ALP LAB

# E1M-X V2N-M1 HW Design Guide

**Document Number:** HG-V2N-M1-001 **Revision:** 0.2 **Date:** June 2026 **Status:** Draft

alplab.ai

## Table of Contents



## List of Figures



## List of Tables



# Release History

| Revision | Changes | Date |
|---|---|---|
| 0.1 | Initial draft. | May 2026 |
| 0.2 | Interface ESD/EMI protection moved to the carrier board (Ethernet, MIPI CSI, MIPI DSI); on-module protection retained only on the SD-card interface. | June 2026 |

**Table 1** Release History

# 1 Introduction

## 1.1 Purpose of This Document

This guide provides practical guidance for developers and system integrators using the **E1M-X V2N-M1 System-on-Module (SoM)**.

It covers:

- Hardware integration and carrier-board design
- Power architecture and system bring-up
- Interface usage and pin functional intent
- DeepX DX-M1 bring-up + PCIe multiplexing
- Common pitfalls during first implementation

This guide complements:

- The **E1M-X V2N-M1 Datasheet** (DS-V2N-M1-001)
- The **E1M-X V2N HW Design Guide** (HG-V2N-001) — the V2N-family rules apply unchanged to V2N-M1 except where §6.7 below diverges
- The **Renesas RZ/V2N** hardware reference manual and PCB design guidelines
- The **DeepX DX-M1** hardware integration manual
- The **E1M Specification** (E1M-STD-1.0) for form-factor and pinout requirements
- The **Alp SDK™** documentation

## 1.2 Product Overview

The **E1M-X V2N-M1** is a compact **45 × 65 mm Edge-AI System-on-Module** that pairs the **Renesas RZ/V2N** vision-AI MPU with the **DeepX DX-M1** standalone AI accelerator on a single SoM. It targets heavy-inference workloads (multi-camera object detection, LLM-class inference, model ensembles) that exceed the 4 TOPS budget of the V2N's internal DRP-AI3.

The module follows the **E1M-X™ open-standard pin layout** and is **pin-compatible with the V2N family**, so customers can prototype on V2N and upgrade to V2N-M1 without redesigning the carrier board.

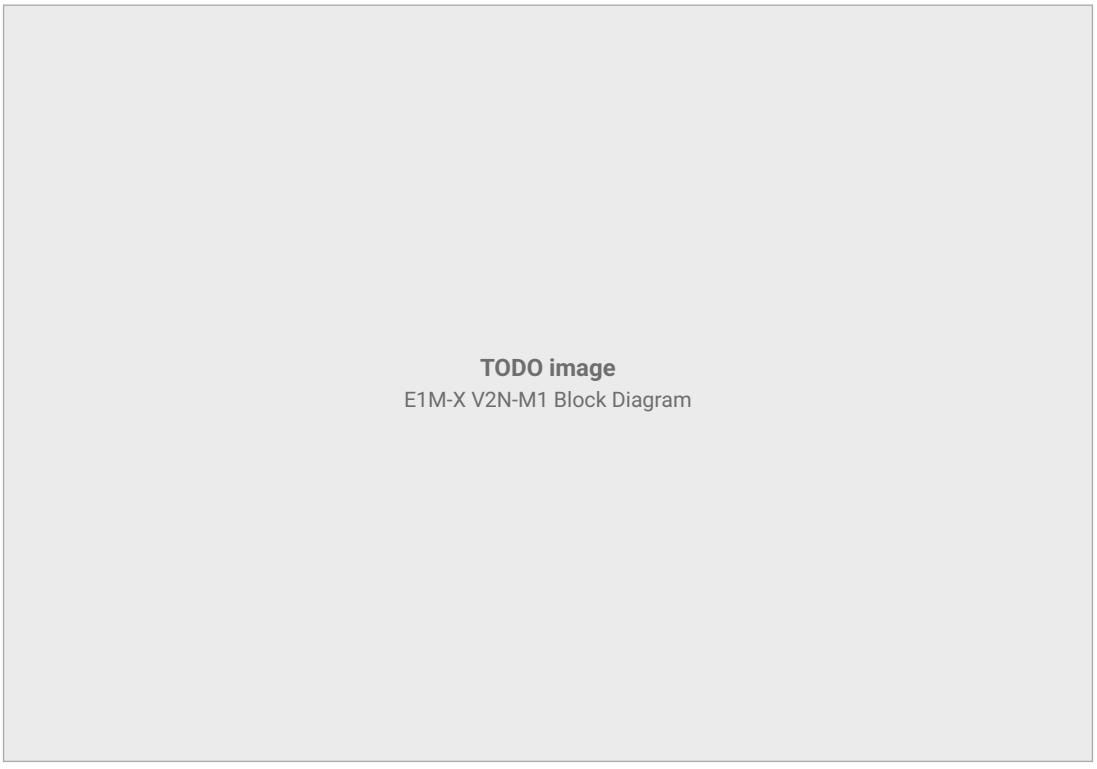


**Figure 1** E1M-X V2N-M1 Block Diagram

**Key characteristics:**

- Unified E1M-X pinout, drop-in compatible with E1M-X V2N
- Quad Cortex-A55 @ 1.8 GHz + Cortex-M33 @ 200 MHz application + real-time cluster
- Renesas DRP-AI3 (4 TOPS) **plus** DeepX DX-M1 (25 TOPS) AI accelerators in parallel
- 2 × LPDDR5X companion memory for the DX-M1
- SPI NAND companion storage for the DX-M1
- Arm Mali-C55 ISP + Mali-G31 GPU
- On-module 1 Gbit Ethernet PHYs (× 2), Wi-Fi 6 + BLE 5.4 combo, CAN-FD transceivers (× 2)
- On-module power management (Renesas DA9292 PMIC) plus GD32 supervisor MCU
- Single external **5 V supply**
- Long-term lifecycle orientation

# 2 Supported SoC Variants

| SKU | Renesas Part | LPDDR4X | eMMC | AI capability |
|---|---|---|---|---|
| E1M-V2M101 | R9A09G056N44GBG#AC0 | 32 Gbit | eMMC 5.1, 32 Gbit | V2N: 4 TOPS<br>DX-M1: 25 TOPS |
| E1M-V2M102 | R9A09G056N44GBG#AC0 | 64 Gbit | eMMC 5.1, 128 Gbit | V2N: 4 TOPS<br>DX-M1: 25 TOPS |

**Table 2** V2N-M1-family SKUs

Both SKUs share identical pinouts and silicon outside of the memory tier; carrier boards designed against one SKU work unmodified with the other.

# 3 Mechanical Information

## 3.1 Dimensions

- **Size:** 45 × 65 mm
- **Form factor:** Connectorless, solder-down LGA SoM
- **Height profile:** Standard-height class per **E1M Spec** §5.2.1

The DX-M1 + LPDDR5X + SPI NAND additions to the V2N stack are all top-side components within the standard-height envelope; no carrier-side cutout is required.

For the full mechanical drawing including land pattern and keep-out zones, see the **E1M-X V2N-M1 Datasheet** (DS-V2N-M1-001) §13.

## 3.2 Handling and Assembly

- ESD-sensitive device. Use grounded handling equipment.
- Avoid mechanical stress after soldering.
- Follow IPC J-STD-020 reflow profile (see DS-V2N-M1-001 §11.2).
- Maintain PCB flatness; warpage > TBD µm may compromise BGA solder joints under both the SoC and the DX-M1.

# 4 Power Architecture

## 4.1 Power Input

- **Input voltage:** 5 V DC, single rail on VDD_5V_IN
- All regulators, the PMIC, and sequencing are on-module
- No additional power rails must be supplied by the carrier

**Warning:** Do **not** attempt to power individual rails (`VIO_OUT`, `SD_VDD_OUT`, etc.) externally. These pads are **outputs** from the module; sourcing them externally will fight the on-module regulator and may damage the PMIC.

## 4.2 Power Design Guidelines

- Place bulk capacitance (≥ 10 µF) near the 5 V input
- Use a solid ground plane covering the entire SoM footprint
- Keep high-current paths short and wide
- Nominal input voltage: 5.0 V; recommended tolerance: ±5%
- Connect **every** GND pad on the SoM to a low-impedance plane (105 pads on E1M-X)
- No external sequencing required; the GD32 supervisor MCU handles the on-module rail order, including bring-up of the DX-M1

**Note:** The V2N-M1 has measurably higher peak power than the V2N at the same workload because the DX-M1 + companion LPDDR5X + SPI NAND draw additional current during inference. Size the carrier supply for ≥ 5 A peak.

## 4.3 Carrier-Board 5 V Generation

A buck converter is the recommended way to generate the 5 V rail from the carrier-board input voltage. Choose a converter rated for transients up to **5 A** to cover worst-case AI (DRP-AI3 + DX-M1 in parallel) + dual-Ethernet + M.2 SSD activity simultaneously.

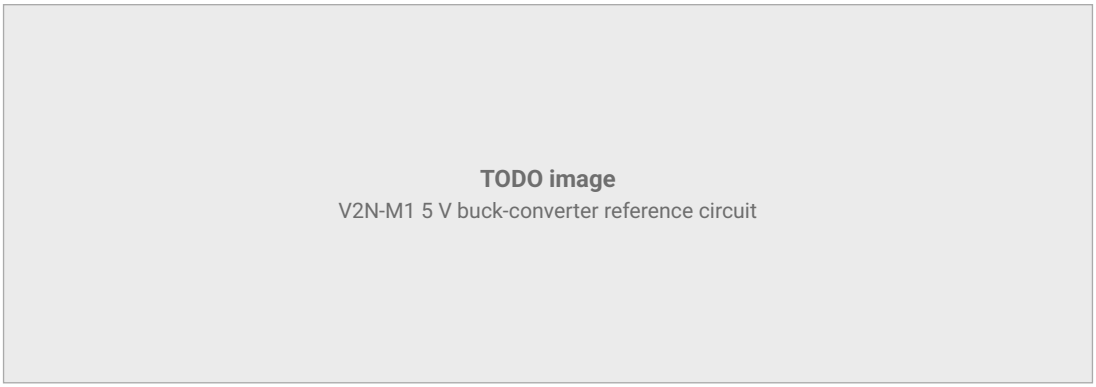


**Figure 2** V2N-M1 5 V buck-converter reference circuit

Power-OR'ing using diodes or eFuses is supported. When multiplexed (e.g. between barrel jack and USB-PD on a carrier), use a unidirectional eFuse to avoid back-feeding the inactive source.

# 5 Memory and Boot Architecture

## 5.1 Internal Memory (SoC-Dependent)

- **On-chip SRAM:** 1.5 MB (RZ/V2N)
- **LPDDR4X DRAM:** 32-bit, 3.2 GT/s, up to 8 GB (per SKU)

## 5.2 On-Module Memory

- **eMMC 5.1** up to 128 Gbit (per SKU)
- **SPI NOR flash:** 128 Mbit, boot-supported (RZ/V2N boot path)
- **2 × LPDDR5X** companion memory for the DeepX DX-M1
- **SPI NAND** companion storage for the DeepX DX-M1 (model weights, neural-network binaries)
- **$I^2C$ EEPROM:** optional, accessible via `I2C0`

**Note:** The DX-M1 has its **own** memory subsystem (LPDDR5X + SPI NAND) that is not accessible from the carrier board. Loading models into the DX-M1 is a software-managed step from the RZ/V2N application cores via the PCIe link; see **§Software Support**.

## 5.3 Boot

The RZ/V2N samples four boot-strap pins at P0Rn release. The V2N-M1 boot flow is **identical to V2N** for the RZ/V2N itself; the DX-M1 brings itself up under control of the GD32 supervisor MCU at the same time as the SoC rails come up.

See the **E1M-X V2N-M1 Datasheet** (DS-V2N-M1-001) §4.1 for the full boot-mode truth table, and the **E1M-X V2N HW Design Guide** (HG-V2N-001) §6 for the carrier-side DIP-switch implementation.

# 6 Carrier-Board Reference Block Diagram

The carrier board has been designed for the **E1M-X** open standard. V2N-M1 routes the **same subset** of E1M-X pads as the V2N family — the DX-M1 is fully encapsulated on the SoM and does not consume any additional carrier-board pads.

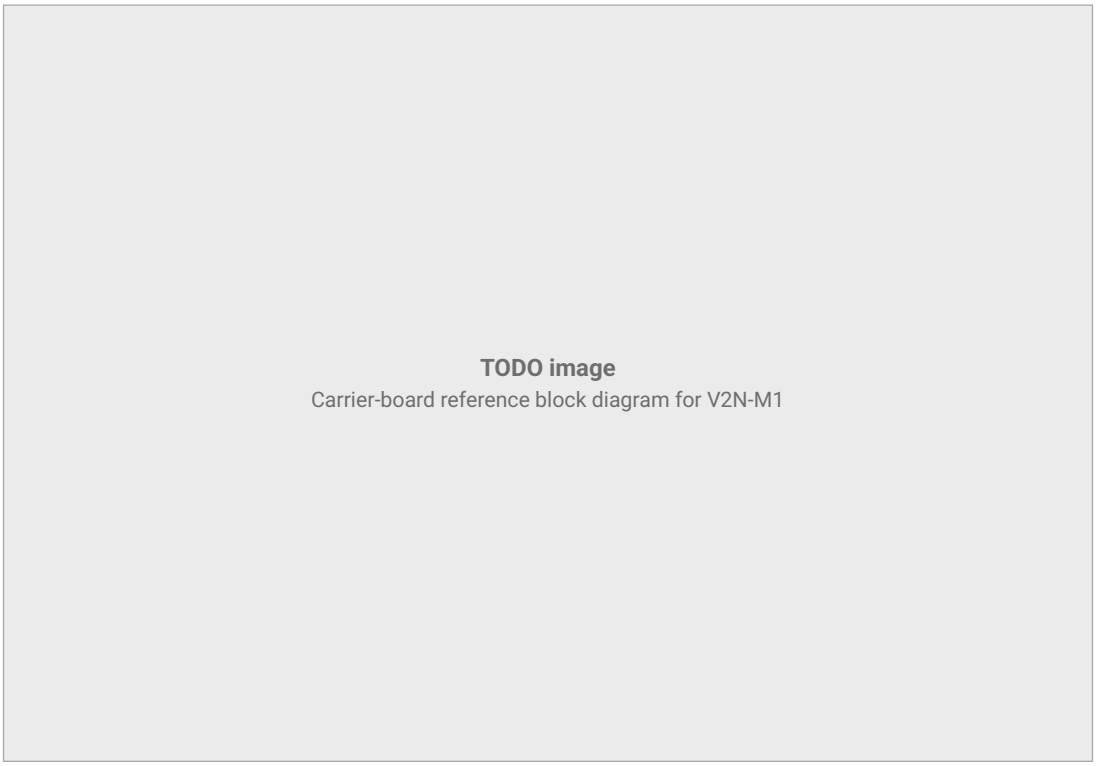


**Figure 3** Carrier-board reference block diagram for V2N-M1

## 6.1 System-Level Overview

Identical to the V2N family. See **HG-V2N-001** §5.

## 6.2 Design Philosophy

- Keep the carrier board simple and application-specific.
- **Reuse the V2N carrier design without changes.** The V2N-M1 is pin-compatible.
- Plan for higher peak power (≥ 5 A on the 5 V rail) and slightly higher thermal dissipation than V2N.

# 7 Interfaces Overview

All interfaces are exposed through the standardised **E1M-X™ pinout**. The carrier-board rules in **HG-V2N-001** §6 apply unchanged for **every interface except PCIe**. This section gives a brief summary; for full details see the V2N HW Design Guide.

> **Note:** See the **E1M-X V2N-M1 Datasheet** (DS-V2N-M1-001) and the **E1M Specification** (E1M-STD-1.0) for the canonical pad-list and pin-function mapping.

## 7.1 Ethernet, Wi-Fi/BLE, MIPI CSI/DSI, USB, CAN, Serial Buses, Audio, Analog, PWM, Encoder

These interfaces are unchanged from V2N. Follow the rules in **HG-V2N-001** §6.1 – §6.10. The on-module components, routing requirements, and carrier-side ESD/EMI protection requirements are all identical (interface ESD/EMI protection is provided on the carrier board, except the SD-card interface which is protected on-module).

## 7.2 PCIe (V2N-M1 specifics)

This is the only interface where V2N-M1 design differs materially from V2N.

The RZ/V2N has **one** PCIe Gen3 × 2 controller. On the V2N-M1 SoM this controller is permanently routed to the on-module DeepX DX-M1, **and** multiplexed out to the carrier on the same `PCIE0_*` pads. The on-module multiplexer is controlled by an I/O-expander bit driven by the GD32 supervisor; software selects whether the PCIe lanes serve the on-module DX-M1 or the external M.2 slot at any given moment.


**Figure 4** V2N-M1 PCIe multiplexer block diagram

### 7.2.1 Routing rules

- Treat `PCIE0_*` exactly as in V2N: 85 Ω differential, lane-to-lane matching ±2 ps, AC-couple every lane.
- The on-module mux adds < TBD ps of skew per lane – design margin TBD ps additional length-tuning allowance on the carrier compared to a direct connection.

### 7.2.2 Power and reset

- The DX-M1 has its own dedicated rails generated on-module; the carrier provides nothing extra.
- The DX-M1 resets when `PORn` is asserted to the SoC; no carrier-side reset is exposed for the DX-M1.

### 7.2.3 Software-selectable routing

The PCIe routing is selected at runtime via an $I^2C$-write to the on-module I/O expander (`I2C0` address `0x21`, register `0x07`, bit 2):

- **bit 2 = 0:** PCIe routed to the carrier board (external M.2 slot is active, DX-M1 is held in reset).
- **bit 2 = 1:** PCIe routed to the on-module DX-M1 (carrier M.2 is disconnected).

> **Warning:** Do **not** assume both endpoints can be active concurrently. The PCIe controller is single-instance; the multiplexer is exclusive. Software must release the M.2 device cleanly before flipping the bit, or the host PCIe stack will trigger an AER fault.

The Alp SDK™ handles this transition through the `alp_pcie_select(target)` API; consult `docs/bring-up-v2n-m1.md` in the SDK for the full sequence.

# 8 Software Support

## 8.1 Alp SDK™

- Open-source, CMSIS-compliant
- Unified HAL across all E1M-X variants (V2N, V2N-M1, future SKUs)
- Vendor-agnostic software architecture

The SDK exposes the same board targets as V2N, with a V2N-M1-specific A55 target:

- `alp_e1m_v2n_m33_io` — on-module GD32 supervisor (216 MHz)
- `alp_e1m_v2n_m33_rt` — RZ/V2N internal Cortex-M33 (200 MHz)
- `alp_e1m_v2m_a55` — RZ/V2N quad Cortex-A55 (1.8 GHz, Zephyr SMP or Linux/Yocto), **DX-M1-aware**

The `alp_e1m_v2m_a55` target enables the DX-M1 driver and the `alp_pcie_select()` API automatically; using the V2N target on a V2N-M1 module leaves the DX-M1 in reset.

## 8.2 AI Acceleration

- Renesas DRP-AI3 (4 TOPS, integrated in the RZ/V2N)
- DeepX DX-M1 (25 TOPS, PCIe-attached on-module)
- Both accelerators can run in parallel; the SDK's scheduler can fan out a single inference graph across both, or pin per-camera streams to dedicated accelerators

### 8.2.1 DX-M1 model deployment

DeepX provides the **DX-COM™** tool which converts PyTorch / ONNX models into DX-M1-optimised executable code. Compiled models are loaded by the RZ/V2N at runtime via the PCIe link into the DX-M1's companion LPDDR5X.

> **Note:** Persisting compiled models in the on-module SPI NAND requires the production-flashing path documented in `docs/bring-up-v2n-m1.md`. For development, models live in the carrier-board SD card or carrier eMMC and are pushed over PCIe on each boot.

# 9 Bring-Up Checklist

## 9.1 Pre-Power Checks

- Confirm the correct SoM variant (`E1M-V2M101` vs `E1M-V2M102`).
- Verify orientation and solder joints (A1 corner against carrier fiducial).
- Ensure only 5 V is connected to `VDD_5V_IN`.
- Check ground continuity across all 105 GND pads.
- Confirm DIP-switch state for the intended boot mode.

## 9.2 First Power-On

- Apply 5 V (recommend barrel jack; USB-PD also works).

- Monitor current consumption (expected idle: TBD mA; expect higher than V2N because of the DX-M1 supervisor power).
- Check for abnormal heating on the SoM — particularly over the DX-M1.
- Confirm PWR and IO_EN LEDs light within 100 ms.
- Open the console UART at 115 200 8N1 and watch for boot output.

## 9.3 Basic Validation

- UART boot / debug output on the SCIF console.
- Both Ethernet links coming up when cables are plugged in.
- GPIO or PWM toggle from a simple Zephyr build.
- `lspci` (on Linux) shows the DX-M1 enumerating at `0x1c3c:0x0001` after `alp_pcie_select(M1)`.
- First DX-M1 inference completes via the SDK's `deepx-hello-world` example.

## 9.4 Common Issues

- Wrong DIP-switch position for the intended boot device (most common first-time failure).
- DX-M1 not enumerating because the PCIe mux is still pointed at the carrier-board M.2 slot (call `alp_pcie_select(M1)` before `lspci`).
- DX-M1 enumerating but inference failing because the model wasn't loaded – check `docs/bring-up-v2n-m1.md` §3.
- Wrong board target (`alp_e1m_v2n_a55` instead of `alp_e1m_v2m_a55`); the DX-M1 driver isn't compiled in.
- Insufficient supply current under DRP-AI3 + DX-M1 + dual-Ethernet load.

# 10 Design Checklist (Summary)

| Item | Status |
|---|---|
| Single 5 V supply, ≥ 5 A rated | OK |
| Bulk + decoupling capacitance close to VDD_5V_IN | OK |
| Ethernet magnetics + RJ45 on both ports | OK |
| Controlled-impedance MIPI and PCIe routing | OK |
| RF-clean antenna routing, ≥ 1 RF connector or PCB antenna | OK |
| CAN-bus termination + ESD on both connectors | OK |
| Boot-mode DIP switch or hard-strap | OK |
| PORn and MODULE_EN accessible (button or test-point) | OK |
| Debug header on JTAG / SWD pins (optional but recommended) | OK |
| PCIe routing supports both on-module DX-M1 and external M.2 (or accepts that external M.2 is disabled when DX-M1 active) | OK |
| Thermal: copper pour or heatsink contact over the DX-M1 region of the SoM | OK |
| Correct firmware per SKU + target core (`alp_e1m_v2m_a55` for full DX-M1 access) | OK |

**Table 3** V2N-M1 carrier-board design checklist

# 11 Ordering Information

E1M-X V2N-M1 modules are ordered based on:

- Memory tier (32 Gbit / 64 Gbit LPDDR4X; 32 / 128 Gbit eMMC)
- Assembly options (lead finish, ball material)
- DX-M1 firmware option (factory-loaded model-zoo vs blank companion NAND)

Refer to the **E1M-X V2N-M1 Datasheet** (DS-V2N-M1-001) §15 for the complete ordering matrix and MPN decoder.

# 12 Notices

- Specifications subject to change without notice.
- Some features may be preliminary at this revision.
- Always verify against the latest **E1M-X V2N-M1 Datasheet** (DS-V2N-M1-001).
- This guide complements but does not replace the Renesas RZ/V2N PCB design guidelines or the DeepX DX-M1 hardware integration manual.