

VMware Project Monterey (vSphere support for DPUs) EAP Begins

Silverton Consulting, Inc. StorInt™ Briefing

At VMworld 2020, VMware announced Project Monterey which was to add support for DPUs and other SmartNICs to VMware vSphere® ESXi™. VMware Early Access Program (EAP) to NVIDIA® BlueField® 2 DPU support on ESXi or Project Monterey begins this month.

VMware Project Monterey

There's a lot more intelligence being supplied outside the CPU these days. Part of this is due to the slowdown in Moore's Law and part of this is due to the expansion of hardware development capabilities which have decreased time to market for new ASICs and other electronic hardware.

Nonetheless, as networking speeds reach rates @200Gb/s or above, the amount of CPU cycles required for packet encryption/decryption, data movement and networking transaction activity is starting to eat into customer workload vCPU and bare metal CPU cycles. In addition, it becomes difficult to hit line rate at bandwidth rates of 100 or 200 Gb/s without some sort of aggressive offload capabilities. Enter DPUs, these devices have specialized hardware accelerators and dedicated compute cores used to offload support for encryption, software defined networking and software defined storage services.

The BlueField 2 DPU is NVIDIA's 2nd generation DPU device and NVIDIA already has on the market BlueField 3 DPUs and are working on BlueField 4 DPUs to be available in the following years. The BlueField DPU technology came from their acquisition of Mellanox in 2020. **Project Monterey will initially support BlueField 2 DPUs on Dell PowerEdge R750 servers and on HPE servers.**

The BlueField 2 DPUs integrates NVIDIA's ConnectX-6 Dx network functions, acceleration engines and 8 ARMv8 gen2 64-bit cores. It supports:

- 25, 100, 200Gb/s Ethernet and InfiniBand networking with PCIe 4.0,
- Secure root-of-trust and secure boot,
- NVIDIA ASAP² (Accelerated Switch and Packet Processing[®]) technology accelerates network performance and the CPU overhead of IP packet transport,
- AES encryption with IPsec and TLS acceleration,
- Storage NVMeoF, VirtIO block, RDMA & Decompression acceleration,
- GPU Direct communication between NVIDIA GPUs in remote systems,
- VXLANs and Geneve overlay network offloading,
- DOCA™ (Data center infrastructure-on-a-Chip Architecture) SDK/API enabling developers to create apps & services to run on BlueField DPUs.

No indication when ESXi support will be available for BlueField 3 DPUs. NVIDIA Gen 3 BlueField adds to the above more and faster cores, higher speed networking and more accelerated services to the mix.

But NVIDIA isn't the only DPU supplier. Both Pensando and Intel were announced as part of Project Monterey. Fungible also has a DPU and BroadCom offers SmartNICs. No telling when ESXi support for Pensando and Intel will be in EAP.

Significance

The emergence of outboard smarts seems to replay what occurred with floating point and other accelerators of the past century. Overtime, Intel (and AMD) just added all this hardware functionality to their core CPUs and killed these accelerators all off over a couple of generations of CPU chips.

What's different this time? For one thing, the industry is running up against some physical speed limits which means while all these activities can still be done at the CPU, doing so adds latency, as it takes time to move data from the network to the CPU to be processed. Nowadays, with network and IO transaction overheads started being measured in 100 nanoseconds to a few microseconds, there's not a lot of time to move data around anymore.

There are many benefits beyond just raw performance. Such as security enclaves can now move outside the CPU and onto a DPU to enhance/isolate security, management of the network moves outside CPU to an independently power cycled DPU, DPUs can expose virtual devices on the PCIe bus, and VMware can run ESXi on the DPU for bare metal Windows/Linux servers to manage bare metal systems like VMs (requires porting ESXi to run on DPUs).

Finally, Moore's law is starting to slow and CPU performance is only increasing linearly, not exponentially as before. This is giving startups and other hardware developers more time to introduce hardware functionality outboard of the CPU. Our guess is that Moore's law will continue to slow down. So, the ability of Intel and AMD to add lots of additional hardware functionality to the next CPU is being delayed, giving more time for others to create new functionality enclaves to increase performance and add value.

But where to draw the line between what's properly done at the CPU vs. the DPU is an open question. There doesn't appear to be any standards body working on this boundary. But barring that, having a major vendor publish an interface (say DOCA), which by default becomes an industry standard interface that other vendors could code APIs to will work fine, iff the interface remains open source.

Welcome to the DPU wars, NVIDIA has a great lead now, but Pensando and Intel are right behind, and other big vendors are not standing still. Then there's this skulk of startups sitting on the horizon all waiting for their chance as well.

Silverton Consulting, Inc., is a U.S.-based Storage, Strategy & Systems consulting firm offering products and services to the data storage community.

