IMXRT1180CCE

Chip Errata

Rev. 3 — 25 June 2025 Errata

1 i.MX RT1180C Chip Errata

1.1 Revision History

This document details all known silicon errata for the:

i.MX RT1180C

Table 1. Revision History

Revision	Release Date	Significant Changes
3	6/2025	The following errata were removed. • ERR052031 • ERR052165 The following errata were revised. • ERR052288
2	4/2025	The following errata were removed. • ERR052134 The following errata were added. • ERR052163 • ERR052358 • ERR052343 • ERR052344 • ERR052206 • ERR052164 • ERR052165 • ERR052167 • ERR052288 • ERR052198
1	12/2024	Initial Revision

1.2 Summary of Silicon Errata

This table summarizes all known errata and lists the corresponding silicon revision level to which they apply. A 'Yes' entry indicates the erratum applies to a particular revision level, and a 'No' entry means it does not apply.

Table 2. Summary of Silicon Errata

Erratum ID	Erratum Title
ADC	
ERR051385	ADC: ADC INLDNL degrade under high ADC clock frequency when VREFH selected as reference
FlexSPI	
ERR011377	FlexSPI: DLL lock status bit not accurate due to timing issue

Table 2. Summary of Silicon Errata*...continued*

Erratum ID	Erratum Title
FlexSPI_follower	
ERR052145	FlexSPI_follower: Block behaviors may mismatch the definitions in register after software reset
GPT	
ERR003777	GPT: Possibility of additional pulse on src_clk when switching between clock sources
I3C	
ERR052086	I3C: RX FIFO may require separate read for the last data in DMA operation
ERR052343	I3C: Target Early Termination Feature not available with DMA controller
ERR052344	I3C: Controller Clock stalling feature not available in I3C Controller
LPSPI	
ERR051588	LPSPI:Reset transmit FIFO after FIFO underrun by LPSPI Slave.
LPUART	
ERR051629	LPUART:Transmit Complete bit (STAT[TC]) is not set.
NETC	
ERR051188	NETC: Register SITSR not accessible for Virtual Station Interfaces
ERR051444	NETC : ICM discards when two priority queues are in use can cause potential memory loss
ERR051460	NETC: Ingress Port Filter (Table ID 13) can overflow rsp buffer during multi-entry query management command
ERR051524	NETC: Ingress Stream Identification Payload construction evaluates incorrectly for frames >1kB
ERR051530	NETC: hash multi-entry search table management continuation command can miss returning entries
ERR051587	NETC: Time gate scheduling update command response can be erroneous if the AdminBaseTime specified is near the current time
ERR051607	NETC: NETC PF may override VF bus master enable and MSI-X function mask for MSI-X generation
ERR051616	NETC: FDB (Table ID 15) always overrides Egress Treatment Table access defined by Ingress Stream (Table ID 31)
ERR051622	NETC: TABLE_BIR of NETC TIMER is wrong
ERR051649	NETC: Switch ports support only one ‘802.1p and DEI to internal QoS’ mapping profile
ERR051651	NETC: Use of 10Mbps restricted to PHYs that support octet alignment for preamble with minimum 1B
ERR051707	NETC: MAC Merge stat inaccurate after lost Rx fragments
ERR051710	NETC: MAC Tx statistic counters inaccurate after half duplex late collision and excessive collision events
ERR051711	NETC: MAC Stats Octet Counter Mismatch after Pause Frame Transmitted with Flexible Preamble Enabled
ERR051887	NETC: VLAN qualifiers missed in ingress classification lookups
ERR051936	NETC: RGMII Half Duplex: MAC Tx FIFO status may not return to empty after FLR and transmitter may become inoperable after FLR

Table 2. Summary of Silicon Errata*...continued*

Erratum ID	Erratum Title
ERR051994	NETC: Energy Efficient Ethernet Tx autonomous Low Power Idle not supported
ERR052024	NETC: Half duplex transmit throughput degraded by up to 2.3%
ERR052129	NETC: Restricted support for short preamble and short IPG
ERR052131	NETC: Half duplex collision not triggered for early CRS in RMII
ERR052133	NETC: MAC does not detect mVerify after start fragment
ERR052163	NETC: RevMII MDIO read data is incorrect for register bit STATUS[PREAMBLE]
ERR052164	NETC: RevMII status bits AN_COMP and LINKST are not cleared when auto-negotiation is disabled
ERR052167	NETC: MAC Tx IPG is longer than expected when configured for MII half duplex
ERR052206	NETC: SG_DROP_COUNT value in the Ingress Stream Count STSE_DATA response begins at an incorrect bit offset, causing it to be read incorrectly.
ERR052288	NETC: The first frame after SequenceRecoveryReset (FRER) is dropped and miscounted
PWM	
ERR051374	PWM fault may work abnormally when the fault signal is very narrow
ERR051989	PWM: output may be abnormal when the value of phase delay register is reduced from a non-zero value to 0.
QDC	
ERR051383	QDC: Quadrature decoder CTRL[DMAEN] bit can not be cleared
SAI	
ERR051405	SAI: Synchronous mode with BYP=1 not supported
SYS_CTR	
ERR051250	SYS_CTR: SYS_CTR_COMPARE register write could be missed
ERR052358	SYS_CTR: CNTCV0 and CNTCV1 register cannot get count value when clocked by 32kHz alternate slow clock
VREF	
ERR052081	VREF: VREF trim value need be loaded with SW before enabling VREF
uSDHC	
ERR052198	uSDHC: eMMC CQE may timeout due to a HW logic issue

2 Known Errata

ERR051524: NETC: Ingress Stream Identification Payload construction evaluates incorrectly for frames >1kB

Description

The Ingress Stream Identification table (Table ID 30) lookup key is built using key construction registers. Payload data can be optionally included to the lookup key.

When present, the payload key field is taken from an offset relative to the Payload Ethertype. Payload Ethertype starts immediately after source MAC address or VLAN tag(s)/RTAG/HSR (if any). The payload key field must be within the frame, and within 116 bytes from the Payload Ethertype. Hardware checks these two conditions before constructing the Ingress Stream Identification table lookup key.

This check may evaluate incorrectly (indicating an invalid key construction) when the frame is received from a pseudo port (internal port) bound to ENETC or the switch, and is 1024 bytes or larger in size.

Workaround

Disable payload as part of the key construction if being used on a pseudo port (bound to ENETC or switch), when maximum frame sizes is expected (or set) to be 1kB or larger.

ERR051649: NETC: Switch ports support only one ‘802.1p and DEI to internal QoS’ mapping profile

Description

One of the methods to set the internal QoS of an incoming frame is to use the '802.1p and DEI to internal QoS' mapping function. The internal QoS represents the local QoS values of frame within the switch, namely the internal priority value (IPV) and drop resilience (DR) of the frame.

The '802.1p and DEI to internal QoS' mapping method can be selected on per switch port, along with specifying the '802.1p and DEI to internal QoS' mapping profile to be used. Two '802.1p and DEI to internal QoS' mapping profiles global to the switch that can be configured. Registers VLANIPVMPaR0/1 and VLANDRMPaR are used to configured the profiles where 'a' identifies the mapping profile instance, and can take the value of 0 or 1.

Only mapping profile instance 0 can be used. Mapping profile instance 1 cannot be used due the decoding of the selected profile is not correct. Therefore, if any switch port is configured to use the second mapping profile (1), all switch ports will not assign the default QoS correctly.

Note, the first profile (0) works correctly. The errata is only when at least one of the switch ports is not using profile 0.

Workaround

All switch ports must use the first VLAN QoS Mapping Profile (PQOSMR[VQMP]=0).

If multiple profiles is critical, then port classification lookups (Ingress Port Filter or Ingress Stream) can be used to provide alternative methods to set internal QoS values

ERR052129: NETC: Restricted support for short preamble and short IPG

Description

This device supports special modes for shorter than standard 7B preamble and shorter than 12 octet IPG on transmit, via PMn_TX_IPG_PREAMBLE. The supported ranges for flexible preamble are 1-7B and for minimum IPG 4-24 octets.

When both values are set to the minimum value (IPG_LEN=4 and FLEX_PREAMBLE_CNT=1, with FLEX_PREAMBLE_EN=1), and the device is connected to a NETC MAC (chip-to-chip or loopback), the receiving NETC MAC cannot reliably detect the frame. As a result, frames may be lost, or partially received with a CRC error (e.g. if there is a data octet later in the frame that matches SFD).

Workaround

When connecting NETC MAC to NETC MAC, for IPG_LEN=4, the minimum supported preamble is 2 bytes.

When connecting NETC MAC to another device that supports shorter-than standard IPG and preamble, ensure that the minimum IPG + preamble is 6 octets.

ERR051188: NETC: Register SITSR not accessible for Virtual Station Interfaces

Description

Register SITSR is a read-only status register which includes the timer function synchronization information. In addition there is a set of interrupt registers that trigger when there is a change in the SYNC bit.

All of these registers are intended to be present in each SI. While the interrupt registers are present, the SITSR register is not available in the VSIs, only in the PSI. The result is that a VSI may know when the SYNC bit changes but will not know the current status.

Workaround

A VSI may send a message to the PSI for general notification or to gain access to hardware resources using the VSIMSG/PSIMSG set of registers. This method can be used to gain access to the PSI SITSR register value when there is a change to the SYNC bit as determined by the VSI interrupt.

ERR051616: NETC: FDB (Table ID 15) always overrides Egress Treatment Table access defined by Ingress Stream (Table ID 31)

Description

Background:

There are 2 methods, via the primary Egress Treatment group assignment, to direct a frame to Egress Sequence Recovery Function used for duplicate frame elimination:

a) using Ingress Stream (IS) table

The ingress stream method to direct to the sequence recovery function is used when stream identifier is encoded in the Ethernet MAC Source address (MAC SA) and VLAN identifier

b) using Bridge method using VLAN and FDB (or IPv4 Multicast Filter) tables.

Can be used in case where the stream identifier is encoded in the Ethernet MAC Destination Address (MAC DA) and the VLAN identifier.

There are 2 methods to forward a frame to egress switch port:

a) Stream forwarding using Ingress Stream and

b) Bridge forwarding using VLAN and FDB (or IPv4 Multicast Filter) tables

Issue:

Using Ingress streams (Table ID 31) to identify Egress Sequence Recovery Function and using Bridging with FDB (Table ID 15) entries (with the Override Egress Treatment Entry ID (OETEID) option field is set to 00b) to forward frame to egress switch ports incorrectly overrides the primary Egress Treatment group assignment (i.e. the Egress Sequence Recovery Function) specified by the Ingress Stream.

This breaks duplicate frame elimination capability provided by Egress Sequence Recovery Function.

This is of particular concern when the Ingress stream is defined using MAC SA and the Bridging method is used to identify egress switch port.

Note: Egress Sequence Recovery is incorrectly overridden only if the bridging (FDB or IPv4 Multicast Filter) lookup finds an entry.

Egress Sequence Recovery is not overridden if FDB (or IPv4 Multicast Filter) entry is not found causing the frame to be flooded.

Workaround

Goal: Avoid having Ingress Stream (IS) specify ET_EID and using FDB entry to forward frame. 2 acceptable workaround.

1) Use IS to specify ET_EID but don’t use FDB entry

For each {VLAN, MAC SA} or {MAC SA} stream, add 2 Ingress Stream (IS) entries:

a) IS for terminating flows: Based on {VLAN, MAC SA, MAC DA} with FA=StreamFwd (010b) and destination port being ENETC’s Host MAC

b) IS for transiting flows: Based on {VLAN, MAC SA} with FA=Bridging (011b).

FDB entry will flood to all ports except the one received.

Another option: specify FA=StreamFwd with SPPD=0 (source port pruning enabled) & fwd to all ports. This eliminate need to do Bridging and use VLAN filter table.

Both IS entries would have same configuration except for FA field. That is, ET_EID & EGRESS_PORT_BITMAP are the same.

Note: IS for terminating flow has a higher precedence than transiting flows during Ingress Stream lookup (using either EM ISID or IPF TCAM lookup).

Pro: simple to implement.

Con: scalability concern if multiple MAC address needs to be migrated from FDB to IS table (#MAC SA streams * # MAC DA)

Note: For HSR, there would be 1 and maybe 2 MAC Addr. So scalability should not a concern.

Note: HSR implementing this workaround.

2) Use IS to identify internal VLAN which specifies ET_EID entry. FDB used to forward.

a) For each {VLAN, MAC SA}, {MAC SA} stream use a reserved VLAN (eg: VLAN 3000-4095)

b) IS adds interim VLAN (Ingress Frame Modification)

c) VLAN specifies the ET_EID entry responsible for sequence recovery function and removes the interim VLAN.

Pro: address scalability concern of Workaround 1

Con: more configuration needed.

ERR051444: NETC : ICM discards when two priority queues are in use can cause potential memory loss

Description

- Issue is related to the Ingress ENETC Ingress Congestion Manager (ICM) which provides the main internal buffering/queuing point on the ENETC receive datapath.

- At the ICM queuing point, frames can be optionally arranged by their schedule priority (priority queuing) and drop resiliency (drop resiliency queuing).

- When priority queuing is used (2 priority levels supported), a frame with high priority is scheduled to HTA Receive before a frame with low priority. Within a priority, first-in-first-out (FIFO) ordering is followed.

- When drop resiliency queuing (4 levels) is used, under congestion (queue buildup), a frame with lower drop resiliency is dropped first before frames with higher drop resiliency. Within a drop resiliency level, first-in-first-out (FIFO) ordering is followed, with dropping occurring at the head of the queue.

- Both priority and drop resiliency queuing are optional and not used when ENETC comes out of reset.

- Issue occurs when priority queueing is used, and ICM becomes congested, and under a combination of rare events, internal buffering memory is depleted/loss when frames are being dropped to relieve congestion.

Workaround

To avoid this defect, ICM priority queueing should not be used; always use single priority by keeping the Receive IPV to ICM priority mapping register 0 (IPV2ICMPMR0) setting to its power on reset value.

ERR051651: NETC: Use of 10Mbps restricted to PHYs that support octet alignment for preamble with minimum 1B

Description

802.3 frames are octet based, but preambles are not necessarily in multiples of octets, depending on interface data width. When operating at 10Mbps, MII or RMII PHYs may remove bits of the preamble in multiples of 4b, leaving the packet non-octet aligned.

This device only supports octet-aligned packets. When running 10Mbps MII or RMII on a PHY that does not support octet-only alignment, the device may discard or flag an error on packets with a non-octet aligned preamble, as it would be unable to recognize the SFD/SMD. For some PHYs, that may impact the majority of packets.

Some 10 Mbps PHYs may also strip the entire preamble. This device does not support that mode of operation, but requires minimum 1B of preamble.

Workaround

1. Use 100Mbps operation instead of 10 Mbps, or

2. Use a PHY that supports octet-only alignment for 10 Mbps and does not strip all preamble bytes

ERR052163: NETC: RevMII MDIO read data is incorrect for register bit STATUS[PREAMBLE]

Description

The RevMII MDIO register bit STATUS[PREAMBLE] reads incorrectly as 1 instead of 0 for both the MAC-side and PHY-side management registers. The incorrect read value of 1 indicates that the preamble can be suppressed for MDIO management frames. The STATUS[PREAMBLE] bit should always read as 0 to indicate preamble is required for all MDIO management frames.

Workaround

When operating in RevMII mode, both the local device and link partner should ignore the actual read value returned for the STATUS[PREAMBLE] and interpret it as zero. This means both devices must configure MDIO accesses to use the normal preamble length of 32 bits as defined in 802.3 Clause 22.

ERR051887: NETC: VLAN qualifiers missed in ingress classification lookups

Description

The Ingress Port Filter (Table ID 13) and Ingress Stream Identification (Table ID 30) lookups can optionally include in the matching key outer and inner VLAN tag information from the received fame.

For the Ingress Port Filter tale, when Hardware constructs the lookup key for a frame header field for which the frame header is not present in the frame, it sets the key field to zero. If zero corresponds to a valid value in a frame header field, then this could lead to a false positive match. To avoid these false positive matches, present frame header (or field) 1-bit key matching fields (or frame attribute flags) are defined in the table entry to explicitly specify whether the presence of a frame header (or field) must be present or not present in a frame header, for an entry to match.

For VLAN headers, an Outer VLAN Tag Present attribute flag and Inner VLAN Tag Present attribute flag have been defined. The expectation is that each of these present matching key flag will be set if parser has found a VLAN tag in their respective position (outer/inner) in the received frame, and it has passed the TPID acceptance criteria for the tag position in the frame (defined in Port TPID acceptance register (PTAR)). Currently the issue (or behaviour) is that each of these present matching key flag will be set if parser has found a VLAN tag in their respective position (outer/inner), regardless of the TPID acceptance criteria for the tag position in the frame. This is applicable for frames received from either the Switch or an ENETC.

For the Ingress Stream Identification table, when Hardware constructs the lookup key for a frame the same behaviour is used as the Ingress Port Filter table - the VLAN tag is found in the frame and has passed the acceptance criteria. Currently the issue (or behaviour) is that if VLAN is not found in the frame or hasn’t passed the acceptance criteria, the default VLAN information is used in the lookup key. This is only applicable for frames received from an ENETC.

Workaround

Option A (Ingress Stream Identification only): Do not use the VLAN tag present/not present attribute in the Ingress Stream Identification lookup.

Option B: Add an entry with higher precedence (Ingress Port Filter) or using the first key construction (Ingress Stream Identification) that contains the VLAN attributes (VID, PCP) of the default VLAN. Configure these entries with the applicable result actions for frames that use the default VLAN information. Frames that do not contain the default VLAN information will match an entry with a lower precedence (Ingress Port Filter) or during the second Ingress Stream Identification lookup.

ERR051936: NETC: RGMII Half Duplex: MAC Tx FIFO status may not return to empty after FLR and transmitter may become inoperable after FLR

Description

MAC Tx FIFO status may not report empty after FLR when operating in RGMII half duplex mode. In some cases, the transmitter may become inoperable and not be able to recover from FLR requiring a full reset instead. The issue can occur when FLR is triggered around the time MAC Tx has started backing off due to a half duplex collision detection.

Workaround

This issue can be avoided by transitioning to full duplex mode just prior to triggering the FLR.

ERR052131: NETC: Half duplex collision not triggered for early CRS in RMII

Description

In RMII operating in half-duplex mode potential collisions will not be detected until after receiving preamble. Pin CRS_DV should assert as soon as preamble is about to begin, but it occurs later resulting in delayed detection of a potential packet collision. In half-duplex mode, pins MII_CRS0 (Carrier sense, only HD) and MII_RX_DV0 (Data valid) rely on assertion of RMII_CRS_DV (Carrier sense, data valid), which is delayed. System recovers normally from this condition.

Workaround

When using RMII in half-duplex mode expect collisions occurring at the beginning of the preamble to be detected with a delay and align the system to react accordingly.

ERR052133: NETC: MAC does not detect mVerify after start fragment

Description

A link partner should be able to initiate preemption verification at any time, e.g. for a port reset or SW reconfiguration of a port. If the MAC receives an mVerify packet after a start fragment, and before the corresponding ending continuation fragment, it will not recognize the mVerify packet, and the link partner verify process will time out.

Workaround

The recommended workaround in a closed system is to disable preemption verification. On NETC, that is done by setting MAC_MERGE_MMCSR[VDIS]=1 before setting [ME]=1.

In an open system where VDIS must be 0, system SW must provide an alternate mechanism for signaling when preemption verify has failed on the link partner. SW can then clear the MAC receiver state by toggling PM1_COMMAND_CONFIG[RX_EN], which will abort the preemptable packet in progress, after which the MAC will be able to receive mVerify packets again.

ERR051460: NETC: Ingress Port Filter (Table ID 13) can overflow rsp buffer during multi-entry query management command

Description

The Ingress Port Filter table which is implemented using a Ternary Content Addressable Memory (TCAM), supports a search query-operation command (ACCESS_METHOD=Search) which is used to return either the “Entry ID” (QUERY_ACTIONS=1) or the full content (QUERY_ACTIONS=0) of every entry in the table. When the full content of every entry in the table is queried (QUERY_ACTIONS=0), the command may return more response data than will fit in the allocated buffer if one or more entry is compressed and the response buffer size is too small to accommodate all query entries. The extra data will be written beyond the end of allocated buffer.

Compression is used to reduce the amount of TCAM space required to store each table entry. A narrow (48-bit wide), but deep TCAM organization is used, where a table entry can occupy multiple sequential lines of the TCAM. TCAM lines for which all fields are masked out (their corresponding mask field set to all 0s) will not actually exist in the TCAM, thus what is being referred above as compression

Workaround

Option A:

Issue a query table management command, with ACCESS_METHOD=Search and QUERY_ACTIONS=1. This returns only the list of entry IDs currently in use.

For each entry ID returned, issue a query table management command, with ACCESS_METHOD=Entry ID Match and QUERY_ACTIONS=0. This returns the full entry data of each entry.

Option B:

Allocate a large enough response buffer to fit all potential entries. The response buffer needs to be 4+(N*232) bytes, where N is the max possible number of entries (96 entries in RT1180, 64 entries in S32Z1/Z2/E2 devices).

ERR052206: NETC: SG_DROP_COUNT value in the Ingress Stream Count STSE_DATA response begins at an incorrect bit offset, causing it to be read incorrectly.

Description

The Ingress Stream Count (Table ID 38) contains multiple statistic counters. When the statistics are queried by the software, one of the counters in the response buffer, namely SG_DROP_COUNT is returned in the response buffer shifted. No information is lost, but the placement of the information in the response buffer is sifted compared to expectation per the reference manual.

Workaround

Adjust the starting offset for the SG_DROP_COUNT statistic in the response buffer from the definedby 7-bits

The actual offset of the SG_DROP_COUNT in the STSE_DATA element is bit 199

Or the offset within the entire response buffer format is bit 231

ERR051994: NETC: Energy Efficient Ethernet Tx autonomous Low Power Idle not supported

Description

The NETC has two methods for entering Tx low power idle state for Energy Efficient Ethernet:

- Direct SW control via PMn_COMMAND_CONFIG[TX_LOWP_ENA]. This is a debug mode not intended for general use

- Autonomous entry/exit based on Tx activity via PMn_SLEEP_TIMER and PMn_LPWAKETIMER

The NETC does not always obey the wakeup time in PMn_LPWAKETIMER. and as a result may transmit frames before the PHY has woken up from low power state. Such frames would be lost.

Workaround

Set PM0_SLEEP_TIMER[SLEEPT]=0 to disable autonomous low power idle on Tx.

ERR052024: NETC: Half duplex transmit throughput degraded by up to 2.3%

Description

When configured for half duplex mode, the theoretical maximum transmit throughput is degraded by up to 2.3%. This issue can occur when transmitting back-to-back frames in the absence of any receive traffic and/or collisions. When transmitting back-to-back frames in half duplex mode, the minimum Inter-packet Gap (IPG) is 14 rather than 12 bytes times, resulting in the degradation. The transmit performance degradation decreases for larger frame sizes as follows:

Frame Length (B) % Degradation

64 2.3

128 1.3

256 0.7

512 0.4

1024 0.2

1518 0.1

Workaround

For devices such as RT1180 supporting flexible preamble, the preamble length can be decreased by 2 if possible to recover the lost IPG time. There is no workaround on other devices without flexible preamble support. The application must tolerate the reduced transmit throughput when configured for half duplex.

ERR051607: NETC: NETC PF may override VF bus master enable and MSI-X function mask for MSI-X generation

Description

Each PCIe function within NETC allocates its MSI-X table resources from a single larger MSI-X table that serves the entire NETC. An ENETC instance acquires a total table size that includes both PF and VF(s)

as defined by IERB register EaMCR[NUM_MSIX], which is split between PF and VF(s) as defined by PSI0CFGR[NUM_MSIX] and PSIaCFGR[NUM_MSIX] (a>0) respectively.

When ENETC PF sets the PCIe MSI-X Function Mask (PCI_CFC_MSIX_MSG_CTL[FUNC_MASK]=1), it will also override the VF function's setting of the same bit, making it appear to the VF MSI-X table as always set. This has the potential side effect of making VF MSI-X interrupt go to the PBA (Pending Bit Array) instead of generating an MSI-X interrupt.

When ENETC PF sets the PCIe Bus Master Enable bit (PCI_CFH_CMD[BUS_MASTER_EN]=1), it will also override the VF function's setting of the same bit, making it appear to the VF MSI-X table as always set. This has the potential side effect of generating an MSI-X interrupt by VF when in fact is has disabled it by setting PCI_CFH_CMD[BUS_MASTER_EN]=0.

Workaround

For an ENETC PF that requires the use of masking an MSI-X interrupt, it should rely on the per-vector masking instead of the global MSI-X table function mask. This is a bit part of the MSI-X table entry VC Control field.

Optionally if the ENETC PF uses the function mask during processing of an interrupt, it should only temporarily set this bit, allowing the VF to log interrupts in the PBA and when PF clear the function mask, VF interrupts will again be generated. This could result in a delay of VF interrupt notification.

ERR052164: NETC: RevMII status bits AN_COMP and LINKST are not cleared when auto-negotiation is disabled

Description

After disablling RevMII auto-negotiation by setting CONTROL[AUTO]=0, the STATUS[AN_COMP] and STATUS[LINKST] status bit are not automatically cleared. Both the AN_COMP and LINKST status bits should be cleared in this case according to 802.3 Figure 37-6 (Auto-Negotiation state diagram).

Workaround

After clearing CONTROL[AUTO] to disable RevMII auto-negotiation, setting CONTROL[RESTART_AN] will then trigger STATUS[AN_COMP] to then be cleared as expected.

ERR052167: NETC: MAC Tx IPG is longer than expected when configured for MII half duplex

Description

Actual MAC Tx IPG is longer than configured when transmitting back-to-back packets in MII half duplex mode by approximately 15 extra bytes. For example, when configured for IPG=12, the actual IPG will be approximately 27.

The net result is that maximum throughput will be reduced also in the absence of half-duplex collision/retry events.

Workaround

Use full duplex mode when possible. If full duplex is not possible and half duplex must be enabled, then RMII interface mode is recommended instead of MII. In cases where MII half duplex must be used, reducing the size of the transmitted preamble using register field PM0_TX_IPG_PREAMBLE[FLEX_PREAMBLE_CNT] may have

the effect of regaining some of the lost bandwidth on the link if the link partner is able to support a reduced-length preamble on Rx.

Reducing NETC IPG by adjusting register field PM0_TX_IPG_PREAMBLE[IPG_LEN] to less than the default 12 does not affect the behavior of this erratum.

ERR051530: NETC: hash multi-entry search table management continuation command can miss returning entries

Description

NETC implements a common hash space that is used for multiple tables, across functions. The hash table management commands support a multi-entry Search command (ACCESS_METHOD = Search). The Search command completes when the full hash space has been searched, or the response buffer has been filled with matching entries. For the latter, Hardware returns a resume entry ID which Software uses to continue the Search command.

When the resume point of a continued Search command is in the middle of a collision chain (i.e. the last entry returned in the previous Search command segment, that filled the response buffer, was not the last entry in the collision chain), and the resume point cannot be found, Hardware may not return the matching entries of the next valid collision chain.

Workaround

OPTION A:

- provide a large enough buffer to hold all entries, avoiding the need to resume.

OPTION B:

- When resuming a search command, zero the lower 4 bits of the returned resume entry ID, and issue a resume search. NOTE: this may return some duplicates overlapping with the end of the previous search, up to HBTCAPR[MAX_COL] entries.

ERR052288: NETC: The first frame after SequenceRecoveryReset (FRER) is dropped and miscounted

Description

There is a slight variation from the standard in the implementation of the IEEE 802.1CB RECOVERY_TIMEOUT event.

IEEE 802.1CB specifies that the SequenceRecoveryReset function is to be executed immediately after a RECOVERY_TIMEOUT event has occurred, meaning that the first packet received after the RECOVERY_TIEMOUT event, is processed using the sequence recovery function reset variables. In other words, the SequenceRecoveryReset should take effect as soon as the timeout expires.

However, the defect (or variation from the standard) is that the SequneceRecoveryReset only takes effect after the first frame arrives after the RECOVERY_TIMEOUT event . This first frame should be processed using the sequence recovery function reset variables, but it is likely to be discarded as it is processed using the sequence recovery function variables values prior to the RECOVERY_TIMEOUT event.

Below is a more detailed description of how the implementation is processing a RECOVERY_TIMEOUT event.

1. When a RECOVERY_TIMEOUT event occurs, the next packet received is processed by the sequence recovery function under normal operation using the current “RecovSeqNum” and “SequenceHistory” values.

2. Once completing the processing of the first packet received after the RECOVERY_TIMEOUT event, the SequenceRecoveryReset function is executed, which sets the “RecovSeqNum” to “RecovSeqSpace - 1”, clears the“SequenceHistory” array and sets “TakeAny” to true. The next packet received (i.e. second packet received after the RECOVERY_TIMEOUT event has occurred) is processed using the sequence recovery function reset variables.

Workaround

If given conditions of the erratum are met (a rare occasion), a workaround should be implemented at the system level e.g., retransmitting lost packets if Rx end misses multiple frames in a sequence, or, tolerate and account for one additional packet packet drop following a Sequence Recovery function reset (multiple frame drop).

ERR051707: NETC: MAC Merge stat inaccurate after lost Rx fragments

Description

When frame preemption is enabled, MAC Merge stat counter MAC_MERGE_MMFAECR may record an inaccurate count of Rx frame assembly errors in some scenarios where continuation fragments are lost and/or received out of order. The MAC_MERGE_MMFAECR stat counter will be lower than expected when the issue occurs.

Workaround

The host that is reading MAC_MERGE_MMFAECR register should check status of PM1_RFCS. If the PM1_RFCS indicates no error then MAC_MERGE_MMFAECR is valid and can be used if on other hand there is an error reported in PM1_RFCS register then MAC_MERGE_MMFAECR might be incorrect and should be treated accordingly

ERR051622: NETC: TABLE_BIR of NETC TIMER is wrong

Description

The NETC PCI MSI-X Table Offset/BIR register (PCI_CFC_MSIX_TABLE_OFF_BIR) and PCI MSI-X PBA Offset/BIR register (PCI_CFC_MSIX_PBA_OFF_BIR) include the field BIR which indicates which one of a Function’s Base Address

registers, located beginning at 10h in Configuration Space, or entry in the Enhanced Allocation capability with a matching BEI, is used to map the Function’s MSI-X Table/PBA into Memory Space.

NETC supports the PCIe Enhanced Allocation (EA) capability which replaces the BARs located in the Configuration Space, starting at 10h. The EA Capability uses fixed base addresses and includes the entries described below. Note that the 4-bit BEI value is supposed to be used to match the BIR values to find an entry.

- EA Entry 0 (BEI 0): Physical Function (PF)

- EA Entry 1 (BEI 2): Physical Function (PF) MSI-X Table/PBA

- EA Entry 2 (BEI 9): Physical Function's first Virtual Function (VF) (ENETC only if applicable)

- EA Entry 3 (BEI 11): Physical Function's first Virtual Function (VF) MSI-C Table/PBA (ENETC only if applicable)

This bug affects all NETC physical/virtual functions specifying an MSI-X table/PBA BIR value of 1 or 3, which are not BEI values found in the EA capability entries, only BIR values of 0 and 2 are valid (matches PF BEI 0/2 and VF BEI 9/11). Any standard PCIe driver that tries to find the correct EA capability entry for the MSI-X table/PBA, will not be able to do so by matching the MSI-X BIR value with an EA entry's BEI value.

Workaround

Software can use the MSI-X Table/PBA BIR value as the index to find the Enhanced Allocation entry. For example, if MSI-X TABLE_BIR=1, EA per-entry registers PCI_CFC_EA_PE{BIR}_* would be used.

ERR051710: NETC: MAC Tx statistic counters inaccurate after half duplex late collision and excessive collision events

Description

The MAC Tx stat counters are incorrect after half duplex late or excessive collision events. When half duplex late or excessive collision events occur, the associated packets will be counted as both good and bad packets causing the following stat counter discrepancies:

1. PMa_TOCTn is incremented for Tx packets with late or excessive collision events

2. PMa_TFRMn is incremented for Tx packets with late or excessive collision events

The specification is that these two counters should only increment for transmission of good packets. These counters incorrectly increment for packets transmitted with half duplex late or excessive collisions.

Workaround

After one or more late collision or excessive collision events, counters PMa_TOCTn and PMa_TFRMn will be higher than expected. The accurate value cannot be recovered for PMa_TOCTn, but PMa_TRFMn can be recovered as follows:

PMa_TFRMn' = PMa_TFRMn - PMa_TLCOLn - PMa_TECOLn

ERR051711: NETC: MAC Stats Octet Counter Mismatch after Pause Frame Transmitted with Flexible Preamble Enabled

Description

MAC statistic counters TEOCT and TOCT are inaccurate after Pause frames are transmitted with flexible preamble enabled (PM0_TX_IPG_PREAMBLE[FLEX_PREAMBLE_EN] = 1) and flexible preamble count (PM0_TX_IPG_PREAMBLE[FLEX_PREAMBLE_CNT]) set to less than 7.

When this condition occurs, the TEOCT and TOCT counters will be higher than expected by an amount of (7-FLEX_PREAMBLE_CNT) for each pause frame transmitted. Additionally, stat PM0_T127n is incorrectly incremented rather than PM0_T64n when the pause frame is transmitted.

Workaround

The corrected TEOCT and TOCT counts can be obtained by multiplying the number of number of pause frames transmitted by the excess octets recorded with each pause frame and subtracting from the actual TEOCT and TOCT counts as follows:

1. Corrected TEOCT = TEOCT - TXPF * (7 - PM0_TX_IPG_PREAMBLE[FLEX_PREAMBLE_CNT])

2. Corrected TOCT = TOCT - TXPF * (7 - PM0_TX_IPG_PREAMBLE[FLEX_PREAMBLE_CNT])

ERR051587: NETC: Time gate scheduling update command response can be erroneous if the AdminBaseTime specified is near the current time

Description

Under the following circumstances, the Time Gate Scheduling Table may not execute configuration commands with the update action correctly:

(1) If an operational gate control list is installed and active (I.E. PTGAGLSR[TG] corresponding to the port is 1), and the ADMIN_BASE_TIME specified in the command is in the range [current time, current time + advance time + the command processing duration + the minimum duration before admin gate list installation] then the command may unexpectedly fail and return the error code “0x0D6” instead of succeeding.

(2) If an operational gate control list is installed and active (I.E. PTGAGLSR[TG] corresponding to the port is 1), and the ADMIN_BASE_TIME specified in the command is in the range (current time – 2*2^30 nanoseconds, current time), then the command may unexpectedly succeeds instead of failing and returning the error code “0x0D6”

Where “advance time” is:

The duration that the time reference used by HW scheduler is moved forward to adjust for latency encountered in the transmit processing pipeline (~ 10 usec for ENETC, 0. 1 usec for the switch)

Where “current time” is:

The timestamp of the 1588 (synchronized) timer at the moment the driver (software) begins to set up and dispatch the command.

Where “the command processing duration" is:

The sum of the duration it takes software to set up and dispatch the command, and the duration it takes hardware to process the command.

Where "the duration it takes software to set up and dispatch the command" is:

The sum of the durations it takes to setup the CBD command, dispatch the CBD command, and complete the CBD command by updating the BDR producer index.

Where "the duration it takes hardware to process the command" is:

(300 + (7 * ADMIN_CONTROL_LIST_LENGTH) + (5 * (ADMIN_CONTROL_LIST_LENGTH)^2)) * the NETC platform clock period

Where “the minimum duration before admin gate installation” is:

The sum of the cycle duration of the active operational gate list and the duration remaining in the current gate list cycle.

Workaround

If an operational gate control list is installed and active (I.E. PTGAGLSR[TG] corresponding to the port is 1), and the ADMIN_BASE_TIME that would be specified in an update command is within the range (current time – 2*2^30 nanoseconds, current time + advance time + the command processing duration + the minimum duration before admin gate list installation), then software must increase the ADMIN_BASE_TIME that would be specified in increments of ADMIN_CYCLE_TIME until it is no longer within that range.

Functionally this workaround has minimum impact on the config change time computation performed by hardware (time at which the admin gate control list is installed to become operational). With the workaround the “minimum time before admin gate control list can be installed as the operation list” is increased

From: time left to end of the current operation cycle + one entire operation cycle time

To: 2 × operational cycle times

ERR011377: FlexSPI: DLL lock status bit not accurate due to timing issue

Description

After configuring DLL and the lock status bit is set, still may get wrong data if immediately read/write from FLEXSPI based external flash due to timing issue

Workaround

Adding a delay time (equal or more than 512 FlexSPI root clock cycle) after the DLL lock status is set.

ERR051383: QDC: Quadrature decoder CTRL[DMAEN] bit can not be cleared

Description

1) Register CTRL[DMAEN] bit can not be cleared by SW.

2) DMA request can not be cleared by manually.

Workaround

Use QDC reset to clear CTRL[DMAEN] bit.

ERR052145: FlexSPI_follower: Block behaviors may mismatch the definitions in register after software reset

Description

When setting the block operations (Block Read, Block Write, Block Next Read and Block Next Write) in MODULE_CONTROL register and implementing the software reset, the block behavior will mismatch the bit definition in the register after the reset.

Workaround

Write 0 to all block operation bits in MODULE_CONTROL register before software reset, then reconfigure them by use case after the reset.

ERR052086: I3C: RX FIFO may require separate read for the last data in DMA operation

Description

When there are more than one data received, the last data (last 1 byte in byte mode, or last 2 bytes in half word mode) received in RX FIFO can not trigger DMA request even DMA is enabled. The last data will be left in the RX FIFO. It impacts both target and controller modes, both I2C and I3C modes.

Workaround

When proceeding the received message, SW need check RX_PEND and do separate read to obtain the last data from RX FIFO if RX_PEND is high. For target mode it need read SRDATAB or SRDATAH depending on

DMAWIDTH configured. For controller mode it need read MRDATAB or MRDATAH depending on DMAWIDTH configured.

ERR052343: I3C: Target Early Termination Feature not available with DMA controller

Description

The indication of Early read termination by I3C Target is reflected by register status and interrupt in I3C Controller. Due to custom interface of DMA controller

this pre-mature termination information is not propagated to it consequently the partial data is not transfer into memory in case of Early termination by Target.

So, Early termination cannot be used with this DMA controller for I3C.

Workaround

1) If DMA is used, do not use Target Early termination feature. Instead use smaller data size which is guaranteed to be arranged by Target always and early termination will never occur.

2) Do not use DMA , use Host to do all transfers and handle Early termination by ISR or by polling of status method.

ERR052344: I3C: Controller Clock stalling feature not available in I3C Controller

Description

The clock stalling feature as per section “5.1.2.5 : Controller Clock Stalling” in MIPI I3C Basic Specification(Improved Inter Integrated Circuit) Specification Version 1.1.1 is not implemented for I3C controller.

For target assuming the stalling of clock between C8 and C9 in I3C SDR write followed with repeated start such as read with register address use case, need to set the appropriate SCL frequency based on the Slave device performance.

Workaround

For target assuming the stalling of clock between C8 and C9 in I3C SDR write followed with repeated start such as read with register address use case, need to set the appropriate SCL frequency based on the Slave device performance.

ERR051250: SYS_CTR: SYS_CTR_COMPARE register write could be missed

Description

In some rare case, SYS_CTR module's SYS_CTR_COMPARE registers might not be updated to the value that was supposed to be written into it.

Workaround

Do read back check after writing to registers of SYS_CTR_COMPARE. If the read back check fails, redo the write until it successes.

ERR052358: SYS_CTR: CNTCV0 and CNTCV1 register cannot get count value when clocked by 32kHz alternate slow clock

Description

SYS_CTR has two selectable clock sources for its counter, base_clk and slow_clk, base_clk is from 24Mhz OSC clock and slow_clk is from 32KHz OSC clock. CNTCV0 and CNTCV1 register cannot get count value when clocked by 32kHz alternate slow clock.

Workaround

Select base_clk for the SYS_CTR counters.

ERR003777: GPT: Possibility of additional pulse on src_clk when switching between clock sources

Description

There is a possibility of an extra pulse on SCLK in the GPT when switching between the clock sources.

Workaround

Changing the clock source should only be done when the GPT is disabled. A way to accomplished this is as follows:

Disable GPT—Write 1'b0 to EN bit of GPTCR

Disable interrupts—Write 6'b000000 in Bits [5:0] of GPTIR

Configure Output Mode to unconnected/ disconnected—Write zeros in OM3, OM2, OM1 in GPTCR

Disable Input Capture Modes—Write zeros in IM1,IM2 in GPTCR

Change clock source CLKSRC in GPTCR

Clear Status register—Write 003F in GPTSR

Set ENMOD in GPTCR

ENABLE GPT—Write 1'b1 to EN bit of GPTCR. The GPTSR should not be read immediately after changing the clock source (a wait of at least one SCLK is required).

ERR051588: LPSPI:Reset transmit FIFO after FIFO underrun by LPSPI Slave.

Description

Transmit FIFO pointers are corrupted when a transmit FIFO underrun occurs (SR[TEF]) in slave mode.

Workaround

When clearing the transmit error flag (SR[TEF] = 0b1) following a transmit FIFO underrun, reset the transmit FIFO (CR[RTF] = 0b1) before writing any new data to the transmit FIFO.

ERR052081: VREF: VREF trim value need be loaded with SW before enabling VREF

Description

VREF trim values are stored in fuse and not loaded to VREF successfully after system reset. SW has to load it after each reset before VREF is enabled for use.

Workaround

SW can follow below sequence to load the trim value to VREF after reset.

1. Unlock VREF trim registers by writing 0x5AA5 into bits 15:1 of TEST_UNLOCK register (VREF address offset 0x1C).

2. Call ELE API to read fuse word 39, WORD39 below refers to the 32bit fuse data.

3. Write WORD39[9:4] into bits 13:8 of TRIM0 register (VREF address offset 0x24).

4. Write WORD39[10] into bits 31 TRIM0 register

5. Call ELE API to read fuse word 38, WORD38 below refers to the 32bit fuse data.

6. Write WORD38[31:28] into bits 27:24 of TRIM0 register

7. Write WORD38[10:7] into bits 19:16 of TRIM0 register

8. Write WORD38[6:4] into bits 15:13 of TRIM0 register

9. Write WORD38[3:0] into bits 18:11 of TRIM0 register

10. Write WORD38[16:14] into bits 7:5 of TRIM0 register

11. Write WORD38[13:11] into bits 2:0 of TRIM0 register

ERR051405: SAI: Synchronous mode with BYP=1 not supported

Description

If the transmitter or receiver is configured for synchronous mode ((SYNC=01) or Bit Clock Swap (BCS=1), and the receiver or transmitter that is the source of the BCLK is configured with BYP=1 and BCD=1, then the transmitter or receiver must set BCI=1.

Workaround

Set BCI=1 when configuring for synchronous mode (SYNC=01) or Bit Clock Swap (BCS=1) and the source of the BCLK has configured BCD=1 and BYP=1.

ERR051374: PWM fault may work abnormally when the fault signal is very narrow

Description

If the fault signal pulse width is narrower than a certain threshold, the protected PWM channels may generate a glitch, which occurs after the PWM channel outputs become inactive.

Workaround

(1) When FCTRL2[NOCOMB] = 0, FFILT [GSTR]= 0, and FFILT[FILT_PER]=0, pulse width of fault signals must be larger than 6 PWM clock periods, otherwise a glitch may be generated on the protected PWM channels.

(2) When FCTRL2[NOCOMB] = 0, FFILT [GSTR]= 1, and FFILT[FILT_PER]=0, pulse width of fault signals must be larger than 3 PWM clock periods, otherwise a glitch may be generated on the protected PWM channels.

(3) When FCTRL2[NOCOMB] = 0, FFILT [GSTR]= 1, and FFILT[FILT_PER] has non-zero values, pulse width of fault signals must be larger than FILT_PER*(FILT_CNT+3)+6 PWM clock periods, otherwise a glitch may be generated on the protected PWM channels.

(4) When FCTRL2[NOCOMB] = 0, FFILT [GSTR]= 0, and FFILT[FILT_PER] has non-zero values, pulse width of fault signals must be larger than FILT_PER*(FILT_CNT+3)+9 PWM clock periods, otherwise a glitch may be generated on the protected PWM channels.

ERR051989: PWM: output may be abnormal when the value of phase delay register is reduced from a non-zero value to 0.

Description

When the value of the SMxPHASEDLY register is reduced from a non-zero value to 0 and the SMxCTRL2[RELOAD_SEL]=1, the submodule x may output an unexpected wide PWM pulse (x=1,2,3).

Workaround

The minimum value of the SMxPHASEDLY register should be set as 1 in this process. To realize no phase delay between the submodule 0 and submodule x in this process, set the SMxPHASEDLY=1, SMxINIT=SM0INIT–1, SMxVALy=SM0VALy–1 (x=1,2,3, y=0,1,2,3,4,5).

ERR051629: LPUART:Transmit Complete bit (STAT[TC]) is not set.

Description

When the CTS pin is negated and the CTS feature is enabled (MODIR[TXCTSE] = 0b1) and the TX FIFO is flushed by software then, the Transmit Complete (STAT[TC]) flag is not set.

Workaround

Clear (MODIR[TXCTSE]) bit and reset the transmit FIFO (FIFO[TXFLUSH] = 0b1) when flushing the FIFO with CTS enabled(MODIR[TXCTSE] = 0b1).

ERR051385: ADC: ADC INLDNL degrade under high ADC clock frequency when VREFH selected as reference

Description

The ADC can only support 24MHz operation for VREFH as reference to meet DNL spec.

Workaround

Select VDDA_ADC_1P8 as the ADC reference instead of VREFH

ERR052198: uSDHC: eMMC CQE may timeout due to a HW logic issue

Description

To facilitate command queuing in eMMC, uSDHC supports CQE (Command Queue) so that the host can queue data transfer tasks. In some cases where the CQDPT (Device Pending Task) bit is being set and cleared at the same time, the CQDPT bit is not cleared even after the task has completed execution, resulting in a CQE timeout.

Workaround

There are two possible workarounds:

1.Disable CQE feature using CQCFG register

2. Before sending a DCMD (Direct Command) request, SW must manually clear the CQDPT[bit n] by setting CQTCLR[bit n] if DPT[bit n] != CQTDBR[bit n], in order to clear the pending task.

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.

Suitability for use — NXP Semiconductors products are not designed, authorized or warranted to be suitable for use in life support, life-critical or safety-critical systems or equipment, nor in applications where failure or malfunction of an NXP Semiconductors product can reasonably be expected to result in personal injury, death or severe property or environmental damage. NXP Semiconductors and its suppliers accept no liability for inclusion and/or use of NXP Semiconductors products in such equipment or applications and therefore such inclusion and/or use is at the customer's own risk.

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.

Limiting values — Stress above one or more limiting values (as defined in the Absolute Maximum Ratings System of IEC 60134) will cause permanent damage to the device. Limiting values are stress ratings only and (proper) operation of the device at these or any other conditions above those given in the Recommended operating conditions section (if present) or the Characteristics sections of this document is not warranted. Constant or repeated exposure to limiting values will permanently and irreversibly affect the quality and reliability of the device.

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.

No offer to sell or license — Nothing in this document may be interpreted or construed as an offer to sell products that is open for acceptance or the grant, conveyance or implication of any license under any copyrights, patents or other industrial or intellectual property rights.

Quick reference data — The Quick reference data is an extract of the product data given in the Limiting values and Characteristics sections of this document, and as such is not complete, exhaustive or legally binding.

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.

Suitability for use in non-automotive qualified products — Unless this document expressly states that this specific NXP Semiconductors product is automotive qualified, the product is not suitable for automotive use. It is neither qualified nor tested in accordance with automotive testing or application requirements. NXP Semiconductors accepts no liability for inclusion and/or use of non-automotive qualified products in automotive equipment or applications.

In the event that customer uses the product for design-in and use in automotive applications to automotive specifications and standards, customer (a) shall use the product without NXP Semiconductors' warranty of the product for such automotive applications, use and specifications, and (b) whenever customer uses the product for automotive applications beyond NXP Semiconductors' specifications such use shall be solely at customer's own risk, and (c) customer fully indemnifies NXP Semiconductors for any liability, damages or failed product claims resulting from customer design and use of the product for automotive applications beyond NXP Semiconductors' standard warranty and NXP Semiconductors' product specifications.

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.

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.

Contents