



# Xsight Softens the DPU

*By Bob Wheeler, Principal Analyst*

*October 2024*

*Powering SmartNICs, the data-processing unit (DPU) has become nearly ubiquitous in the leading public clouds. Existing designs maximize power efficiency for a constrained feature set, and they require proprietary software tools. Xsight Labs aims to break this paradigm with its new E1 DPU, which promises the openness of an Arm server CPU. Xsight Labs sponsored the creation of this white paper, but the opinions and analysis are those of the author.*

### *Cloud Infrastructure Demands Programmability*

In the world of computing, there's a constant tension between generalization and specialization. Nowhere is this more evident than in AI, where everything from general-purpose CPUs to highly specialized accelerators can run established models. Feature velocity favors the former, whereas the latter maximize performance and power efficiency. Market forces indicate that GPUs are in the sweet spot, delivering the right blend of flexibility and performance for AI workloads. Cloud service providers must perform similar calculus when it comes to various other components of their infrastructure. An added dimension, however, is that the provider's architecture should be as transparent as possible to the end customer.

The world's leading public-cloud provider, Amazon has been advancing its Nitro system for more than a decade, delivering increasing network and storage bandwidth combined with advanced security and near-bare-metal compute performance for AWS customers. What began with a smart network interface card (SmartNIC), the Nitro system now comprises a network card, a storage card, a security card, and a highly optimized hypervisor. The hardware delivers transparent virtualization of network and storage interfaces to the customers' compute instances. Nitro stands apart not only for its completeness but also because Amazon has continued to execute across several generations of custom silicon (ASICs). Competitors have attempted to imitate Nitro, but none has delivered feature parity across multiple generations.

With the benefit of hindsight, we can now credit Amazon with establishing the market for data-processing units (DPUs). Although the acronym is newer, the company began deploying DPU-based SmartNICs for AWS infrastructure a decade ago. In the intervening years, both cloud providers and merchant-chip vendors have developed DPUs using a wide variety of architectures. AMD and Nvidia acquired DPU designs through their respective acquisitions of Pensando and Mellanox, whereas Intel codeveloped its DPU with Google. Microsoft developed generations of SmartNICs using field-programmable gate arrays (FPGAs) before acquiring DPU-startup Fungible in early 2023. Other merchant-DPU vendors have come and gone, having failed to secure ongoing business at a leading cloud provider.

Against this backdrop, it's fair to ask whether the market needs another DPU vendor. AMD, Intel, and Nvidia all provide DPUs that combine a programmable packet processing pipeline with a complex of Arm cores for exception and control-plane processing. Due to ship by 2025, AMD's "Salina" and Intel's "Mount Morgan" will support 400G Ethernet, keeping pace with both front-end networks for AI and compute-server demands. In a back-to-the-future moment, however, Xsight Labs will offer a highly programmable DPU reminiscent of the original SoCs used in Nitro SmartNICs a decade ago.

### *A Different Kind of DPU*

Figure 1 compares the architecture of Xsight's new E1 DPU with that of available merchant DPUs. The AMD and Intel designs include a packet processing pipeline that is programmable using the P4 language, whereas Nvidia reuses the microcoded pipeline from its ConnectX NIC chips. Although details vary, all three vendors' DPUs implement similar functions in their respective pipelines, such as packet parsing, match-action stages, and stateless offloads. All three also include Arm cores that,

logically, sit to the side of the data path. In the latest shipping generation, these DPUs include 16 cores that are either Cortex-A72/78 or Neoverse N1 and backed by multiple external-DRAM channels.

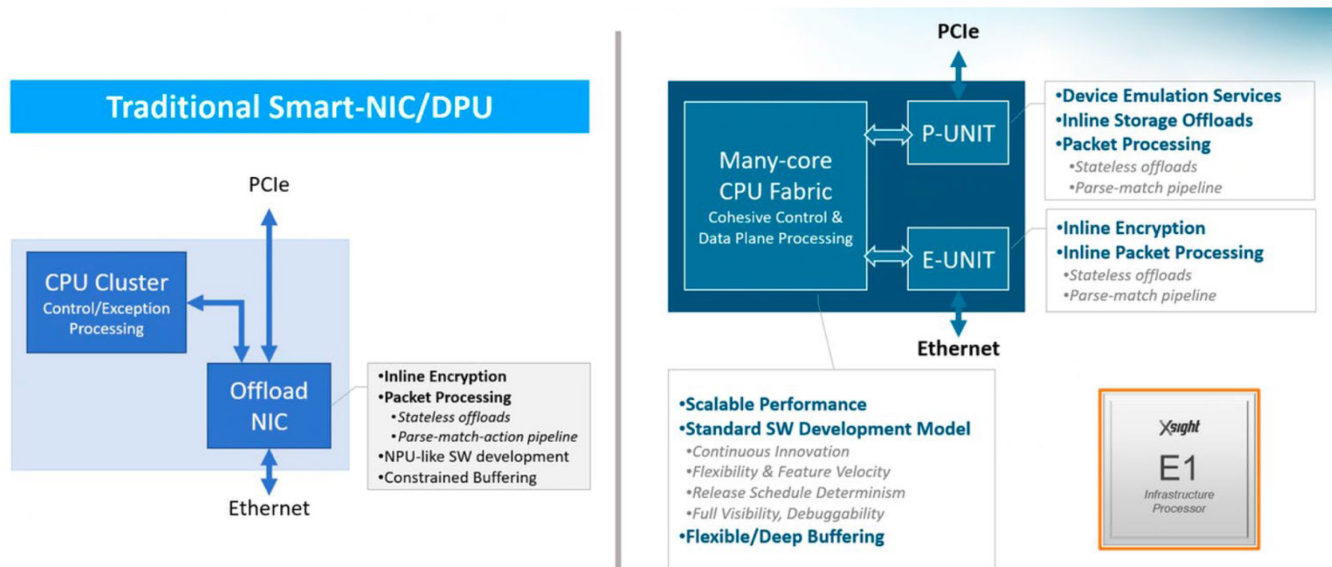


FIGURE 1. DPU ARCHITECTURE COMPARISON  
(Source: Xsight Labs)

The general theory of operation for these DPUs is that the Arm cluster touches only new packet flows, whereas the pipeline handles established flows. This approach is well proven for certain applications such as OVS offload, but new applications may exceed the pipeline's capabilities. In particular, stateful applications can quickly exceed whatever limited capacity the pipeline may have for tracking flows and other state metadata. When the CPU cluster must handle all traffic rather than just known flows, performance drops dramatically. Another customer concern with the available merchant DPUs is that they require either vendor-supplied data-plane code or programming using unfamiliar languages and tools. AMD and Intel's support for P4 reduces but does not eliminate this burden, as few engineers are familiar with the language. In addition, the P4 programmer must still understand the resource limitations of the underlying hardware.

Xsight's new DPU takes a different approach, one that emphasizes the use of standard software while still supporting 400G Ethernet port speeds. The E1 is a unique design that represents a hybrid of a DPU and an Arm server processor, implementing many DPU features but placing the Arm cluster directly in the data path. The SoC is designed to boot and run any OS that's Arm SystemReady compliant, primarily various Linux distributions like Sonic and Debian. Xsight will supply standard network drivers for its integrated Ethernet controllers (E-Unit), enabling open-source applications to run with little porting effort. The company's pre-silicon validation shows the E1 should run open-source SDN applications at 800Gbps even though its Arm cores will process every packet. This software-centric design will certainly be less power efficient than a purpose-built pipeline, but the E1 should still dissipate less power than one of the shipping 400Gbps DPUs built in an older process technology.

## E1 is Armed for SDN

The E1 SoC design starts with 64 Arm Neoverse N2 cores, as Figure 2 shows. In evaluating CPU cores, Xsight considered using the Neoverse E series commonly found in embedded designs, but the less powerful cores would have limited single-flow performance. Instead, the company optimized the N2 core's cache size and clock speed to suit the DPU application. The E1 will offer 512KB of L2 cache per

core for a total of 32MB. Xsight's performance estimates assume a 2.0GHz clock speed, whereas Arm specifies the N2 for operation up to 3.6GHz in a 5nm process. Arm's CMN-700 mesh fabric provides the coherent interconnect between the 64 cores, 32MB of system-level cache (SLC), and I/O. The SoC's external-memory subsystem comprises four 80-bit-wide DDR5-5200 interfaces with inline encryption. There's nothing extraordinary about the CPU cluster, which is a good thing when you want to boot and run standard OS distributions.

What sets the E1 apart from an Arm server processor is its E-Units and P-Units, which handle the Ethernet and PCI Express sides of the SoC, respectively. The E-Units are essentially sophisticated dual 400G Ethernet controllers that support RDMA as well as IPsec and PSP using AES-GCM encryption. Their external interface consists of 8x112Gbps PAM4 serdes using the silicon-proven design from Xsight's X2 switch chip. Each 400Gbps MAC can handle 1x400GbE, 2x200GbE, or 4x100GbE ports. The ingress packet processor includes programmable header parsing and a three-stage lookup engine for flow tables. DMA engines handle data movement into and out of memory, typically the system-level cache. The E-Unit also supports virtual functions so that multiple drivers – for example, a Linux kernel driver and a DPDK driver – can share a physical port.




FIGURE 2. E1 BLOCK DIAGRAM  
(Source: Xsight Labs)

The dual P-Units each start with 16 Gen5 serdes for the physical PCIe interface, providing 2Tbps of aggregate bidirectional bandwidth. Each P-Unit has four PCIe controllers that handle 1x16, 4x4, or 1x8+2x4 logical ports, and each port can be configured as a PCIe root (host) or endpoint (device). Most E1 applications will configure at least one x16 port as an endpoint. In a multihost SmartNIC design, the E1 can appear as a x16 device to two independent hosts, such as a pair of GPUs.

The DPU's PCIe-device-emulation hardware enables virtual-network and virtual-storage interfaces to the host, similar to how Amazon's Nitro virtualizes these resources. The hardware includes PCIe

configuration-space and MMIO-space emulation as well as transaction-layer packet (TLP) steering. The P-Unit also handles inline processing in the DMA path, such as packet processing for emulated network devices (e.g., VirtIO-net). The inline packet processing can offload parsing and lookups, whereas the inline storage processing can offload AES-XTS block encryption and block protection information (PI). These functions offload the Arm cluster, maximizing transaction rate while making emulation more efficient.

Xsight developed the E1 using 5nm process technology and plans to offer two versions (SKUs). The E1-64 will offer the full complement of 64 cores, 32MB of SLC, and 4xDDR5 channels at an estimated power of 90W (TDP). A 32-core version (E1-32) with half the memory resources should reduce power to 65W (TDP). The company announced the E1 in October 2024 as it achieved tape out, with customer sampling expected in 2Q25.

## Maximizing Off-the-Shelf Software

As discussed above, Xsight designed the E1 to run existing software, including open-source network operating systems. As Figure 3 shows, bump-in-the-wire applications, such as network appliances, can run an existing NOS such as SONiC and access the E1's E-Unit using standard DPDK and Linux drivers supplied by Xsight. The DPU should run any application software that runs on an Arm server without any modifications. In fact, Ampere's Altra CPU can serve as a development vehicle with similar processing power.

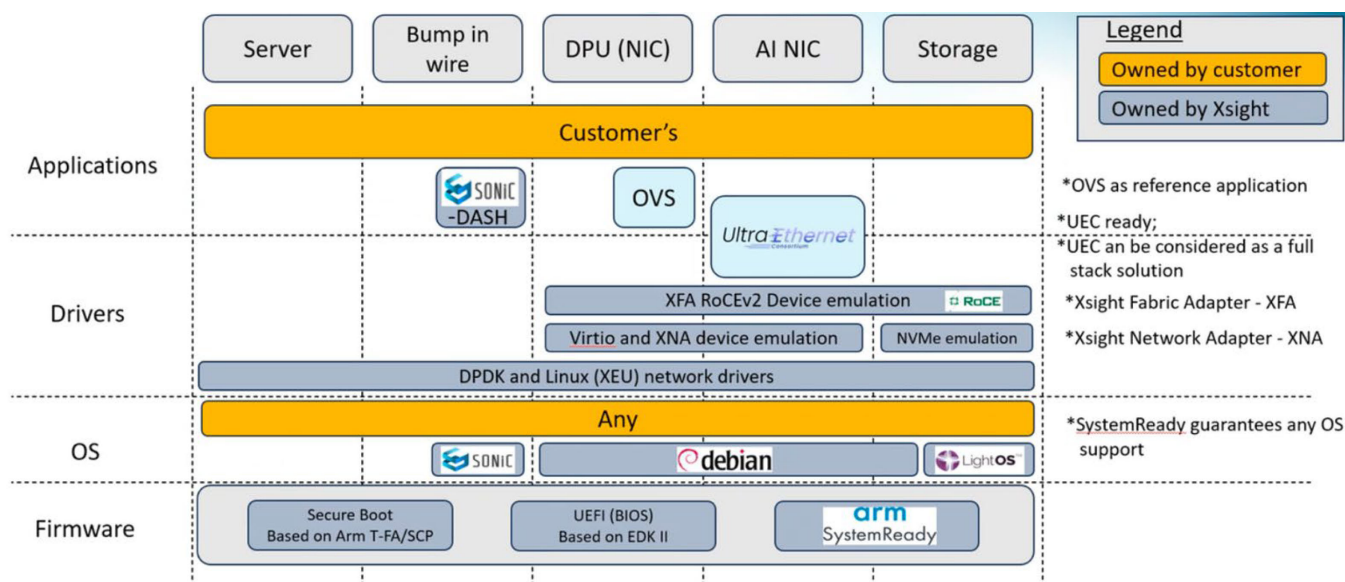


FIGURE 3. E1 SOFTWARE STACK  
 (Source: Xsight Labs)

For NIC applications, Xsight adds emulated-network-device drivers to virtualize the network interfaces presented to the server CPU (or host). In addition to standard network drivers, the E1 supports RDMA (RNIC) emulation, enabling data transfers directly into host memory. This capability is particularly important for latency-sensitive AI applications, as it avoids double buffering received packets. Xsight will initially support the RoCEv2 protocol, but it's also a member of the Ultra Ethernet Consortium and plans to support that body's new RDMA transport protocol (UET). Given UET is an immature protocol, the company believes it's uniquely positioned to adapt its design without silicon

revisions as the protocol evolves. For storage virtualization, Xsight will supply standard NVMe-over-Fabrics drivers for TCP and RDMA that emulate a local SSD and initiate NVMe-oF transactions.

An example workload that demonstrates the E1's capabilities is SONiC-DASH (Disaggregated API for SONiC Hosts), an open-source project started by Microsoft for use in Azure. DASH is intended to improve the performance and scale of flow processing for SDN using SmartNIC hardware. It defines a set of overlay and underlay APIs using the Switch-Abstraction Interface (SAI). Using hardware emulation, Xsight has run the DASH Hero benchmark and shown the E1 design passes the test at 800Gbps. It says a single E1 can run the DASH data plane at this speed while also running the SONiC NOS. Customers could run the same software on different platforms, such as an E1-based smart switch or a standard Arm or x86 server.

Xsight offered some low-level benchmarks to help customers estimate the performance of their specific workloads. For DPDK applications, it rates the E1's TestPMD packet rate at 43Mpps per core, whereas Virtio comes in at 15Mpps per core. Adding RDMA-device emulation reduces performance to 5Mpps per core, which is adequate for workloads with large messages, such as AI training, but probably not suitable for HPC. In terms of raw CPU performance, Xsight rates the SPECint2017\_rate of the E1 at 175, which places it in between an Intel Xeon D-27xx below and Ampere's 80-core Altra above. As with any highly programmable platform, customers' "mileage" will vary, but they can use existing Arm servers to estimate baseline performance before any additional offloads provided by the DPU's hardware.

### *Lowering the DPU-Adoption Barrier*

The challenge for Xsight, as well as other merchant-DPU vendors, is that the market is concentrated around a small set of large data-center operators. Amazon and Microsoft have internal SmartNIC teams, and Google is working with Intel on its second-generation DPU. AMD (Pensando) won SmartNIC business at IBM and Oracle as well as a lower-volume network virtual appliance (NVA) design for Azure. One advantage of Xsight's architecture is that the E1 can serve a diversity of designs, including network appliances and storage systems.

With its E1 DPU, the company is betting that the architectural pendulum is swinging back toward programmability and away from more-rigid pipelines. Pipelined processing guarantees packet rate so long as the required feature set fits the available resources. As soon as features exceed the capabilities of the pipeline, however, packets must be forward to the CPU cluster. For example, a pipeline may not support a new protocol needed for AI workloads, such as UET. The available AMD, Intel, and Nvidia DPUs all offer 16 Arm cores with per-core performance roughly similar to that of Xsight's DPU. Thus, the E1 should deliver about four times the CPU performance of existing DPU designs. Among SoCs of this class, the E1 is unique in integrating 400G Ethernet ports and PCIe x16 ports with endpoint support, suiting it to SmartNIC designs.

While the E1's processing power enables more CPU cycles per packet for a given packet rate, it also opens the programming model. Customers needn't rely on proprietary software stacks or tools as provided by other DPU vendors. As a startup, Xsight benefits too, as it doesn't require the large software-development team needed to develop these components. Instead, it can focus on delivering driver-level code that's compliant with NOSs and networking stacks, leaving application code to customers and open-source communities. With ties to Azure, SONiC-DASH is a good proof-of-concept application for Xsight's approach. The company's immediate task is to deliver working silicon and associated drivers, proving the E1 lives up to its design concept. Once it does, design wins should follow as customers discover the ease of use that a different kind of DPU can deliver.

*Bob Wheeler is an independent industry analyst covering semiconductors and networking for more than two decades. He is currently principal analyst at Wheeler's Network, established in 2022. Previously, Wheeler was a principal analyst at The Linley Group and a senior editor for Microprocessor Report. Joining the company in 2001, he authored articles, reports, and white papers covering a range of chips including Ethernet switches, DPUs, server processors, and embedded processors, as well as emerging technologies. Wheeler's Network offers white papers, strategic consulting, roadmap reviews, and custom reports. Our free blog is available at [www.wheelersnetwork.com](http://www.wheelersnetwork.com).*