ALP LAB

SoM Selection Guide

Alp Lab E1M / E1M-X System-on-Module Catalogue

Document Number: SG-SOM-001 **Revision:** 0.11 **Date:** November 2026 **Status:** Preliminary

alplab.ai

© 2026 Alp Lab AB. All rights reserved.

Table of Contents

1 Introduction . . . 3
2 At a Glance . . . 3
3 Quick Decision Tree . . . 3
4 E1M-AEN Family (Alif Ensemble) . . . 4
4.1 Pick the AEN tier . . . 4
5 E1M-X V2N Family (Renesas RZ/V2N) . . . 4
5.1 When to pick V2N over V2N-M1 . . . 5
6 E1M-X V2N-M1 Family (RZ/V2N + DeepX DX-M1) . . . 5
6.1 When to pick V2N-M1 over V2N . . . 6
7 Software Portability . . . 6
8 Roadmap (Informative) . . . 6
9 Cross-Form-Factor Compatibility . . . 6
10 Where to Next . . . 6
11 Revision History . . . 7

List of Tables

Table 1 Top-level positioning . . . 3
Table 2 E1M-AEN SKU comparison . . . 4
Table 3 E1M-X V2N SKU comparison . . . 4
Table 4 E1M-X V2N-M1 SKU comparison . . . 5
Table 5 Software target matrix . . . 6
Table 6 Next-step references . . . 6
Table 7 Revision History . . . 7

Copyright © 2026 Alp Lab AB
Preliminary. Subject to change without notice.

1 Introduction

This guide compares the Alp Lab System-on-Module catalogue side-by-side so engineers and procurement can pick a SoM in one sitting. It covers every SKU currently in **production** or **preliminary** status; future SoMs are listed under **Roadmap** when their form factor and silicon are public.

The catalogue is split across two form factors:

- **E1M** (35 × 35 mm, 312 pads) — compact, lower-cost, MCU/MPU silicon. Targets industrial edge-AI, smart appliances, IoT endpoints.
- **E1M-X** (45 × 65 mm, 496 pads) — larger, higher-throughput, full MPU silicon. Targets robotics, multi-camera vision, automotive driver-monitoring, server-class edge inference.

Both form factors share the same family of signal classes per **E1M Spec** (E1M-STD-1.0), so a carrier board can host modules across both form factors with a footprint swap and no functional re-architecture.

2 At a Glance

Family	E1M-AEN	E1M-X V2N	E1M-X V2N-M1
Form factor	E1M (35 × 35 mm)	E1M-X (45 × 65 mm)	E1M-X (45 × 65 mm)
Primary silicon	Alif Ensemble (E3– E8)	Renesas RZ/V2N	Renesas RZ/V2N + DeepX DX-M1
Peak AI	up to 454 GOPS (2 × U55 + U85)	4 TOPS (DRP-AI3)	29 TOPS (DRP-AI3 + DX-M1)
Application CPU	up to 2 × A32 @ 800 MHz	4 × A55 @ 1.8 GHz	4 × A55 @ 1.8 GHz
Real-time CPU	up to 2 × M55	M33 @ 200 MHz	M33 @ 200 MHz
Linux capable	Yes (E5+)	Yes	Yes
Wi-Fi 6 + BLE 5.4	Yes (TI CC3501E)	Yes (Murata LBEE5HY2FY)	Yes (Murata LBEE5HY2FY)
Ethernet	1 × 100 Mbit	2 × 1 Gbit	2 × 1 Gbit
PCIe	—	Gen3 × 2	Gen3 × 2 (DX-M1 + external)
Operating temp	– 40 to +75 °C	– 40 to +85 °C	– 40 to +85 °C

Table 1 Top-level positioning

3 Quick Decision Tree

Three questions usually pick the SoM:

1. **Do you need to run Linux?**
   - No (bare-metal / Zephyr only) ⟶ consider **E1M-AEN301** / **AEN401** (E3 / E4) for the lowest BOM cost.
   - Yes ⟶ continue.
2. **What's your peak AI budget?**
   - <= 4 TOPS, vision-friendly workloads ⟶ **E1M-X V2N101** / **V2N102** / **V2N103**.
   - > 4 TOPS, or you want headroom for model ensembles / LLM-class inference ⟶ **E1M-X V2M101** / **V2M102** / **V2M103** (adds DX-M1 25 TOPS).
3. **Are size or BOM cost dominant?**
   - Yes, and AI workload fits the Alif Ensemble NPU (Ethos-U55/U85, 250 GOPS on E5/E7, up to 454 GOPS on E6/E8) ⟶ **E1M-AEN501**– **AEN803**.
   - Yes, and the workload is "always-on listen + occasional vision" ⟶ **E1M-AEN401** (E4, M55-only with U55+U85 NPUs and on-die ISP/JPEG).
   - No, full E1M-X feature surface is desirable ⟶ **V2N** or **V2N-M1**.

Copyright © 2026 Alp Lab AB

4 E1M-AEN Family (Alif Ensemble)

The AEN family pairs the **Alif Semiconductor Ensemble** MCU/MPU series with the E1M (35 × 35 mm) form factor. All seven SKUs use the **same module PCB** — routing, power controller and regulators, Wi-Fi combo, Ethernet PHY, and CAN transceiver are identical. SKUs differ in the Alif silicon tier (cores + NPU + memory) and in external-memory population.

SKU	Variant	Application CPU (A32)	Real-time CPU (M55)	NPU	ISP / JPEG
E1M-AEN301	E3	—	2 × M55	2 × U55	—
E1M-AEN401	E4	—	2 × M55	2 × U55 + 1 × U85	Yes
E1M-AEN501	E5	1 × A32 @ 800 MHz	2 × M55	2 × U55	—
E1M-AEN601	E6	1 × A32 @ 800 MHz	2 × M55	2 × U55 + 1 × U85	Yes
E1M-AEN701	E7	2 × A32 @ 800 MHz	2 × M55	2 × U55	—
E1M-AEN801	E8	2 × A32 @ 800 MHz	2 × M55	2 × U55 + 1 × U85	Yes
E1M-AEN803	E8	2 × A32 @ 800 MHz	2 × M55	2 × U55 + 1 × U85	Yes

Table 2 E1M-AEN SKU comparison

All seven SKUs share:

- **Memory:** 13.5 MB (E3 / E5 / E7) or 9.75 MB (E4 / E6 / E8) SRAM on-die, 5.5 MB MRAM. External memory differs by SKU: **E1M-AEN803** fits 512 Mbit HyperRAM + 256 Mbit xSPI NOR flash (both on OSPI0); **E1M-AEN801** fits neither and runs from on-die MRAM; on AEN301– AEN701 the external HyperRAM / NOR are BOM options.
- **Wireless:** TI CC3501E Wi-Fi 6 (2.4 GHz, 20 Mbps) + BLE 5.4.
- **Ethernet:** TI DP83825 100 Mbit PHY on-module.
- **Operating temperature:** – 40 to +75 °C.
- **On-module:** programmable power controller with always-on 3.3 V / 1.8 V regulators, buck regulators for the SoC and the 1.8 V I/O rail, and a brownout supervisor; RTC (RV-3028-C7), temperature sensor (TMP112), EEPROM (N24S128). No E1M-AEN module fits a secure element.

> **Note:** For the full AEN technical reference see **DS-AEN-001** and **HG-AEN-001**.

4.1 Pick the AEN tier

- **Smallest budget, bare-metal:** AEN301 (E3) — M55-only, two NPUs, no Linux.
- **Vision-light + always-on:** AEN401 (E4) — adds on-die ISP/JPEG + the U85 NPU.
- **Linux-capable, single A32:** AEN501 / AEN601.
- **Linux-capable, dual A32 (recommended balanced pick):** AEN701 / AEN801. The 801 adds the U85 NPU and on-die ISP/JPEG.
- **E8 with external memory:** AEN803 — the AEN801 silicon plus 512 Mbit HyperRAM and 256 Mbit xSPI NOR flash fitted (preliminary).

5 E1M-X V2N Family (Renesas RZ/V2N)

The V2N family pairs the **Renesas RZ/V2N** vision-AI MPU with the E1M-X (45 × 65 mm) form factor. Three SKUs differ only in memory tier:

SKU	Renesas part	LPDDR4X	eMMC
E1M-V2N101	R9A09G056N44GBG#AC0	4 GB	eMMC 5.1, 4 GB
E1M-V2N102	R9A09G056N44GBG#AC0	8 GB	eMMC 5.1, 16 GB

Copyright © 2026 Alp Lab AB

SKU	Renesas part	LPDDR4X	eMMC
E1M-V2N103	R9A09G056N44GBG[1]	4 GB	eMMC 5.1, 16 GB

Table 3 E1M-X V2N SKU comparison

[1] Same RZ/V2N silicon variant (R9A09G056N44GBG) as V2N101 / V2N102; the ordering suffix for E1M-V2N103 is not stated in this edition.

All three SKUs share:

- **Application cluster:** 4 × Arm Cortex-A55 @ 1.8 GHz.
- **Real-time core:** Arm Cortex-M33 @ 200 MHz (in the RZ/V2N).
- **I/O MCU:** GigaDevice GD32G553 Cortex-M33 @ 216 MHz (separate chip, on-module supervisor).
- **AI accelerator:** Renesas DRP-AI3, 4 TOPS dense.
- **ISP / GPU:** Arm Mali-C55 ISP + Mali-G31 GPU.
- **Codec:** H.264 1920 × 1080 @ 60 fps; H.265 3840 × 2160p @ 30 fps.
- **Connectivity:** 2 × 1 Gbit Ethernet (RTL8211FI-VD-CG), Wi-Fi 6 + BLE 5.4 (Murata LBEE5HY2FY), 2 × CAN-FD, PCIe Gen3 × 2.
- **Operating temperature:** − 40 to +85 °C.

> **Note:** For the full V2N technical reference see **DS-V2N-001** and **HG-V2N-001**.

5.1 When to pick V2N over V2N-M1

V2N is the right pick when:

- Peak inference fits within the 4 TOPS DRP-AI3 budget.
- Power envelope is tight (the DX-M1 on V2N-M1 adds TBD W under load).
- BOM cost matters more than AI headroom.

6 E1M-X V2N-M1 Family (RZ/V2N + DeepX DX-M1)

V2N-M1 **is** the V2N SoM plus an on-module **DeepX DX-M1** 25-TOPS AI accelerator, PCIe-attached. Three SKUs match the V2N memory tiers.

SKU	DX-M1 status	Memory tier	Combined AI capability
E1M-V2M101	Populated	4 GB LPDDR4X + 4 GB eMMC	DRP-AI3 (4 TOPS) + DX-M1 (25 TOPS)
E1M-V2M102	Populated	8 GB LPDDR4X + 16 GB eMMC	DRP-AI3 (4 TOPS) + DX-M1 (25 TOPS)
E1M-V2M103	Populated	4 GB LPDDR4X + 16 GB eMMC	DRP-AI3 (4 TOPS) + DX-M1 (25 TOPS)

Table 4 E1M-X V2N-M1 SKU comparison

V2N-M1 adds, on top of every V2N feature:

- **DeepX DX-M1** 25-TOPS AI accelerator, PCIe-attached on-module.
- **2 × LPDDR5X** companion memory for the DX-M1.
- **SPI NAND** companion storage for the DX-M1 (for model weights).

The DX-M1 is invisible to the carrier board: it sits behind the V2N's PCIe controller via an on-module multiplexer. Software selects whether the PCIe lanes serve the DX-M1 or the external M.2 slot at runtime (see **HG-V2N-M1-001 §6.7**).

> **Note:** For the full V2N-M1 technical reference see **DS-V2N-M1-001** and **HG-V2N-M1-001**.

Copyright © 2026 Alp Lab AB

6.1 When to pick V2N-M1 over V2N

V2N-M1 is the right pick when:

- AI workload exceeds 4 TOPS or has burst-y peaks.
- Multiple inference graphs run in parallel (DRP-AI3 + DX-M1 can both fire).
- LLM-class inference, semantic-search embedding, or large model ensembles are part of the workload.

7 Software Portability

All SoMs are supported by the **Alp SDK™**, which presents a uniform HAL across families. Most peripheral examples (`gpio-button-led`, `i2c-scanner`, `pwm-led-fade`, `audio-wake-word`, `iot-connected-camera`, etc.) build **without source changes** on every SKU.

Target / runtime	AEN301/401	AEN501-803	V2N101/102/103	V2M101/102/103
Zephyr on M-class core	Yes	Yes	Yes	Yes
Bare-metal C on M-class core	Yes	Yes	Yes	Yes
Linux (Yocto / buildroot)	—	Yes	Yes	Yes
DRP-AI3 inference	—	—	Yes	Yes
DX-M1 inference	—	—	—	Yes
Ethos-U55 inference	Yes	Yes	—	—

Table 5 Software target matrix

The full per-target board strings are documented in the SDK's `docs/firmware-quickstart.md`.

8 Roadmap (Informative)

Additional SoMs planned for upcoming editions of this guide:

- **E1M-X (TBD)** with secondary AI accelerator alternatives beyond the DX-M1.
- **E1M-X HIGH-CORE** (TBD) — more A55 cores, more PCIe lanes, targeted at server-class edge.

Roadmap items will appear in this document once their pin-out and electrical envelope are signed off; for the live status of each item see the alp-sdk firmware quickstart §1 or contact Alp Lab.

9 Cross-Form-Factor Compatibility

A carrier board that follows the **E1M Spec** (E1M-STD-1.0) is electrically and mechanically compatible with **every conformant SoM in its form factor**. A V2N101 and a V2M102 can be swapped on the same E1M-X carrier without re-routing; an AEN301 and an AEN801 can be swapped on the same E1M carrier without re-routing.

Software portability is similarly first-class: per **E1M Spec** §8, every digital pad can be repurposed as GPIO on any conformant SoM, so a carrier-board design that uses only the default-function pinout is hardware-compatible with any future E1M / E1M-X SoM.

10 Where to Next

Once you've narrowed the catalogue to a family:

If you're picking…	Read next
E1M-AEN301..803	**DS-AEN-001** (Datasheet), **HG-AEN-001** (HW Design Guide), **QS-E1M-EVK-001** (Getting Started)

Copyright © 2026 Alp Lab AB

If you're picking…	Read next
E1M-X V2N101/102/103	**DS-V2N-001** (Datasheet), **HG-V2N-001** (HW Design Guide), **QS-E1M-X-EVK-001** (Getting Started)
E1M-X V2M101/102/103	**DS-V2N-M1-001** (Datasheet), **HG-V2N-M1-001** (HW Design Guide), **QS-E1M-X-EVK-001** (Getting Started)
Any SoM (form-factor reference)	**E1M-STD-1.0** (E1M Specification)
Software	`github.com/alplabai/alp-sdk` (Alp SDK™) — `docs/firmware-quickstart.md` and the matching `docs/soms/<family>.md` one-pager

Table 6 Next-step references

For volume-pricing or custom-variant requests, contact Alp Lab at contact@alplab.ai.

11 Revision History

Revision	Changes	Date
0.5	• Re-synced to alp-sdk v0.17.0. • Removed a withdrawn E1M-form-factor SoM from the roadmap: it was never produced and is no longer supported by the SDK. • Added **E1M-AEN803** (E8 with 512 Mbit HyperRAM + 256 Mbit xSPI NOR fitted), **E1M-V2N103** and **E1M-V2M103** (32 Gbit LPDDR4X + 128 Gbit eMMC) to the SKU tables, decision tree, software matrix and reference table. • Ethernet PHY corrected from `DP83825I` to `DP83825` (the order-code suffix is not established). • Stated that the OPTIGA Trust M is not fitted on E1M-AEN801 / E1M-AEN803 and corrected the AEN external-memory description per SKU. • First edition with a revision-history table.	October 2026
0.6	• E1M-AEN operating temperature corrected to − 40 to +75 °C (E1M-AEN hardware revision R3, standard module) in the AEN family section.	October 2026

0.7	• E1M-AEN family section aligned with DS-AEN-001:
	  ‣ on-module list no longer includes the OPTIGA Trust M (no E1M-AEN module fits a secure element)
	  ‣ "PMIC" replaced by the R3 power controller with its always-on regulators, buck regulators and brownout supervisor
	  ‣ on-die memory corrected to 13.5 MB (E3 / E5 / E7) or 9.75 MB (E4 / E6 / E8) SRAM and 5.5 MB MRAM (was 128 KB– 13.5 MB SRAM, 256 KB– 5.5 MB MRAM).
	• E1M-X V2N family connectivity: Ethernet PHY `RTL8211FI-VD-CG` (was RTL8211FDI), as in DS-V2N-001.
		October 2026

0.8	• Release history: every entry formatted as a bullet list.	October 2026
0.9	• Contact: contact@alplab.ai is the only contact address. • Memory sizes given in GB (E1M-X V2N and V2N-M1 SKU tables): `101` = 4 GB LPDDR4X + 4 GB eMMC, `102` = 8 GB + 16 GB, `103` = 4 GB + 16 GB.	October 2026
0.10	• Table of Contents: each entry links to its section; the front-matter lists are set tighter.	October 2026
0.11	• Product name on page 1 links to the product page on alplab.ai.	November 2026

Table 7 Revision History

Copyright © 2026 Alp Lab AB