VELSANET ARCHITECTURE GROUP

Global Industry Platform Mapping

A Cross-Platform Technical Analysis of Velsanet Hardware Implementation

ARM + Samsung · NVIDIA + TSMC · Intel + Broadcom

Independent Technical Analysis by Velsanet Architecture Group

v1.0 · March 2026 · Independent Analysis

About This Report

This report consolidates three independent hardware mapping analyses — ARM + Samsung, NVIDIA + TSMC, and Intel + Broadcom — into a single cross-platform reference. Unlike the individual reports, which examine each platform in isolation, this document is organized by Velsanet architectural layer. For each layer, all three platform options are presented and compared side by side.

The goal is not to declare a winner. Each platform combination represents a distinct implementation path with different strengths, timelines, and strategic contexts. The comparison is provided to enable informed engagement — whether for prototyping, partnership discussion, or standardization contribution.

Document Structure

Chapter 1 — Velsanet Architecture

Velsanet is a next-generation network architecture that replaces packet switching with structure-first, AI-native connectivity. There are no routing tables, no switching ASICs, and no address-based forwarding. Connectivity is formed deterministically through physical parallel optical paths, and intelligence is embedded natively within the network nodes — not bolted on top.

1.1 The Five Layers

Node	Shape	Role	AI Layer	Scale
T4	Tetrahedron	Access / device entry	None (passive)	Device / User
H6	Hexahedron	Channel alignment / RAN	None (passive)	Access point
O8	Octahedron	Personal E2E domain	PAI — Personal AI	Neighborhood
D12	Dodecahedron	Regional mediation	AAI — Agent AI	District / City
I20	Icosahedron	Global governance	AsAI — Assistant AI	National / Global

1.2 MOCT — The Physical Substrate

The Multi-Optical-Core Transceiver (MOCT) is the physical foundation of every Velsanet node. It is a fixed, non-pluggable photonic module permanently integrated with the Matrix. Core properties: 100–1000+ parallel optical cores per node, 8 internal channels per core (Channel 1 reserved for sensing, control, and identity), MEMS-based alignment, pre-attached fibers (5m+) fused at manufacture. There are no connectors, no packet switching, no routing logic at the physical layer.

1.3 DIKWEI Cognitive Cycle

All AI processing in Velsanet follows the DIKWEI cycle: Data → Information → Knowledge → Wisdom → Execution → Intent. This cycle runs continuously across PAI → AAI → AsAI layers. The Dedicated Equipment is the hardware gateway that binds the optical physical layer to this cognitive cycle — converting Cube optical lanes into PAI cognitive channels and writing AsAI outputs back into Cube memory.

1.4 Dedicated Equipment — The Synaptic Transformer

Function	Operation
Cube → PAI Demux	Converts 384 Cube optical lanes into 8-channel PAI cognitive format

PAI → AAI Mux	Combines PAI outputs into 24-core bundles for AAI ingress
AAI ↔ AsAI Conversion	Structural mapping: 12-face (AAI) topology ↔ 20-face (AsAI) topology
AsAI → Cube Write-back	Converts AsAI results back into Cube-addressable DIKWEI memory packets

Chapter 2 — MOCT / Matrix: Photonic Physical Layer

The MOCT is not a conventional optoelectronic component. It is a semiconductor-grade photonic substrate requiring sub-micron lithographic precision, MEMS process integration, and multi-material co-fabrication. This is the most demanding manufacturing requirement in the Velsanet hardware stack — and the one where the three platform paths diverge most clearly.

PATH A — SAMSUNG FOUNDRY

SF3 / SF4 Process + MEMS PDK + H-Cube Packaging

Samsung Foundry's advanced GAA processes provide the lithographic precision required for MOCT polymer waveguide structures. Samsung offers an integrated path covering silicon photonics PDK, MEMS foundry capability for MAI alignment fixtures, and H-Cube / X-Cube heterogeneous packaging for multi-die optical+electronic integration.

- Silicon photonics PDK: integrated waveguide structures on SF4 GAA process
- MEMS process integration: MAI static alignment fixture fabrication
- H-Cube / X-Cube: multi-die stacking of photonic and electronic layers
- Samsung MEMS foundry: separate process line for MEMS structures if needed
- Strategic note: only Korean domestic partner covering entire hardware stack in a single engagement

PATH B — TSMC

SiPho PDK + N3P + CoWoS-L Packaging

TSMC's silicon photonics PDK is production-proven at scale in commercial co-packaged optics (CPO) applications for hyperscaler data centers. TSMC offers photonics-on-CMOS co-fabrication, where optical and electronic functions share a single wafer process. CoWoS-L enables 2.5D multi-die integration of photonic, compute, and memory dies on a single silicon interposer.

- TSMC SiPho PDK: production-grade — waveguides, modulators, Ge photodetectors on CMOS
- N3P (3nm): dimensional control for integrated MEMS alignment structures
- CoWoS-L: 2.5D co-packaging — photonic + compute + HBM3e on single interposer
- InFO: thin-form-factor integration for MOCT polymer-form optical core layer
- Strategic note: CPO production experience directly maps to MOCT; current industry frontier for SiPho at scale

PATH C — INTEL FOUNDRY

18A Process + Intel SiPho PDK + EMIB / Foveros

Intel is the originator of silicon photonics as a manufacturing discipline, having produced the world's first silicon Raman laser (2005) and the first GHz-speed silicon modulator (2004). Intel Foundry's SiPho PDK is built on decades of production experience, now available to external customers through Intel Foundry Services. Intel's EMIB and Foveros packaging technologies provide unique die-to-die integration paths for the MOCT's multi-layer structure.

- Intel SiPho PDK: deepest heritage — originator of silicon photonics manufacturing

- Intel 18A (1.8nm-class RibbonFET GAA): highest-precision logic for MOCT control circuits
- EMIB: 2.5D die-to-die bridge interconnect — avoids full interposer cost
- Foveros 3D: vertical stacking for MOCT polymer optical core + control circuit integration
- 800G CPO in production today: proven fabrication at MOCT-adjacent density
- Strategic note: Intel CPO and Velsanet MOCT share the same principle — MOCT is an extension of Intel's CPO roadmap, not a departure from it

2.1 MOCT Fabrication Path Comparison

Criterion	Samsung Foundry	TSMC	Intel Foundry
SiPho PDK maturity	Growing — SF4 photonics emerging	Production-proven (CPO at scale)	Originator — 20+ years heritage
MEMS integration	Dedicated MEMS foundry available	CMOS-compatible passive alignment	MEMS structures on 18A compatible
Multi-die packaging	H-Cube / X-Cube	CoWoS-L (industry leading)	EMIB + Foveros
Process node	SF3 / SF4 (3–4nm GAA)	N3P (3nm FinFET+)	18A (1.8nm RibbonFET)
CPO production	Limited (research phase)	High volume (hyperscalers)	800G in production
Strategic fit	Domestic Korean partner	Global CPO standard platform	Deepest photonics expertise

Chapter 3 — Dedicated Equipment: Optical↔AI Channel Bridge

The Dedicated Equipment is Velsanet's synaptic transformer — the real-time hardware bridge between the MOCT optical physical layer and the PAI/AAI/AsAI cognitive layers. It must perform 384-lane optical demultiplexing, 8-channel PAI format conversion, 24-bundle AAI multiplexing, 12-to-20-face structural topology mapping, and DIKWEI gateway processing simultaneously at line rate. This is the highest-bandwidth, most structurally unique compute element in the architecture.

PATH A — SAMSUNG SYSTEM LSI (EXYNOS CUSTOM SOC)

Exynos Platform + PUF Identity IP + LPDDR5X PHY

A custom Exynos SoC integrating ARM Cortex-A720 compute, Ethos-N78 NPU (or Samsung NPU), CMN-700 interconnect, and Samsung's PUF (Physical Unclonable Function) for Channel 1 manufacturing-time identity. This is a fully custom silicon path — highest design effort but maximum integration with Samsung's full manufacturing stack.

- ARM Cortex-A720: deterministic structural topology mapping logic
- Samsung NPU / Ethos-N78: DIKWEI cognitive cycle gateway inference
- Samsung PUF: manufacturing-time identity — maps to Channel 1 physical-layer authentication
- LPDDR5X PHY: 384-lane Cube state memory bandwidth
- Samsung 5G modem IP: wireless interface path for H6/T4 node connectivity

PATH B — NVIDIA BLUEFIELD-3 DPU

BlueField-3 + DOCA Programmable Data Plane

The BlueField-3 DPU is the closest existing product to Velsanet's Dedicated Equipment in terms of structural role: it sits between network fabric and host compute, performing offload, transformation, and protocol processing. The DOCA (Data-center On-a-Chip Architecture) platform allows custom Velsanet channel mapping logic to be implemented in the programmable data plane without custom ASIC development.

- 400Gb/s fabric interface: sufficient for 384-lane Cube parallel data ingress
- DOCA programmable pipeline: Cube→PAI demux and PAI→AAI mux implementable in P4
- Embedded ARM Cortex-A78: AAI structural topology conversion (12-face ↔ 20-face)
- Hardware crypto engine: Channel 1 physical-layer authentication
- Available today: production hardware — fastest prototype path
- ConnectX-8 roadmap: 800Gb/s — matches MOCT Phase 3 core density

PATH C — INTEL IPU (MOUNT EVANS)

E2000 Infrastructure Processing Unit + P4 Pipeline

Intel's IPU (Infrastructure Processing Unit), codenamed Mount Evans, was purpose-built to offload infrastructure processing from host compute — sitting structurally between the network fabric and server CPUs. This is the most precise architectural match for the Dedicated Equipment's role. Intel TDX (Trust Domain Extensions) provides the hardware-enforced isolation equivalent to Velsanet's

PAI sovereignty boundary.

- P4-programmable data plane: Velsanet channel mapping logic without custom ASIC
- ARM Neoverse-class cores: DIKWEI gateway and topology conversion logic
- 100Gb/s+ network interface: 384-lane Cube ingress capable
- Intel TDX: hardware domain isolation — Channel 1 identity and PAI sovereignty
- E2000 SoC: production silicon — available for prototype development today

3.1 Dedicated Equipment Comparison

Criterion	Samsung Exynos SoC	NVIDIA BlueField-3	Intel IPU (E2000)
Availability	Custom design (12–18 months)	Production today	Production today
Programmability	Full custom — maximum control	DOCA / P4 — high flexibility	P4 — high flexibility
Bandwidth	Custom spec (design target)	400Gb/s (800 roadmap)	100Gb/s+
Identity / security	Samsung PUF — mfg-time identity	HW crypto engine	Intel TDX — domain isolation
Structural fit	Highest integration potential	Closest structural role match	Most precise architectural match

Chapter 4 — PAI Node (O8): Personal AI, Always-On Edge

PAI nodes are the distributed nervous system of Velsanet — potentially thousands of always-on instances running the DIKWEI cognitive cycle continuously at neighborhood scale. Each PAI represents one user's personal context, intent, and history. Requirements: ultra-low power (always-on), dedicated inference acceleration, 8-channel multimodal input processing, hardware identity isolation, local DIKWEI metadata storage.

PATH A — ARM CORTEX-M55 + ETHOS-U65

Helium SIMD + Dedicated NPU + TrustZone-M

The Cortex-M55 + Ethos-U65 combination is designed precisely for embedded always-on AI inference. The Cortex-M55's Helium (M-Profile Vector Extension) provides SIMD acceleration for the 8-channel PAI input stream, while the Ethos-U65 NPU handles DIKWEI inference cycle computation at sub-1W power. TrustZone-M provides PAI identity isolation at the hardware level.

- Helium SIMD: 8-channel parallel cognitive input processing
- Ethos-U65: dedicated DIKWEI inference without CPU involvement — always-on capable
- TrustZone-M: hardware-enforced PAI identity isolation (Channel 1 sovereignty)
- Sub-500mW active inference — feasible for always-on distributed deployment
- Lowest power profile of the three paths — volume deployment advantage

PATH B — NVIDIA JETSON ORIN NX

Ampere GPU + Dual DLA + Orin Security Engine

The Jetson Orin NX combines 1024 Ampere CUDA cores with dual DLA (Deep Learning Accelerator) engines within a 10–25W power envelope. The dual DLA enables always-on DIKWEI inference without activating the main GPU — preserving GPU resources for burst multimodal processing. NVIDIA TensorRT provides optimized inference runtime for the DIKWEI cognitive cycle.

- Dual DLA engines: always-on DIKWEI inference, no GPU activation required
- Ampere GPU: burst multimodal processing across 8 cognitive channels
- Orin hardware security engine: root of trust for PAI identity
- NVIDIA TensorRT: optimized DIKWEI model inference runtime
- 10–25W TDP: higher than ARM path but with significantly more burst capability

PATH C — INTEL CORE ULTRA (METEOR LAKE)

Dedicated NPU Tile + Heterogeneous Compute + TME

Intel Core Ultra introduces a dedicated NPU tile — physically separate from CPU and GPU — designed specifically for always-on ambient intelligence workloads. This is the most direct hardware match for PAI's always-on sequential inference requirement: the NPU runs the DIKWEI cycle without activating CPU or GPU cores. Intel Thread Director optimizes workload placement across heterogeneous tiles.

- Dedicated NPU tile: always-on DIKWEI inference, independent of CPU and GPU

- P-core + E-core + NPU + SoC tile: heterogeneous parallel processing of 8 cognitive channels
- Intel Thread Director: intelligent workload routing to NPU for DIKWEI
- Intel TME (Total Memory Encryption): PAI personal context protection
- Intel 4 process: compact SoC for distributed deployment

4.1 PAI Node Comparison

Criterion	ARM Cortex-M55 + Ethos-U65	NVIDIA Jetson Orin NX	Intel Core Ultra NPU
Power (always-on)	Sub-500mW — lowest	10–25W — highest	2–9W — mid range
Inference engine	Ethos-U65 NPU dedicated	Dual DLA dedicated	NPU tile dedicated
Always-on capability	Native — designed for this	DLA enables it	NPU tile enables it
Multimodal burst	Limited (Helium SIMD)	Highest (Ampere GPU)	Mid (GPU tile + NPU)
Identity isolation	TrustZone-M (hardware)	HW security engine	TME + TDX (hardware)
Volume cost	Lowest (embedded class)	Highest (GPU SoC)	Mid (client PC SoC)

Chapter 5 — AAI Node (D12): Regional Multi-Agent Coordination

AAI is the social validation and intent mediation layer — operating at D12 (Dodecahedron) nodes at district and city scale. It receives intent flows from regionally distributed PAI nodes, evaluates them against normative constraints (ethical, legal, policy), reconciles heterogeneous intents into Structured Intent Packets, and routes validated outputs to AsAI. AAI does not generate intent; it governs it.

PATH A — ARM CORTEX-X4 + MALI-G720 + CMN-700

Big-Core Reasoning + Parallel Matrix GPU + Coherent Mesh

The Cortex-X4 + Mali-G720 combination on a CMN-700 interconnect mesh provides the peak single-thread reasoning performance and parallel matrix capability needed for AAI's normative reasoning and multi-agent intent consolidation workloads. CMN-700 scales to the AAI node's multi-PAI ingress bandwidth across all D12 N-faces.

- Cortex-X4 big-core: complex normative reasoning — social/ethical/legal validation logic
- Mali-G720: parallel matrix operations for multi-agent intent consolidation
- CMN-700: coherent mesh — scales to N-face multi-PAI ingress bandwidth
- TrustZone: governance boundary enforcement between AAI jurisdiction domains
- Custom SoC path: manufacturable by Samsung System LSI with ARM IP license

PATH B — NVIDIA GRACE-HOPPER GH200

Neoverse V2 CPU + Hopper GPU + NVLink-C2C 900GB/s

The Grace-Hopper Superchip combines ARM Neoverse V2 with an H200 Hopper GPU via NVLink-C2C at 900GB/s bidirectional bandwidth — eliminating the PCIe bottleneck for multi-agent workloads. The Transformer Engine handles multi-PAI intent consolidation, while Confidential Computing enforces hardware-level AAI governance jurisdiction isolation.

- Grace CPU (Neoverse V2): deterministic normative reasoning logic
- H200 Hopper Transformer Engine: multi-agent intent consolidation at scale
- NVLink-C2C 900GB/s: eliminates CPU↔GPU memory bandwidth bottleneck
- 141GB HBM3e: sufficient for full regional PAI intent context storage
- Hopper Confidential Computing: hardware-enforced AAI jurisdiction isolation
- MIG: partitioned GPU resources for independent parallel PAI intent streams

PATH C — INTEL XEON + GAUDI 3

Enterprise CPU + Open Ethernet AI Fabric + SGX Enclaves

Intel's Xeon + Gaudi 3 provides AAI-scale compute with a key differentiator: Gaudi 3 uses standard Ethernet for inter-accelerator communication rather than proprietary NVLink. This open fabric is strategically relevant for multi-jurisdiction AAI deployments where no single governance domain should control the interconnect fabric. Intel SGX provides hardware enclave isolation for AAI governance boundaries.

- Intel Xeon: enterprise-grade CPU for deterministic normative reasoning
- Gaudi 3: 128GB HBM2e, optimized for LLM inference — AAI intent consolidation
- Open Ethernet fabric: no proprietary lock-in — strategically relevant for AAI governance
- Intel SGX: hardware enclave isolation for jurisdiction-level AAI governance boundaries
- Intel AMX: CPU-side matrix operations for Structured Intent Packet formation

5.1 AAI Node Comparison

Criterion	ARM Cortex-X4 + Mali	NVIDIA GH200	Intel Xeon + Gaudi 3
Reasoning performance	High (Cortex-X4 big-core)	Highest (Neoverse V2 + HBM3e)	High (Xeon enterprise)
Multi-agent parallelism	Mali-G720 GPU	Hopper Transformer Engine	Gaudi 3 HBM2e
CPU↔GPU bandwidth	CMN-700 coherent mesh	NVLink-C2C 900GB/s	PCIe 5.0
Governance isolation	TrustZone (hardware)	Confidential Computing	SGX enclaves (hardware)
Interconnect	Proprietary CMN	Proprietary NVLink	Open Ethernet
Availability	Custom SoC required	Production today	Production today

Chapter 6 — AsAI Node (I20): Infrastructure-Scale Wisdom Execution

AsAI is the highest layer of the Velsanet AI architecture, operating at I20 (Icosahedron) nodes at national and global scale. It executes wisdom-layer reasoning across aggregated AAI-validated intent patterns. No individual access path exists — AsAI receives only AAI-authorized requests. Structurally, this impossibility is enforced by architecture, not policy. AsAI requires the highest compute density and memory bandwidth available.

PATH A — ARM NEOVERSE V2

Server-Class SVE2 + High Core Count + RAS

The Neoverse V2 is ARM's highest-performance server processor, used in cloud infrastructure deployments (AWS Graviton4, Google Axion). Its SVE2 vector engine provides the large-scale vector inference capability needed for DIKWEI wisdom-layer reasoning. RAS (Reliability, Availability, Serviceability) features ensure continuous operation for infrastructure-level AI.

- SVE2: large-scale DIKWEI wisdom-layer inference — scalable vector architecture
- High core count: parallel processing of multi-AAI input streams
- CCIX/PCIe 5.0: high-bandwidth interface to Samsung HBM3E memory subsystem
- RAS features: continuous operation — infrastructure AI cannot have planned downtime
- Custom SoC path: implementable by Samsung System LSI with Neoverse V2 license

PATH B — NVIDIA BLACKWELL GB200 NVL72

72 Blackwell GPUs + NVSwitch 1.8TB/s All-to-All + CoWoS-L

The Blackwell GB200 NVL72 is a rack-scale AI supercomputer: 72 Blackwell GPUs connected via NVSwitch at 1.8TB/s all-to-all bandwidth. This is the closest existing hardware implementation of Velsanet's parallel E2E connectivity principle at the compute layer. NVSwitch's non-blocking, deterministic, fixed-topology fabric mirrors the MOCT's parallel optical core architecture — the same design principle at different scale.

- NVSwitch all-to-all 1.8TB/s: non-blocking, deterministic — structural analogue of MOCT parallel E2E
- Blackwell Transformer Engine FP4: 30× inference throughput vs. previous generation
- NVIDIA Magnum IO: AsAI collective validation memory — accumulated AAI-validated patterns
- Hardware RAS: continuous operation for national/global AsAI infrastructure
- TSMC CoWoS-L: compute + HBM3e co-integrated — same packaging path as MOCT integration

PATH C — BROADCOM XPU CUSTOM AI ASIC

Purpose-Built AI ASIC + HBM3 + Custom Scale-Out Fabric

Broadcom's Custom Silicon Solutions division designed Google TPU v4/v5, Meta MTIA, and ByteDance AI accelerators — the hyperscale AI ASIC market. For Velsanet's AsAI, a purpose-built XPU ASIC would be optimized specifically for DIKWEI wisdom-layer inference, with HBM3 memory

integration and a custom scale-out fabric for multi-chip I20 node formation. This path requires the longest development timeline but yields the highest optimization.

- Purpose-built ASIC: DIKWEI wisdom-layer inference optimized — not general-purpose
- HBM3 integration: highest available memory bandwidth for multi-AAI aggregation
- Custom scale-out fabric: multi-chip AsAI cluster — full control over topology
- Proven delivery: Google TPU, Meta MTIA at production volume
- Intel Gaudi 3 alternative: near-term deployable AsAI node pending custom ASIC

6.1 AsAI Node Comparison

Criterion	ARM Neoverse V2	NVIDIA Blackwell GB200	Broadcom XPU ASIC
Inference throughput	High (SVE2 vector)	Highest (FP4 Transformer Engine)	Highest when built (custom)
Memory bandwidth	HBM3E (via Samsung)	HBM3e CoWoS-L integrated	HBM3 integrated (custom)
Fabric topology	Custom (PCIe / CCIX)	NVSwitch all-to-all	Custom — full control
MOCT topology parallel	Partial (parallel processing)	Strongest (NVSwitch = MOCT principle)	Partial (custom fabric)
Availability	Custom SoC (12–18 months)	Production today	Custom ASIC (18–24 months)
Optimization level	High (Velsanet-specific SoC)	General AI — high	Highest (purpose-built)

Chapter 7 — Cross-Platform Comparison

This chapter consolidates the layer-by-layer comparisons into a single reference view and introduces two cross-cutting analyses: the structural parallel between NVSwitch and MOCT, and the Intel silicon photonics heritage analysis.

7.1 Master Mapping Table

Velsanet Layer	ARM + Samsung	NVIDIA + TSMC	Intel + Broadcom
MOCT / Matrix	Samsung Foundry SF3/SF4	TSMC SiPho PDK + N3P + CoWoS-L	Intel Foundry 18A + SiPho PDK
Dedicated Equipment	Exynos Custom SoC + PUF	NVIDIA BlueField-3 DPU + DOCA	Intel IPU (Mount Evans) + P4
PAI Node (O8)	Cortex-M55 + Ethos-U65	Jetson Orin NX (Dual DLA)	Core Ultra NPU tile
AAI Node (D12)	Cortex-X4 + Mali-G720 + CMN-700	Grace-Hopper GH200	Xeon + Gaudi 3
AsAI Node (I20)	Neoverse V2 + Samsung HBM3E	Blackwell GB200 NVL72	Broadcom XPU ASIC

7.2 Platform Character Summary

Dimension	ARM + Samsung	NVIDIA + TSMC	Intel + Broadcom
Path type	IP licensing + foundry	System platform + foundry	Network infra-native
MOCT strength	Korean vertical integration	CPO production scale	Deepest SiPho heritage
Prototype speed	Slower (custom SoC required)	Fastest (all available today)	Fast (IPU + Core Ultra today)
Power efficiency	Best (Cortex-M55 PAI)	Highest performance	Balanced
Topology parallel	No direct analogue	NVSwitch ≈ MOCT principle	CPO ≈ MOCT principle
Governance fit	TrustZone + PUF	Confidential Computing	SGX enclaves + TDX
Strategic context	Korean domestic leadership	AI infra standard platform	Network infra incumbent

7.3 The NVSwitch ↔ MOCT Structural Parallel

The deepest cross-platform insight in this report is the structural parallel between NVIDIA NVSwitch and Velsanet MOCT. Both implement the same design principle: replace routing and switching with fixed, deterministic, parallel physical connectivity.

Property	NVIDIA NVSwitch / NVLink	Velsanet MOCT
Topology	All-to-all fixed physical links	Parallel fixed optical cores
Blocking	Non-blocking — always full BW	Non-blocking — dedicated cores
Latency	Deterministic — no routing	Deterministic — physical optical
Packet switching	None — point-to-point fixed	None — abandoned entirely
Scale	Rack scale (NVL72: 72 GPUs)	Network scale (city to global)
Origin	Solving AI compute bandwidth	Solving AI-native network

The engineering intuitions that produced NVLink — developed to solve the bandwidth and latency problems of large-scale AI training — are the same intuitions that produced the MOCT architecture for network-scale AI. The convergence is structural, not coincidental. NVSwitch does this at rack scale. Velsanet MOCT does this at network scale. The principle is the same. The scale is the difference.

7.4 The Intel CPO ↔ MOCT Fabrication Parallel

Intel's Co-Packaged Optics (CPO) and Velsanet MOCT share the same fabrication principle: silicon photonics on CMOS, factory-attached fiber, integrated with compute. The MOCT is an extension of Intel's CPO direction — not a departure from it.

Property	Intel CPO (Current)	Velsanet MOCT (Target)
Optical cores	8–16 parallel lanes	100–1000+ parallel cores
Module type	Co-packaged (fixed)	Fixed, non-pluggable
Fiber attachment	Factory-attached	Factory-fused (5m+)
MEMS alignment	Passive alignment structures	MEMS Alignment Interface (MAI)
AI channel control	Not native	Native (Channel 1 per core)
Fabrication process	Intel SiPho PDK (production)	Intel SiPho PDK (target)

Chapter 8 — Implementation Paths

This chapter maps each platform combination to a concrete prototype and deployment timeline, enabling a direct comparison of when each path can deliver working hardware.

8.1 ARM + Samsung — Custom Silicon Path

TIMELINE: 12–24 MONTHS TO FIRST PROTOTYPE

- Month 1–3: ARM IP licensing — Cortex-M55, Ethos-U65, Cortex-A720, Ethos-N78, CMN-700, Neoverse V2
- Month 3–6: Samsung System LSI engagement — Dedicated Equipment SoC specification
- Month 4–8: Samsung Foundry SiPho PDK feasibility study for MOCT
- Month 6–12: PAI node prototype on ARM Corstone reference platform
- Month 12–18: Dedicated Equipment SoC tape-out on SF4
- Month 18–24: Full PAI + Dedicated Equipment hardware demonstration
- Advantage: highest integration, Samsung vertical stack, Korean domestic leadership
- Challenge: no off-the-shelf path — everything requires custom silicon development

8.2 NVIDIA + TSMC — System Platform Path

TIMELINE: 1–3 MONTHS TO FIRST PROTOTYPE (TODAY'S HARDWARE)

- Week 1–2: Jetson Orin NX Developer Kit — PAI node prototype on existing hardware
- Week 2–4: BlueField-3 DPU developer kit — Dedicated Equipment DOCA prototype
- Month 1–2: Grace-Hopper GH200 (cloud) — AAI node DIKWEI validation
- Month 2–3: Full PAI → Dedicated Equipment → AAI chain demonstration
- Month 3–6: TSMC OIP engagement for MOCT SiPho feasibility
- Month 6–12: AsAI demonstration on GB200 NVL72
- Advantage: fastest path — all compute hardware available today
- Challenge: MOCT still requires custom fabrication; NVSwitch fabric is rack-scale not network-scale

8.3 Intel + Broadcom — Network Infra-Native Path

TIMELINE: 1–6 MONTHS TO FIRST PROTOTYPE (TODAY'S HARDWARE)

- Week 1–2: Intel Core Ultra developer platform — PAI NPU prototype
- Week 2–4: Intel IPU (Mount Evans) dev kit — Dedicated Equipment P4 prototype
- Month 1–3: Intel Gaudi 3 (Intel Developer Cloud) — AAI validation
- Month 3–6: Intel Foundry Services SiPho PDK engagement for MOCT
- Month 6–12: Broadcom Custom Silicon Solutions engagement — AsAI ASIC scoping
- Month 12–24: Broadcom XPU AsAI ASIC development (if selected)
- Advantage: IPU is the best structural match for Dedicated Equipment; Intel has deepest SiPho heritage for MOCT

- Challenge: AsAI requires either Gaudi 3 (near-term) or long Broadcom ASIC development cycle

8.4 Implementation Path Comparison

Milestone	ARM + Samsung	NVIDIA + TSMC	Intel + Broadcom
PAI prototype	6–12 months (Corstone)	1–2 weeks (Jetson kit)	1–2 weeks (Core Ultra)
Dedicated Equipment	12–18 months (SoC)	2–4 weeks (BlueField kit)	2–4 weeks (IPU kit)
AAI prototype	Custom SoC needed	1–2 months (GH200 cloud)	1–3 months (Gaudi 3 cloud)
AsAI prototype	Custom SoC needed	Available (GB200)	Gaudi 3 near-term / XPU later
MOCT feasibility	Samsung Foundry study	TSMC OIP engagement	Intel IFS engagement
Full chain demo	18–24 months	2–3 months	3–6 months

References

Velsanet White Papers — https://joa337.github.io/velsanet-whitepapers/

- WP03 — Multi-Optical-Core Transceiver (MOCT) Architecture
- WP08 — AI Architecture: DIKWEI Cognitive Cycle
- WP09 — Three-Layer AI System: PAI / AAI / AsAI
- WP10 — Network AI
- WP13 — Node Color and Identity Semantics
- WP15 — Network-Native Identity and Path Authority
- WP18 — Network AI and Global Governance Architecture
- WP20 — Node Intelligence and Optical Parallelism
- WP21 — AI Safety by Architecture

ARM References

- ARM Cortex-M55, Ethos-U65 NPU, Cortex-X4, Mali-G720, Neoverse V2, CoreLink CMN-700

Samsung References

- Samsung Foundry SF3/SF4 Process Technology, System LSI Exynos Platform, HBM3E, Samsung Advanced Package (SAP), PUF IP

NVIDIA References

- Jetson Orin NX, BlueField-3 DPU + DOCA SDK, Grace-Hopper GH200, Blackwell GB200 NVL72, NVSwitch 4.0 / NVLink 4.0

TSMC References

- TSMC Silicon Photonics PDK, N3P / N4P Process, CoWoS-L, InFO Packaging, Open Innovation Platform (OIP)

Intel References

- Intel Foundry 18A, Silicon Photonics PDK, IPU (Mount Evans / E2000), Core Ultra (Meteor Lake), Xeon Scalable, Gaudi 3, EMIB, Foveros 3D, TDX, SGX

Broadcom References

- Broadcom Custom Silicon Solutions (XPU Platform), Jericho3-AI, Google TPU / Meta MTIA (public references)