VELSANET ARCHITECTURE GROUP

Technical Platform Mapping Report

Velsanet Hardware Implementation on Intel + Broadcom Platform

Independent Technical Analysis by Velsanet Architecture Group

v1.0 · March 2026 · Independent Analysis

Executive Summary

This is the third in a series of independent technical mapping reports examining how Velsanet's architecture maps to major semiconductor and systems platforms. The first report covered ARM + Samsung (IP licensing + foundry path). The second covered NVIDIA + TSMC (system platform + photonics foundry path). This report covers Intel + Broadcom — a network infrastructure-native path with a distinct and historically significant connection to silicon photonics.

Intel's relevance to Velsanet is unique among all platform options: Intel is the originator of silicon photonics as a manufacturing discipline, having developed the first silicon photonic modulators and waveguide structures in the early 2000s. Intel's IPU (Infrastructure Processing Unit) is structurally the closest existing product to Velsanet's Dedicated Equipment. Broadcom, as the dominant supplier of network switching silicon and custom AI ASIC platforms (Google TPU, Meta MTIA), addresses the network fabric and AsAI compute layers — though with important structural caveats around its packet-switching heritage.

The core conclusion:

Velsanet Layer	Hardware Need	Platform Fit
MOCT / Matrix	Silicon photonics, MEMS, polymer waveguide fab	Intel Foundry 18A + Intel Silicon Photonics PDK
Dedicated Equipment	Optical↔AI channel bridge, DIKWEI gateway	Intel IPU (Mount Evans) + P4 programmable pipeline
PAI Node (O8)	Always-on edge AI, 8-ch multimodal inference	Intel Core Ultra (Meteor Lake) — Intel NPU
AAI Node (D12)	Regional multi-agent social validation	Intel Xeon + Gaudi 3 AI Accelerator
AsAI Node (I20)	Infrastructure-scale wisdom execution	Broadcom XPU custom AI ASIC + HBM3

Network Fabric	Inter-node connectivity (with structural caveat)	Broadcom Jericho3-AI (packet-switching caveat applies)

1. Why Intel + Broadcom Is a Distinct Mapping

The ARM + Samsung mapping is an IP licensing path. The NVIDIA + TSMC mapping is a GPU-native parallel inference path. The Intel + Broadcom mapping is a network infrastructure-native path — rooted in the two companies that have historically dominated network silicon and data center interconnect.

Three observations define this mapping's character:

Intel is the originator of silicon photonics

Intel's silicon photonics research began in 2000 and produced the world's first continuous-wave silicon Raman laser in 2005. Intel's current silicon photonics products — used in 400G and 800G optical transceivers — are manufactured on a dedicated SiPho process that predates TSMC's commercial offering. For Velsanet's MOCT, Intel Foundry's silicon photonics heritage is the deepest available in the industry.

Intel IPU is structurally the closest product to Velsanet's Dedicated Equipment

Intel's IPU (Infrastructure Processing Unit), codenamed Mount Evans, was designed to offload, accelerate, and isolate infrastructure services from host compute. It sits between the network fabric and host processors — exactly the structural position of Velsanet's Dedicated Equipment between the MOCT optical layer and the PAI/AAI/AsAI cognitive layers.

Broadcom's network switching heritage is both an asset and a structural tension

Broadcom supplies the switching silicon inside most of the world's data center networks (Tomahawk, Trident, Jericho series). However, all of these products are packet-switching architectures — fundamentally incompatible with Velsanet's packet-free E2E model. The mapping therefore distinguishes between Broadcom's AI ASIC capabilities (XPU platform, used for Google TPU and Meta MTIA) which are directly applicable, and its switching products which are structurally incompatible with Velsanet's physical layer.

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. Intel Platform Mapping

2.1 MOCT / Matrix — Intel Foundry Silicon Photonics

Intel's silicon photonics heritage makes it uniquely positioned for MOCT fabrication. Intel has been manufacturing integrated photonic circuits since the mid-2000s and currently produces 400G/800G co-packaged optics (CPO) modules using its SiPho process. Key capabilities:

Intel Foundry 18A + Intel SiPho PDK

- Intel 18A (1.8nm-class RibbonFET GAA): Intel's most advanced logic process — provides the dimensional control required for MOCT polymer waveguide and MEMS alignment structure fabrication
- Intel Silicon Photonics PDK: dedicated process design kit for integrated optical circuits — waveguides, ring modulators, Ge photodetectors, grating couplers — all co-fabricated with CMOS electronics
- Intel's CPO (Co-Packaged Optics) program: production experience integrating photonic dies with compute packages — directly applicable to MOCT's requirement for optical + compute co-integration
- Intel EMIB (Embedded Multi-die Interconnect Bridge): 2.5D die-to-die interconnect for MOCT multi-layer matrix-transceiver integration without full interposer
- Intel Foveros 3D: vertical die stacking for MOCT's layered polymer-form optical core + electrical control circuit integration

Intel's silicon photonics process is notable for one specific property: it is the only major SiPho platform developed by a company that also operates a leading-edge logic foundry, manufactures CPO products at volume, and designs the compute silicon that will interface with those photonic components. This vertical integration is directly relevant to MOCT fabrication.

2.2 Dedicated Equipment — Intel IPU (Mount Evans)

The Intel IPU (Infrastructure Processing Unit) is the most architecturally aligned existing product to Velsanet's Dedicated Equipment concept. It was designed specifically to offload infrastructure processing from host compute — performing network, storage, and security functions in a dedicated

silicon layer between the fabric and the server.

Intel IPU (Mount Evans) — structural fit analysis

- Fabric-to-compute bridge: IPU sits between network fabric and host CPU — same structural position as Dedicated Equipment between MOCT optical layer and PAI/AAI cognitive layers
- P4-programmable data plane: custom packet/data processing logic programmable in P4 — Velsanet channel mapping (Cube→PAI demux, PAI→AAI mux) implementable as P4 pipeline
- ARM Neoverse cores embedded: host-side compute for DIKWEI gateway logic — structural topology conversion (12-face AAI ↔ 20-face AsAI)
- 100Gb/s+ network interface: handles 384-lane Cube parallel data ingress
- Intel TDX (Trust Domain Extensions): hardware-enforced isolation — maps to Channel 1 identity and PAI sovereignty boundary
- Intel E2000 (Mount Evans SoC): purpose-built for infrastructure offload — available today as production silicon

2.3 PAI Node (O8) — Intel Core Ultra (Meteor Lake)

The PAI layer requires always-on, low-power, multimodal inference at neighborhood scale. Intel's Core Ultra (Meteor Lake) introduced a dedicated NPU (Neural Processing Unit) tile — Intel's first client-side AI accelerator, designed precisely for always-on ambient intelligence workloads.

Intel Core Ultra — NPU tile

- Intel NPU (Neural Processing Unit): dedicated AI inference engine separate from CPU and GPU — enables always-on DIKWEI cognitive cycle without activating main CPU cores
- Heterogeneous tile architecture (CPU + GPU + NPU + SoC tile): maps to PAI's requirement for parallel multimodal input processing across 8 cognitive channels
- Intel Thread Director: workload-aware scheduling between P-cores, E-cores, and NPU — optimizes DIKWEI inference loop placement
- Intel TDX / TME (Total Memory Encryption): PAI identity isolation and personal context protection
- Intel 4 process (TSMC N3 equivalent class): compact, power-efficient SoC for distributed PAI node deployment
- Low TDP variants (9–28W): feasible for always-on neighborhood-scale deployment

2.4 AAI Node (D12) — Intel Xeon + Gaudi 3

AAI requires high-performance multi-agent reasoning at regional scale. Intel's Xeon + Gaudi 3 combination provides the CPU-AI accelerator pairing needed for AAI's social validation and intent mediation workloads.

Intel Xeon Scalable (Sapphire Rapids / Emerald Rapids) + Gaudi 3

- Intel Xeon: enterprise-grade CPU for AAI's deterministic normative reasoning logic — social/ethical/legal validation of PAI-interpreted intent
- Intel Gaudi 3 AI Accelerator: 128GB HBM2e per card, 8× 200Gb/s Ethernet fabric — designed for LLM inference at scale, maps to AAI multi-agent intent consolidation
- Intel Gaudi's Ethernet-based fabric: unlike NVLink, Gaudi uses standard Ethernet for inter-accelerator communication — lower proprietary lock-in, relevant for distributed AAI deployments across governance jurisdictions
- Intel AMX (Advanced Matrix Extensions): CPU-side matrix operations for Structured Intent Packet formation
- Intel SGX (Software Guard Extensions): hardware-enforced enclaves for AAI governance boundary isolation — maps to Velsanet sovereignty model (WP18)

2.5 Intel Platform Summary

Velsanet Layer	Node	Intel Platform	Key Capability Match
MOCT	Physical substrate	Intel Foundry 18A + SiPho PDK	Originator SiPho heritage, CPO production experience, EMIB/Foveros
Dedicated Equip.	O8 bridge	Intel IPU (Mount Evans)	P4 programmable fabric-to-compute bridge, TDX isolation
PAI	O8 (Octahedron)	Core Ultra (Meteor Lake) NPU	Always-on dedicated NPU tile, heterogeneous tiles, TME
AAI	D12 (Dodecahedron)	Xeon + Gaudi 3	Open Ethernet fabric, Gaudi 3 HBM2e, SGX enclave isolation

3. Broadcom Platform Mapping

Broadcom's role in this mapping is more nuanced than Intel's. As the dominant supplier of network switching silicon, Broadcom's core products are packet-switching architectures — structurally incompatible with Velsanet's packet-free E2E model. However, Broadcom also operates one of the most sophisticated custom AI ASIC platforms in the industry, responsible for fabricating Google's TPU series and Meta's MTIA. This dual character requires a careful separation of applicable and inapplicable capabilities.

3.1 The Structural Caveat: Broadcom Switching vs. Velsanet MOCT

Broadcom's Tomahawk, Trident, and Jericho switching ASICs are the backbone of most data center and carrier networks. They are designed for packet forwarding: ingress → lookup → egress, with buffering, scheduling, and routing logic at every stage.

Velsanet abandons packet switching entirely. The MOCT provides deterministic, parallel, physically isolated E2E optical paths — no lookup tables, no buffers, no routing decisions. Broadcom's switching products are therefore structurally inapplicable to Velsanet's physical layer.

This is not a limitation of Broadcom as a company. It is a fundamental architectural distinction. Velsanet and packet-switching networks solve different problems at the physical layer.

3.2 AsAI Node (I20) — Broadcom XPU Custom AI ASIC

Broadcom's custom AI ASIC platform (internally referred to as XPU) is a different business from its switching silicon. Through its Custom Silicon Solutions division, Broadcom has designed and manufactured the Google TPU v4/v5, Meta MTIA, and ByteDance custom AI accelerators — all large-scale, high-bandwidth inference and training silicon with HBM memory integration.

Broadcom XPU (Custom AI ASIC) — AsAI fit

- Custom silicon design: Broadcom can design a purpose-built AsAI accelerator optimized for Velsanet's DIKWEI wisdom-layer inference — not a general-purpose GPU but a task-specific AI ASIC
- HBM3 integration: Broadcom's XPU platform integrates HBM stacks directly — AsAI requires the highest memory bandwidth available for aggregated multi-AAI wisdom inference
- Scale-out fabric: Broadcom's custom AI ASICs use high-bandwidth die-to-die interconnects for multi-chip AsAI cluster formation — analogous to NVSwitch but implementable as custom logic
- TSMC / Intel Foundry fabrication: Broadcom's XPU products are typically fabricated on TSMC N3/N5 or Intel Foundry — complementary to the photonics fabrication path
- Proven delivery track record: Google TPU, Meta MTIA — Broadcom has delivered hyperscale AI silicon at production volume

3.3 Jericho3-AI — Limited Applicability with Caveat

Broadcom's Jericho3-AI is a network routing ASIC designed for AI traffic patterns — supporting RDMA, RoCE, and high-bandwidth elephant flows in AI training clusters. It represents Broadcom's attempt to bridge network switching and AI fabric.

Jericho3-AI — applicable scope and hard limit

- Applicable: Jericho3-AI's AI-optimized traffic scheduling could handle the inter-node communication between AAI and AsAI layers in a transitional deployment scenario where full MOCT implementation is not yet available
- Hard limit: Jericho3-AI is a packet-switching ASIC. It cannot replace Velsanet's MOCT physical layer. Any deployment using Jericho3-AI for physical connectivity would be a transitional implementation — not the target Velsanet architecture

- Value in transition: as a bridge technology enabling AAI↔AsAI communication during prototype phases before full MOCT deployment, Jericho3-AI has practical utility

3.4 Broadcom Platform Summary

Velsanet Component	Broadcom Product	Assessment
MOCT / Physical Layer	Tomahawk / Trident	NOT APPLICABLE — packet-switching architecture is structurally incompatible with Velsanet's packet-free E2E model
AsAI Compute (I20)	XPU Custom AI ASIC	DIRECTLY APPLICABLE — custom AI ASIC + HBM3, proven at Google TPU / Meta MTIA scale
Inter-node Fabric (transitional)	Jericho3-AI	LIMITED — AI-optimized packet switching, applicable only as transitional bridge before full MOCT deployment
Custom Node SoC	Custom Silicon Solutions	APPLICABLE — Broadcom can design purpose-built Velsanet node ASICs (PAI, AAI, AsAI) on TSMC / Intel Foundry

---

4. The Intel Silicon Photonics Advantage

Across all four platform mapping reports in this series, Intel's silicon photonics heritage represents the single most historically significant connection to Velsanet's MOCT architecture. It deserves dedicated analysis.

4.1 Intel's Silicon Photonics Timeline

- 2000: Intel Research begins silicon photonics program
- 2004: First silicon-based optical modulator operating at GHz speeds
- 2005: World's first continuous-wave silicon Raman laser
- 2010: 50Gbps silicon photonic link demonstrated
- 2016: Intel ships 100G PSM4 silicon photonics transceivers at volume
- 2020: Intel introduces co-packaged optics (CPO) research platform
- 2023: Intel 800G silicon photonics transceiver in production
- 2025+: Intel Foundry SiPho PDK available for external customers via Intel Foundry Services

No other company in the semiconductor industry has a deeper or longer silicon photonics manufacturing history. For Velsanet's MOCT — which requires production-grade silicon photonics fabrication at densities (100–1000+ cores) that exceed current commercial transceivers — Intel

Foundry's SiPho expertise is the most relevant available.

4.2 MOCT vs. Intel CPO: Same Principle, Different Scale

Property	Intel CPO (Current)	Velsanet MOCT (Target)
Optical cores	8–16 parallel lanes	100–1000+ parallel cores
Pluggable	No (co-packaged)	No (fixed, non-pluggable)
Fiber attachment	Factory-attached	Factory-fused (5m+)
MEMS alignment	Passive alignment structures	MEMS Alignment Interface (MAI)
AI control	Not native	Native (Channel 1)
Fabrication	Intel SiPho PDK	Intel SiPho PDK (target)

The MOCT is not a departure from Intel's CPO direction — it is an extension of it. Intel CPO and Velsanet MOCT share the same fabrication principle (silicon photonics on CMOS), the same integration approach (co-packaged with compute), and the same reliability philosophy (factory-attached, no field connectors). The MOCT's innovation is scale and AI-native control — not a different manufacturing paradigm.

---

5. Partnership Model

5.1 Division of Roles

Layer	Intel Role	Broadcom Role	Joint Opportunity
MOCT / Physical	SiPho PDK + 18A + EMIB/Foveros	—	Intel-exclusive: deepest SiPho heritage, CPO production experience
Dedicated Equipment	Intel IPU (Mount Evans)	—	IPU as DIKWEI gateway prototype — available today
PAI (O8)	Core Ultra NPU tile	—	NPU-based PAI node — widely available, low cost
AAI (D12)	Xeon + Gaudi 3	—	Open Ethernet fabric enables multi-jurisdiction AAI deployment
AsAI (I20)	Gaudi 3 cluster	XPU Custom AI ASIC	Broadcom XPU for purpose-built AsAI ASIC; Intel Gaudi 3 for near-term deployment

Transitional Fabric	—	Jericho3-AI (limited)	Bridge technology only — pre-MOCT transitional inter-node connectivity

5.2 Near-Term Prototype Path

Like the NVIDIA + TSMC path, the Intel + Broadcom path offers near-term prototype opportunities using existing production hardware:

1. **PAI node:** Intel Core Ultra developer platform (available today) — implement DIKWEI cognitive cycle on Intel NPU tile using OpenVINO inference runtime. Validate 8-channel multimodal input and Channel 1 identity model.
2. **Dedicated Equipment:** Intel IPU (Mount Evans) development kit — implement Velsanet channel mapping and DIKWEI gateway in P4 programmable pipeline. Validate Cube→PAI demux and PAI→AAI mux logic.
3. **AAI node:** Intel Gaudi 3 cluster (available via Intel Developer Cloud) — implement AAI social validation and multi-agent intent consolidation. Validate DIKWEI AAI layer at regional simulation scale.
4. **MOCT feasibility:** Intel Foundry Services engagement for SiPho PDK access — assess MOCT polymer waveguide and MEMS alignment requirements against Intel's production SiPho process design rules.
5. **AsAI ASIC scoping:** Broadcom Custom Silicon Solutions engagement — define AsAI ASIC requirements (DIKWEI wisdom-layer inference, HBM3 bandwidth, multi-chip fabric) for a purpose-built Velsanet AsAI chip.

The Intel + Broadcom 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

6. **Technical Review Meeting —** Joint session with Intel Foundry and Broadcom Custom Silicon Solutions to validate the hardware mapping and identify specification gaps, particularly around Intel SiPho PDK compatibility with MOCT polymer waveguide requirements and Broadcom XPU design parameters for AsAI.
7. **Intel IPU Prototype —** Develop Dedicated Equipment prototype on Intel IPU (Mount Evans): implement DIKWEI gateway and Velsanet channel mapping in P4 data plane. Validate Cube→PAI→AAI transformation pipeline on production hardware.
8. **Intel SiPho Feasibility Study —** Intel Foundry Services engagement: assess MOCT fabrication requirements against Intel SiPho PDK — waveguide density, MEMS integration path, EMIB/Foveros packaging for multi-layer MOCT structure.

9. **Broadcom AsAI ASIC Scoping —** Engage Broadcom Custom Silicon Solutions for AsAI ASIC architecture definition: DIKWEI wisdom-layer inference requirements, HBM3 memory subsystem, multi-chip scale-out fabric for I20 node.
10. **ITU-T SG13 Joint Contribution —** Submission referencing Intel + Broadcom as a candidate implementation path for Velsanet's IMT-2030 AI-native network node architecture, with specific emphasis on Intel's silicon photonics heritage as a MOCT fabrication foundation.

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

Intel References

- Intel Silicon Photonics Technology Overview
- Intel Foundry 18A / Intel 4 Process Technology
- Intel IPU (Mount Evans) Architecture — Infrastructure Processing Unit
- Intel Core Ultra (Meteor Lake) NPU Technical Brief
- Intel Gaudi 3 AI Accelerator Architecture
- Intel EMIB and Foveros 3D Packaging Technology

Broadcom References

- Broadcom Custom Silicon Solutions — XPU Platform
- Broadcom Jericho3-AI Routing ASIC
- Broadcom Tomahawk 5 / Trident 5 — Network Switching ASIC
