EL2Monitor

Partitioning and isolation solution for Arm® Cortex®-R52 cores

Overview

EL2Monitor is an NXP in-house solution providing partitioning and isolation at the Arm Cortex-R52 core level for NXP S32 automotive processing platform devices. It was born from the need to address hardware limitations of the Cortex-R52 generic interrupt controller (GIC) distributor on S32Z/E and S32N devices.

Running at Arm EL2, EL2Monitor provides each core with a virtual GIC distributor and virtual message receive unit (MRU), enabling true partitioning and isolation for Cortex-R52 cores in the real-time unit (RTU) cluster with zero interrupt latency penalty.

EL2Monitor is developed according to ISO 26262 (ASIL D) and ISO 21434 standards and is delivered as part of the NXP S32N SDK and GreenVIP bundles.

When would you need EL2Monitor?

EL2Monitor is the right choice when you need to run multiple independent operating systems or bare-metal applications on separate R52 cores of the same system in chip (SoC), each requiring strict isolation from the others.

Typical scenarios

- Multiple RTOSes (e.g. NXP RTOS, FreeRTOS, Zephyr®) running independently, each at EL1 on a dedicated R52 core
- EL2Monitor sits at EL2 on each core, virtualising the shared GIC Distributor and MRU so each EL1 OS sees its own isolated peripheral view
- Safety partitioning: isolating ASIL D workloads from QM workloads on the same RTU cluster
- Any use case requiring freedom of interference between cores sharing hardware peripherals (GIC, MRU)

The EL2Monitor story

Challenge 1: Shared GIC distributor

- HW provides shared GIC distributor; SW needs isolated resources
- EL2Monitor provides a virtual GIC distributor per core
- Trap-and-emulate: 1.48 µs average latency per event
- Paravirtualised* CMSIS API: reduced to 0.78 µs average

Challenge 2: Shared MRU IP

- MRU IP cannot be partitioned by multiple software entities
- EL2Monitor provides a virtual MRU per software entity
- Guarantees freedom from interference

Main characteristics

- Partitioning and isolation at Cortex-R52 core level
- Manages shared peripherals: GIC distributor, MRU
- Virtual GIC distributor per core
- Virtual MRU per software entity
- Paravirtualised* GIC interface (CMSIS API)
- EL2 MPU memory protection
- 1:1 VM-to-CPU allocation
- Device virtualization: always passthrough
- Memory footprint: Only 10 KB
- NXP S32 design studio configurable
- EB tresos Studio configurable

Supported hardware

SoC families

- NXP S32Z/E: Arm Cortex-R52 RTU
- NXP S32N: Arm Cortex-R52 RTU cluster
- NXP S32K5: Arm Cortex-R52
- NXP S32J100: Arm Cortex-R52

Execution level

- Runs at Arm EL2 (per core)
- RTOS/apps run at EL0/EL1

Business and customers

- Delivered as part of NXP S32N SDK and GreenVIP bundles

Authors

- SW Engineering
- Software Engineer
- SW Engineering

Software offering

- Virtual GIC distributor emulation
- Paravirtualised* GIC interface (CMSIS API)
- Virtual MRU emulation and allocation
- EL2 MPU memory region configuration
- EL2 sync for shared peripheral access
- NXP S32 design studio integration
- EB tresos studio plugin

Key use cases

S32N R52 split lock: 4x single-core RTOS

- Each R52 core runs its own RTOS independently
- Virtual GIC and MRU per core for full isolation

Safety and security

- ASIL D—highest automotive safety level
- ISO 26262 safety standard compliance
- ISO 21434 cybersecurity compliance
- EL2 MPU for memory protection and isolation

Performance vs. hypervisors

- Zero IRQ latency penalty (vs. 13.3–19.99 µs)
- Zero interrupt jitter (vs. significant)
- 10 KB footprint (vs. 250–500+ KB)
- Paravirt* CMSIS GIC API: 0.78 µs average
- Purpose-built for NXP S32 Arm Cortex-R52

* Paravirtualisation: a technique where the guest OS is modified to call a hypervisor API directly (here: CMSIS-based), avoiding costly trap-and-emulate overhead and reducing GIC emulation latency from 1.48 µs to 0.78 µs average

10 KB	Zero	Zero	ASIL D	0.78 µs
Memory footprint vs 250–500+ KB (Third party)	**IRQ latency penalty** vs 13.3–19.99 µs (Third party)	**Interrupt jitter** vs 250–500+ KB (Third party)	**Safety certified** ISO 26262 & ISO 21434-compliant	**GIC Emulation** Paravirtualised* CMSIS API

EL2Monitor vs. Third party hypervisors

	EL2Monitor	Third party hypervisor
Memory footprint	10 KB	250–500+ KB
IRQ latency penalty	Zero	13.3–19.99 µs average
Interrupt jitter	Zero	Significant
Virtual machines	No	Yes
VM:CPU allocation	1:1	N:1
Device virtualization	Passthrough	Yes
GIC and MRU eemulation	Yes	Yes
GIC Emul. latency	0.78 µs (paravirt.*)	Not measured
EL2 MPU protection	Yes	Yes
Safety compliance	ASIL-D	Vendor-dependent
ISO 26262/21434	Yes	Vendor-dependent

* 1. Trap-and-emulate GIC.

* 2. Paravirtualised CMSIS API. 3rd-party: Vendor 1 = 13.3 µs avg., Vendor 2 = 19.99 µs avg.

* Paravirtualisation: guest OS calls hypervisor API directly, avoiding trap-and-emulate overhead.

Feature matrix

Feature	✓
Virtual GIC distributor per core	✓
Virtual MRU per software entity	✓
Paravirtualised GIC interface (CMSIS API)	✓
EL2 MPU memory protection	✓
Zero interrupt latency and jitter penalty	✓
Partitioning and isolation at R52 core level	✓
ISO 26262 (ASIL D) safety compliance	✓
ISO 21434 security standard compliance	✓
NXP S32 Design studio configuration	✓
EB tresos Studio configuration	✓
Optimised GIC driver for lower latency	✓
Delivered as part of VIP bundle	✓

Additional information

Demos and integration

EL2Monitor is available as part of the NXP S32N SDK and GreenVIP bundles. Live demos are available showcasing EL2Monitor running multiple RTOSes on NXP S32N and S32Z/E devices. On request, NXP can provide support for integration into customer projects.

For more information and demo access, visit: **nxp.com/el2m**

nxp.com/el2m

Document Number: EL2MONITORFSA4 REV 0