VELSANET ARCHITECTURE GROUP

Technical Platform Mapping Report

Velsanet Hardware Implementation on NVIDIA + TSMC Platform

Independent Technical Analysis by Velsanet Architecture Group

v1.0 · March 2026 · Independent Analysis

Executive Summary

This report maps the Velsanet network architecture — its five polyhedral node types, three AI layers (PAI, AAI, AsAI), photonic physical layer (MOCT), and Dedicated Equipment — to the NVIDIA compute platform and TSMC fabrication capabilities.

The NVIDIA + TSMC path is distinct in character from the ARM + Samsung mapping. Where ARM provides a power-efficient distributed edge compute path via IP licensing, NVIDIA provides a GPU-native parallel inference path with an interconnect fabric (NVLink / NVSwitch) that structurally mirrors Velsanet's polyhedral topology. Where Samsung Foundry provides MEMS and silicon photonics integration, TSMC's silicon photonics PDK and CoWoS packaging represent the current industry frontier for MOCT-class fabrication.

The core conclusion:

Velsanet Layer	Hardware Need	Platform Fit
MOCT / Matrix	Silicon photonics, MEMS, polymer waveguide fab	TSMC SiPho PDK + N3P + CoWoS-L
Dedicated Equipment	Optical↔AI channel bridge, DIKWEI gateway	NVIDIA BlueField-3 DPU + DOCA
PAI Node (O8)	Always-on edge AI, 8-ch multimodal inference	NVIDIA Jetson Orin NX (TSMC 8nm)
AAI Node (D12)	Regional multi-agent social validation	NVIDIA Grace-Hopper GH200 (NVLink-C2C)
AsAI Node (I20)	Infrastructure-scale wisdom execution	NVIDIA Blackwell GB200 NVL72 (CoWoS-L)

1. Why NVIDIA + TSMC Is a Distinct Mapping

Two structural observations make this mapping particularly compelling.

NVSwitch topology mirrors Velsanet parallel E2E

NVLink-C2C and NVSwitch form a fixed, high-bandwidth, low-latency fabric between compute dies — not packet switching. It is deterministic parallel fabric. The topology of an NVSwitch-connected Blackwell cluster (all-to-all, deterministic bandwidth) is a hardware-level analogue of Velsanet's parallel E2E path model through the MOCT. Both replace routing and switching with fixed, deterministic, parallel physical connectivity.

BlueField DPU is the closest existing product to Velsanet's Dedicated Equipment

The BlueField-3 DPU sits between the network fabric and the compute layer, performing offload, transformation, and protocol processing. Velsanet's Dedicated Equipment sits between the optical physical layer and the AI cognitive layers, performing channel mapping, format conversion, and DIKWEI gateway functions. The structural role is identical; the implementation context differs.

1.1 Velsanet Architecture Recap

Node	Shape	Role	AI Layer	Scale
T4	Tetrahedron	Access / device entry	None (passive)	Device
H6	Hexahedron	Channel alignment / RAN	None (passive)	Access
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

2. NVIDIA Platform Mapping

2.1 PAI Node (O8) — NVIDIA Jetson Orin NX

PAI nodes are always-on personal AI instances — potentially thousands running the DIKWEI cognitive cycle continuously at neighborhood scale. Requirements: ultra-low power, always-on sequential inference, 8-channel multimodal input, local DIKWEI metadata storage, hardware identity

isolation.

NVIDIA Fit: Jetson Orin NX

- 12-core ARM Cortex-A78AE CPU + Ampere GPU (1024 CUDA cores) within 10–25W
- Dual DLA (Deep Learning Accelerator) engines: always-on DIKWEI inference without GPU involvement — maps to PAI's sequential loop requirement
- NVIDIA TensorRT: optimized DIKWEI cognitive cycle model inference
- Orin SoC hardware security engine: root of trust for PAI identity — maps to Channel 1 authentication
- Unified memory architecture: CPU + GPU share memory pool — simplifies Cube state management
- TSMC 8nm fabrication: cost-effective, proven process for volume PAI node deployment

2.2 Dedicated Equipment — NVIDIA BlueField-3 DPU

The Dedicated Equipment must perform: 384-lane Cube demux, 8-channel PAI format conversion, 24-bundle AAI multiplexing, 12-face to 20-face topology mapping, and DIKWEI gateway processing — all in real time at line rate.

NVIDIA Fit: BlueField-3 DPU + DOCA platform

- 400Gb/s network interface: sufficient for 384-lane Cube parallel data ingress
- NVIDIA DOCA programmable data plane: implements custom Velsanet channel mapping and DIKWEI gateway logic without custom ASIC
- Embedded ARM Cortex-A78 compute complex: runs AAI structural topology conversion (12-face ↔ 20-face)
- Hardware crypto engine: Channel 1 physical-layer authentication
- P4 programmability: Velsanet-specific PAI→AAI→AsAI forwarding logic in data plane
- ConnectX-8 (next-gen): 800Gb/s throughput — matches MOCT Phase 3 density roadmap (1000+ cores)

2.3 AAI Node (D12) — NVIDIA Grace-Hopper GH200

AAI performs social validation and multi-agent intent mediation across flows from many distributed PAI nodes. It needs: high single-thread performance for normative reasoning, parallel processing for multi-agent intent consolidation, and high-bandwidth CPU↔GPU communication.

NVIDIA Fit: Grace-Hopper GH200 Superchip

- Grace CPU (Neoverse V2): deterministic normative reasoning — social/ethical/legal validation of PAI-interpreted intent
- H200 Hopper GPU (Transformer Engine): multi-agent intent consolidation — reconciling heterogeneous PAI streams into Structured Intent Packets

- NVLink-C2C at 900GB/s bidirectional: eliminates CPU↔GPU memory bottleneck for parallel multi-agent workloads
- 141GB HBM3e: sufficient for all active PAI intent contexts within an AAI jurisdiction
- Hopper Confidential Computing: hardware-enforced isolation between AAI governance jurisdictions — maps to Velsanet sovereignty model (WP18)
- MIG (Multi-Instance GPU): partitioned resources for independent parallel PAI intent streams

2.4 AsAI Node (I20) — NVIDIA Blackwell GB200 NVL72

AsAI is the wisdom layer — infrastructure-scale execution across I20 nodes, coordinating multiple AAI domains. No individual access path exists by architecture. This is the most compute-intensive layer, requiring the industry's highest-density inference platform.

NVIDIA Fit: Blackwell GB200 NVL72

- 72 Blackwell GPUs connected via NVSwitch at 1.8TB/s all-to-all — non-blocking, deterministic fabric
- 5th-generation NVLink: 1.8TB/s aggregate bandwidth for AsAI's multi-AAI input streams at national scale
- Blackwell Transformer Engine FP4: 30× inference throughput improvement — AsAI wisdom-layer reasoning at infrastructure scale
- NVIDIA Magnum IO: high-bandwidth storage for AsAI collective validation memory — accumulated AAI-validated patterns shaping AsAI evolution
- Hardware RAS (Reliability, Availability, Serviceability): continuous operation for infrastructure AI — no planned downtime
- TSMC CoWoS-L packaging: compute + HBM3e co-integrated on silicon interposer — same packaging trajectory as MOCT photonic integration

The NVSwitch topology deserves specific attention. In a GB200 NVL72 cluster, every GPU communicates with every other at full bandwidth simultaneously — no routing, no arbitration, no packet switching. This is deterministic parallel fabric at rack scale. It is the closest existing hardware implementation of the connectivity principle that Velsanet's MOCT defines at network scale.

2.5 NVIDIA Platform Summary

Velsanet Layer	Node	NVIDIA Platform	Key Capability Match
PAI	O8 (Octahedron)	Jetson Orin NX	Dual DLA always-on inference, HW security, 10-25W
Dedicated Equip.	O8 bridge	BlueField-3 DPU	400Gb/s fabric, DOCA programmable DIKWEI gateway
AAI	D12 (Dodecahedron)	Grace-Hopper GH200	NVLink-C2C 900GB/s, Transformer Engine, MIG isolation

AsAI	I20 (Icosahedron)	Blackwell GB200 NVL72	NVSwitch 1.8TB/s all-to-all, FP4, CoWoS-L, RAS

3. TSMC Platform Mapping

TSMC's relevance to Velsanet spans two domains: advanced logic process nodes for AI compute SoCs, and silicon photonics fabrication for the MOCT physical layer. In the photonics domain, TSMC has a specific advantage: its silicon photonics process is production-proven at scale, used for commercial photonic integrated circuits (PICs) in hyperscaler data center co-packaged optics (CPO) applications.

3.1 MOCT / Matrix — TSMC Silicon Photonics PDK

The MOCT requires sub-micron lithographic precision to integrate polymer waveguide structures, MEMS alignment fixtures, and electrical control circuits on a single substrate. TSMC's SiPho PDK directly addresses this:

- TSMC SiPho PDK: production-ready process for integrated photonic circuits — waveguides, modulators, photodetectors, grating couplers on CMOS substrate
- Photonics-on-CMOS: optical and electronic functions co-fabricated on the same wafer — maps to MOCT's requirement for integrated MEMS alignment and optical core structures
- TSMC N3P (3nm): highest-density logic process for electrical control circuits co-integrated with photonic structures
- TSMC InFO (Integrated Fan-Out): thin-form-factor multi-layer integration for MOCT polymer-form optical core layer
- TSMC CoWoS: 2.5D packaging for MOCT multi-layer matrix-transceiver integrated structure

TSMC's SiPho process has been validated in production CPO modules. The MOCT extends this domain: CPO typically supports 8–16 parallel lanes; Velsanet MOCT Phase 1 targets 100 cores, Phase 3 targets 1000+. The fabrication principle is the same; the scale is Velsanet's innovation.

3.2 AI Node SoCs — TSMC Advanced Nodes

- Jetson Orin NX (PAI node): TSMC 8nm — mature, cost-effective for volume edge AI deployment
- Grace CPU (AAI/AsAI): TSMC N4P — high-performance process for Neoverse V2 cores
- Blackwell GPU (AsAI): TSMC N4P + CoWoS-L — highest-density GPU with 2.5D HBM3e integration
- BlueField-3 DPU: TSMC 5nm — optimized for high-bandwidth network processing with embedded compute

3.3 CoWoS-L and the MOCT Integration Path

TSMC's CoWoS-L enables co-packaging of compute dies with HBM3e memory and photonic dies on a single silicon interposer. The evolution from current CoWoS-L (used in GB200) to full photonic integration is the same trajectory that leads from today's co-packaged optics to Velsanet's MOCT-integrated node.

A Velsanet node in its most integrated form would be a CoWoS assembly with compute dies, HBM memory, and MOCT photonic dies co-packaged on one interposer — exactly the trajectory TSMC's advanced packaging roadmap is developing.

3.4 TSMC Platform Summary

Velsanet Component	TSMC Technology	Specific Capability
MOCT Photonic Layer	SiPho PDK + N3P	Production silicon photonics — waveguides, MEMS, photodetectors on CMOS
MOCT Multi-layer Integration	CoWoS + InFO	2.5D co-packaging of photonic + electronic + MEMS dies
PAI Node SoC	N8 (8nm FinFET)	Jetson Orin process — volume AI SoC with proven yield
AAI / AsAI Compute	N4P (4nm)	Grace CPU + Blackwell GPU — highest-performance AI compute
AsAI Memory Integration	CoWoS-L	GB200 HBM3e co-packaging — 1.8TB/s memory bandwidth per node

4. Structural Parallel: NVSwitch vs. Velsanet MOCT

The deepest architectural parallel between NVIDIA's platform and Velsanet is the topology of their respective interconnect fabrics.

Property	NVSwitch / NVLink	Velsanet MOCT
Blocking	Non-blocking all-to-all	Non-blocking — each optical core is dedicated
Latency	Deterministic — no routing decisions	Deterministic — physical optical propagation
Topology	Fixed physical links — does not reconfigure	Fixed optical cores — permanent structures
Packet switching	None — NVLink is point-to-point fixed fabric	None — Velsanet abandons packet switching entirely

Scale	Rack scale (NVL72: 72 GPUs)	Network scale (city to global)
Purpose	AI training / inference bandwidth	AI-native network E2E connectivity

Both NVSwitch and Velsanet MOCT embody the same architectural principle: replace routing and switching with fixed, deterministic, parallel physical connectivity. The engineering intuitions behind NVLink — developed to solve AI compute bandwidth problems — are the same intuitions that produced the MOCT architecture for network-scale AI. The convergence is structural, not superficial.

---

5. Partnership Model

5.1 Division of Roles

Layer	NVIDIA Role	TSMC Role	Joint Opportunity
MOCT / Physical	—	SiPho PDK + CoWoS	TSMC-exclusive photonics fab path to MOCT
Dedicated Equipment	BlueField-3 DPU (DOCA)	5nm fabrication	DOCA-programmed DIKWEI gateway — near-term prototype
PAI (O8)	Jetson Orin NX	N8 volume production	Jetson PAI node — available today for prototype
AAI (D12)	Grace-Hopper GH200	N4P fabrication	GH200 regional AAI node — NVLink-C2C maps PAI→AAI BW
AsAI (I20)	Blackwell GB200 NVL72	CoWoS-L packaging	GB200 NVL72 as AsAI rack — NVSwitch = structural E2E

5.2 Near-Term Prototype Path (Today's Hardware)

The NVIDIA + TSMC path has a significant near-term advantage: NVIDIA's products exist today and are production-proven. A Velsanet prototype can start now:

1. **PAI node:** Jetson Orin NX Developer Kit — available today. Implement DIKWEI cognitive cycle as CUDA/DLA inference pipeline. Validate 8-channel multimodal input and Channel 1 identity model.
2. **Dedicated Equipment:** BlueField-3 DPU developer kit — available today. Implement channel mapping and DIKWEI gateway logic in DOCA. Validate Cube→PAI demux and PAI→AAI mux in programmable data plane.
3. **AAI node:** Grace-Hopper GH200 available via cloud (CoreWeave, Lambda). Implement AAI social validation and multi-agent intent consolidation as Grace + Hopper pipeline.

4. **MOCT feasibility:** TSMC OIP (Open Innovation Platform) engagement for SiPho PDK access — assess MOCT polymer waveguide requirements on production TSMC processes.

The NVIDIA + TSMC path demonstrates that Velsanet nodes can be implemented using existing AI infrastructure technologies, enabling near-term prototyping without requiring new semiconductor architectures.

---

6. Proposed Next Steps

5. **Technical Review Meeting —** Joint session with NVIDIA Korea and TSMC to validate the hardware mapping and identify specification gaps, particularly around BlueField DPU programmability for DIKWEI gateway implementation.
6. **PAI + Dedicated Equipment Prototype —** Develop functional prototype using Jetson Orin NX + BlueField-3 DPU: Cube data ingress → PAI DIKWEI inference → AAI forwarding. Demonstrate the full cognitive chain on existing NVIDIA hardware.
7. **TSMC SiPho Feasibility Study —** TSMC OIP engagement: SiPho PDK compatibility with MOCT polymer waveguide requirements, MEMS integration path, CoWoS packaging design for multi-layer MOCT structure.
8. **NVSwitch ↔ MOCT Topology Study —** Joint technical paper documenting the architectural parallel between NVSwitch all-to-all fabric and Velsanet MOCT parallel E2E — both as instances of the deterministic parallel fabric principle for AI-native compute and networks.
9. **ITU-T SG13 Joint Contribution —** Submission referencing NVIDIA + TSMC as a candidate demonstration platform for Velsanet's IMT-2030 AI-native network node architecture.

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)
- WP18 — Network AI and Global Governance Architecture
- WP20 — Node Intelligence and Optical Parallelism
- WP21 — AI Safety by Architecture

NVIDIA References

- NVIDIA Jetson Orin NX Product Brief & Developer Kit

- NVIDIA BlueField-3 DPU Architecture Overview + DOCA SDK
- NVIDIA Grace-Hopper GH200 Superchip Architecture
- NVIDIA Blackwell GB200 NVL72 Technical Brief
- NVIDIA NVSwitch 4.0 and NVLink 4.0 Technical Overview

TSMC References

- TSMC Silicon Photonics (SiPho) Process Technology
- TSMC N3P / N4P Process Node Overview
- TSMC CoWoS and CoWoS-L Advanced Packaging
- TSMC Open Innovation Platform (OIP) — SiPho PDK
