compute density. FPGAs that rely only on traditional programmable logic are challenged to satisfy these competing requirements, making them undeployable for many next-generation systems.

SOLUTION: VERSAL PRIME SERIES FOR COMPUTATIONAL STORAGE

The Versal™ Prime adaptive compute acceleration platform (ACAP) is a highly integrated, multicore, heterogeneous device that balances hardened IP cores with adaptable hardware to deliver the necessary compute capability, power efficiency, and flexibility needed to implement a wide range of storage workloads while conforming to EDSFF and other common storage form factor standards.

Highest Compute Density Storage Accelerator in its Class

Versal Prime series devices provide 2.3X compute density[1] vs. competing 10nm FPGAs as well as breakthrough integration of hardened IP, including connectivity cores and high-bandwidth interconnect, delivering superior performance/watt in a small form factor while enabling future hardware adaptability.

The Only 7nm/10nm Hardware Adaptable Device for EDSFF Deployment

With package dimensions as small as 31mm x 31mm, Versal Prime ACAPs meet the form factor constraints prescribed by EDSFF specifications—in contrast to competing 10nm FPGAs whose smallest package dimensions exceed board dimensions. Coupled with dense I/O for DDR4 and NAND flash connectivity along with PCIe Gen4 compliance, Versal Prime devices are ideal accelerators for computational storage.

Scalable Acceleration and Custom Datapath Features

Capable of diverse computational storage workloads due to their heterogeneous compute engines, Versal Prime ACAPs also scale to Versal AI Edge devices for AI/ML inference and image processing functionality.
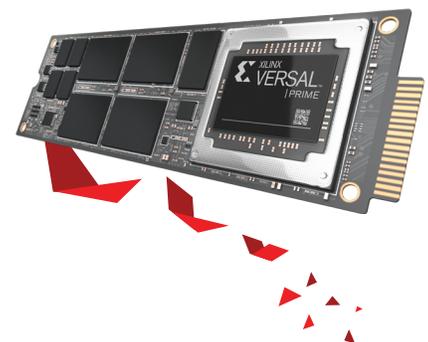
2.3X
Compute Density[1]

6.9GB/s
Read Rate per SSD

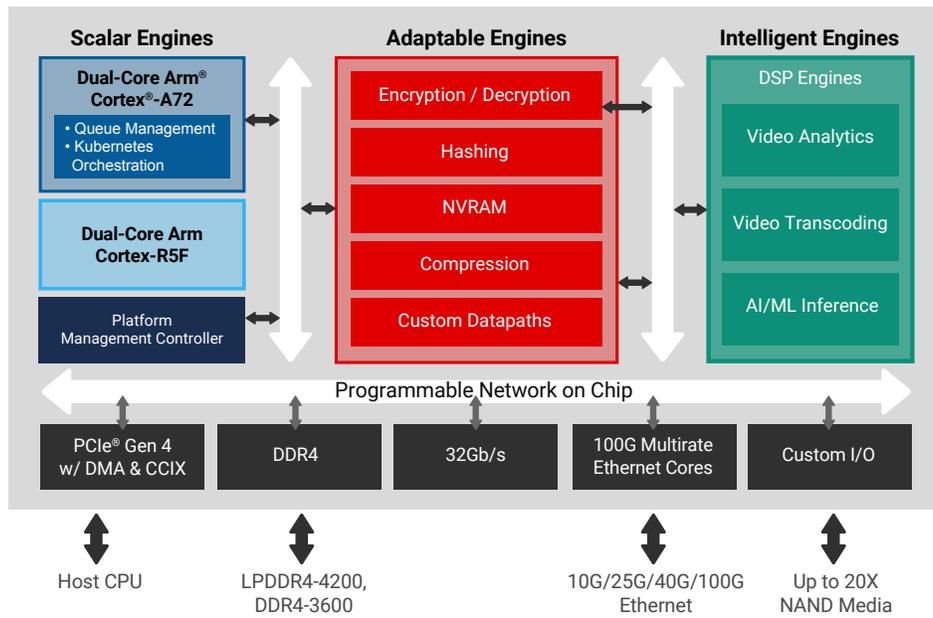The Only 7nm/10nm Hardware Adaptable Device Deployable for EDSFF E1 SSDs

1: Versal Prime VM1402 ACAP vs. Intel Agilex AGF014 FPGA (logic density / mm^2)

XILINX.

VERSAL ACAP IMPLEMENTATION

A Device to Enable Superior Inline Acceleration for Next-Generation Computational Storage

A Versal Prime ACAP provides heterogeneous compute engines for diverse storage computation, a hardened shell for off-the-shelf connectivity to compute infrastructure, hardware adaptability to support evolving algorithms and custom datapaths, and a form factor and I/O density ideal for EDSFF and other standard storage form factor deployments. A storage accelerator implemented with a Versal Prime VM1402 device delivers read rates of up to 6.9GB/s, consumes 17 watts for a typical storage workload, and provides connectivity to DDR memory, NAND flash, or other components.
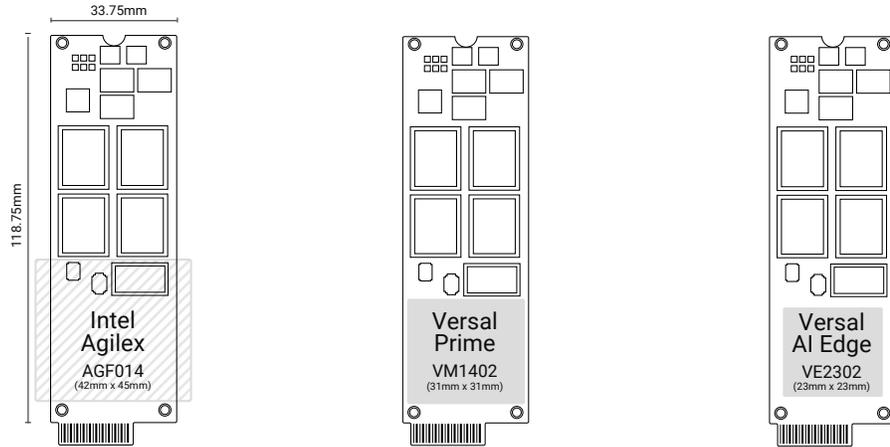
Versal Prime ACAP

Scalar Engines	Adaptable Engines	Intelligent Engines
Dual-Core Arm® Cortex®-A72 • Queue Management • Kubernetes Orchestration	Encryption / Decryption	DSP Engines
Dual-Core Arm Cortex-R5F	Hashing NVRAM Compression Custom Datapaths	Video Analytics
Platform Management Controller		Video Transcoding AI/ML Inference

Programmable Network on Chip

PCIe® Gen 4 w/ DMA & CCIX	DDR4	32Gb/s	100G Multirate Ethernet Cores	Custom I/O
Host CPU	LPDDR4-4200, DDR4-3600		10G/25G/40G/100G Ethernet	Up to 20X NAND Media

PLATFORM HIGHLIGHTS

Adaptable Engines	> 565k look-up-tables (LUTs) in a small 31mm x 31mm package > Adaptable to diverse storage workloads, including compression, hashing, NVRAM, and more > Enable custom encoding schemes and other datapath features with low latency
Intelligent (DSP) Engines	> Variable fixed- and floating-point DSP compute > Ideal for video analytics, video file transcoding, and AI/ML inference workloads
Scalar Engines	> Arm processing subsystem for queue management and Kubernetes orchestration > Platform management controller for security, power management, and bitstream management
Programmable Network on Chip (NoC)	> Seamlessly integrates all engines and key interfaces > Simplifies kernel and IP placement, reducing soft logic needed for connectivity > Streamlines programming experience for software and hardware developers
Integrated Shell	> Comprises hardened host interface, programmable NoC, and Scalar Engines > Ensures streamlined device bring-up and connectivity to off-chip interfaces—making the platform available at boot > Delivers pre-engineered timing closure, power savings, and logic resource savings
I/O and DDR4	> 648 single-ended I/O in a 35mm x 35mm package > 324 single-ended I/O in a 31mm x 31mm package > Scalable connectivity to DDR memory, NAND flash, or other components

XILINX

Adaptable. Intelligent.

BENCHMARK

Enterprise DC Storage Form Factor (EDSFF) Comparison

As shown below, the closest competing adaptive hardware device (Intel Agilex AGF014 FPGA) is not offered in package dimensions smaller than 42mm x 45mm, making it undeployable for many enterprise data center storage form factors. Through the integration of hard IP (integrated DMA engine, programmable network on chip, PCIe Gen4, and DDR4 memory controllers), Versal Prime devices deliver superior performance/watt in significantly smaller form factors, with scalability to Versal AI Edge series for machine learning acceleration.

	INTEL AGILEX AGF014[2]	VERSAL PRIME VM1402[3]	VERSAL AI EDGE VE2302[5]
Logic Density[1]	487K ALMs + DDR	565K LUTs + CPM[4] + NoC + DDR	328K LUTs + NoC + DDR
Device Package Size	42mm x 45mm	31mm x 31mm	23mm x 23mm
EDSFF Form Factor	NOT DEPLOYABLE	DEPLOYABLE	DEPLOYABLE
COMPUTE STORAGE FUNCTIONS			
Encryption	NOT DEPLOYABLE	■	■
Compression	NOT DEPLOYABLE	■	■
Hashing	NOT DEPLOYABLE	■	■
NVRAM Management	NOT DEPLOYABLE	■	■
ML Acceleration	NOT DEPLOYABLE	1,696 DSP Engines	34 AI Engine-ML Tiles[5]

1: Storage acceleration functions at 6.9GB/s typically require ~300K LUTs or more

2: Intel Agilex AGF014-2340A FPGA package

3: See Versal ACAP Prime Series Product Selection Guide for full product specifications

4: VM1402 features a CPM4, offering integrated PCIe Gen4 with hardened DMA, eliminating the need to implement DMA in programmable logic

5: See Versal ACAP AI Edge Series Product Selection Guide for more details

Versal Prime VMK180 Evaluation Kit
www.xilinx.com/vmk180

TAKE THE NEXT STEP

> For more information on the Versal Prime series, visit www.xilinx.com/versal-prime

> To start designing with a Versal Prime Evaluation Kit, visit www.xilinx.com/vmk180

> To learn more about Xilinx computational storage solutions, visit www.xilinx.com/computational-storage

> To contact your local Xilinx sales representative, visit Contact Sales