

Serveritehnoloogia tipptase ja tulevik 2025

Madis Pärn
IT Arhitekt
ATEA

madis.parn@atea.ee

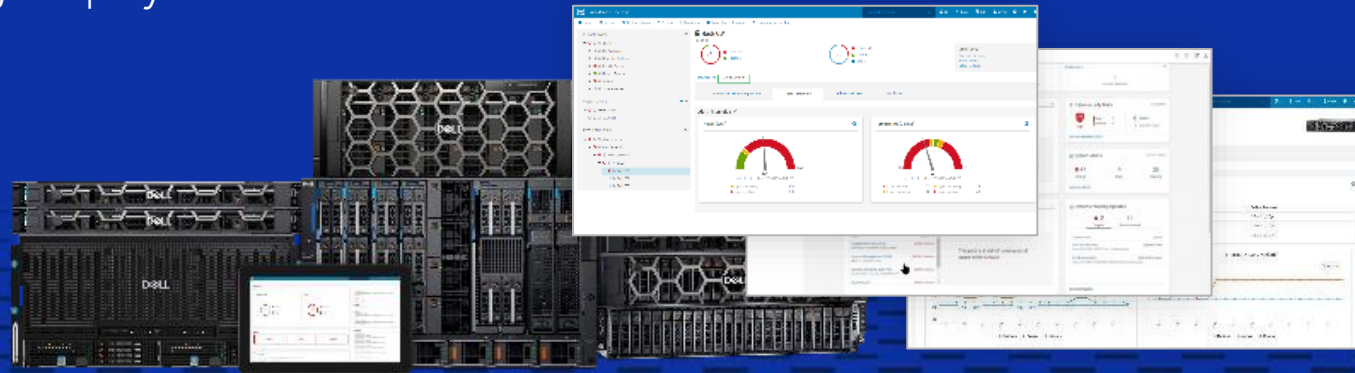
ATEA

Agenda

- Technology Trends
- Key .Next Technology Enablers
 - Intel Xeon 6 detour
 - A Look at DPU
- Cooling Technologies
- OCP Open Architecture

PowerEdge Servers

Purpose-built | Intelligent | Cyber Resilient



Purpose-built

Scale AI, Edge & Performance Anywhere



Intelligent

Accomplish more with Automation & Improve Operational Efficiencies

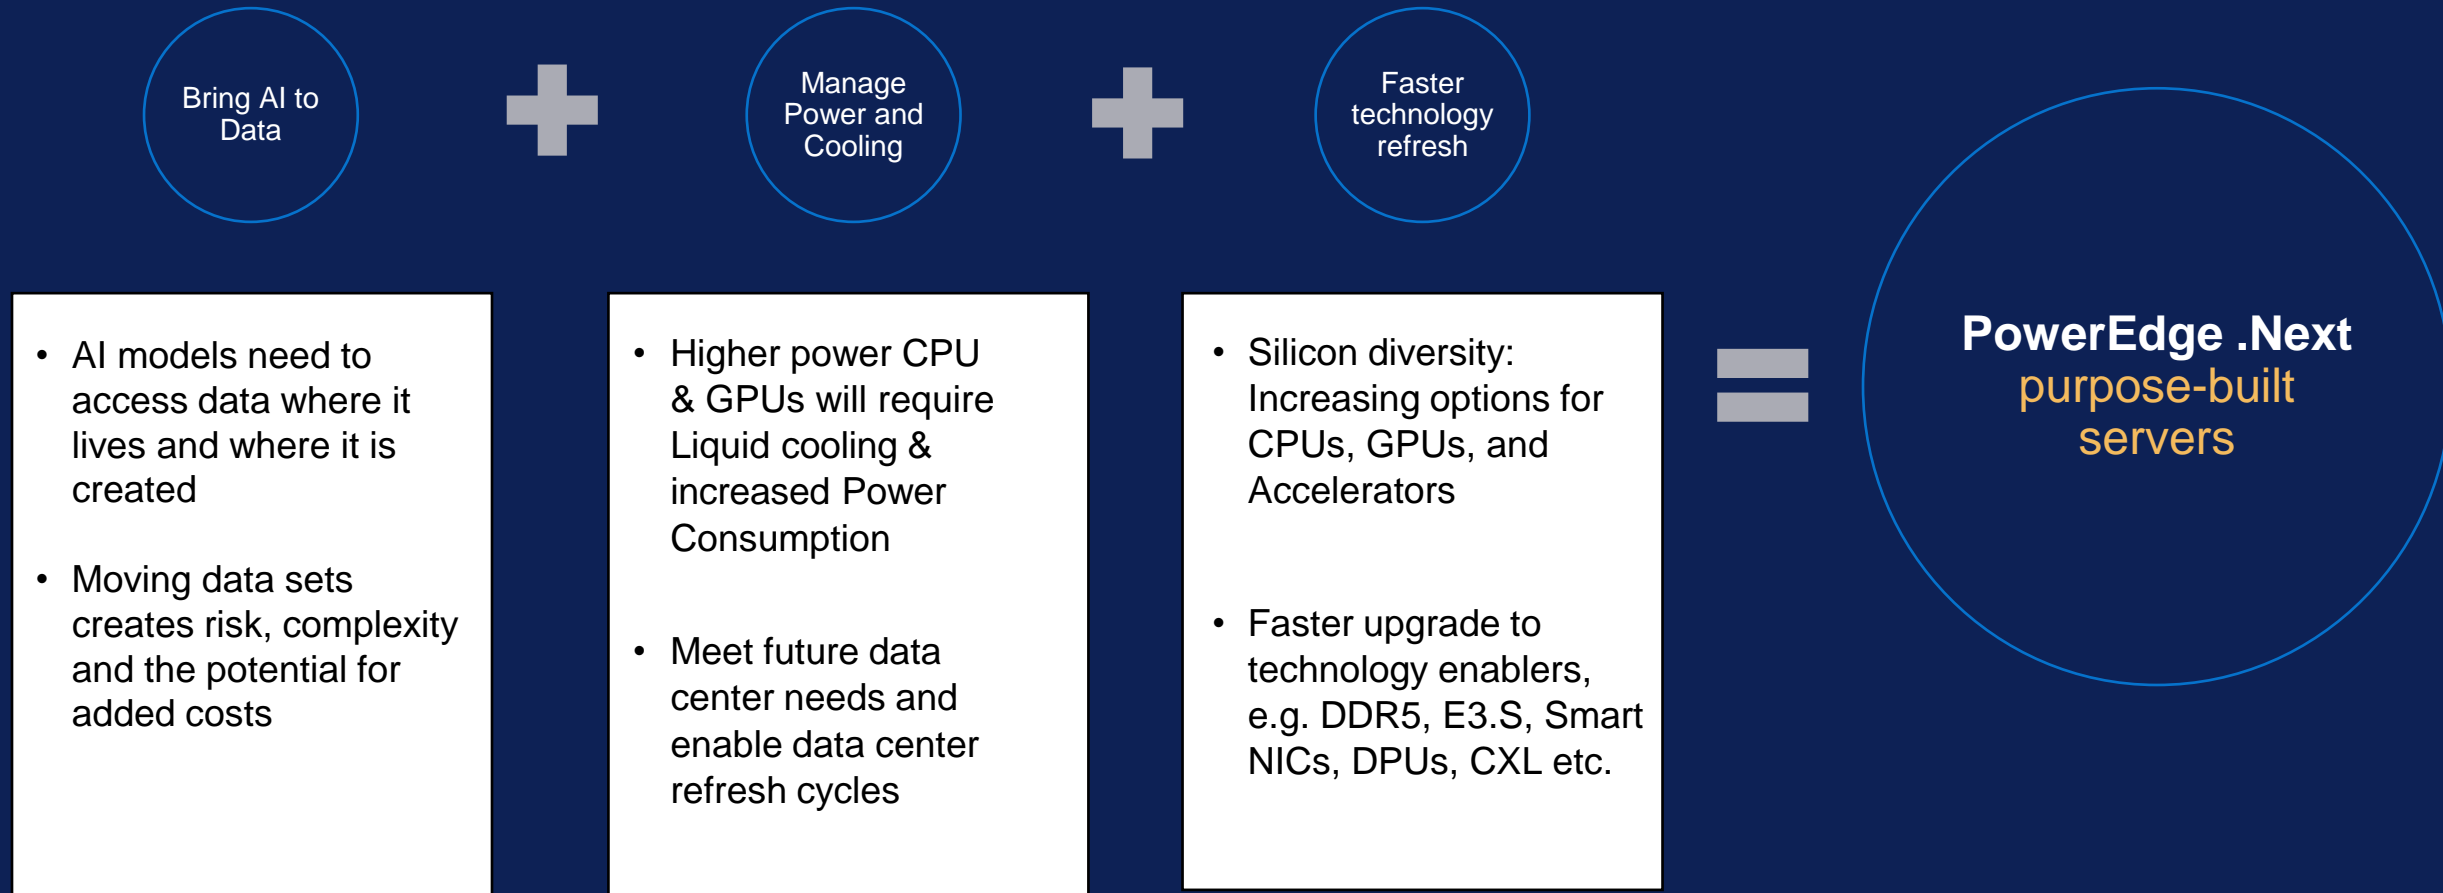


Cyber Resilient


Accelerate Zero Trust Adoption

Subscribe or Consume aaS with APEX

Evolving customer needs met with purpose-built servers



THE WORLD HAS 11,800 DATA CENTERS



Who is driving the design?

~2/3 of Global DC IT spend is
Hyperscale CSPs

Rest is Enterprise with share
decreasing

The next generation PowerEdge server portfolio

Purpose-built to address evolving customer needs

Data Center

Optimized



R470



R570



R6715



R7715

Performance



R670



R770



R6725



R7725

Modular



PowerEdge Dense Server

AI/ML




XE9680L/XE9685L



XE7740/XE7745
PCIe GPUs



XE9712*
GB200 NVL72



XE8712*
GB200 NVL4 Dense

*subject to change

Technology Enablers

Next-Gen CPU Overview

Multi-vendor chip manufactures offering the best solutions for all workloads.

Solutions for workloads in Mainstream Enterprise, AI/ML inferencing, Dense Virtualization, Data Analytics, HPC/HFT, Design and Simulation balanced by versatile and entry-level offerings.

Leading-edge architectures offer higher performance, improved power efficiency, expanded memory and richer storage configurations.

Trusted performance, exceptional efficiency.

Next-Gen Processors, the Best Solutions for all Workloads

AMD EPYC

Turin

- 1S & 2S
- Up to 192 cores
- Up to 500W TDP
- DDR5: Up to 6000 MT/s
- Advanced I/O: Up to 128 lane PCIe 5.0/CXL 2.0
- Multi-threaded

Intel® Xeon® 6

Efficiency E-Core SP (Sierra Forest)

- 1S & 2S
- Up to 144 cores
- Up to 330W TDP
- DDR5: Up to 6400 MT/s
- Advanced I/O: Up to 88 lanes PCIe 5.0/CXL 2.0
- Single threaded

Performance P-Core SP (Granite Rapids)

- 1S & 2S
- Up to 128 cores
- Up to 500W TDP
- DDR5: Up to 6400 MT/s
- Advanced I/O: Up to 88 lanes PCIe 5.0/CXL 2.0; Rich I/O SKUs up to 136 lanes
- Multi-threaded

New capabilities & technologies for the next gen...

Help customers with no compromise native computing and higher cache for technical compute with expanded AMD 9004 series

No Compromise Native Computing with Zen4c and next gen 3D die stacking technology

- AMD 97X4 “Zen4c” (Bergamo) processors optimized for scale-out performance, significant power efficiency and density-optimized cache hierarchy Increased socket level performance driving lower TCO
 - Option as alternative to ARM with dense core counts, better memory density and performance with 12 ch DDR5 4800MT/s, support multi-threading apps, and leverage existing x86 software applications with little to no code changes
- AMD 96X4X “Zen4” (Genoa-X) processors optimized for Technical Compute and High Cache per core workloads
 - Leadership Socket and Per-Core Performance for EDA, CFD and FEA workloads
- Socket, infrastructure, BIOS and software compatible with “Genoa”



1. More cores

With Zen 4c “Bergamo”, 100% more core count over previous 8003 generation, Up to 128 SMT-capable (Simultaneous Multithreading) supported and up to 96 cores with Zen 4 with 3D V-Cache

2. Thread Density

AMD “Bergamo” highest Thread Density for highest HPL/FLOP performance and with up to 128C focused on perf/watt.
2.5X-3.9X improvements on key cloud native workloads

3. Cache

AMD “Genoa-X” >1GC of L3 Cache providing highest Perf/Core Focused on highly cache bound workloads, relieves memory bandwidth pressure and reduces latency

PowerEdge with AMD Processors



- Next gen iDRAC 10, Latest memory, IO
- Our technology ecosystem for the future (DLC Solution, GPU, DPU, Storage, Memory Solutions)
- Enables CXL 2.0 features
- Limited configs at launch with feature complete planned for 2025Q2

Higher performance, greater cores, faster memory & IO

Single Socket

PowerEdge R6715

- Powered by one AMD Turin proc, up to 160 cores
- 2DPC with up to 24x DDR5 up to 5200 MT/s
- 3.5", 2.5", E3.S drive options (up to 22x)
- 100% Gen5 IO, Dual OCP with most configs





PowerEdge R7715

- Powered by one AMD Turin proc, up to 160 cores
- 2DPC with up to 24x DDR5 up to 5200 MT/s
- 3.5", 2.5", E3.S drive options (up to 40x)
- 100% Gen5 IO, Dual OCP with most configs

Two Socket

PowerEdge R6725

- Powered by two AMD Turin proc, up to 192 cores
- 1DPC with up to 24x DDR5 up to 6000MT/s
- 3.5", 2.5", E3.S drive options (up to 22x)
- 100% Gen5 IO, Dual OCP with most configs



PowerEdge R7725

- Powered by two AMD Turin proc, up to 192 cores
- 1DPC with up to 24x DDR5 up to 6000MT/s
- 3.5", 2.5", E3.S drive options (up to 40x)
- 100% Gen5 IO, Dual OCP with most configs

Comprehensive portfolio to meet diverse customer needs

PowerEdge with Intel Processors

- Focused CSP Edition launch in July '24 with E-core CPUs and OSM 3.0
- Additional configs available in 2024Q4 & 2025Q2
- Technology ecosystem for the future with E-Core & P-Core procs, DLC, GPU, DPU, E3 Storage & DDR5-6400
- Advanced management with iDRAC 10 and OSM 3.0
- CXL 2.0 Enabled

Full Performance



R670
2S / 32 DIMMs / 1U



R770
2S / 32 DIMMs / 2U

Optimized



R470
1S / 16 DIMMs / 1U



R570
1S / 16 DIMMs / 2U

Next-gen features

Next-gen Xeon 6 E-core and P-core SP processors

Cold aisle support

iDRAC 10* with expanded functions

Industry Standard Firmware

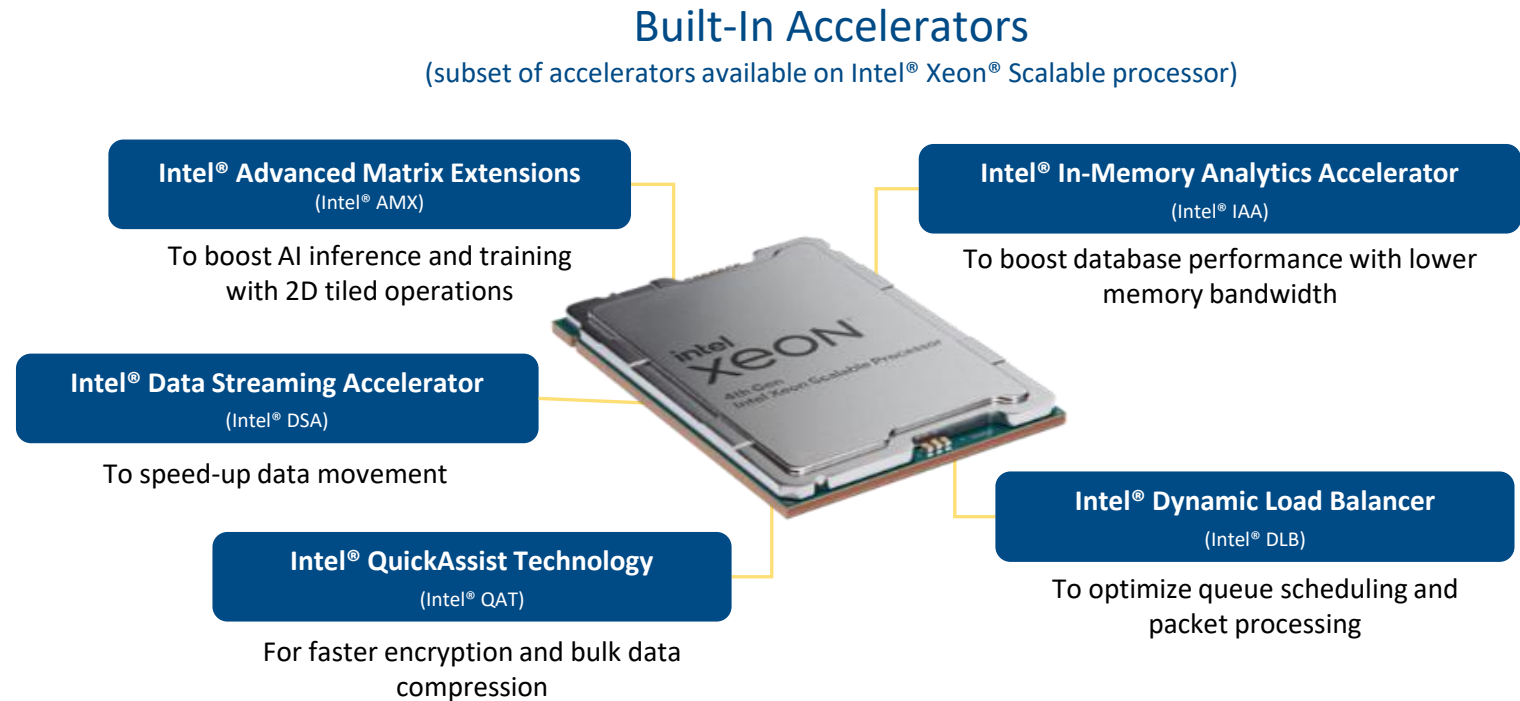
DC-MHS Compliant

*CSP Edition with OSM launches in July '24, and the mainstream edition with iDRAC follows in 2024Q4

Intel Xeon 6 detour

Get to New Levels of Efficiency and Performance with Built-In Accelerators

Intel® Accelerator Engines are built-in accelerators on Intel® Xeon® Scalable processors that increase ROI and open the door to new degrees of system power efficiency and performance that cannot be achieved by simply adding more cores.



<https://www.intel.com/content/www/us/en/products/docs/accelerator-engines/overview.html>

Make the best use of CPU core resources with built-in accelerators



AMX for AI

Accelerate AI workloads **3x to 5x** for deep learning inference on SSD-ResNet34 and up to **2x** for training on ResNet50 v1.5 with Intel® Advanced Matrix Extensions (Intel® AMX) compared with the previous generation.¹

IAA for data analytics

Improve database and analytics performance with **1.91x** higher throughput for data decompression in the open source RocksDB engine, using Intel® In-Memory Analytics Accelerator (Intel® IAA) compared to software compression on cores without acceleration.⁴

DSA for 5G/networks

Get better networking performance and efficiency with built-in accelerators. Gain up to **1.8x** higher throughput for packet processing on Open vSwitch (OVS)4 with Intel® Data Streaming Accelerator (Intel® DSA) compared to software on cores without acceleration.³

DLB for storage

Get up to **3.27x** performance on packet forwarding with Intel® Dynamic Load Balancer (Intel® DLB) vs. software queue management on cores without acceleration.³

QAT for cloud

Run cloud and networking workloads using fewer cores with faster cryptography. Increase client density by up to **4.35x** on an open source NGINX web server with Intel® QuickAssist Technology (Intel® QAT) using RSA4K compared to software running on CPU cores without acceleration.²

Intel® Xeon® 6 (@2025-01)

12 Products [COMPARE ALL](#)

Product Name	Launch Date	Total Cores	Max Turbo Frequency	Processor Base Frequency	Cache	TDP
<input type="checkbox"/> Intel® Xeon® 6952P Processor	Q3'24	96	3.9 GHz	2.1 GHz	480 MB	400 W
<input type="checkbox"/> Intel® Xeon® 6960P Processor	Q3'24	72	3.9 GHz	2.7 GHz	432 MB	500 W
<input type="checkbox"/> Intel® Xeon® 6972P Processor	Q3'24	96	3.9 GHz	2.4 GHz	480 MB	500 W
<input type="checkbox"/> Intel® Xeon® 6979P Processor	Q3'24	120	3.9 GHz	2.1 GHz	504 MB	500 W
<input type="checkbox"/> Intel® Xeon® 6980P Processor	Q3'24	128	3.9 GHz	2 GHz	504 MB	500 W
<input type="checkbox"/> Intel® Xeon® 6710E Processor	Q2'24	64	3.2 GHz	2.4 GHz	96 MB	205 W
<input type="checkbox"/> Intel® Xeon® 6731E Processor	Q2'24	96	3.1 GHz	2.2 GHz	96 MB	250 W
<input type="checkbox"/> Intel® Xeon® 6740E Processor	Q2'24	96	3.2 GHz	2.4 GHz	96 MB	250 W
<input type="checkbox"/> Intel® Xeon® 6746E Processor	Q2'24	112	2.7 GHz	2 GHz	96 MB	250 W
<input type="checkbox"/> Intel® Xeon® 6756E Processor	Q2'24	128	2.6 GHz	1.8 GHz	96 MB	225 W
<input type="checkbox"/> Intel® Xeon® 6766E Processor	Q2'24	144	2.7 GHz	1.9 GHz	108 MB	250 W
<input type="checkbox"/> Intel® Xeon® 6780E Processor	Q2'24	144	3 GHz	2.2 GHz	108 MB	330 W



Back to Technology Enablers

Next-Gen Memory Overview

- Higher speed offering with DDR5 technology along with the following improvements
 - Memory bus speeds of 6400MT/s and above expected,
 - Single bit error correction in DRAM die
 - Support for up to 32Gb density DRAM
 - Enhanced data integrity feature
- Scalability for multi-core workload
- Post package repair now available

Latest Memory Technology, higher speed & bandwidth

DIMMs Capacities

DDR5* (16Gb/ 24Gb/ 32Gb density, 6400MT/s)

- 16GB RDIMMs
 - (Mar CY25)
- 32GB RDIMMs
 - (June CY24)
- 64GB RDIMMs
 - (June CY24)
- 96GB RDIMMs
 - (Nov CY24)
- 128GB RDIMMs
 - (Dec CY24)
- 256GB RDIMMs
 - (Mar CY25)

DIMMs per System

- Intel Sierra Forest/Intel Granite Rapids
 - 2 socket servers up to 32 DIMMs
 - 1 socket servers: up to 16 DIMMs
- AMD Turin
 - 2 socket servers: 24 DIMMs
 - 1 socket servers: 24 DIMMs

6400MT/s Max Memory Bus

- Higher memory bus speed
- Intel Xeon 6
 - Max CPU memory bus speed is up to **6400MT/s**
- AMD Turin
 - Max CPU memory bus speed is up to **6000MT/s**

*Dell does not support DIMM capacity mixing on 17th Generation

**17th Generation memory will have staggered launch; exact launch dates will be communicated through 411

PCIe Gen 5

Technologies taking advantage of PCIe Gen 5


- NVMe Drives
- GPUs
- DPUs
- NICs
- CXL

Database and AI/ML workloads will benefit from the new bandwidth

PCIe Gen 5: Future proofing servers for the new ecosystem



- PCIe Gen 5 is twice the speed of PCIe Gen 4 with backward compatibility
- PCIe Gen 5 32 GTs data rate vs PCIe Gen 4 16 GTs data rate
- PCIe Gen 5 has full duplex bandwidth for x16 interface at 128 GB/s vs Gen 4 at 64 GB/s




Generation	Raw Bit Rate	Interconnect Bandwidth	Bandwidth Lane Direction	Total Bandwidth for x16 Link
Gen1	2.5 GT/s	2 Gb/s	~250 MB/s	~8 GB/s
Gen2	5 GT/s	4 Gb/s	~500 MB/s	~16 GB/s
Gen3	8 GT/s	8 Gb/s	~1 GB/s	~32 GB/s
Gen4	16 GT/s	16 Gb/s	~2 GB/s	~64 GB/s
Gen5	32 GT/s	32 Gb/s	~4 GB/s	~128 GB/s

CXL Memory Expansion Overview

Compute Express Link (CXL) Memory

- Seamlessly integrates into supported PCIe slots for enhanced memory capacity.
- Facilitates high-speed, low-latency communication between CPUs and devices for optimized performance.



<https://computeexpresslink.org/>

CXL adds channels to attach memory



More memory bandwidth

Up to 25% more vs 8 channels of native DDR5 6400 MT/s alone



More memory capacity


Up to 50% more vs sixteen native DDR5 128GB DIMMs alone



Right-sizing memory capacity

Add memory capacity and bandwidth to existing system DDR configs

DDR5+CXL - Cost-effective memory expansion



CXL Memory Devices	R770	R670	R7725	R7715
CXL DIMM Add-In-Card	•	•	•	•

- PCIe connection: x16 PCIe single wide slot. This device can also connect on x8 PCIe slot
- Up to two AICs per CPU
- Max capacity with AIC per CPU 1TB

GPU Accelerators

Broad multi-vendor portfolio catering to applications ranging from the Cloud to Core to Edge

Solutions for targeted workloads in Gen AI, ML/DL, HPC, VDI, Data Analytics balanced by versatile and entry-level offerings to boost utilization, help the AI journey

Leading edge technology ingredients in core & memory architecture, fabrication technology, air and liquid cooling, interconnect bandwidths to deliver breakthrough performance


Growing ecosystem of frameworks, GPU- accelerated libraries that are optimized & ready-to-deploy and the necessary development tools

Accelerate insight and innovation with Dell's GPU portfolio on PowerEdge servers


Brand	GPU Model	GPU Memory	Max Power Consumption	Form-factor	2-way Bridge	Recommended Workloads
PCIe Adapter form-factor						
Nvidia	A2	16 GB GDDR6	60W	SW, HHHL or FHHL	n/a	AI Inferencing, Edge, VDI
Nvidia	L4	24 GB GDDR6	72W	SW, HHHL or FHHL	n/a	AI Inferencing, Edge, VDI
Intel	Flex 140	12 GB GDDR6	75W	SW, HHHL or FHHL	n/a	AI Inferencing, Edge, VDI
Nvidia	A16	64 GB GDDR6	250W	DW, FHFL	n/a	VDI
Nvidia	L40	48 GB GDDR6	300W	DW, FHFL	N	Performance graphics, Multi-workload
Nvidia	L40S	48 GB GDDR6	350W	DW, FHFL	N	AI Inferencing, Multi-workload
AMD	MI210	64 GB HBM2e	300W	DW, FHFL	Y	HPC, AI Training
Nvidia	H100 NVL	94 GB HBM3	400W	DW, FHFL	Y	AI Training, HPC, AI Inferencing
Nvidia	H200 NVL *	141 GB HBM3e	450W-600W	DW, FHFL	Y	AI Training, HPC, AI Inferencing
SXM / OAM form-factor						
Nvidia	HGX H100	80 GB HBM3 or 94 GB HBM2e	700W	SXM w/ NVLink	n/a	AI Training, Inferencing, HPC
Nvidia	HGX H200	141 GB HBM3e	700W	SXM w/ NVLink	n/a	AI Training, Inferencing, HPC
Nvidia	HGB B200 **	Up to 192GB HBM3e	1000W	SXM w/ NVLink	n/a	AI Training, Inferencing
AMD	MI300X	192 GB HBM3	750W	OAM w/ XGMI	n/a	AI Training, Inferencing
Intel	Gaudi3	128 GB HBM2e	850W	OAM w/ RoCE	n/a	AI Inferencing, Training

* - Development or under evaluation ** - See your local 411 updates

PCIe Adapter



PCIe with 2-way Bridge



4-way SXM / OAM Baseboard



8-way SXM / OAM Baseboard



EDSFF-E3

Increased Performance

- Supports PCIe Gen5; 100% increase in Sequential Reads, 62% increase in Sequential Writes, 60% improvement in Random Reads, and 33% improvement in Random Writes

Greater Storage Density

- 60% increase on 1U and 33% increase on 2U
- Total capacity increase:

	15G	16G
1U	154TB	245TB
2U	368TB	491TB

Improved Thermals


- Airflow can be optimized through the server due to the smaller drive size

EDSFF and E3.S, a form factor optimized for SSDs and the future of Server Storage


- EDSFF is a new family of form factors optimized for Flash storage devices designed to support high frequency interfaces like PCIe Gen5 and Gen6
- PowerEdge will utilize the E3.S form factor and it will be the launch vehicle for PCIe Gen5 NVMe
- E3.S is roughly half the size of a 2.5" SSD benefitting density, thermals, and improved packaging in space constrained servers
- E3.S SSDs will have the same Serviceability and Manageability as our current 2.5" SSDs

(EDSFF E3.S 2T will not be supported)


15G



MX750c: Up to 6 x 2.5" SSD




R650/R6515: Up to 10 x 2.5" SSD




R7525/R7515: Up to 24 x 2.5" SSD


16G



MX760c: Up to 8 x E3.S NVMe



R660/R6615: Up to 16 E3.S NVMe




R7625/R7615: Up to 32 x E3.S NVMe

BOSS N1 for Next-Gen Servers

Boot Optimized Storage Solution (BOSS) uses 1 or 2 M.2 devices to store the OS, allowing customers the ability to segregate their boot devices from primary storage

- BOSS N1 for 17th Generation leverages 16th Generation BOSS chip, Marvell NevoX with some additional features and a new form factor
 - Root of Trust – Authenticates FW and BIOS
 - Flat Bread Form Factor – BOSS Flat Bread form factor to support DC-MHS, modular design able to support multiple applications
- Key Features:
 - Supports UEFI Boot Only
 - Secure Firmware update
 - Marvell NevoX chip
 - Supports 480GB and 960GB M.2 2280 NVMe
 - LKM and SEKM support for SED and FIPS M.2 NVMe
 - SMCU for Root of Trust
 - Supports HW RAID1 and RAID0
 - Hot Serviceable
 - DC-MHS - Modularity



Platinum and Titanium efficiency

- PSU efficiency is dictated by the [80 Plus certification program](#)
- Testing and certification is performed at specific voltages: 115Vac and 230Vac
 - 80 Plus standard does not cover DC voltage inputs or higher AC voltages such as 277Vac
- Performance testing is done at various load levels
 - For example, an 1100W Titanium PSU is 96% efficient at 550W power draw (50% load)

• EU ErP Lot9

- Since January 1st, 2024, Titanium PSU are required to ship into CE countries

• Energy Star 4.0

- From January 12th, 2024, Titanium PSU are required for 750W or greater, in order to comply with Energy Star

Required Efficiency depending on % of Rated Load

80 PLUS Certification	115V Internal Non-Redundant			230V Internal Redundant				
	% of Rated Load	20%	50%	100%	10%	20%	50%	100%
80 PLUS		80%	80%	80%	N/A			
80 PLUS Bronze		82%	85%	82%	---	81%	85%	81%
80 PLUS Silver		85%	88%	85%	---	85%	89%	85%
80 PLUS Gold		87%	90%	87%	---	88%	92%	88%
80 PLUS Platinum		90%	92%	89%	---	90%	94%	91%
80 PLUS Titanium		---	---	---	90%	94%	96%	91%



Jumper Cords

Connector (female)	Appliance inlet (male)	Configuration Female/Male	Earth contact	International		North America		Max. pin temp. (°C)
				Max. current (A)	Voltage (V)	Max. current (A)	Voltage (V)	
C13	C14		Yes	10	250	15	125/250	70
C15	C16		Yes	10	250	15	125/250	120
C19	C20		Yes	16	250	20	125/250	70
C21	C22		Yes	16	250	20	125/250	155

Standard PDUs are C13 & C19

PSU

Power Cord examples



C14 ← **C13-C14** → C13



C16 ← **C15-C14** → C13



C20 ← **C19-C20** → C19



C22 ← **C21-C20** → C19

PDU




Next-Gen M-CRPS PSU

- 1500W Titanium
- 1100W Platinum
- 800W Platinum
- Following wattages:
 - 3200W Titanium
 - 1100W Titanium
 - 800W Titanium
 - 1500W 277Vac
 - 2400W Titanium
 - 1400W -48Vdc

M-CRPS PSUs, a smaller form factor to free-up valuable real estate in Servers and Storage

- Next-Gen M-CRPS Power supplies, as part of the DC-MHS project, will come in 2 form factors : 73.5mm and 60mm (width), respectively 30% and 18% smaller than the 15th/16th generation PSUs :




- Live FW Update - no need to turn off the server to update the PSU Firmware
- Signed FW images (Firmware security)

DPU (Data Processing Unit)

Silicon Chipset Options in Modern Servers

A quick primer



Central Processing Unit
(CPU)

Runs Workloads


General purpose x86 cores running application software



Graphical Processing Unit
(GPU)

Accelerates Specific Workloads

Optimized for parallel processing, enabling specific workloads, AI/ML, to be accelerated in hardware

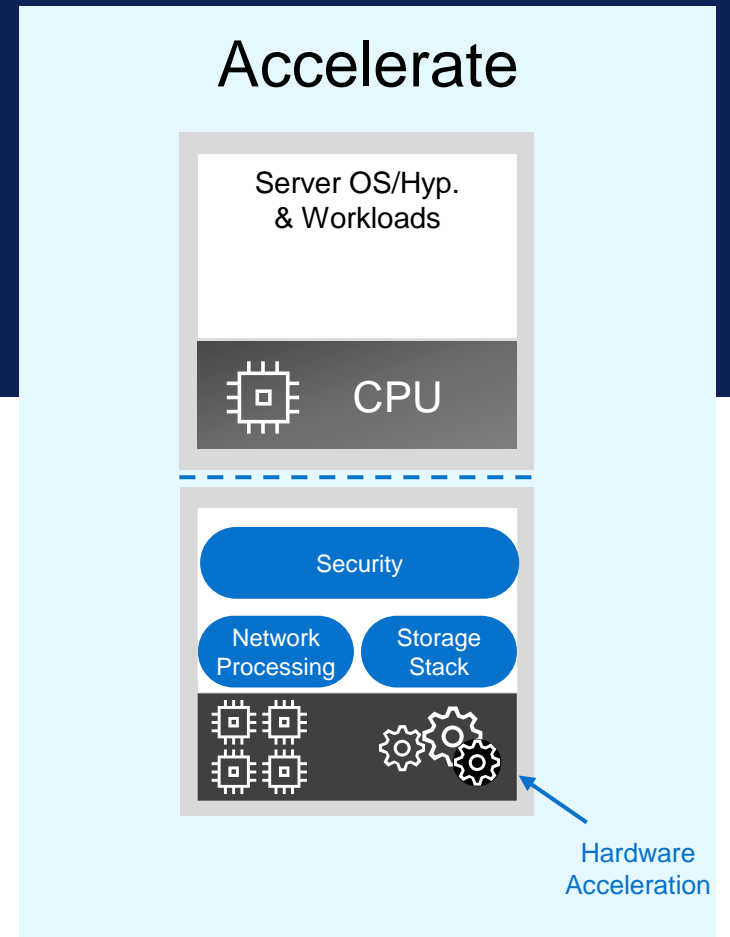
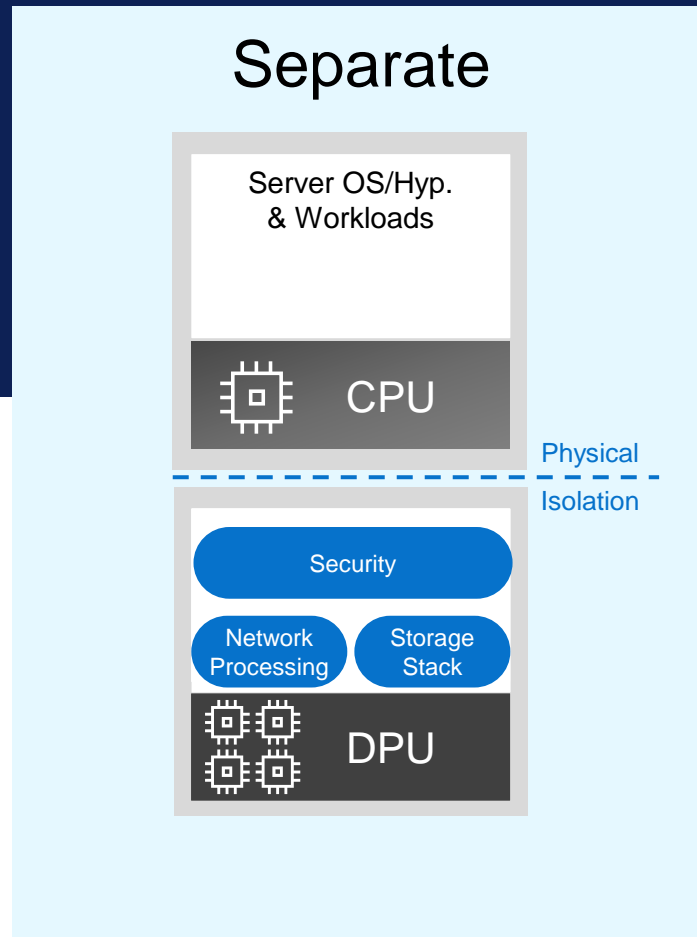
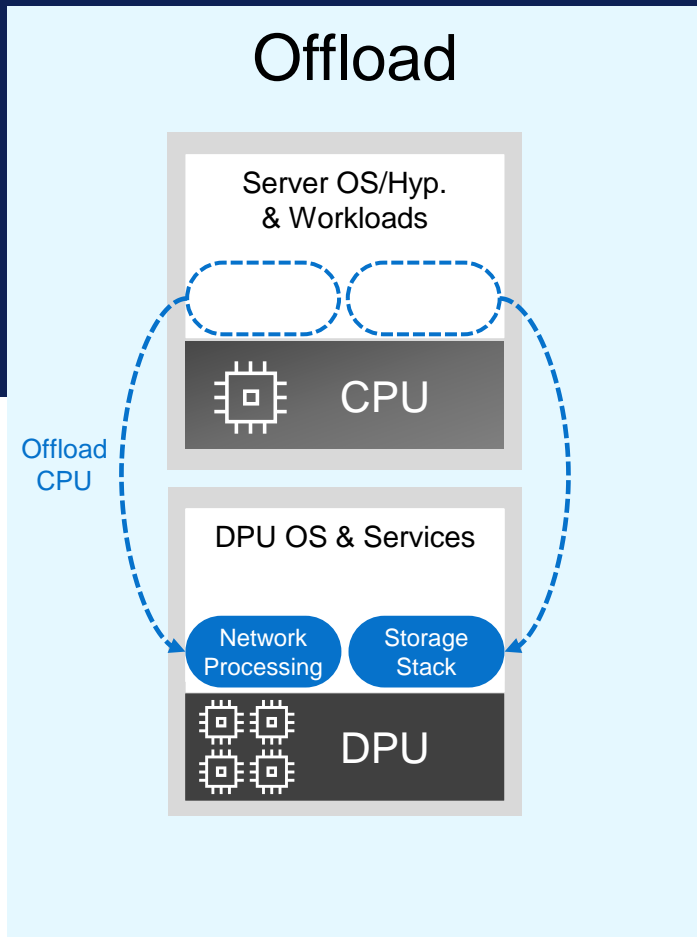


Data Processing Unit
(DPU)

Accelerates Infrastructure Services


Combination of a NIC ASIC and ARM cores and it is a server within a server. Runs and accelerates infrastructure services such as networking, storage, and security.

DPU Capabilities



Main DPU Value Propositions

Networking Stack



- Switch packets in hardware instead of in software
- Achieve line rate performance at 25GbE and 100GbE

Storage Stack



- Hardware accelerate NVMe/TCP and SDS
- Reduce storage performance jitter
- Hardware accelerated compression

Security Stack



- Isolate workloads from security services because they run on different processors
- Line rate packet inspection done in hardware
- In-flight encryption

Simplify Host Deployment

Compute/
Workloads

Infrastructure

- Moving infrastructure services to the DPU simplifies workload deployment

Dell Has Two Sets of DPUs

Dell DPUs (vSphere on DPU)

Features

- Hardware accelerated NSX
- Hardware accelerated Distributed Switch
- Near Future: NVMe/TCP

Advantages

- Ease of management and integration into vSphere and iDRAC
- VxRail support
- DPU services are provided by VMware

DPU Vendors

- AMD/Pensando
- NVIDIA



Partner DPUs (Linux Stack)

Features

- Hardware accelerated firewall (Palo Alto, F5)
- OVN/OVS acceleration
- In-Flight encryption and TLS/SSL acceleration

Advantages

- Flexibility
- Multiple 3rd party ISVs are creating DPU services
- Broad set of use cases

DPU Vendors

- NVIDIA
- Intel



Hardware cannot switch between Dell DPU and Partner DPUs.

These are different part numbers, and the server recognizes them and treats them differently.

Partner DPUs currently does not support host servers running VMware vSphere OS

Data Processing Units

A DPU is a network adapter with ARM Cores integrated with hardware acceleration blocks, memory and storage that runs its own operating system and applications.

DPU helps free up the cores on the host system by offloading and running infrastructure services (networking, storage and security) and improve performance.

Offload and Accelerate Infrastructure services with DPUs



Services that can run on the DPUs

- HW accelerated Firewall
- In-Flight Encryption, TLS/SSL acceleration
- Packet Processing/Switching/Inspection
- NVMeoTCP
- Compression
- UPF Acceleration

Customers can develop their solution using Vendor SDKs or work with 3rd Party Software vendors to deploy solutions suited to their needs on the DPUs

Brand	DPU Model Number	DPU ASIC	Core Type	Host OS Supported	Port Speed	Number of Ports
PCIe Adapter form-factor FHHL						
Nvidia	BlueField-2	BlueField-2	8 ARMv8 A72	Linux-Based	25GbE	2
Nvidia	B3210E	BlueField-3	16 Arm Cortex-A78	Linux-Based	100GbE	2
Nvidia	B3220	BlueField-3	16 Arm Cortex-A78	Linux-Based	200GbE	2
Nvidia	B3140H^	BlueField-3	8 Arm Cortex-A78	Linux-Based	400GbE	1
Intel	E2100-CCQDA2*	Mt Evans	16 ARM N1	Linux-Based	100GbE	2
Intel	E2100-CCQDA2HL	Mt Evans	16 ARM N1	Linux-Based	100GbE	2

*only FH ¾ length Form factor, will transition to the Half Length Intel DPU E2100-CCQDA2HL by mid-CY25

^Marketed as SuperNIC (Cores can be turned off)

Cooling Technologies

Growing Space constraints and Power requirements

- CPU size increase and additional number of Memory channels results in higher CPU TDPs, and space constraints within servers
- Growth in GPU TDPs and with higher CPU TDPs necessitate new solutions for cooling
- Rack level infrastructure and super pods expected to be future of infrastructure solutions

Over Next 2 generation of servers



Intel CPU size increase



AMD CPU size increase



Memory channels growing from 8 to 16


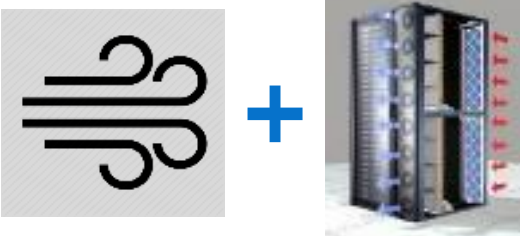

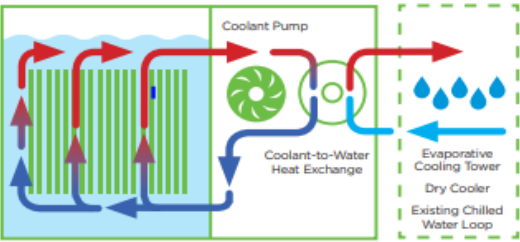
↑ Max CPU TDP expected to grow to ~500W from ~350W

↑ Max GPU PCIe CEM TDP expected to grow to ~900W from ~350W

↑ Max GPU SXM/OAM TDP expected to grow to ~1400W from ~750W


- Direct Liquid to Chip cooling necessary for highest TDP GPUs
- Workloads for 8-way SXM/OAM super pods eventually move to multi-node scale-up solutions
- ORv3-based rack infrastructure becomes direction for some workloads:
 - Very Large AI Training
 - Dense High TDP GPU and CPU Compute

Cooling Technology Comparisons


	Air cooling	Air + Supplemental	Direct Liquid Cooling (DLC)	Immersion
Cooling Solution Options				
Products	<ul style="list-style-type: none"> • Traditional air-cooling & air-handling equipment • Containment 	<ul style="list-style-type: none"> • In-row coolers • Rear Door Heat Exchangers (RDHx) • Containment (hot & cold aisle) 	<ul style="list-style-type: none"> • CPU/GPU Cold-plate loops • Rack/facility level DLC products required 	Single-phase (1P) and Two-phase (2P) Immersion tank solutions
Environments	Traditional data centers	Traditional data centers, with facility water	Traditional data centers, with facility water	<ul style="list-style-type: none"> • Non-traditional spaces, no conditioned air required (ex. - warehouse) • Note: facility water required
Main usage model	<ul style="list-style-type: none"> • Low to Mid-density racks • Up to ~ 15kW/rack 	<ul style="list-style-type: none"> • Mid to High-density racks • Up to ~30kW/rack 	<ul style="list-style-type: none"> • Systems with high TDP parts • High-density racks, up to ~80kW/rack 	<ul style="list-style-type: none"> • Limited/no air cooling available • High-density racks, or high TDP parts
Typical Cost Adder	NA	+	++	Single phase (1P): ++ Two-phase (2P): +++
Availability	Standard cooling	Standard server cooling + 3 rd party supplemental cooling solutions	Dell factory supported configurations	Dell OEM project engagement

The Dell Cooling Strategy

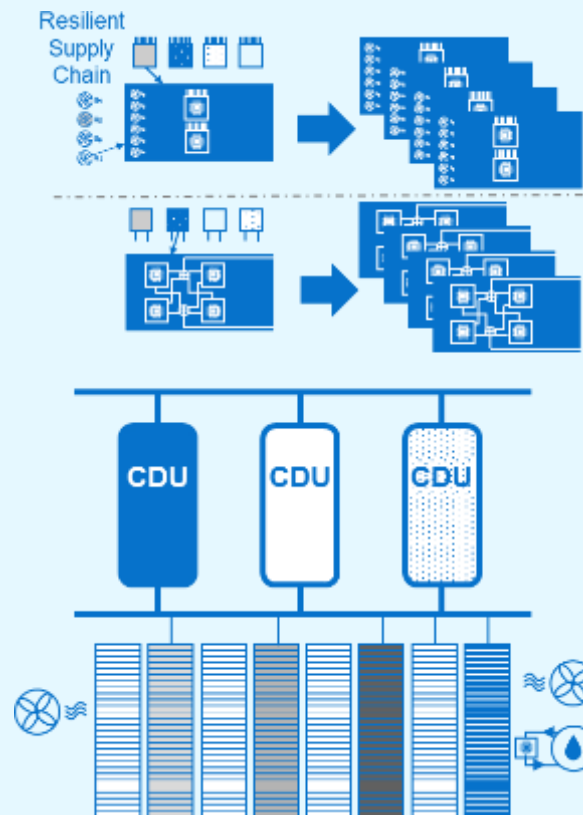
Continued innovation with **Air** and **Liquid** cooled solutions







Standards enabling interoperability & interchangeability

























With standards and commonality, **Scale** the supply chain



Cooling technology available to serve the market

-  - Enables high density racks
-  - Enables high power chips
-  - Requires facility water to rack/row
-  - 100% heat capture to facility water

“At-the-Server” Cooling Hardware	Capacity	Attributes	Overview
Air	1 kW/U		<ul style="list-style-type: none"> Traditional air-cooling utilizing fans and heatsinks. Heated air is released into the datacenter environment.
Immersion	1 kW/U	 	<ul style="list-style-type: none"> Equivalent in cooling performance to traditional air-cooling. Server immersed in a bath of oil-based liquid for heat transfer.
In-server Liquid-Assisted Air Cooling (iLAAC)	1.2 kW/U		<ul style="list-style-type: none"> Utilizes liquid for heat transfer off the chip. Radiator enclosed in the server chassis to facilitate cooling the liquid.
Direct-to-Chip Liquid Cooling, Liquid-to-Air (L2A DLC)	3 kW/U	 	<ul style="list-style-type: none"> Utilizes liquid for heat transfer off the chip. Liquid transfers heat to facility air in an external coolant distribution unit (CDU).
Direct-to-Chip Liquid Cooling, Liquid-to-Liquid (L2L DLC)	10 kW/U	  	<ul style="list-style-type: none"> Utilizes liquid for heat transfer off the chip. Liquid transfers heat to facility water in an external coolant distribution unit (CDU).

“At-the-Rack” Cooling Hardware	Capacity	Attributes	Overview
Air / Immersion	20 kW/Rack	Same as above, respectively	<ul style="list-style-type: none"> Air: Perimeter air handler (CRAH) to cool heated air. Immersion: Facility water plumbed to immersion tank which cools heated liquid.
Enclosed Cooling (EC)	30 kW/Rack	  	<ul style="list-style-type: none"> Isolates hot air with rack dedicated air-cooling units, requires additional floor space. Heated air does not release into the datacenter environment. Warm water
Heat Capture Cabinet (HCC, concept)	70 kW/Rack	  	<ul style="list-style-type: none"> Isolated hot air with rack dedicated air-cooling units integrated into rack. Heated air does not release into the datacenter environment. Warm water
Rear Door Heat Exchanger (RDHx)	70 kW/Rack	  	<ul style="list-style-type: none"> Air-cooling unit attached to rear of rack. Cold water Heated air is cooled and released into data center environment.
Coolant Distribution Unit, Liquid-to-Air (L2A CDU)	80 kW/Rack	 	<ul style="list-style-type: none"> Liquid cooling unit either “in-rack” or as a “in-row” appliance. Heated liquid is cooled and released into data center environment.
Coolant Distribution Unit, Liquid-to-Liquid (L2L CDU)	100 kW/Rack	  	<ul style="list-style-type: none"> Liquid cooling unit either “in-rack” or as a “in-row” appliance. Heated liquid is cooled by facility water which is then.

OCP Open Architecture



OPEN

Compute Project®



Server

Sub-Projects:

- High Performance Computing
- NIC
- Open Accelerator Infrastructure
- Open Chiplet Economy
- AI HW SW CoDesign
- Composable Memory System
- MHS



Cooling Environments



Data Center Facility



Rack & Power



Storage


Alibaba Cloud (since 2017) 	AMD (since 2012) 	AMI (since 2018) 	Ampere Computing (since 2019) 	Arista Networks (since 2019) 	ARM (since 2018)
Auras Technology (since 2024) 	Baidu (since 2019) 	Broadcom Inc (since 2022) 	ByteDance Technology (since 2021) 	Castrol Lmtd (BP) (since 2021) 	Cisco (since 2014)
Dell Technologies (since 2015) 	Delta Electronics (since 2016) 	Edgecore Networks (since 2016) 	Google (since 2015) 	HPE (since 2015) 	Hyve Solutions (since 2012)
IEIT SYSTEMS (since 2016) 	Intel (since 2011) 	Inventec (since 2014) 	Lattice Semiconductor (since 2017) 	Meta (since 2011) 	Micas Networks (since 2023)
Microsoft (since 2014) 	NVIDIA (since 2017) 	Quanta Cloud Technology (since 2012) 	Rackspace (since 2011) 	Rittal (since 2017) 	Samsung Electronics (since 2019)
Sanmina (since 2023) 	Seagate (since 2017) 	Shell (since 2020) 	Sims Lifecycle Services (since 2020) 	Submer (since 2018) 	Supernmicro (since 2019)
Tencent (since 2018) 	Vertiv (since 2019) 	Wiwynn (since 2014) 			

OCP Open Architecture

DC-MHS (Datacenter ready Modular Hardware System) key tenets

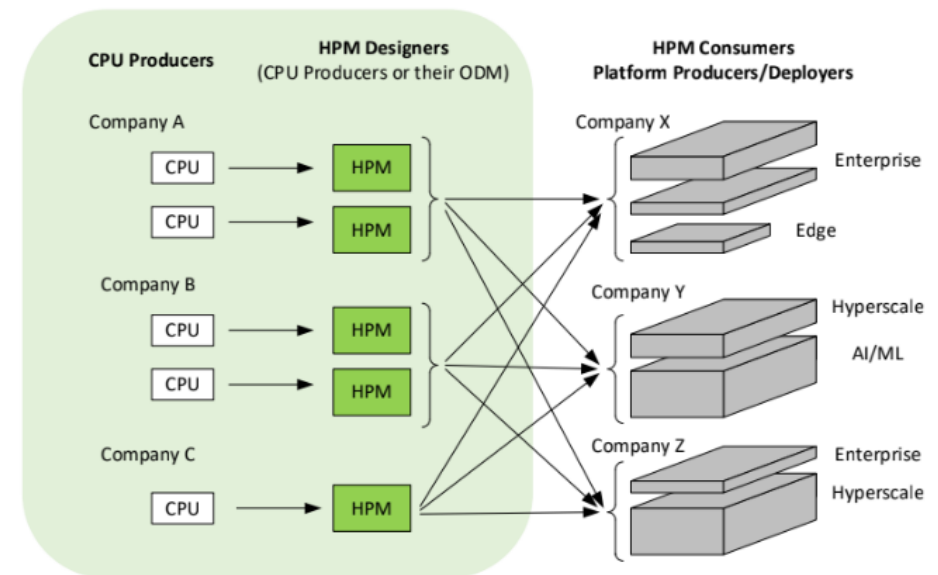
- Drive interoperability between key elements of datacenter, edge and enterprise infrastructure
- Industry standard building blocks and supporting ingredients for greater agility, speed, and leverage

Today's model



Unique company platforms that deliver on specific workloads and/or requirements

DC-MHS model



Industry leveraged platforms that reduce validation scope and with industry wide interoperability

DC-MHS



The DC-MHS R1 Mission

- **What:** Data Center – Modular Hardware System Revision 1.0

DC-MHS R1 envisions interoperability between key elements of datacenter, edge and enterprise infrastructure by providing consistent interfaces and form factors among modular building blocks.

DC-MHS R1 standardizes a collection of HPM (Host Processor Modules) form-factors and supporting ingredients to allow interoperability of HPMs and platforms.

- **Why**
 1. DC-MHS R1 aims to ultimately improve industry efficiency and innovation.
 - Enable the CPU Suppliers to design and validate the circuit board under their CPUs
 - While preserving the ability for the rest of the supply chain to innovate beyond the CPU
 2. CPU Suppliers are enabled to innovate without barriers to adoption.
 3. Platform Suppliers may innovate without burden of redesigning HPMs
- **When:** Enabling for producing solutions late 2023, early 2024.
- **Who:**



Dell	Shawn Dube	shawn.dube@dell.com
Google	Siamak Tavallaei	tavallaei@google.com
HPE	Jean-Marie Verdun	verdun@hpe.com
Intel	Brian Aspnes	brian.d.aspnes@intel.com
Meta	Todd Westhauser	twesthauser@meta.com
Microsoft	Mark A. Shaw	mashaw@microsoft.com

<https://www.opencompute.org/wiki/Server/Working>

DC-MHS: specifications



The Data Center – Modular Hardware System (DC-MHS) family of specifications are written to enable interoperability between key elements of datacenter and enterprise infrastructure by providing consistent interfaces and form factors among modular building blocks. As of September 2022, DC-MHS includes the following specification workstreams:

- **M-FLW (Modular Hardware System Full Width Specification)** – Host Processor Module (HPM) form factor specification optimized for using the full width of a Standard EIA-310-D Rack mountable server. The specification is not limited to use within the EIA-310 Rack but is used to serve as a template for a common target where the design is expected to be utilized.
- **M-DNO (Modular Hardware System Partial Width Density Optimized Specification)** – Host Processor Module (HPM) specification targeted to partial width (i.e. $\frac{1}{2}$ width or $\frac{3}{4}$ width) form factors. Such form factors are often depth challenged and found not only in enterprise applications but also in Telecommunications, Cloud and Edge Deployments. While the EIA-310 Rack implementation is chosen as a key test case for use, the specification is not limited to use within the EIA-310-D Rack but is used to serve as a template for a common target where the design is expected to be utilized.
- **M-CRPS (Modular Hardware System Common Redundant Power Supply Specification)** – Specifies the power supply solutions and signaling expected to be utilized by DC-MHS compatible systems.
- **M-PIC (Modular Hardware System Platform Infrastructure Connectivity Specification)** – Specifies common elements needed to interface a Host Processor Module (HPM) to the platform/chassis infrastructure elements/subsystems. Examples include power management, control panel and cooling amongst others.
- **M-XIO (Modular Hardware System Extensible I/O)** – Specifies the high-speed connector hardware strategy. An M-XIO source connector enables entry and exit points between sources such as Motherboards, Host Processor Modules & RAID Controllers with peripheral subsystems such as PCIe risers, backplanes, etc. M-XIO includes the connector, high speed and management signal interface details and supported pinouts.
- **M-PESTI (Modular Hardware System Peripheral Sideband Tunneling Interface)** – Specifies a standard method for discovery of subsystems, self-describing attributes, and status (e.g., versus a priori knowledge/hard coding firmware and BIOS for fixed/limited configurations). Examples: vendor/module class, physical connectivity descriptions, add-in card presence, precise source to destination cable coupling determination.

DC-MHS: ingredients



- HPM Form-factors:

- Dimensions, mounting, KOs, Connectors

- Power Supply:


- Form-factor
- Electrical operation



- System Connectivity (conn's, pins-outs, signaling)

- PCIe/CXL, cabled and riser
- Sideband Virtualization
- Power Distribution
- Control Panel

- Utilization of OCP NIC R3 and DC-SCM R2



DC-MHS v1.0 specifications

[M-FLW R1 v1p0](#)

[M-DNO R1 v1p0](#)

[M-PESTI R1 v1p0](#)

[M-PIC R1 v1p0](#)


[M-XIO R1 v1p0](#)

[M-CRPS R1 v1p0](#)


Modular Form Factors

Modular Building Blocks


- **Chassis**
 - Full length (M-FLW)
 - Dense Optimized (M-DNO)
- **Planar (HPM)**
 - Full width (M-FLW) – 2S, 2DPC
 - ¾ width (M-DNO ¾) – 2S, 1DPC
 - ½ width (M-DNO ½) – 1S, 2DPC
- **Components**
 - BOSS
 - Power distribution
 - DC-SCM (IDRAC & OSM SW)
 - PSU's (M-CRPS)
 - Fan solution
 - Riser (TBD)
 - South IO placement
 - Backplanes (TBD)



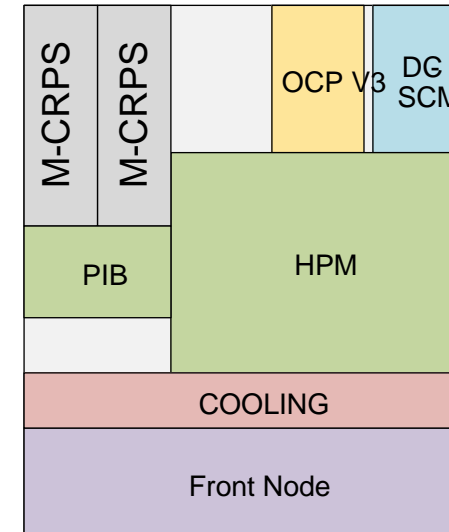
M-FLW – Chassis 1U/2U




SRF 2S
GNR SP
2DPC



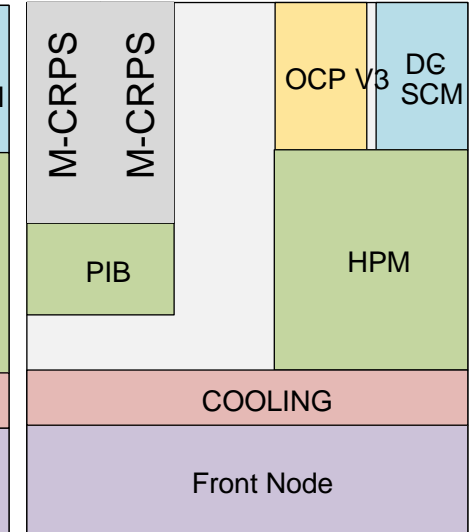
M-DNO – Chassis 1U/2U/3/4



Future
Option



SRF 1S
GNR SP
GNR RIO
2DPC



Dell Leads with OCP Solutions for Large-scale AI & Dense Compute

Forward-thinking infrastructure that will scale to support multiple generations of the latest CPU and GPU technologies, with optimized power and cooling for energy efficiency

Dell Integrated Rack 7000



Dell PowerEdge M7725



Dell PowerEdge XE9712



← Delivered by expanded rack scale integration services →

Efficient

Dense

Scalable

Dell Integrated Rack 7000



OCP Standards-based Rack Scale Infrastructure for Large-Scale AI and Dense Compute

Unparalleled simplicity

Cable-free liquid and power delivery.

Flexibility

Support multiple architectures in one rack.

Scalable

Grow as your compute demand grows. Disaggregated power for seamless scaling

Rapid deployment

Your entire HPC cluster at-scale with a white glove experience

Future-ready design

Support up to 480kW in each rack. Multigenerational compute support

Efficient

Integrated DLC for energy efficiency

Rainwater Infrastructure

(Dell Integrated Rack 7000)

- ORv3 Standard based Infrastructure for Large AI & Dense Compute
- Designed for multigeneration & heterogenous technology (CPU, GPU & CPU+GPU)


Rack level solution for next-gen Dense Compute infrastructure

Optimized for Mainstream & Large-Scale deploys

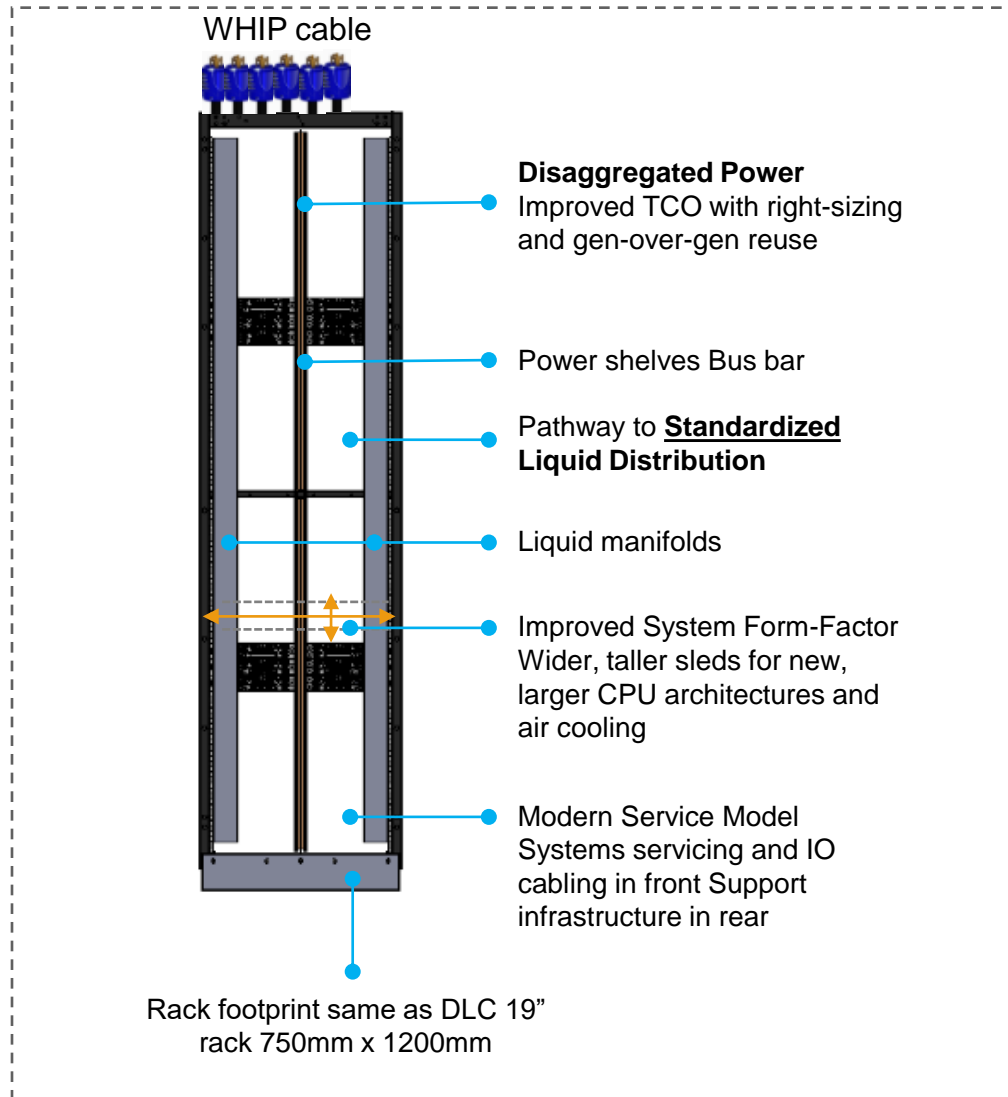
- Rack scale clusters
- Multi-generational – future proofing infrastructure

Standards based infrastructure that scales



- ORv3 racks supporting 44OU, 50OU
- Future proofed to scale up to 480kW with future platforms
- 18kW single AC input, 33kW dual AC input options in 1HCY25
- Up to 10kW per OU heat capture
- Compatible with Off the shelf CDUs



Rainwater Infrastructure View




Power shelves




Product Concepts- Subject to Change

DELL Technologies

High Density ORv3 + DLC 100% Heat Capture Concept



Dell PowerEdge M7725



The future of dense high-performance compute

Dense form factor

10U with 2x 2S server nodes, up to 72 nodes per rack

Uncompromised performance

The latest AMD EPYC 5th Gen CPUs with up to 27,000 cores per rack

Energy efficient

Hybrid cooling with air + liquid for optimized power utilization

Easy to deploy and manage

Cold aisle serviceability. Quick disconnects for cable free liquid connectivity. Front I/O cabling.



72

72 nodes per rack

27K

Up to 27,000 cores per rack

32%

32% higher performance than previous AMD processor

PowerEdge Dense Compute Server Brief

- Dense Compute offering with EPYC 5 2S CPUs

Leading the future of Dense Compute

Highest performance offering – 2x 2S / 10U

- Up to 500W & 192 cores with EPYC 5
- 50 – 70% higher performance vs. Genoa
- 1 DPC DDR5 Memory at 6000MT/s

Optimized for cold Aisle serviceability

- Front I/O, 100% Gen5 enabled
- Two x16 IO Slots (Eth, IB & DPUs)
- Optional E3 & BOSS Module

Eliminates complex PSU cabling

- Blind-mates to DC Bus Bar

DLC to enable easier operation

- Quick disconnect
- Hybrid Cooling – liquid to CPUs, air for rest
- Near 100% heat capture for sustainable deployments



M7725
up to 72 nodes per rack

Dell PowerEdge XE9712 with NVIDIA GB200 NVL72



The future of dense acceleration for real time inference

Powering Generative AI

Up to 30x faster real-time LLM Performance

Energy efficient

Liquid cooled to maximize your datacenter power utilization

Lightning-fast connectivity

72 connected GPUs acting as one with NVLink technology

Rapid deployment

Your entire AI cluster at-scale with a white glove experience



25X

25 times more efficient than H100

8K

For LLM Training, highest performance delta at 8k+ GPU Clusters

30X

faster real-time trillion-parameter LLM inference than H100


GB200 NVL72 scale-up architecture

XE9712 GB200 NVL72

Targeted towards the largest AI GPU clusters

Performance-at-scale

Part of a multi-rack reference network topology for modular growth and fast deployment



Grace Blackwell

- 2 Grace CPUs (72 cores each) per node
- 4 B200 GPUs per node

Storage

- 8x E1.S NVMe drives
- Internal M.2



72x B200 GPUs acting as one

- Rack-scale, multi-node NVLink @ 1.8TB/s
- Offered as single 136kW or 2x 73kW racks

I/O

- 1:1 NIC to GPU ratio for east west GPU traffic
- Bluefield3 DPUs for storage and in-band traffic

ORv3 Rack








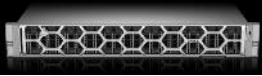
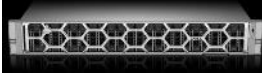




- Bus bar & power Shelves
- ~136kW per GB200 NVL72

- New scale-up NVLink solution based on Grace Blackwell
- Very large AI at scale – intended for “superpod” large model clusters of thousands of GPUs
- Leverages ORv3 rack & power infrastructure, eliminating PDUs from the solution (Rainwater Infrastructure)
- Direct liquid cooled – in rack or in row

Faster technology refresh

- Faster pace of newer generation CPUs
- Introduction of CPUs and GPUs by multiple vendors
- Adoption of new technology enablers
 - DDR5
 - CXL Memory
 - PCIe Gen 5
 - EDSFF E3.S
 - Smart NICs
 - M-CRPS Power supplies
- DC-MHS enables faster time to market

Purpose-built to provide choices with Silicon Diversity

					
	R470	R670	R6715	R6725	PowerEdge Dense Server
					
	R570	R770	R7715	R7725	
CPU	Intel	Intel	AMD	AMD	AMD
GPU	NVIDIA/AMD/Intel	NVIDIA/AMD/Intel	NVIDIA/AMD/Intel	NVIDIA/AMD/Intel	NVIDIA/AMD/Intel
					
	XE9680L/XE9685L	XE7740/XE7745	XE9712* GB200 NVL72	XE8712* GB200 NVL4 Dense	
CPU	Intel/AMD	Intel/AMD	NVIDIA	NVIDIA	
GPU	NVIDIA/AMD/Intel	PCIe NVIDIA/AMD/Intel	NVIDIA	NVIDIA	

*subject to change

