AN14758

Cycle Count Calculator and Guideline for i.MX RT700

Rev. 1.0 — 10 November 2025 Application note

Document information

Information	Content
Keywords	AN14758, i.MX RT700, Cycle Count Calculator
Abstract	This document provides an overview of the NXP i.MX RT700 architecture. It focuses on SRAM accesses, features available to enhance the performance or power reduction, and the memory accesses affected by the application choices.

1 Introduction

This document provides an overview of the NXP i.MX RT700 architecture. It focuses on SRAM accesses, features available to enhance performance or power reduction, and the memory accesses affected by the application choices.

A cycle count calculator, which provides an estimated cycle performance for a specific memory access, is available. For more information, see *Memory partitions calculator*.

2 General architecture

i.MX RT700 has five independent chip domains: Compute, Sense, Common, DSP, and Media.

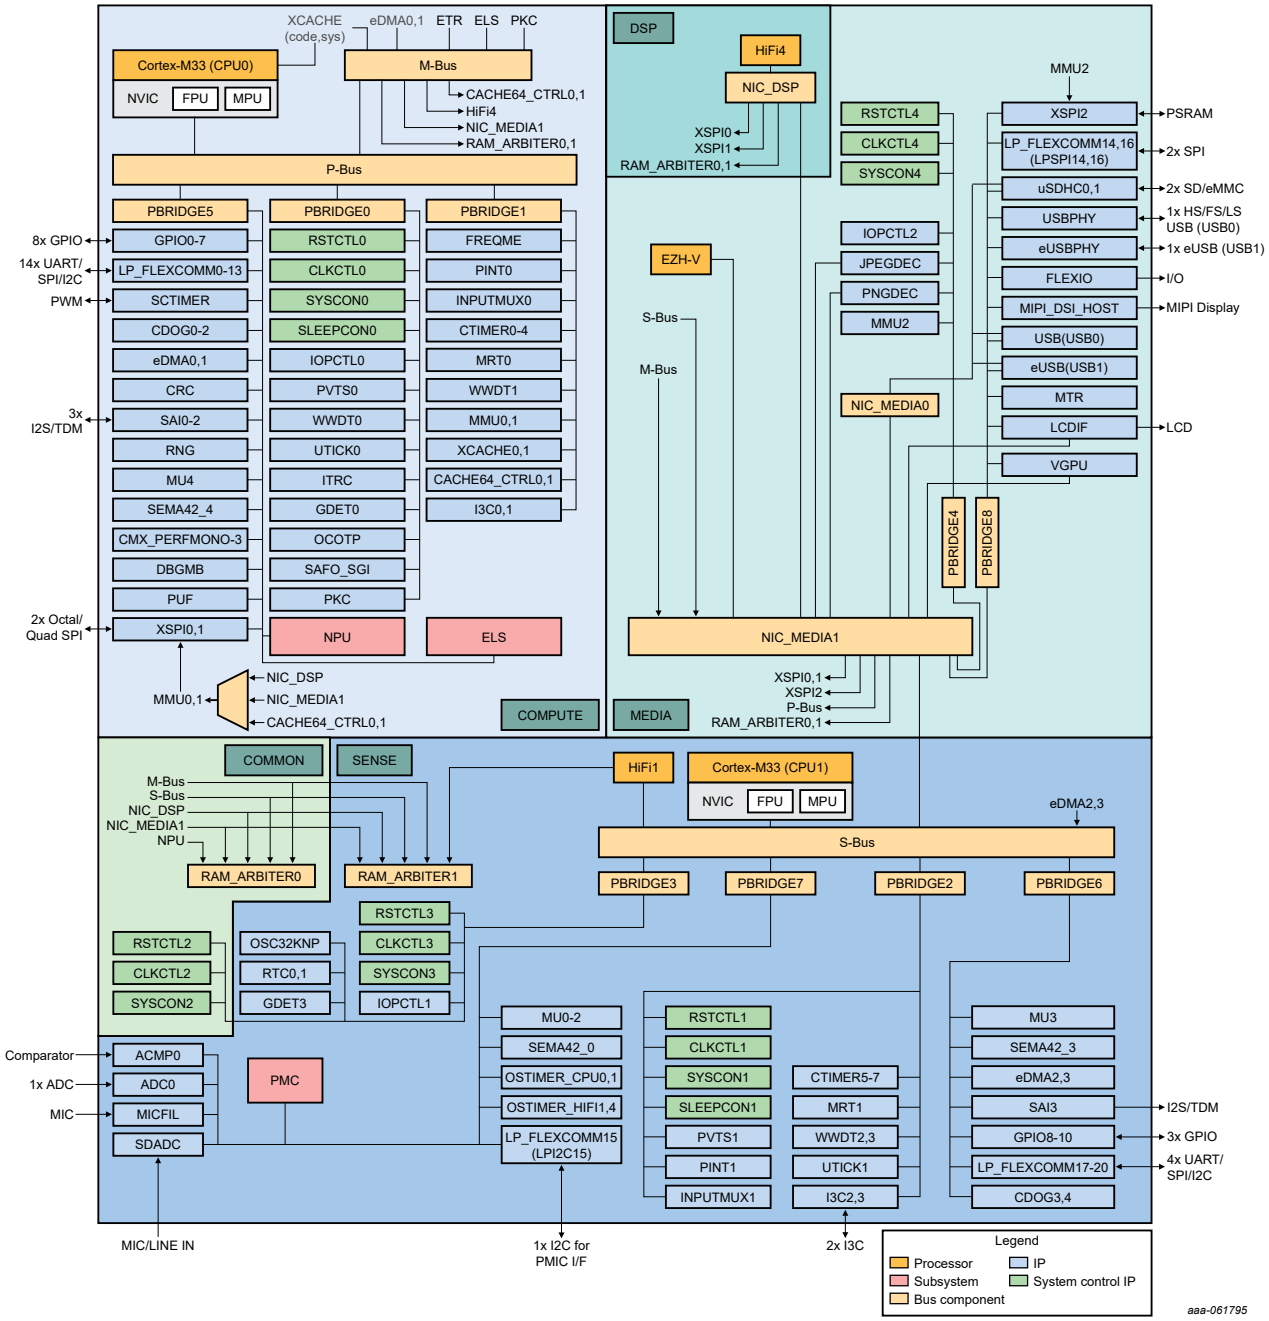

Figure 1. i.MX RT700 architecture

2.1 Compute Domain

This chip domain implements a full-featured Cortex M33 complex and a HiFi4 DSP. It features a split bus architecture to minimize the power consumption.

The bus is split into a memory bus (M-Bus) that connects the core Initiator to (shared) memory and a periphery bus (P-Bus) interfaces to peripherals to keep the toggling activity minimized to only required memory accesses. The M-Bus and P-Bus are connected through a Target port on the M-Bus to an Initiator port on the P-Bus.

Another architectural feature introduced in the Compute Domain is the integration of a cache controller associated with the M33. This provides a quasi-single-cycle access to the shared SRAM and helps eliminate power consumption dependence on the physical position of the SRAM.

2.2 Sense domain

This chip domain has its own bus, S-Bus. This bus connects the other Cortex-M33 core (CPU1), a HiFi1 DSP, and the two DMA controllers as core Initiators.

The CPU1 of the Sense domain can access SRAM for code and data via S-Bus. The HiFi1 DSP does not have a private cache memory, instead its dual-memory TCM bus interfaces are connected to shared memory.

2.3 Common domain

This chip domain holds modules related to system control, clock generation, reset control, pin characteristic control, and internal power regulators, as well as the shared SRAM (P0-P17) and peripherals, such as the memory interface of XSPI0-2.

2.4 Media domain

The purpose of the Media domain is to hold peripherals that require high performance. The Media domain includes a NIC interface (NIC_MEDIA0, NIC_MEDIA1 controllers), the uSDHC (NAND) controllers, the USB (USB0) and eUSB (USB1) IPs, and the graphics subsystem integrated by the JPEG decoder, PNG decoder, VGPU, LCDIF, DSI host interface, as well as a MIPI PHY. The Media domain includes high-speed SPIs, and the smart DMA EZH-V. The Media domain has cross-links to the Compute and Sense domains, and direct access to the shared memory.

2.5 DSP domain

The DSP domain holds the HIFI4 and includes a NIC interface. It has direct access to the shared memory, but also to peripherals in the compute, media and sense domains.

Future cycles values and penalties mentioned in this document are with CACHE disabled or when a cache miss happens. It is the minimum value and can greatly vary depending on many factors, such as RAM activity, RAM clocking, and bus arbitration. It also demonstrates the impact of available features on memory accesses and must be used for guidance only.

3 RAM partitions and RAM arbiters

The 7.5 MB of SRAM is available on the i.MX RT700 and it is divided into two different domains:

- The 5.5 MB of RAM residing in the compute domain is split in 18 individual partitions of different sizes (P0 to P17)
- The 2 MB of RAM residing in the sense domain is split in 12 individual partitions (P18 to P29)

They are in a separate voltage domain and can operate at different voltages.

The whole SRAM is accessible by the compute and sense domain. However, NPU can only access SRAM partition in the compute domain, and HIFI1 can only access SRAM in the sense domain.

The memory is organized in to groups that are composed of partitions. Each partition connects to the underlying memory with an interface called a memory arbiter. The arbiter interfaces the memory to the various buses, that is initiators. Indeed, multiple initiators ports can be connected to a single memory arbiter. Arbiters support AHB and AXI protocols.

Also, memories and arbiters are endpoints, that is they do not have their own dedicated clocks. Instead, they use the clocks of the transaction initiators accessing the shared memory.

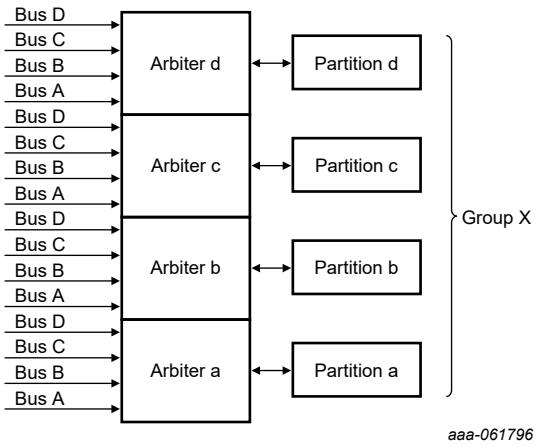

Figure 2. RAM arbiters and partitions

4 Initiators

Initiators to the RAM arbiters are the 6 buses of the i.MX RT700.

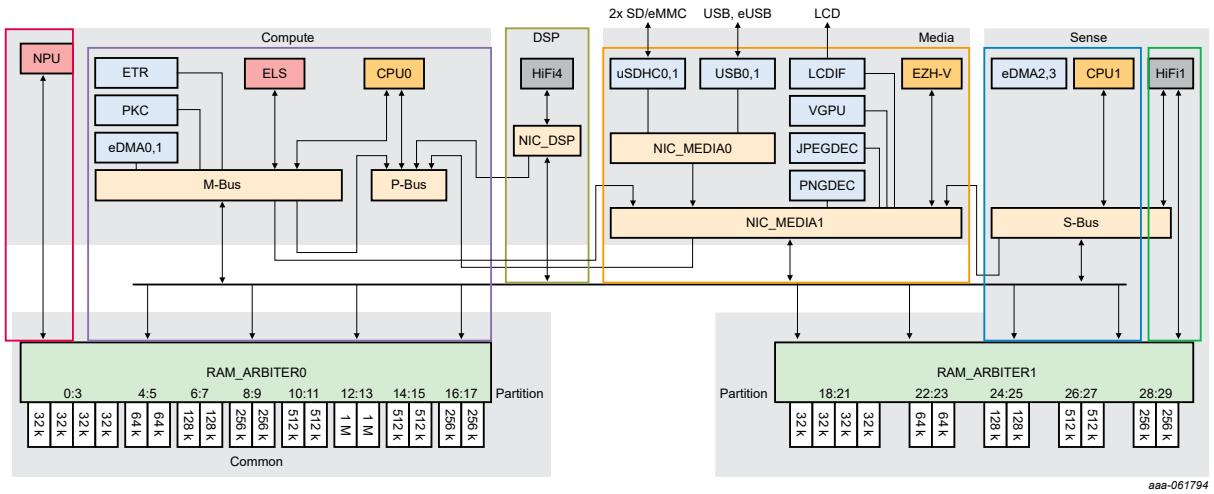

Figure 3. i.MX RT700 initiators architecture

From left to right in the Figure 3:

- **NPU** uses AXI protocol to communicate with the SRAM.
- **M-BUS is** an AHB fabric with multiple controllers connected to different ports as the CPU0, eDMA0, and EDMA1, ELS, PKC, ETR. Resides in the compute domain and uses the AHB protocol to communicate with the memory.
- **NIC_DSP** is a NIC that is a specific initiator to the NIC_MEDIA. NIC_DSP has two separate ports to communicate with SRAM in the compute domain and in the sense domain. The communication is done using the AXI protocol.
- **NIC_MEDIA1** is a NIC that connects multiple initiators to targets, as the SRAM in both the compute and sense domain. The AXI protocol is used to communicate with the different initiators and targets. AHB-to-AXI bridges can be used to convert initiator protocol to an NIC-compliant protocol.

- **S-BUS** is an AHB fabric with multiple controllers connected to different ports as the CPU1, eDMA2, and EDMA3. Resides in the sense domain and uses the AHB protocol to communicate with the memory.
- TCM initiators include TCM-I or TCM-D and are used to receive read/write transactions from the **HIFI1 TCM** interface, and then forwards them to the connected memory partitions.

Each bus is connected to memory via dedicated ports. Here is a description of the Initiator connections to the SRAM partitions via different ports.

Table 1. Initiators port for SRAM partition access

Group	Partitions	Size (kB)	MBUS	NIC_MEDIA1	S-BUS	NIC_DSP	HIFI1_TCM	NPU
G0	0	32	P_A_4	P_C_3	P_E_0	s_d_3	N/A	npu_axi
	1	32						
	2	32						
	3	32						
G1	4	64	P_A_5					
	5	64						
G2	6	128	P_A_6					
	7	128						
G3	8	256	P_A_7					
	9	256						
G4	10	512	P_A_8	P_C_4	P_E_1			
	11	512						
G5	12	1024	P_A_9	P_C_5	P_E_2			
	13	1024						
G6	14	512	P_A_10	P_C_6	P_E_3			
	15	512						
G7	16	256	P_A_11	P_C_7	P_E_4			
	17	256						
G8	18	32	P_A_12	P_C_8	P_E_5	s_d_4	TCM_0	N/A
	19	32						
	20	32						
	21	32						
G9	22	64			P_E_6			
	23	64						
G10	24	128			P_E_7			
	25	128						
G11	26	512	P_A_13	P_C_9	P_E_8		TCM_1	
	27	512						
G12	28	256		P_C_10	P_E_9			
	29	256						

5 Mechanisms to consider

Different mechanisms and features must be considered as they can drastically impact a memory access performance. Here are a few key elements to consider.

The access time from an Initiator to the SRAM partition can vary depending on many factors. Indeed, how often the SRAM partition is accessed and if accessed by multiple initiators, it can impact the performance.

5.1 Multiple accesses to a single SRAM partition

A port on buses or NICs may connect to one or more memory partitions, but it can access one partition at a time. If one memory partition is accessed by several Initiators, it reduces the performance because of latency.

When multiple controllers connected to the same bus request a memory access at the same time, a bus arbitration occurs. Since only one controller can use the bus at a time, the arbitration ensures orderly and fair access depending on arbitration schemes based. For example, on priority conditions. If a low-priority controller requests access to a memory partition at the same time as a high-priority controller, its performance will be impacted. Because the access is stalled until the other high-priority controller releases the bus.

5.2 Initiators clock switch

As explained before, the memory and arbiter use the clock from the initiator. If the same partition was last accessed by an Initiator using a different clock, then a switch to the new initiator clock happens and can impact the overall performance. These clocks can potentially be turned OFF to maximize power savings.

Table 2 and Table 3 summarizes which register controls which initiator clock.

Table 2. Clocks required when accessing SRAM partition in the Compute domain

Clock	Register
COMPUTE_MAIN_CLK for M-BUS	Automatically gated when no access
COMMON_RAM_CLK for S-BUS	CLKCTL1.PSCCTL1[SENSE_ACCESS_RAM_ARBITER0]
Hifi4 clock	Clock off HiFi4 to turn it off
NPU clock	CLKCTL0.PSCCTL5[NPU0]
MEDIA_MAIN_CLK for NIC_MEDIA1	CLKCTL0.PSCCTL5[MEDIA_ACCESS_RAM_ARBITER0]

Table 3. Clocks required when accessing SRAM partition in the Sense domain

Clock	Register
SENSE_RAM_CLK for M-BUS	CLKCTL0.PSCCTL5[COMP_ACCESS_RAM_ARBITER1]
SENSE_MAIN_CLK for S-BUS	Automatically gated when no access
Hifi4 clock	CLKCTL0.PSCCTL5[HiFi4_ACCESS_RAM_ARBITER1]
Hifi1 clock	Clock off HiFi1 to turn it off
SENSE_RAM_CLK for NIC_MEDIA1	CLKCTL3.PSCCTL0_COMP[MEDIA_ACCESS_RAM_ARBITER1] or CLKCTL3.PSCCTL0_SENS[MEDIA_ACCESS_RAM_ARBITER1]

The clock switching requires that both the last clock and the new clock are available, so prematurely turning off the last access clock prevents the clock switching and blocks any new access to that partition.

5.3 Automatic clock gating

Automatic clock gating feature is available. This feature turns OFF clocks to each internal SRAM after 5 bus clocks with no activity. This feature allows to reduce the power consumption; however, it impacts the memory access performance.

The automatic clock gating for SRAM partitions in the compute domain is controlled by the register `SYSCON0->AUTOCLKGATEOVERRIDE0[x]`, with x the memory partitions P0 to P17 in the compute domain.

`SYSCON3->AUTOCLKGATEOVERRIDE0[x]`, with x 0 to 11 corresponding to P18 to P29 in the sense domain, controls the automatic clock gating for the SRAM partitions in the sense domain.

Clearing the bit enables the clock gating for this partition while setting the bit set the clock partition to continuous clocking.

5.4 Cross-domain memory access

As mentioned in the above sections, the COMPUTE, and SENSE domain can access the whole 7.5 MB of SRAM distributed across the two domains. Also, able to operate at different voltages and frequencies. Synchronizers between the domains are required to ensure proper communication and access between domains. The synchronizer takes both the domains clocks to synchronize it and the frequency of these clocks can impact the overall SRAM access from one domain to another. Depending on which domains are crossed during the access, different clocks are involved. For instance, for accessing from the COMPUTE domain to the SENSE domain, COMPUTE_MAIN_CLK and SENSE_RAM_CLK are involved.

For the respective clocks, see Table 2 and Table 3.

5.5 Potential penalties

Here is a summary of the potential penalties to add to the typical access time depending on the configuration:

If auto clock gating is enabled, there is a potential of 10 cycles penalty to wake up the clock.

If the previous and current initiator uses different clocks, there is a potential of 2 cycles penalty to switch the arbiter clock.

If cross-domain access requiring synchronizer, there is a potential of in the best case is 5 controller clock cycles penalty + 5 follower clock cycles penalty. However, it may vary depending on the clock phase relationship between the two clocks, with a potential penalty of 6 controller clock cycles + 6 follower clock cycles.

The cycle count calculator provides an estimation of the number of cycles required for a specific memory access. For more information, see Memory partitions calculator.

You can select the previous initiator, current initiator, partition group to access and if the auto clock gating feature is enabled. The result of the calculator is the minimum number of cycles required for the initiator to access a partition in the selected group. However, the number of cycles varies depending on the different elements explained in the previous section.

6 Conclusion

- For better performance, ensure that a shared memory partition is accessed with only one Initiator at a time.
- Use cross-domain RAM partition access only if it is mandatory.
- If power reduction is more important than performance, enable the auto clock-gating feature, and/or turn OFF the Initiator clocks, when not required.

7 Acronyms and abbreviations

Table 4 provides the acronyms and abbreviations used in this document.

Table 4. Acronyms and abbreviations

Acronym	Abbreviation
SRAM	Static Random Access Memory
DSP	Digital Signal Processing
DMA	Direct Memory Access
TCM	Tightly Coupled Memory
XSPI	External Serial Peripheral Interface
NIC	Network Interconnect Component
AHB	Advanced High-performance Bus
AXI	Advanced eXtensible Interface
ELS	EdgeLock Secure Subsystem
PKC	Public Key Cryptography Co-processor
ETR	Arm Embedded Trace Router

8 Revision history

Table 5 summarizes the revisions done to this document.

Table 5. Revision history

Document ID	Release date	Description
AN14758 v.1.0	10 November 2025	Initial public release

Legal information

Definitions

Draft — A draft status on a document indicates that the content is still under internal review and subject to formal approval, which may result in modifications or additions. NXP Semiconductors does not give any representations or warranties as to the accuracy or completeness of information included in a draft version of a document and shall have no liability for the consequences of use of such information.

Disclaimers

Limited warranty and liability — Information in this document is believed to be accurate and reliable. However, NXP Semiconductors does not give any representations or warranties, expressed or implied, as to the accuracy or completeness of such information and shall have no liability for the consequences of use of such information. NXP Semiconductors takes no responsibility for the content in this document if provided by an information source outside of NXP Semiconductors.

In no event shall NXP Semiconductors be liable for any indirect, incidental, punitive, special or consequential damages (including - without limitation - lost profits, lost savings, business interruption, costs related to the removal or replacement of any products or rework charges) whether or not such damages are based on tort (including negligence), warranty, breach of contract or any other legal theory.

Notwithstanding any damages that customer might incur for any reason whatsoever, NXP Semiconductors' aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms and conditions of commercial sale of NXP Semiconductors.

Right to make changes — NXP Semiconductors reserves the right to make changes to information published in this document, including without limitation specifications and product descriptions, at any time and without notice. This document supersedes and replaces all information supplied prior to the publication hereof.

Applications — Applications that are described herein for any of these products are for illustrative purposes only. NXP Semiconductors makes no representation or warranty that such applications will be suitable for the specified use without further testing or modification.

Customers are responsible for the design and operation of their applications and products using NXP Semiconductors products, and NXP Semiconductors accepts no liability for any assistance with applications or customer product design. It is customer's sole responsibility to determine whether the NXP Semiconductors product is suitable and fit for the customer's applications and products planned, as well as for the planned application and use of customer's third party customer(s). Customers should provide appropriate design and operating safeguards to minimize the risks associated with their applications and products.

NXP Semiconductors does not accept any liability related to any default, damage, costs or problem which is based on any weakness or default in the customer's applications or products, or the application or use by customer's third party customer(s). Customer is responsible for doing all necessary testing for the customer's applications and products using NXP Semiconductors products in order to avoid a default of the applications and the products or of the application or use by customer's third party customer(s). NXP does not accept any liability in this respect.

Terms and conditions of commercial sale — NXP Semiconductors products are sold subject to the general terms and conditions of commercial sale, as published at https://www.nxp.com/profile/terms, unless otherwise agreed in a valid written individual agreement. In case an individual agreement is concluded only the terms and conditions of the respective agreement shall apply. NXP Semiconductors hereby expressly objects to applying the customer's general terms and conditions with regard to the purchase of NXP Semiconductors products by customer.

Export control — This document as well as the item(s) described herein may be subject to export control regulations. Export might require a prior authorization from competent authorities.

HTML publications — An HTML version, if available, of this document is provided as a courtesy. Definitive information is contained in the applicable document in PDF format. If there is a discrepancy between the HTML document and the PDF document, the PDF document has priority.

Translations — A non-English (translated) version of a document, including the legal information in that document, is for reference only. The English version shall prevail in case of any discrepancy between the translated and English versions.

Security — Customer understands that all NXP products may be subject to unidentified vulnerabilities or may support established security standards or specifications with known limitations. Customer is responsible for the design and operation of its applications and products throughout their lifecycles to reduce the effect of these vulnerabilities on customer's applications and products. Customer's responsibility also extends to other open and/or proprietary technologies supported by NXP products for use in customer's applications. NXP accepts no liability for any vulnerability. Customer should regularly check security updates from NXP and follow up appropriately.

Customer shall select products with security features that best meet rules, regulations, and standards of the intended application and make the ultimate design decisions regarding its products and is solely responsible for compliance with all legal, regulatory, and security related requirements concerning its products, regardless of any information or support that may be provided by NXP.

NXP has a Product Security Incident Response Team (PSIRT) (reachable at PSIRT@nxp.com) that manages the investigation, reporting, and solution release to security vulnerabilities of NXP products.

Suitability for use in industrial applications (functional safety) — This NXP product has been qualified for use in industrial applications. It has been developed in accordance with IEC 61508, and has been SIL-classified accordingly. If this product is used by customer in the development of, or for incorporation into, products or services (a) used in safety critical applications or (b) in which failure could lead to death, personal injury, or severe physical or environmental damage (such products and services hereinafter referred to as "Critical Applications"), then customer makes the ultimate design decisions regarding its products and is solely responsible for compliance with all legal, regulatory, safety, and security related requirements concerning its products, regardless of any information or support that may be provided by NXP. As such, customer assumes all risk related to use of any products in Critical Applications and NXP and its suppliers shall not be liable for any such use by customer. Accordingly, customer will indemnify and hold NXP harmless from any claims, liabilities, damages and associated costs and expenses (including attorneys' fees) that NXP may incur related to customer's incorporation of any product in a Critical Application.

NXP B.V. — NXP B.V. is not an operating company and it does not distribute or sell products.

Trademarks

Notice: All referenced brands, product names, service names, and trademarks are the property of their respective owners.

NXP — wordmark and logo are trademarks of NXP B.V.

AMBA, Arm, Arm7, Arm7TDMI, Arm9, Arm11, Artisan, big.LITTLE, Cordio, CoreLink, CoreSight, Cortex, DesignStart, DynamIQ, Jazelle, Keil, Mali, Mbed, Mbed Enabled, NEON, POP, RealView, SecurCore, Socrates, Thumb, TrustZone, ULINK, ULINK2, ULINK-ME, ULINK-PLUS, ULINKpro, µVision, Versatile — are trademarks and/or registered trademarks of Arm Limited (or its subsidiaries or affiliates) in the US and/or elsewhere. The related technology may be protected by any or all of patents, copyrights, designs and trade secrets. All rights reserved.

EdgeLock — is a trademark of NXP B.V.

Contents