# AN13684

## Avoiding WRIOP FIFO Depletion for 100 Gbps Bandwidth Configuration on LX2 Platforms

Rev. 0 — 5 July 2022 Application note

**Document information**

| Information | Content |
|---|---|
| Keywords | WRIOP port, WRIOP FIFO depletion, Avoid WRIOP FIFO depletion, 100 Gbps bandwidth configuration, WRIOP 100 Gbps, RECYCLE ports speed change impact, WRIOP FIFO size distribution, WRIOP total FIFO size |
| Abstract | This document describes the WRIOP port FIFO limitation that may occur on LX2160 when multiple physical ports are configured and their total bandwidth exceeds 100 Gbps. The document also presents a solution to avoid this situation. |

# 1 Introduction

This document describes the WRIOP port FIFO limitation that may occur on LX2160 when multiple physical ports are configured and their total bandwidth exceeds 100 Gbps. The document also presents a solution to avoid this situation.

# 2 Context description

The Management Complex (MC) firmware configures by default two RECYCLE ports on LX2 (DPAA2) platform. These ports are not accessible and are used for virtual connections (DPNI to DPNI, DPNI to DPSW, DPNI to DPDMUX) and frame replication. RECYCLE port is a WRIOP port that is not connected to a physical MAC. It takes frames from QMAN as if they were coming on ingress path. For each RECYCLE port, MC allocates FIFO memory for INGRESS and EGRESS directions.

Each RECYCLE port is rated as a 50 Gbps physical port which means that the reserved FIFO is identical with the FIFO reserved for a 50 Gbps physical port.

The examples given below will focus on EGRESS direction.

On EGRESS side, two RECYCLE ports will consume:

*2 * (0x7F + 1) * 0x100Bytes = 0x10000 Bytes*

If a SERDES protocol with multiple physical interfaces that exceed 100 Gbps throughput is configured by the user, there is a high probability that there will not be sufficient FIFO memory for all the physical ports.

Two protocols selections that can trigger this situation are:

- **8-7-3**

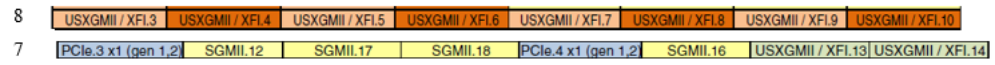


- **18-6-2**

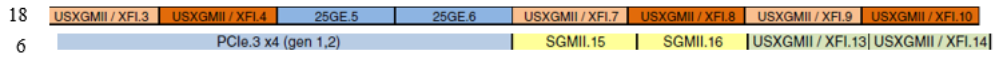


***Note:*** *SERDES 3 is not listed, as only SERDES 1 and 2 are used for networking.*

If SERDES protocol **18-6-2** is chosen, there will not be enough FIFO memory for all the 14 physical ports from the configuration:

- 8 × 10 Gbps ports
- 2 × 25 Gbps ports
- 4 × 1 Gbps port

The FIFO consumption on egress side equals to:

*(0x3F + 1)* 8 * 0x100 + (0x3F + 1)* 2 * 0x100 + (0x2F + 1)* 4 * 0x100 = 0x34000 Byte*

This value is added to the RECYCLE ports FIFO on EGRESS and the final result equals to 0x44000 Bytes.

The output exceeds the total available FIFO on EGRESS direction which is 0x42000 Bytes (270336 decimal) for LX2. For the value provided, see LX2 DPAA2 Reference Manual.

## 3 Manifestation of the limitation

The outcome of the FIFO consumption is that the last remaining ports (starting from the 13th in the current example) that will be configured, will not function properly: TX packets are discarded due to physical errors and RX side will not work (it is very likely that FIFO will be consumed on this direction as well).

Other symptoms of the issue are the messages that will appear in the MC console:

```
cat /dev/dpaa2_mc_console

[W]   [PFS    ] val=0x0000003f got=0x00000001
[W]   [PFS    ] val=0x00000037 got=0x00000001
[W]   [PFS    ] val=0x0000003f got=0x00000001
```

The keyword that points to the limitation in the above snippet is Port FIFO Size (PFS).

According to the Reference Manual:

- The total sum of ports' FIFOs must not exceed the total FIFO size available. The FIFO size is device specific. Check **Device specific features**.

A write to the FIFO size register that would exceed the total FIFO size Is ignored. This explains the prints from the MC console.

## 4 How to avoid the overflow

User must ensure that a firmware version greater or equal to 10.33.0 is loaded on the LX2 platform. The MC version can be checked either in the UBOOT console when MC starts or in Linux:

```
restool -m
```

Besides the MC firmware, the Data Path Configuration (DPC) file, which is loaded by the MC at startup, should contain explicitly at least one of the RECYCLE ports:

```
[…]

board_info {
        /* Insert @recycle_ports node together with the 2
 recycle ports.
        In this case both are needed to reduce the FIFO such
 that 14 ports
        will be available
        */
                        recycle_ports {
                                recycle@2 {
                                        max_rate = "10G";
                                };
                                recycle@1 {
                                        max_rate = "10G";
                                };
                        };
                        ports {
                              /*the usual MACs are here*/
                                                    [..]
                }
```

```
}
```

The `max_rate` parameter determines the FIFO consumption for each RECYCLE port. In the current example, each RECYCE will be rated as a 10 Gbps physical port and the FIFO consumed by both of them will be equal to:

*(0x3F + 1) * 2 * 0x100 = 0x8000 Bytes*

If the above value is added to the FIFO consumed by the 14 ports (0x34000 Bytes), the total FIFO used will be equal to 0x3C000 Bytes which is less than the FIFO available on egress direction 0x42000 Bytes (270336 decimal).

In this way, the 14 physical ports selected by SERDES 18-6-2 protocol combination, will work properly.

## 5 FIFO distribution based on physical ports speeds

Table 1 lists the values written in FIFO size register for each WRIOP physical port depending on its speed selected by the SERDES protocol.

**Table 1. WRIOP ports FIFO size distribution**

| Ingress side | | Egress side | |
|---|---|---|---|
| **Port rate** | **FIFO size** | **Port rate** | **FIFO size** |
| 1 G | 0x2F | 1 G | 0x2F |
| 2.5 G | 0x2F | 2.5 G | 0x2F |
| 5 G | 0x33 | 5 G | 0x3F |
| 10 G | 0x37 | 10 G | 0x3F |
| 20 G | 0x3F | 20 G | 0x3F |
| 25 G | 0x3F | 25 G | 0x3F |
| 40 G | 0x7F | 40 G | 0x7F |
| 50 G | 0x7F | 50 G | 0x7F |
| 100 G | 0xFF | 100 G | 0xFF |

## 6 RECYCLE ports speed change impact

RECYCLE ports are WRIOP internal ports not directly accessible that are used in several situations:

- DPNI to DPNI connection
- DPNI to DPSW connection
- DPNI to DPDMUX connection
- Frame replication

If the setup configured by the user does not contain only physical ports, changing the speed of RECYCLE ports will impact all the above use cases. The impact will be observed at performance level. For example, if the RECYLE port speed is set at 10 G, then all the connections that go through RECYCLE will be limited to maximum 10 Gbps.

# 7 Revision history

| Rev. | Date | Description |
|---|---|---|
| 0 | 8 July 2022 | Initial release |

# 8 Legal information

## 8.1 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.

## 8.2 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.

**Terms and conditions of commercial sale** — NXP Semiconductors products are sold subject to the general terms and conditions of commercial sale, as published at http://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.

**Suitability for use in non-automotive qualified products** — Unless this data sheet 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.

**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.

## 8.3 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.

# Avoiding WRIOP FIFO Depletion for 100 Gbps Bandwidth Configuration on LX2 Platforms

## Contents





For more information, please visit: http://www.nxp.com
For sales office addresses, please send an email to: salesaddresses@nxp.com

Date of release: 5 July 2022
Document identifier: AN13684