VELSANET ARCHITECTURE GROUP

# Technical Platform Mapping Report

## Velsanet Hardware Implementation on ARM + Samsung Platform

Independent Technical Analysis by Velsanet Architecture Group

v1.0 · March 2026 · Independent Analysis

## Executive Summary

Velsanet is a next-generation network architecture built on three structural foundations: a topology-first polyhedral node hierarchy, a deterministic parallel optical physical layer (MOCT), and a three-layer AI system embedded natively within the network. Each of these components has specific, mappable hardware requirements.

This report defines how each layer of the Velsanet architecture maps to available hardware platforms — specifically ARM IP families and Samsung semiconductor and manufacturing capabilities — and identifies the partnership opportunities that arise from those mappings.

**The core conclusion is this:**

| Velsanet Layer | Hardware Need | Platform Fit |
|---|---|---|
| MOCT / Matrix (Physical) | Silicon photonics, MEMS, polymer waveguide fab | Samsung Foundry (SF3 / SF4) |
| Dedicated Equipment | High-bandwidth optical↔AI channel bridge SoC | ARM Cortex-A + Ethos NPU + Samsung Exynos |
| PAI Node (O8) | Ultra-low-power distributed edge AI processor | ARM Cortex-M55 + Ethos-U65 cluster |
| AAI Node (D12) | Regional multi-agent coordination processor | ARM Cortex-X4 + Mali GPU + CMN mesh |
| AsAI Node (I20) | High-throughput wisdom-layer inference engine | ARM Neoverse V2 cluster / Samsung HBM |

# 1. Velsanet Architecture Overview

Velsanet replaces packet-switched networking with a structure-first, AI-native architecture. 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 not bolted onto the network — it is the network.

## 1.1 The Five-Node Polyhedral Hierarchy

Every Velsanet node corresponds to a specific polyhedron and occupies a defined structural role:

| 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 |

## 1.2 The Physical Layer: MOCT

The Multi-Optical-Core Transceiver (MOCT) is the physical substrate of Velsanet. It is a fixed, non-pluggable photonic module permanently integrated with the Matrix. Key properties:

- 100 to 1000+ parallel optical cores per node
- 8 internal channels per core — Channel 1 reserved for sensing, control, and identity
- MEMS-based alignment — micron-level precision, no active switching
- Pre-attached fibers (5m+) fused at manufacture — zero connector loss
- Semiconductor lithography alignment — reliability profile of a packaged IC, not a telecom connector

## 1.3 The Intelligence Layer: PAI → AAI → AsAI

Intelligence in Velsanet follows the DIKWEI cognitive cycle: Data → Information → Knowledge → Wisdom → Execution → Intent. This cycle is executed across three AI layers, each physically bound to a specific node type:

- PAI (O8 node): personal intent interpretation. Accumulates individual history, values, behavioral patterns. Does not execute — only interprets and represents. Runs in parallel across all O8 nodes simultaneously.

- AAI (D12 node): social validation gateway. Evaluates PAI-interpreted intent against ethical, legal, and policy constraints. No intent reaches AsAI without AAI authorization.
- AsAI (I20 node): network-scale execution. Receives only AAI-validated, authorized requests. No direct individual access path exists — structurally impossible by architecture.

### 1.4 The Dedicated Equipment

The Dedicated Equipment is the synaptic transformer of Velsanet — the hardware bridge between the optical physical layer and the AI cognitive layers. Its four core responsibilities:

| 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 between 12-face (AAI) and 20-face (AsAI) topologies |
| AsAI → Cube Write-back | Converts AsAI results back into Cube-addressable DIKWEI memory packets |

---

## 2. ARM Platform Mapping

ARM's IP portfolio covers the full compute spectrum from sub-milliwatt embedded cores to data-center-grade server processors. Velsanet's distributed, heterogeneous intelligence architecture maps precisely across this spectrum.

### 2.1 PAI Node — ARM Cortex-M55 + Ethos-U65

The PAI layer runs on O8 nodes — potentially thousands of units operating in parallel at personal/neighborhood scale. The hardware requirements are:

- Ultra-low power: PAI nodes are always-on, accumulating individual behavioral context
- Always-on sequential inference: DIKWEI cycle runs continuously, not on-demand
- 8-channel parallel input processing from the Dedicated Equipment interface
- Multimodal signal handling: audio, vision, sensor, intent streams
- Local memory for DIKWEI metadata and Cube-plane state

#### ARM Fit: Cortex-M55 + Ethos-U65 ML Accelerator

The Cortex-M55 is ARM's most ML-capable Cortex-M processor, featuring Helium (M-Profile Vector Extension) for DSP and ML workloads. The Ethos-U65 NPU provides dedicated neural network

inference at sub-1W power envelopes. Together, this combination provides:

- Helium SIMD: parallel processing of the 8-channel PAI cognitive input stream
- Ethos-U65: DIKWEI inference cycle acceleration without main CPU involvement
- TrustZone-M: structural isolation of identity and personal context — maps directly to Velsanet's PAI sovereignty model
- Sub-500mW operation in active inference mode — feasible for always-on distributed deployment

## 2.2 Dedicated Equipment — ARM Cortex-A + Ethos-N78 + CMN

The Dedicated Equipment is the highest-bandwidth compute element in the Velsanet node. It must perform real-time optical↔cognitive channel mapping across 384 Cube lanes simultaneously. Requirements:

- 384-lane optical core demultiplexing at line rate
- 8-channel PAI format conversion (Cube → PAI path)
- 24-bundle multiplexing (PAI → AAI path)
- 12-face to 20-face structural topology mapping (AAI ↔ AsAI)
- DIKWEI gateway: first and last gate of the cognitive cycle

### ARM Fit: Cortex-A720 cluster + Ethos-N78 + CoreLink CMN-700

The Cortex-A720 provides high single-thread performance for the deterministic channel mapping logic. The Ethos-N78 handles the neural inference components of the DIKWEI gateway. The CoreLink CMN-700 interconnect fabric provides the high-bandwidth, low-latency mesh needed to move data between the optical interface and the AI processing pipeline at the required throughput.

- CMN-700 mesh: handles 384-lane parallel data movement without bottlenecks
- Ethos-N78: DIKWEI cognitive cycle gate inference — Cube-to-PAI meaning extraction
- Cortex-A720: deterministic structural mapping logic (topology face conversion)
- GIC-700: interrupt management for per-core optical channel events

## 2.3 AAI Node — ARM Cortex-X4 + Mali-G720 + CMN

AAI operates at the D12 (Dodecahedron) node — a regional coordination layer serving multiple O8/PAI nodes. It performs social validation, intent mediation, and multi-agent cognitive consolidation. Compute requirements are significantly higher than PAI:

- Multi-agent parallel processing: receives intent flows from regionally distributed PAI nodes simultaneously
- Normative reasoning engine: evaluates intent against social, ethical, legal, and policy constraints
- Structural alignment: reconciles heterogeneous PAI intents into coherent Structured Intent Packets
- AAI ↔ AAI prohibition enforcement: lateral links require mediated verification

### ARM Fit: Cortex-X4 + Mali-G720 GPU + CMN-700

The Cortex-X4 provides the peak single-thread performance needed for complex normative reasoning. The Mali-G720 GPU handles the parallel matrix operations required for multi-agent intent consolidation. CMN-700 interconnect scales to the AAI node's multi-PAI ingress bandwidth requirements.

- Cortex-X4 big-core: complex judgment logic for social/ethical validation
- Mali-G720: parallel processing of multi-origin PAI intent streams
- CMN-700: scalable mesh for N-face multi-channel ingress from O8 nodes
- TrustZone: governance boundary enforcement between AAI jurisdiction domains

## 2.4 AsAI Node — ARM Neoverse V2

AsAI is the highest layer of the Velsanet AI architecture, operating at I20 (Icosahedron) nodes at national and global scale. It executes at infrastructure scale, coordinating across multiple AAI instances. This is the most compute-intensive layer:

- Wisdom-layer inference: highest-order reasoning across validated intent patterns
- Network-scale execution: coordinates resources and actions across multiple AAI domains
- Collective evolution: updates itself based on aggregated AAI-validated request patterns
- No individual access path: structurally receives only AAI-authorized requests

### ARM Fit: Neoverse V2 (server-class)

The Neoverse V2 is ARM's highest-performance server processor, designed for cloud and infrastructure workloads. AsAI's requirements match data-center-grade compute: high core count, high memory bandwidth, SVE2 for large-scale vector inference, and RAS (Reliability, Availability, Serviceability) for continuous operation.

- SVE2 vector engine: large-scale DIKWEI wisdom-layer inference
- High core count: parallel processing of multi-AAI input streams
- RAS features: continuous operation required for infrastructure-level AI
- CCIX/PCIe 5.0: high-bandwidth interface to Samsung HBM memory subsystem

## 2.5 ARM Platform Summary

| Velsanet Layer | Node | ARM IP | Key Capability Match |
|---|---|---|---|
| **PAI** | O8 (Octahedron) | Cortex-M55 + Ethos-U65 | Always-on edge inference, TrustZone identity isolation |
| **Dedicated Equip.** | O8 bridge | Cortex-A720 + Ethos-N78 + CMN-700 | 384-lane optical↔AI mapping, DIKWEI gateway |
| **AAI** | D12 (Dodecahedron) | Cortex-X4 + Mali-G720 + CMN-700 | Multi-agent intent mediation, social validation |
| **AsAI** | I20 (Icosahedron) | Neoverse V2 | Infrastructure-scale wisdom inference, SVE2 |

# 3. Samsung Platform Mapping

Samsung's capabilities span the full vertical stack from semiconductor fabrication through system integration. This makes Samsung the only Korean partner capable of addressing both the photonic manufacturing layer (MOCT/Matrix) and the compute silicon layer (Dedicated Equipment, AI nodes) within a single engagement.

## 3.1 MOCT and Matrix — Samsung Foundry

The MOCT is not a conventional optoelectronic component. It is a semiconductor-grade photonic substrate. Its manufacturing requirements are:

- Silicon photonics fabrication: polymer waveguide structures require sub-micron lithographic precision
- MEMS fabrication: the MEMS Alignment Interface (MAI) requires surface MEMS process integration
- Multi-material integration: polymer optical cores, MEMS structures, and electrical control circuits on a single substrate
- Permanent fiber fusion: 5m+ pre-attached fibers require precision fusion splicing at wafer level

### Samsung Foundry Fit: SF3 / SF4 (3nm / 4nm GAA/FinFET)

Samsung Foundry's advanced node processes provide the lithographic precision required for MOCT fabrication. Specifically:

- Samsung's silicon photonics PDK (Process Design Kit) supports integrated waveguide structures
- SF4 GAA process offers the dimensional control required for MEMS alignment structure fabrication
- Samsung's heterogeneous integration packaging (H-Cube, X-Cube) enables multi-die stacking of photonic and electronic layers
- Samsung's MEMS foundry capability provides process integration for the MAI static alignment fixtures

## 3.2 Dedicated Equipment SoC — Samsung System LSI

The Dedicated Equipment is the highest-value custom silicon opportunity in the Velsanet hardware stack. It requires a custom SoC that integrates:

- High-bandwidth optical interface IP: 384-lane parallel optical core termination
- ARM compute cluster: Cortex-A720 + Ethos-N78 (licensable from ARM)
- CMN-700 interconnect fabric (ARM licensable)
- LPDDR5X / HBM memory interface for DIKWEI state storage

- Hardware security module: identity key storage for Channel 1 authentication

### Samsung System LSI Fit: Exynos custom SoC platform

Samsung System LSI has an established track record of integrating ARM compute IP, custom NPU blocks (Samsung NPU), and high-bandwidth memory interfaces into production SoCs. The Dedicated Equipment SoC would leverage:

- Exynos SoC platform: ARM Cortex-A720 + Samsung NPU (or Ethos-N78 license) integration
- Samsung HBM3 / LPDDR5X PHY: provides the memory bandwidth required for 384-lane Cube state management
- Samsung PUF (Physical Unclonable Function) IP: maps to Velsanet's manufacturing-time identity model for Channel 1 authentication
- Samsung 5G modem integration path: Dedicated Equipment → wireless interface for H6/T4 node connectivity

## 3.3 AI Node SoCs — Samsung HBM + Packaging

The AAI and AsAI nodes require high-bandwidth memory subsystems to support multi-agent cognitive workloads. Samsung is the world's leading HBM manufacturer:

- HBM3 / HBM3E: provides the memory bandwidth required for AsAI wisdom-layer inference at I20 scale
- Samsung V-NAND: Cube memory layer storage — DIKWEI metadata persistence across cognitive cycles
- 2.5D HBM packaging (HBM + Logic die): co-packaging of ARM Neoverse V2 with Samsung HBM3E for AsAI nodes
- Samsung Advanced Package (SAP): heterogeneous integration of ARM compute + Samsung memory + photonic interface

## 3.4 Samsung Platform Summary

| Velsanet Component | Samsung Capability | Specific Technology |
|---|---|---|
| MOCT / Matrix | Samsung Foundry | SF3/SF4 silicon photonics + MEMS PDK + H-Cube packaging |
| Dedicated Equipment SoC | System LSI | Exynos platform + PUF identity IP + LPDDR5X PHY |
| PAI Node SoC | System LSI + Foundry | Ultra-low-power SoC (ARM M55 + NPU) on SF4 process |
| AAI / AsAI Memory | Samsung Memory | HBM3E (AsAI) + LPDDR5X (AAI) + V-NAND (Cube layer) |
| Multi-die Integration | Samsung Advanced Package | 2.5D / 3D heterogeneous packaging for AI + photonic dies |

# 4. Partnership Model

The ARM + Samsung combination addresses distinct but complementary layers of the Velsanet hardware stack. Neither partner alone covers the full requirement; together they provide complete vertical coverage.

## 4.1 Division of Roles

| Layer | ARM Role | Samsung Role | Joint Opportunity |
|---|---|---|---|
| MOCT / Physical | — | Foundry fabrication | Samsung-exclusive manufacturing + ARM IoT SoC for Channel 1 control |
| Dedicated Equipment | Cortex-A + Ethos + CMN IP licensing | Exynos SoC integration + PUF | Co-designed SoC: ARM IP inside Samsung fab on SF4 |
| PAI (O8) | Cortex-M55 + Ethos-U65 IP | Ultra-low-power SoC fab | Mass-produced PAI node SoC — potential volume play |
| AAI (D12) | Cortex-X + Mali + CMN IP | System LSI integration | Regional hub SoC targeting 6G base station form factors |
| AsAI (I20) | Neoverse V2 cluster | HBM3E + 2.5D packaging | Infrastructure AI node — addressable by Samsung-ARM server partnership |

## 4.2 Why This Matters Now

Three external factors create urgency for this partnership:

- **6G standardization timeline —** ITU-T SG13 engagement:
- Velsanet has submitted a formal contribution to ITU-T SG13 referencing all 21 white papers. The IMT-2030 standardization window for architectural proposals is open now. A hardware partnership with ARM and Samsung would accelerate the path from paper architecture to demonstrated prototype.
- **WP21 (AI Safety by Architecture) —** AI safety regulatory pressure:
- The EU AI Act, NIST AI RMF, and emerging global AI governance frameworks are all rule-based. Velsanet's structural safety model — where direct individual access to AsAI is architecturally impossible — represents a differentiated compliance story that ARM and Samsung could co-brand as a hardware-enforced AI safety platform.

- **Samsung is already a Tier-1 5G network equipment vendor. The Velsanet architecture, if prototyped on ARM + Samsung silicon, would give Samsung Network a path to position next-generation equipment beyond 5G NR — with AI and optical integrated at the node level rather than as separate subsystems.** Samsung Network division:

---

# 5. Proposed Next Steps

This report establishes the technical mapping. The following steps are proposed to move from mapping to engagement:

1. **Technical Review Meeting** — Joint session with ARM Korea and Samsung Research to validate the hardware mapping presented in this report and identify any specification gaps.
2. **Dedicated Equipment SoC Specification** — Co-develop a formal hardware specification for the Dedicated Equipment SoC, defining ARM IP selection, memory interface, optical interface requirements, and Samsung fabrication process target.
3. **MOCT Feasibility Study** — Samsung Foundry assessment of MOCT fabrication requirements: silicon photonics PDK compatibility, MEMS process integration, and polymer waveguide layer feasibility on SF4/SF3 process nodes.
4. **PAI Node Prototype** — Development of a PAI node reference design using Cortex-M55 + Ethos-U65 on an ARM Corstone platform — demonstrating the DIKWEI cognitive cycle in hardware as a proof-of-concept ahead of custom SoC development.
5. **ITU-T SG13 Joint Contribution** — Submission of a joint technical contribution to ITU-T SG13 positioning the ARM + Samsung + Velsanet hardware architecture as a candidate implementation reference for IMT-2030 AI-native network nodes.

---

# References

Velsanet White Papers (all available at 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 IP References**

- ARM Cortex-M55 Technical Reference Manual
- ARM Ethos-U65 NPU Product Brief
- ARM Cortex-X4 Technical Reference Manual
- ARM Neoverse V2 Platform Architecture
- ARM CoreLink CMN-700 Coherent Mesh Network

**Samsung References**

- Samsung Foundry SF4 / SF3 Process Technology Overview
- Samsung System LSI Exynos Platform
- Samsung HBM3E Product Specification
- Samsung Advanced Package (SAP) Technology

---