

REALTIMEEDGEUG

Real-time Edge Software User Guide

Rev. 3.1 — 26 March 2025

User guide

Document information

Information	Content
Keywords	REALTIMEEDGEUG, Real-time Edge Software, Real-time Networking, Real-time System, Protocols, i.MX boards, QorIQ (Layerscape) boards, i.MX 6ULL EVK, i.MX 8DXL EVK, i.MX 8M Mini LPDDR4 EVK, i.MX 8M Plus LPDDR4 EVK, i.MX 93 EVK, i.MX RT1180 EVK, i.MX 93 9x9 QSB, i.MX 93 14x14 EVK, i.MX 91 9x9 LPDDR4 QSB, i.MX 91 11x11 LPDDR4 EVK, i.MX 95 15x15 LPDDR4 EVK, i.MX 95 19x19 EVK, LX2160ARDB Rev2, NXP hardware platforms
Abstract	This document describes the features and implementation of Real-time Edge Software on NXP hardware platforms. The key technology components include Real-time System, Real-time Networking, Heterogeneous Multicore Framework, Heterogeneous Multi-SoC Framework, and Protocols.



1 Introduction

1.1 Real-time Edge software

Real-time Edge software is an evolved version of Open Industrial Linux (OpenIL) for real-time and deterministic systems in different fields. The key technology components include Real-time System, Heterogeneous Multicore Framework, Heterogeneous Multi-SoC Framework, Real-time Networking, and Protocols.

- Real-time System includes Preempt-RT Linux, Native RTOS on Cortex-A, Jailhouse, U-Boot based BareMetal framework, RTOS and BareMetal on Cortex-M, and different combinations of these systems.
- Heterogeneous Multicore Framework provides a general software framework to support Heterogeneous AMP. It enables AMP to be interconnected and provides a unified resource management and life-cycle management.
- Heterogeneous Multi-SoC Framework enables the usage of a combination of MPU and MCU. It extends the hardware components for both, the MCU and MPU.
- Real-time Networking includes TSN technology, TSN standards, management, configuration, and applications. This component also supports networking and redundancy features.
- Protocols component includes support for industry standard protocols such as EtherCAT, CoE, FlexCan, OPC-UA, and others.

This document describes the features and implementation of Real-time Edge Software on NXP hardware platforms.

1.2 Feature support matrix

[Table 1](#) shows the features that are supported in this release.

Table 1. Key features

Feature				i.MX 6ULL 14x14 EVK	i.MX 8DXL LPDDR4 EVK	i.MX 8M Mini LPDDR4 EVK	i.MX 8M Plus LPDDR4 EVK	i.MX 93 EVK	i.MX 93 9x9 LPDDR4 QSB	i.MX 93 14x14 LPDDR4x EVK	i.MX 95 19x19 LPDDR5 EVK	i.MX 95 15x15 LPDDR4 EVK	i.MX 91 11x11 LPDDR4 EVK	i.MX 91 9x9 LPDDR4 QSB	i.MX RT1180 EVK	LS1028 ARDB	LS1043 ARDB	LS1046 ARDB	LX2160 ARDB		
Boot mode	SD			Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y		Y	Y	Y	Y		
	eMMC					Y	Y	Y	Y	Y	Y	Y	Y	Y		Y		Y			
Real-time System	Preempt-RT Linux			Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y		Y	Y	Y	Y		
	BareMetal	ICC				Y	Y	Y	Y								Y	Y	Y	Y	
		PCIe																Y	Y	Y	
		Ethernet				Y	Y	Y										Y	Y	Y	
		GPIO				Y	Y													Y	
		IPI				Y	Y	Y	Y									Y	Y	Y	
		UART				Y	Y	Y	Y									Y	Y	Y	
		USB																	Y	Y	
		SAI																Y			
		CAN																			
		I2C																Y	Y	Y	
		QSPI																		Y	
		IFC																		Y	
		Flextimer																		Y	
		Linux (communication with BareMetal)		ICC			Y	Y	Y	Y								Y	Y	Y	Y
				IPI			Y	Y	Y	Y									Y	Y	Y
		Single HW Interrupt to multiple cores																		Y	
		Newlib Math library					Y	Y	Y	Y									Y	Y	Y
		All Cortex-A cores running under Baremetal																	Y		
	Native RTOS on Cortex-A	FreeRTOS				Y	Y	Y			Y	Y	Y	Y	Y						
		Zephyr				Y	Y	Y			Y	Y	Y	Y	Y						
	Jailhouse					Y	Y			Y		Y						Y	Y	Y	
	Harpoon RTOS	FreeRTOS					Y	Y	Y			Y									
		Zephyr					Y	Y	Y			Y									
Heterogeneous Multicore Frame-work	Flexible Real-time System	Flexible Real-time System	Free RTOS			Y	Y	Y		Y	Y	Y									
			Zephyr			Y	Y	Y		Y	Y	Y									

Table 1. Key features...continued

Feature				i.MX 6ULL 14x14 EVK	i.MX 8DXL LPDDR4 EVK	i.MX 8M Mini LPDDR4 EVK	i.MX 8M Plus LPDDR4 EVK	i.MX 93 EVK	i.MX 93 9x9 LPDDR4 QSB	i.MX 93 14x14 LPDDR4x EVK	i.MX 95 19x19 LPDDR5 EVK	i.MX 95 15x15 LPDDR4 EVK	i.MX 91 11x11 LPDDR4 EVK	i.MX 91 9x9 LPDDR4 QSB	i.MX RT1180 EVK	LS1028 ARDB	LS1043 ARDB	LS1046 ARDB	LX2160 ARDB
		RAM Console	Free RTOS			Y	Y	Y		Y	Y	Y							
			Zephyr			Y	Y	Y		Y	Y	Y							
		Networking stack on A-Core RTOS	Free RTOS			Y	Y	Y		Y									
	Unified Life Cycle Management	U-Boot booting Native RTOS A-Core Image	Free RTOS			Y	Y	Y		Y	Y	Y	Y	Y					
			Zephyr			Y	Y	Y		Y	Y	Y	Y	Y					
		Linux booting Native RTOS A-Core Image	Free RTOS			Y	Y	Y		Y	Y	Y							
			Zephyr			Y	Y	Y		Y	Y	Y							
		U-Boot booting Native RTOS M-Core Image	Free RTOS			Y	Y	Y	Y	Y									
			Free RTOS			Y	Y	Y	Y	Y									
	RPMSG	RPMSG between Linux and M-Core RTOS	Free RTOS			Y	Y	Y	Y	Y									
			Free RTOS			Y	Y	Y		Y									
		RPMSG between 2 A-Core RTOS	Free RTOS				Y												
			Free RTOS			Y													
		RPMSG Performance Evaluation	Free RTOS				Y												
			Free RTOS			Y		Y	Y	Y									
	Heterogeneous Multicore VirtIO	Heterogeneous Multicore VirtIO Performance Evaluation	Free RTOS			Y	Y												
			Free RTOS			Y	Y	Y		Y									
		Heterogeneous Multicore VirtIO Network Sharing	Free RTOS			Y	Y	Y		Y									
	Industrial Applications	SOEM running on A-core or M-core	Free RTOS			Y	Y	Y											
Heterogeneous Multi-SoC Framework	NETC TSN Switch	DSA single port mode					Y	Y							Y				
		DSA bridge mode					Y	Y							Y				

Table 1. Key features...continued

Feature			i.MX 6ULL 14x14 EVK	i.MX 8DXL LPDDR4 EVK	i.MX 8M Mini LPDDR4 EVK	i.MX 8M Plus LPDDR4 EVK	i.MX 93 EVK	i.MX 93 9x9 LPDDR4 QSB	i.MX 93 14x14 LPDDR4x EVK	i.MX 95 19x19 LPDDR5 EVK	i.MX 95 15x15 LPDDR4 EVK	i.MX 91 11x11 LPDDR4 EVK	i.MX 91 9x9 LPDDR4 QSB	i.MX RT1180 EVK	LS1028 ARDB	LS1043 ARDB	LS1046 ARDB	LX2160 ARDB
		MTU configuration				Y	Y							Y				
		VLAN configuration				Y	Y							Y				
		FDB configuration				Y	Y							Y				
		Port statistics				Y	Y							Y				
		PTP stack				Y	Y							Y				
		802.1CB				Y	Y							Y				
		Qci				Y	Y							Y				
		Qbv				Y	Y							Y				
		Qbu				Y	Y							Y				
	Virtual Switch					Y	Y							Y				
Real-time Networking	TSN Standards	Qbv		Y		Y	Y	Y	Y						Y			
		Qbu		Y		Y	Y	Y	Y						Y			
		Qci													Y			
		Qav		Y	Y	Y	Y	Y	Y						Y			
		802.1AS		Y	Y	Y	Y	Y	Y						Y	Y	Y	Y
		802.1CB													Y			
		VCAP chain mode													Y			
		802.1 Q-in-Q													Y			
	TSN Configurations	Linux tc command		Y		Y	Y	Y	Y						Y			
		TSN tool													Y			
		NETCONF/YANG	Qbv	Y		Y	Y	Y	Y						Y			
			Qbu	Y		Y	Y	Y	Y						Y			
			Qci												Y			
			IP	Y		Y	Y	Y	Y						Y			
			MAC	Y		Y	Y	Y	Y						Y			
			VLAN config	Y		Y	Y	Y	Y						Y			
		Web-based configuration	Qbv	Y		Y	Y	Y							Y			
			Qbu	Y		Y	Y	Y							Y			
			Qci												Y			
		Dynamic topology discovery		Y		Y	Y	Y							Y			
		Dynamic TSN configuration	Qci												Y			
			CQF												Y			
			Qbv	Y		Y	Y	Y							Y			

Table 1. Key features...continued

Feature			i.MX 6ULL 14x14 EVK	i.MX 8DXL LPDDR4 EVK	i.MX 8M Mini LPDDR4 EVK	i.MX 8M Plus LPDDR4 EVK	i.MX 93 EVK	i.MX 93 9x9 LPDDR4 QSB	i.MX 93 14x14 LPDDR4x EVK	i.MX 95 19x19 LPDDR5 EVK	i.MX 95 15x15 LPDDR4 EVK	i.MX 91 11x11 LPDDR4 EVK	i.MX 91 9x9 LPDDR4 QSB	i.MX RT1180 EVK	LS1028 ARDB	LS1043 ARDB	LS1046 ARDB	LX2160 ARDB
	GenAVB/TSN Stack	gPTP 802.1AS-2020 with multi-domain support		Y	Y	Y	Y	Y	Y	Y					Y	Y	Y	Y
		SRP IEEE 802.1Q-2022	Y	Y	Y	Y	Y	Y	Y	Y					Y	Y	Y	Y
		AVB Standards	AVTP IEEE 1722-2016	Y	Y	Y	Y	Y	Y									
			ATDECC 1722.1-2021	Y	Y	Y	Y	Y	Y									
			Milan v1.2	Y	Y	Y	Y	Y	Y									
			MAAP IEEE 1722-2016	Y	Y	Y	Y	Y	Y									
		AVB Endpoint	Y	Y	Y	Y	Y		Y									
		AVB Bridge		Y ^[1]					Y ^[1]						Y			
		AVB Hybrid							Y ^[1]									
		TSN Endpoint		Y		Y	Y	Y	Y	Y								
		TSN Bridge													Y			
		Media clock recovery	Y	Y	Y ^[2]	Y	Y ^[2]		Y ^[2]									
Industrial Protocol	EtherCAT Main Device	IGH EtherCAT MainDevice stack	Y	Y	Y	Y	Y								Y	Y	Y	Y
		IGH native Ethernet device driver		Y	Y	Y	Y								Y	Y	Y	
		SOEM			Y	Y	Y											
		CodeSYS EtherCAT MainDevice stack	Y		Y	Y	Y											
	FlexCAN														Y			
	CANopen																	
	OPC UA	open62541	Y	Y	Y	Y	Y	Y	Y						Y	Y	Y	Y
		OPC UA PubSub over TSN		Y		Y	Y	Y	Y						Y			
	BEE (Mikroe Click board)														Y			
	Bluetooth Low Energy (Mikroe Click board)														Y			
	NFC (Mikroe Click board)														Y			
	Modbus	Modbus-RTU	Y	Y	Y	Y	Y	Y	Y									
		Modbus-TCP	Y	Y	Y	Y	Y	Y	Y						Y	Y	Y	Y

[1] Using a SJA1105Q-EVB

[2] Media clock recovery is implemented through a software-based sampling.

1.3 Supported NXP platforms

The [Table 2](#) lists the NXP hardware SoCs and boards that support the Real-time Edge software.

Table 2. Supported NXP platforms

Platform	Architecture	Boot
i.MX 6ULL EVK	Arm v7	SD
i.MX 8DXL LPDDR4 EVK	Arm v8	SD, eMMC
i.MX 8M Mini LPDDR4 EVK	Arm v8	SD, eMMC
i.MX 8M Plus LPDDR4 EVK	Arm v8	SD, eMMC
i.MX 93 EVK	Arm v8	SD, eMMC
i.MX 93 9x9 QSB	Arm v8	SD, eMMC
i.MX 93 14x14 EVK	Arm v8	SD, eMMC
i.MX 95 19x19 LPDDR5 EVK	Arm v8	SD, eMMC
i.MX 95 15x15 LPDDR4X EVK	Arm v8	SD, eMMC
i.MX 91 11x11 LPDDR4X EVK	Arm v8	SD, eMMC
i.MX 91 9x9 LPDDR4 QSB	Arm v8	SD, eMMC
LS1028ARDB	Arm v8	SD, eMMC
LS1043ARDB	Arm v8	SD
LS1046ARDB	Arm v8	SD, eMMC
LX2160ARDB Rev 2	Arm v8	SD

1.3.1 Switch settings

The [Table 3](#) lists and describes the switch configuration for the platforms supported by Real-time Edge software.

Table 3. Switch setting for various NXP platforms

Platform	Boot source	Switch setting
i.MX 6ULL EVK	Internal Boot / MicroSD	SW602 = 0b'10 (internal boot) and SW601[1:4] = 0b'0010 (MicroSD)
i.MX 8DXL LPDDR4 EVK	SD, eMMC	<ul style="list-style-type: none"> SD: SW1[1:4] = 0b'1100 eMMC: SW1[1:4] = 0b'0100
i.MX 8M Mini LPDDR4 EVK	MicroSD / uSDHC2	<ul style="list-style-type: none"> SW1101[1:10] = 0b' 0110110010 SW1102[1:10] = 0b' 0001101000
i.MX 8M Plus LPDDR4 EVK	MicroSD / SDHC2	SW4[1:4] = 0b'0011
i.MX 93 EVK	MicroSD / uSDHC2	SW1301[1:4] = 0b'0100
i.MX 93 9x9 QSB	MicroSD / uSDHC2	SW601[1:4] = 0b'0011
i.MX 93 14x14 EVK	SD, eMMC	<ul style="list-style-type: none"> SD: SW5[1:4] = 0b'0100 eMMC: SW5[1:4] = 0b'0000
i.MX 95 19x19 LPDDR5 EVK	SD, eMMC	<ul style="list-style-type: none"> SD: SW7[1:4] = 0b'1011 eMMC: SW7[1:4] = 0b'1010
i.MX 95 15x15 LPDDR4X EVK	SD, eMMC	<ul style="list-style-type: none"> SD: SW7[1:4] = 0b1011 eMMC: SW7[1:4] = 0b1010
i.MX 91 11x11 LPDDR4X EVK	SD, eMMC	<ul style="list-style-type: none"> SD: SW1301[1:4] = 0b0100 eMMC: SW1301[1:4] = 0b0000
i.MX 91 9x9 LPDDR4 QSB	SD, eMMC	<ul style="list-style-type: none"> SD: SW3[1:4] = 0b0011 eMMC: SW3[1:4] = 0b0010
LS1028ARDB	SD, eMMC	<ul style="list-style-type: none"> SD: SW2[1:8] = 0b'10001000 eMMC: SW2[1:8] = 0b'10011000
LS1043ARDB	SD	SW4[1:8] + SW5[1] = 0b'00100000_0
		UART1 output select <ul style="list-style-type: none"> SW3[3] = 0b'0: RJ45 SW3[3] = 0b'1: CMSIS-DAP (MiniUSB)
LS1046ARDB	SD, eMMC	SW5[1:8] + SW4[1] = 0b'00100000_0
		UART1 output select <ul style="list-style-type: none"> SW4[4] = 0b'0: RJ45 SW4[4] = 0b'1: CMSIS-DAP (MicroUSB)
LX2160ARDB Rev2	SD	SW1[1:8] = 0b'10001000

1.3.2 Flashing pre-built images

Pre-built images for platforms supported by Real-time Edge software can be downloaded from NXP website from the below URL:

<https://www.nxp.com/design/software/development-software/real-time-edge-software:REALTIME-EDGE-SOFTWARE>

Download the image required and extract it by using the commands below: (The code below shows the commands used for i.MX93 EVK as an example)

```
$ unzip Real-time_Edge_v3.1_IMX93EVK.zip
$ cd Real-time_Edge_v3.1_IMX93EVK/real-time-edge/
$ ls
Image-imx93evk.bin                imx93-11x11-evk-igh.dtb
imx93-11x11-evk-avb.dtb           imx93-11x11-evk-inmate.dtb
imx93-11x11-evk-dpdk.dtb          imx93-11x11-evk-multicore-rtos.dtb
imx93-11x11-evk-dsa.dtb           imx93-11x11-evk-root.dtb
imx93-11x11-evk-dsa-enetc.dtb     imx93-11x11-evk-rpmsg.dtb
imx93-11x11-evk.dtb              imx93-11x11-evk-uart-sharing-cm33.dtb
imx93-11x11-evk-ecat.dtb          imx93-11x11-evk-virtio-net-ca55.dtb
imx93-11x11-evk-ecat-userspace.dtb imx93-11x11-evk-virtio-net-cm33.dtb
imx93evk.rootfs.manifest          nxp-image-real-time-edge-
imx93-11x11-evk-harpoon.dtb       nxp-image-real-time-edge-
imx93evk.rootfs.tar.zst           nxp-image-real-time-edge-
imx93-11x11-evk-harpoon-industrial.dtb
imx93evk.rootfs.wic.zst
imx93-11x11-evk-harpoon-virtio-net.dtb
```

Extract the compressed wic image

```
$ zstd -d nxp-image-real-time-edge-imx93evk.rootfs.wic.zst
```

Insert SD card, device node “sdx” (for example: sdc) is created in directory “/dev/” with USB reader, flash file “nxp-image-real-time-edge-imx93evk.rootfs.wic” to SD card:

```
$ sudo dd if=./nxp-image-real-time-edge-imx93evk.rootfs.wic of=/dev/sdc bs=1M
conv=fsync
```

After flashing this image to SD card, insert this SD card into i.MX93 EVK board, connect DBG port and open it. Then, powering on the i.MX93 EVK board displays the message.

```
U-Boot SPL 2024.04+gee8fce7cee5+p0 (Dec 05 2024 - 02:53:42 +0000)
SOC: 0xa1009300
LC: 0x2040010
PMIC: PCA9451A
PMIC: Over Drive Voltage Mode
DDR: 3733MTS
M33 prepare ok
Normal Boot
Trying to boot from BOOTROM
Boot Stage: Primary boot
image offset 0x0, pagesize 0x200, ivt offset 0x0
Load image from 0x4f400 by ROM_API
NOTICE: TRDC init done
NOTICE: BL31: v2.10.0 (release):lf-6.6.36-2.1.0-dirty
NOTICE: BL31: Built : 10:35:32, Aug 29 2024
```

```
U-Boot 2024.04+gee8fce7cee5+p0 (Dec 05 2024 - 02:53:42 +0000)

Reset Status: POR

CPU:      NXP i.MX93(52) Rev1.1 A55 at 1700 MHz
CPU:      Industrial temperature grade (-40C to 105C) at 36C

Model: NXP i.MX93 11X11 EVK board
DRAM: 2 GiB
TCPC: Vendor ID [0x1fc9], Product ID [0x5110], Addr [I2C2 0x52]
TCPC: Vendor ID [0x1fc9], Product ID [0x5110], Addr [I2C2 0x51]
TCPC: Vendor ID [0x1fc9], Product ID [0x5110], Addr [I2C2 0x50]
Core: 235 devices, 37 uclasses, devicetree: separate
MMC: FSL_SDHC: 0, FSL_SDHC: 1
Loading Environment from MMC... OK
.....
```

1.4 Real-time Edge Software Yocto Project

For using the Yocto build environment, refer to the *Real-time Edge Yocto Project User Guide*. This document describes the steps to build Real-time Edge images using a Yocto Project build environment for both i.MX and QorIQ (Layerscape) boards.

1.5 Related documentation

All documentation related to Real Time Edge software is available on the link: [REALTIME EDGE Documentation](#). The following documents are available:

- **Real-time Edge Release Notes (RN00161)** (provides release information)
- **Real-time Edge User Guide (REALTIMEEDGEUG)** (provides detailed user guide)
- **Real-time Edge Yocto Project User Guide (RTEDGEYOCTOUG)** (provides information for using the Yocto build environment)
- **GenAVB/TSN Stack Evaluation User Guide (GENAVBTSNUG)** (provides information on how to set up Audio Video Bridging evaluation experiments of the GenAVB/TSN Stack on NXP platforms)
- **Harpoon User's Guide (HRPNUG)** (provides information to build Harpoon Yocto images)
- **i.MX6ULL EVK GenAVB/TSN Rework Application Note (AN13678)**
- **NXP MPU Cortex-A Core Zephyr User Guide (UG10199)**
- For details about the graphics feature available in i.MX 8M Plus and i.MX 8M Mini boards, refer to the [i.MX Graphics User's Guide](#).

To boot up and set up the boards mentioned in this document, refer to the instructions available in the following user guides:

- [i.MX 6ULL EVK Quick Start Guide](#)
- [i.MX 8M Mini LPDDR4 EVK Quick Start Guide](#)
- [i.MX 8M Plus LPDDR4 EVK Quick Start Guide](#)
- [i.MX 8X Lite EVK Quick Start Guide](#)
- [LS1028ARDB Quick Start Guide](#)
- [LS1043ARDB Getting Started Guide](#)
- [LS1046ARDB Getting Started Guide](#)
- [LX2160A/LX2160A-Rev2 RDB Quick Start Guide](#)
- [i.MX 93 EVK Quick Start Guide](#)

- [i.MX 91 EVK Quick Start Guide](#)
- [i.MX RT1180 Getting Started Guide](#)

Note: For access to these documents, contact your local NXP Field Application Engineer (FAE) or support executive.

2 Real-time system

2.1 Overview

Real-time System addresses several real-time requirements on multicore platforms. These include schedule latency, inter-core communication, hardware resource sharing, unified life-cycle management, and unified building and deployment mechanism.

To support different schedule latency requirements, the Real-time System provides:

- Preempt-RT Linux
- native RTOS on Cortex-A core and Cortex-M core
- RTOS on Cortex-A core with Jailhouse and
- BareMetal framework.

It also provides a flexible combination of different cores running Preempt-RT Linux and RTOS/BareMetal to meet different real-time requirements for different use cases.

The heterogeneous multicore framework feature of Real-time Edge software provides different intercore communication mechanisms and hardware-resource sharing mechanisms between different CPU cores and different OSes to cover high-performance communication and real-time requirements. Unified CPU Core life cycle management provides a unified mechanism to bootstrap the Cortex-A core and Cortex-M core on the heterogeneous MPU system. Refer to chapter [Section 3](#) for details.

The unified software building and deploying mechanism provides easy building and deployment for the software components running on the Cortex-A core and Cortex-M core. These components and features of this mechanism are depicted in [Figure 1](#).

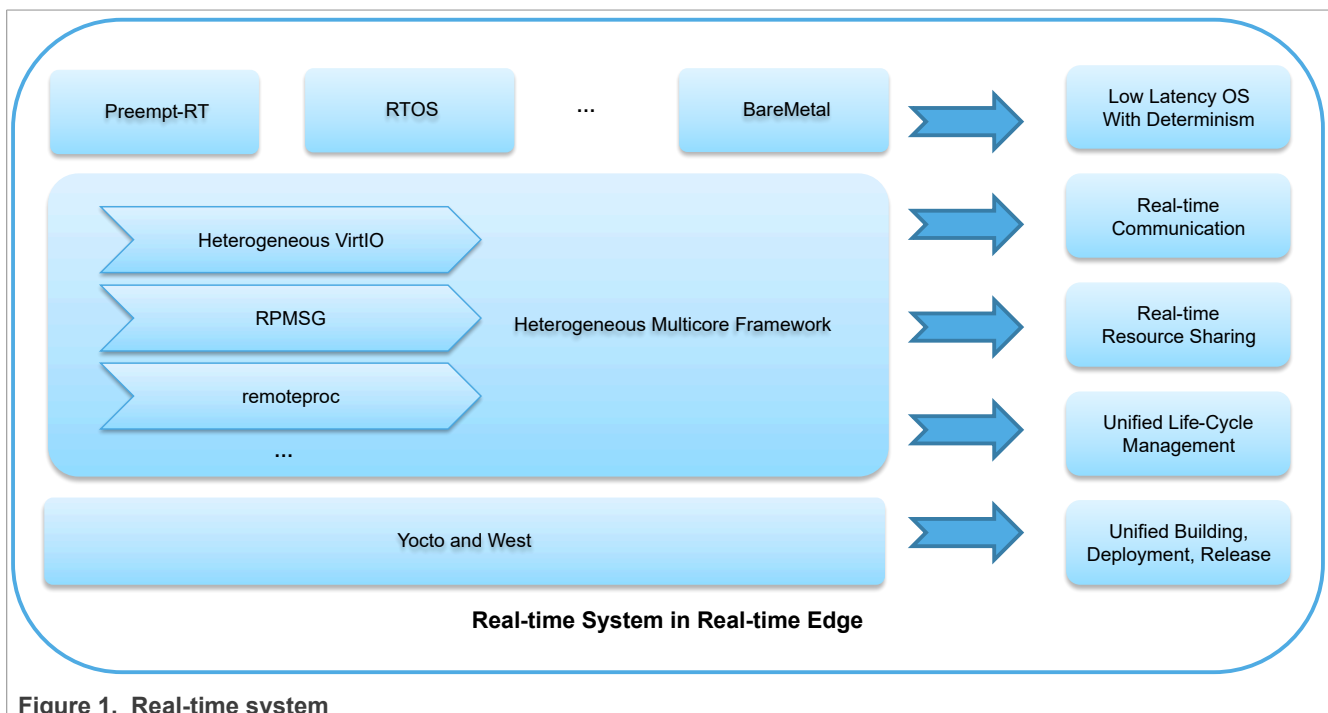


Figure 1. Real-time system

Real-time Edge software provides a general software architecture to run Real-time Systems on MPU platforms with the following features:

- **Different frameworks and flexible combinations for different schedule latency requirements**

The [Figure 2](#) shows all the OSES/BareMetal supported in Real-time Edge on MPU platforms with different-scale schedule latency:

– **Preempt-RT Linux on Cortex-A Core**

Real-time Linux kernel provides deterministic low latency compared to Linux.

– **BareMetal on Cortex-A Core**

Single or multiple BareMetal instance run Cortex-A Cores with zero schedule latency.

– **Native RTOS SMP/AMP on Cortex-A Core**

Native RTOS (FreeRTOS or Zephyr) is kicked to one or more Cortex-A Core from U-Boot. Jailhouse is not used and this combination targets lower latency and higher performance as compared to RTOS with Jailhouse.

– **RTOS SMP/AMP on Corex-A Core with Jailhouse**

RTOS (FreeRTOS or Zephyr) runs in Jailhouse inmate with hardware resource isolation on Cortex-A Core.

– **RTOS and BareMetal on Cortex-M Core**

This combination is generally used for Real-time Control systems, but has less CPU computing ability than Cortex-A Core.

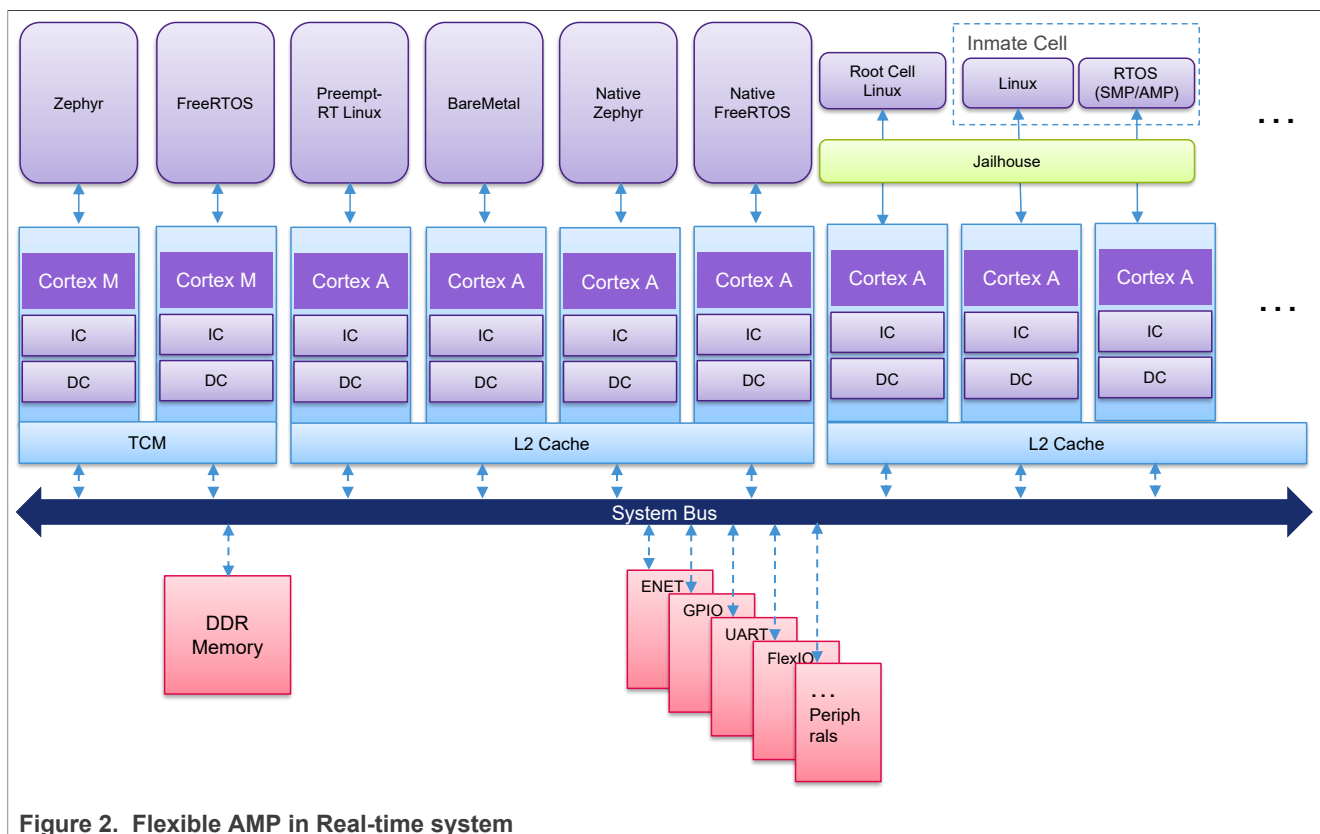


Figure 2. Flexible AMP in Real-time system

All these Real-time operating systems or BareMetal can be combined to be a flexible AMP system on multicore system. For example, the i.MX 8M Plus platform has four Cortex-A53 cores and one Cortex-M7 core. The Real-time Edge software supports the flexible AMP system to run these OS/BareMetal combinations:

- Four Cortex-A53 cores run SMP Preempt-RT Linux, Cortex-M7 core run RTOS.
- Four Cortex-A53 cores run four Baremetal/RTOS instances, Cortex-M7 core run RTOS.
- One or more Cortex-A53 cores run Preempt-RT Linux, the other Cortex-A53 cores run one or more Baremetal/RTOS instances, Cortex-M7 core run RTOS.

- One or more Cortex-A53 cores run Preempt-RT Linux as Jailhouse Root Cell, the other Cortex-A53 cores run as one or more inmate cell(s) with RTOS.
- **Unified Software Building, Deploy, and Release**
 - All different OS/application running on different cores are built via Yocto.
 - A bitbake command is used to create all images on different cores.
 - Single Flash Image includes all OS/applications running on all CPU cores.
- **Heterogeneous Multicore Framework**

A common framework with the following key features and functions:

 - **Inter-Core Data Communication and Resource Sharing**

Common Heterogeneous Multicore Framework provides data communication and resource sharing between M-Core and A-Core(s) or between different A-Cores simultaneously: RPMSG provides standard message communication for low bandwidth use cases, and Heterogeneous Multicore Virtio provides high performance data path and resource sharing to meet high bandwidth requirement.
 - **Unified Life-Cycle management for flexible AMP**

Real-time Edge software supports Preempt-RT Linux, FreeRTOS, Zephyr, and BareMetal running on different processors with Heterogeneous Multicore Framework.

2.2 Real-time-system support matrix

The [Table 4](#) shows the support matrix on NXP platforms:

Table 4. Real-time-system support matrix

Real-time System		i.MX 8M Mini LPDDR4 EVK			i.MX 8DXL			i.MX 93 EVK		i.MX 93 EVK (14 x 14)		i.MX RT1180-EVK	i.MX 91 9x9 QSB	i.MX 91 11x11 EVK	i.MX 95 15x15 EVK		i.MX 95 19x19 EVK		LS1028ARDB	LS1043ARDB	LS1046ARDB	LX2160ARDB
Cores		4 x A53	1 x M4	4 x A53	1 x M7	2 x A35	1 x M4	2 x A55	1 x M33	2 x A55	1 x M33		1 x A55	1 x A55	6 x A55	1 x M7	6 x A55	1 x M7	2 x A72	4 x A53	4 x A72	16 x A72
Preempt-RT Linux		Y		Y		Y		Y		Y			Y	Y	Y		Y		Y	Y	Y	Y
BareMetal		Y	Y	Y	Y			Y	Y										Y	Y	Y	Y
Native RTOS	FreeRTOS	Y	Y	Y	Y		Y	Y	Y	Y	Y	Y	Y	Y	Y		Y					
	Zephyr	Y	Y	Y	Y			Y	Y	Y			Y	Y	Y		Y					
Jailhouse	Baremetal	Y		Y				Y											Y	Y	Y	Y
	FreeRTOS	Y		Y				Y														
	Zephyr	Y		Y				Y														
	Harpoon	Y		Y				Y														

2.3 Building, deploying, and releasing unified software

The Yocto project is an open source collaboration project that helps developers create custom Linux-based systems regardless of the hardware architecture. The project provides a flexible set of tools and a space where embedded developers worldwide can share technologies, software stacks, configurations, and best practices. These features can further be used to create tailored Linux images for embedded and IoT devices, or for a customized Linux OS. Moreover, the Linux factory selects Yocto as the building tool. Real-time Edge also selects Yocto as the unified SW release tool. [Figure 3](#) shows the unified Yocto structure for Heterogeneous AMP.

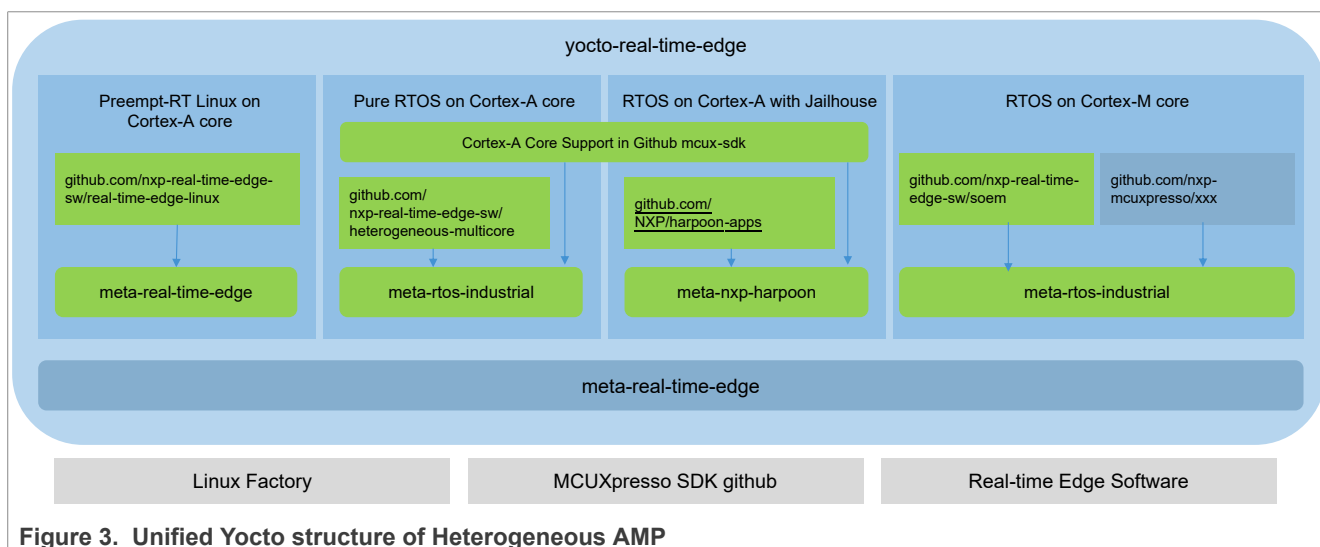


Figure 3. Unified Yocto structure of Heterogeneous AMP

The different Yocto Layers support different functions:

- The `meta-real-time-edge` layer focuses on building Linux and BareMetal application on Cortex-A core.
- The `meta-nxp-harpoon` layer focuses on building RTOS on Cortex-A core.
- The `meta-rtos-industrial` layer focuses on building RTOS running on Cortex-M core.

Only one build command is required to generate the complete image including all binaries running on core A and core M.

For example:

```
# setup yocto environment for imx8mp-lpddr4-evk board
$ DISTRO=nxp-real-time-edge MACHINE=imx8mp-lpddr4-evk source real-time-edge-setup-env.sh
# build all images for imx8mp-lpddr4-evk board
$ bitbake nxp-image-real-time-edge
```

2.3.1 Yocto layer for Cortex-A core

The Cortex-A core allows users to run Linux, Jailhouse, Baremetal, and RTOS. The corresponding Yocto layer description is as follows:

1. Linux and Rootfs

The Yocto layer `meta-real-time-edge` focuses on Linux building on Cortex-A cores. This layer is based on Linux factory and describes the process for building all applications for Linux and rootfs on Cortex-A core.

2. Jailhouse

The scripts under `meta-real-time-edge/recipes-extended/real-time-edge-jailhouse` describe how to build Jailhouse running on Cortex-A core.

3. Baremetal application

The scripts under `meta-real-time-edge/recipes-extended/real-time-edge-baremetal` describe how to build baremetal application on Cortex-A core. Refer to [Section 2.5](#) for details.

4. Harpoon (RTOS on A core)

Harpoon provides an environment for developing real-time demanding applications on an RTOS running on one (or several) Cortex-A core(s) in parallel of a Linux distribution, leveraging the 64-bit Arm (R) architecture for higher performance. The system starts on Linux and the Jailhouse partitions the hardware to run both Linux and the guest RTOS in parallel. The hardware partitioning is configurable and depends on the use case. The Yocto layer `meta-nxp-harpoon` describes how to build these applications on Cortex-A core. For more information, refer to Harpoon User's Guide. See [Section 1.5](#).

2.3.2 Yocto layer for Cortex-M core

When the application runs on the Cortex-M core, it uses different toolchain and source code. For a unified compilation interface, Yocto meta layer `meta-rtos-industrial` is introduced into Real-time Edge project. The `meta-rtos-industrial` layer provides the build environment to create MCUX SDK application for Cortex-M cores.

2.3.2.1 Introduction to meta-rtos-industrial

The [Figure 4](#) shows the `meta-rtos-industrial` file structure.

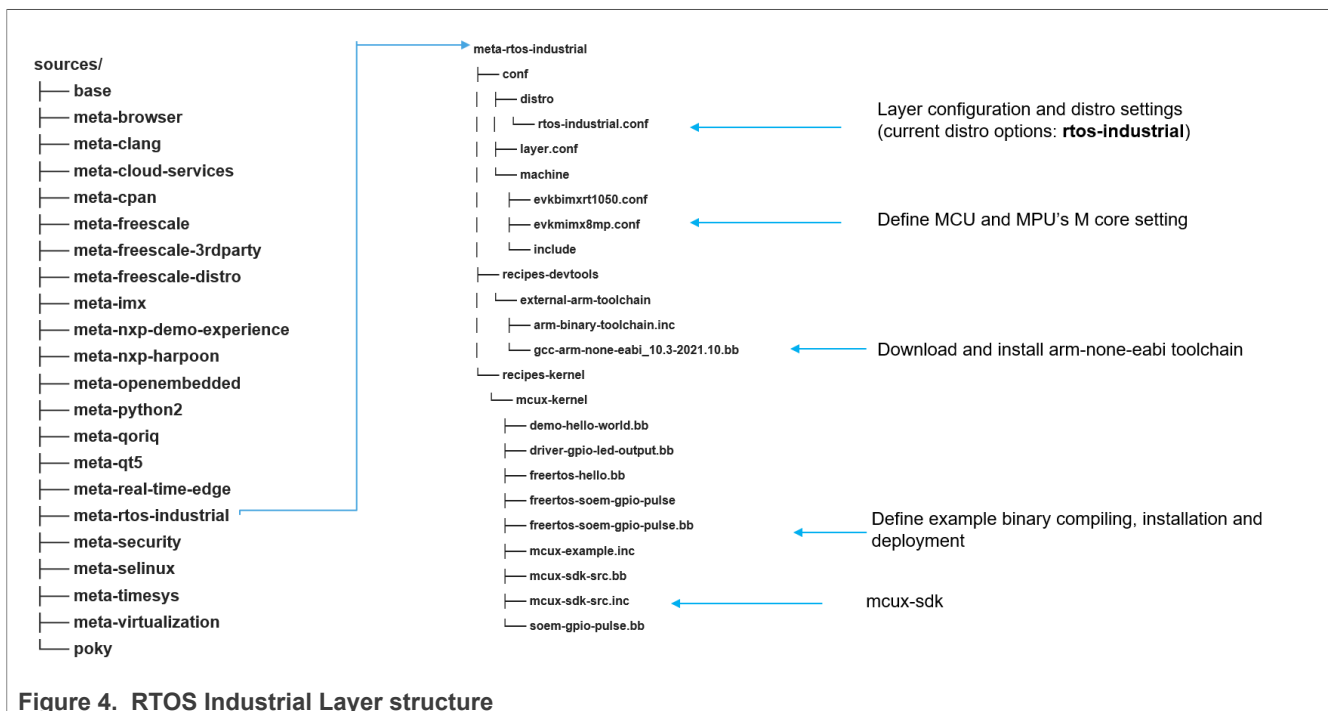


Figure 4. RTOS Industrial Layer structure

2.3.2.1.1 Source code definition

All source code related definition is under `recipes-kernel/mcux-kernel` folder.

The file `mcux-sdk-src.inc` defines all the repos of ([NXPmicro/mcux-sdk: MCUXpresso SDK \(github.com\)](#)) and the new repos.

If a new repo needs to be downloaded, append a new line to "SRC_URI" with the URL and location of the required repo.

For example, use the below code to download 'SOEM' stack to `git/core/components/SOEM` folder:

```
git://${NXPMICRO_BASE}/soem.git;protocol=https;nobranch=1;destsuffix=git/core/
components/SOEM;name=SOEM \
```

The `mcux-sdk-src-XXX.inc` defines the MCUX SDK repo commit ID for the release XXX. For example, `mcux-sdk-src-2.11.0.inc` contains all the commit IDs for the release 2.11.0 repository.

The parameter `PREFERRED_VERSION_MCUX-SDK` defines the default version in `mcux-sdk-src.inc`. If you want to compile a different version, overwrite this parameter in the `local.conf`.

For example:

```
# Add the below line into local.conf
PREFERRED_VERSION_MCUX-SDK = "2.10.0"
```

2.3.2.1.2 Example definition

The file `mcux-examples.inc` describes the common method to compile, install, and deploy examples. Each example `bb` file must include this file and then specify the folder of the example.

Use the command below to add a new example:

```
include mcux-example.inc
MCUX_EXAMPLE_DIR = "examples/${RTOS-INDUSTRIAL-BOARD}/demo_apps/hello_world"
```

2.3.2.1.3 Toolchain definition

The file `recipes-devtools/external-arm-toolchain/gcc-arm-none-eabi_VERSION.bb` describes how to download, install, and deploy `gcc-arm-none-eabi` toolchain of the specific VERSION.

This layer also supports external toolchain. Parameter "ARMGCC_DIR" can be overwritten to point the external toolchain.

For example:

```
ARMGCC_DIR = "/MYPATH/arm-none-eabi"
```

2.3.2.2 Integration of meta-rtos-industrial

To integrate `meta-rtos-industrial` into the Real-time Edge project, you need to specify the board and examples.

The board name is different between i.MX SDK and MCUX SDK. For example, in order to compile Cortex-M application for i.MX 8M Mini EVK with LPDDR4, use the board name `evkmimx8mm` instead of `imx8mm-lpddr4-evk`. The file `rtos-industrial-examples.inc` is created under `meta-real-time-edge/distro/include` to map the board names. The board name used by MCUX SDK should be set to parameter `RTOS-INDUSTRIAL-BOARD`.

In the path `meta-real-time-edge/recipes-nxp/packagegroups/packagegroup-real-time-edge-rtos.bb` is used for examples that are compiled. These examples should be installed into rootfs.

2.3.2.3 Building meta-rtos-industrial

As the meta-rtos-industrial is already integrated into Real-time Edge, we do not need any special commands or settings to enable building the rtos application. When building `nxp-image-real-time-edge` image, all examples defined in `packagegroup-real-time-edge-rtos.bb` are built and installed into “/examples” folder in rootfs.

Use the below commands to create `nxp-image-real-time-edge` image for `imx8mm-lpddr4-evk` board.

```
$ mkdir yocto-real-time-edge
$ cd yocto-real-time-edge
$ repo init -u https://github.com/nxp-real-time-edge-sw/yocto-real-time-edge.git
  -b real-time-edge-scarthgap -m real-time-edge-3.1.0.xml
$ repo sync
$ DISTRO=nxp-real-time-edge MACHINE=imx8mm-lpddr4-evk source real-time-edge-
  setup-env.sh -b build-real-time-edge
$ bitbake nxp-image-real-time-edge
```

The sample binary files are located under the `tmp/deploy/images/imx8mm-lpddr4-evk/examples` and `examples` directories of rootfs.

```
examples/
|-- heterogeneous-multicore
|   |-- hello-world-ca
|   |   |-- ddr_release
|   |   |   |-- hello_world_ca53.bin
|   |   |   |-- hello_world_ca53.elf
|   |-- virtio-net-backend-ca
|   |   |-- ddr_release
|   |   |   |-- virtio_net_backend_ca53.bin
|   |   |   |-- virtio_net_backend_ca53.elf
|   |-- virtio-net-backend-cm
|   |   |-- release
|   |   |   |-- virtio_net_backend_cm4.bin
|   |   |   |-- virtio_net_backend_cm4.elf
|   |-- virtio-perf-ca
|   |   |-- ddr_release
|   |   |   |-- virtio_perf_ca53.bin
|   |   |   |-- virtio_perf_ca53.elf
|   |-- virtio-perf-cm
|   |   |-- release
|   |   |   |-- virtio_perf_cm4.bin
|   |   |   |-- virtio_perf_cm4.elf
|-- mcux-sdk
|   |-- 9bit-uart-interrupt-transfer
|   |   |-- ddr_release
|   |   |   |-- 9bit_uart_interrupt_transfer.bin
|   |   |   |-- 9bit_uart_interrupt_transfer.elf
|   |   |-- release
|   |   |   |-- 9bit_uart_interrupt_transfer.bin
|   |   |   |-- 9bit_uart_interrupt_transfer.elf
|   |-- 9bit-uart-polling
|   |   |-- ddr_release
|   |   |   |-- 9bit_uart_polling.bin
|   |   |   |-- 9bit_uart_polling.elf
|   |   |-- release
|   |   |   |-- 9bit_uart_polling.bin
|   |   |   |-- 9bit_uart_polling.elf
|-- demo-hello-world
```

```
| | -- ddr_release  
| | | -- hello_world.bin  
| | `-- hello_world.elf  
| | `-- release  
| | | -- hello_world.bin  
| | `-- hello_world.elf  
| -- driver-gpio-led-output  
| | -- ddr_release  
| | | -- igpio_led_output.bin  
| | `-- igpio_led_output.elf  
| | `-- release  
| | | -- igpio_led_output.bin  
| | `-- igpio_led_output.elf  
| -- freertos-hello  
| | -- ddr_release  
| | | -- freertos_hello.bin  
| | `-- freertos_hello.elf  
| | `-- release  
| | | -- freertos_hello.bin  
| | `-- freertos_hello.elf  
| -- freertos-soem-gpio-pulse  
| | -- ddr_release  
| | | -- soem_gpio_pulse_freertos.bin  
| | `-- soem_gpio_pulse_freertos.elf  
| | `-- release  
| | | -- soem_gpio_pulse_freertos.bin  
| | `-- soem_gpio_pulse_freertos.elf  
| -- rpmsg-lite-str-echo-rtos-8m-cm  
| | -- ddr_release  
| | | -- rpmsg_lite_str_echo_rtos.bin  
| | `-- rpmsg_lite_str_echo_rtos_imxcm4.elf  
| | `-- release  
| | | -- rpmsg_lite_str_echo_rtos.bin  
| | `-- rpmsg_lite_str_echo_rtos_imxcm4.elf  
| -- rpmsg-lite-str-echo-rtos-ca  
| | -- ddr_release  
| | | -- rpmsg_lite_str_echo_rtos.bin  
| | `-- rpmsg_lite_str_echo_rtos_imxca53.elf  
| -- rpmsg-lite-uart-sharing-rtos  
| | -- release  
| | | -- rpmsg_lite_uart_sharing_rtos.bin  
| | `-- rpmsg_lite_uart_sharing_rtos.elf  
| -- soem-gpio-pulse  
| | -- ddr_release  
| | | -- soem_gpio_pulse.bin  
| | `-- soem_gpio_pulse.elf  
| | `-- release  
| | | -- soem_gpio_pulse.bin  
| | `-- soem_gpio_pulse.elf
```

If you just want to compile a special example, you can use the following command:

For example:

```
$ DISTRO=nxp-real-time-edge MACHINE=imx8mm-lpddr4-evk bitbake demo-hello-world
```

2.4 Preempt-RT Linux

The Preempt-RT Linux option turns the kernel into a real-time kernel. It does so by replacing various locking primitives (for example, spinlocks and rwlocks) with preemptible priority-inheritance aware variants. The Preempt-RT Linux option also enforces interrupt threading and introduces mechanisms to break up long non-preemptible sections. This makes the kernel fully preemptible and brings most execution contexts under scheduler control. However, very low level and critical code paths (entry code, scheduler, low level interrupt handling) remain non-preemptible.

2.4.1 System Real-time Latency tests

The basic measurement tool for Real-time Linux is `cyclicttest`.

2.4.1.1 Running Cyclicttest

Cyclicttest provides statistics about the latencies of the system. It accurately and repeatedly measures the difference between the intended wake-up time of a thread and the time at which it actually wakes up. It can measure latencies in real-time systems caused by the hardware, the firmware, and the operating system.

Thomas Gleixner (tglx) wrote the original test, but several people had later contributed modifications. Cyclicttest is part of the test suite, [rt-tests](#). Clark Williams and John Kacur currently maintain Cyclicttest.

cyclicttest:

- Use the below command to perform Latency Test:

```
$ cyclicttest -p90 -h50 -D30m
```

Note: For detailed parameters of Cyclicttest, refer to [Cyclicttest Web Page](#).

2.4.2 Real-time application development

This section describes the steps for developing the Real-time application.

Real-time Application: API, Basic Structure, Background:

- Basic Linux application rules are the same; Use the POSIX API.
- There is still a division of Kernel Space and User Space.
- Linux applications run in User Space.
- For details, refer to: https://wiki.linuxfoundation.org/realtime/documentation/howto/applications/application_base

Real-time Application: Users can build it using the steps below:

- **Using the cross-compiler example:**

```
$ arm-linux-gnueabihf-gcc <filename>.c -o <filename>.out -lrt -Wall
```

- **Using the native compiler on a target example:**

```
$ gcc <filename>.c -o <filename>.out -lrt -Wall
```

Scheduling policies have two classes:

1. Completely Fair Scheduling (CFS)

- **SCHED_NORMAL** (traditionally called SCHED_OTHER): The scheduling policy that is used for regular tasks. Every task gets a so called 'nice value'. It is a value between -20 for the highest nice value and 19 for

the lowest `nice` value. The average value of execution time of the task depends on the associated '`nice` value'.

- **SCHED_BATCH**: Does not preempt nearly as often as regular tasks. Hence, it allows tasks to run longer and make better use of caches, but at the cost of interactivity. This is well suited for batch jobs and optimized for throughput.
- **SCHED_IDLE**: This policy is even weaker than a `nice` value of 19. However, it is not a true idle timer scheduler in order to avoid getting into priority inversion problems, which would deadlock the machine.

2. Real-time policies

- **SCHED_FIFO**: Tasks have a priority between 1 (low) and 99 (high). A task running under this policy is scheduled until it finishes or a higher prioritized task preempts it.
- **SCHED_RR**: This policy is derived from SCHED_FIFO. The difference with respect to SCHED_FIFO policy is that a task runs during a defined time slice (if it is not preempted by a higher prioritized task). It can be interrupted by a task with the same priority once the time slice is used up. The time slice definition is exported in `procfs` (`/proc/sys/kernel/sched_rr_timeslice_ms`).
- **SCHED_DEADLINE**: This policy implements the Global Earliest Deadline First (GEDF) algorithm. Tasks scheduled under this policy can preempt any task scheduled with SCHED_FIFO or SCHED_RR.

2.5 Baremetal on Cortex-A core

The following sections provide an overview of the Real-time Edge BareMetal framework on A core including:

- Features supported
- Getting started with the BareMetal framework using the supported platforms:
 - NXP Layerscape platforms
 - i.MX 8M / i.MX 93 platforms

It also describes how to run a sample BareMetal framework on the host environment and develop customer-specific applications based on this framework.

2.5.1 BareMetal framework

The BareMetal framework supports the scenarios that need low latency, real-time response, and high performance. There is no OS running on the cores and customer-specific application runs on the core directly. The [Figure 5](#) depicts the BareMetal framework architecture.

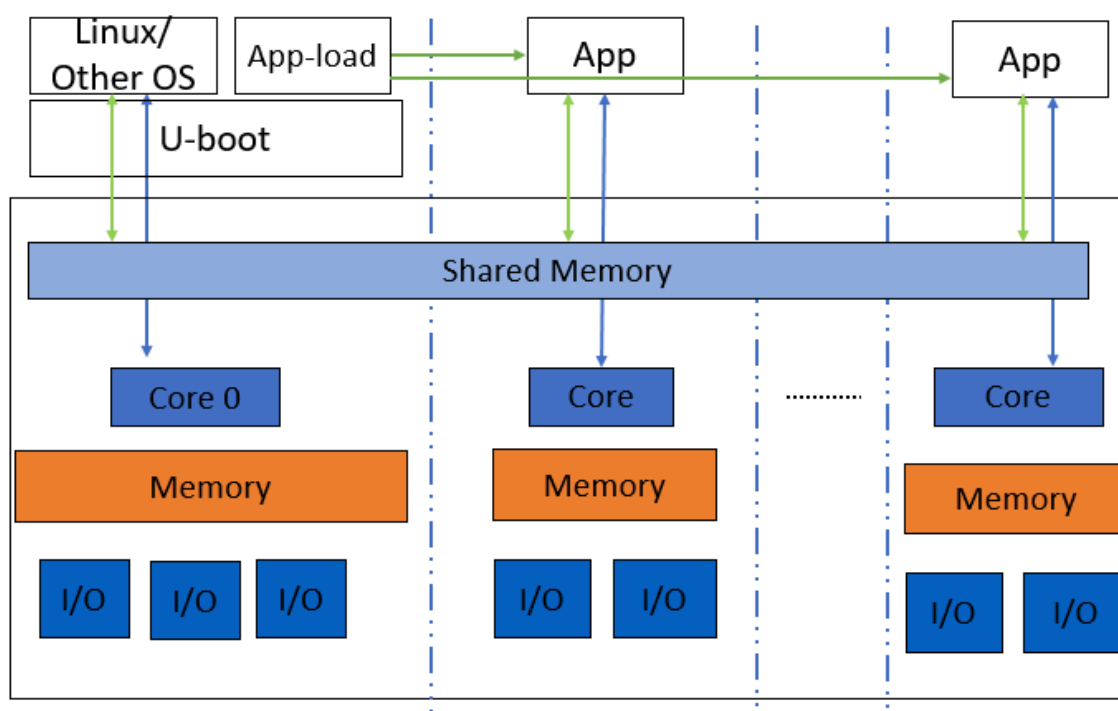


Figure 5. BareMetal framework architecture

The main features of the BareMetal framework are as follows:

- Core0 runs as a controller and runs the BareMetal or the operating system such as Linux, Vxworks.
- Responder cores run the BareMetal application.
- Easy assignment of different IP blocks to different cores.
- Interrupts between different cores and a high-performance mechanism for data transfer.
- Different UART for core0 and slave cores for easy debug.
- Communication via shared memory.

The master core0 runs the Baremetal under master mode. It then loads the BareMetal application to the slave cores and starts the BareMetal application. The [Figure 6](#) depicts the boot flow diagram:

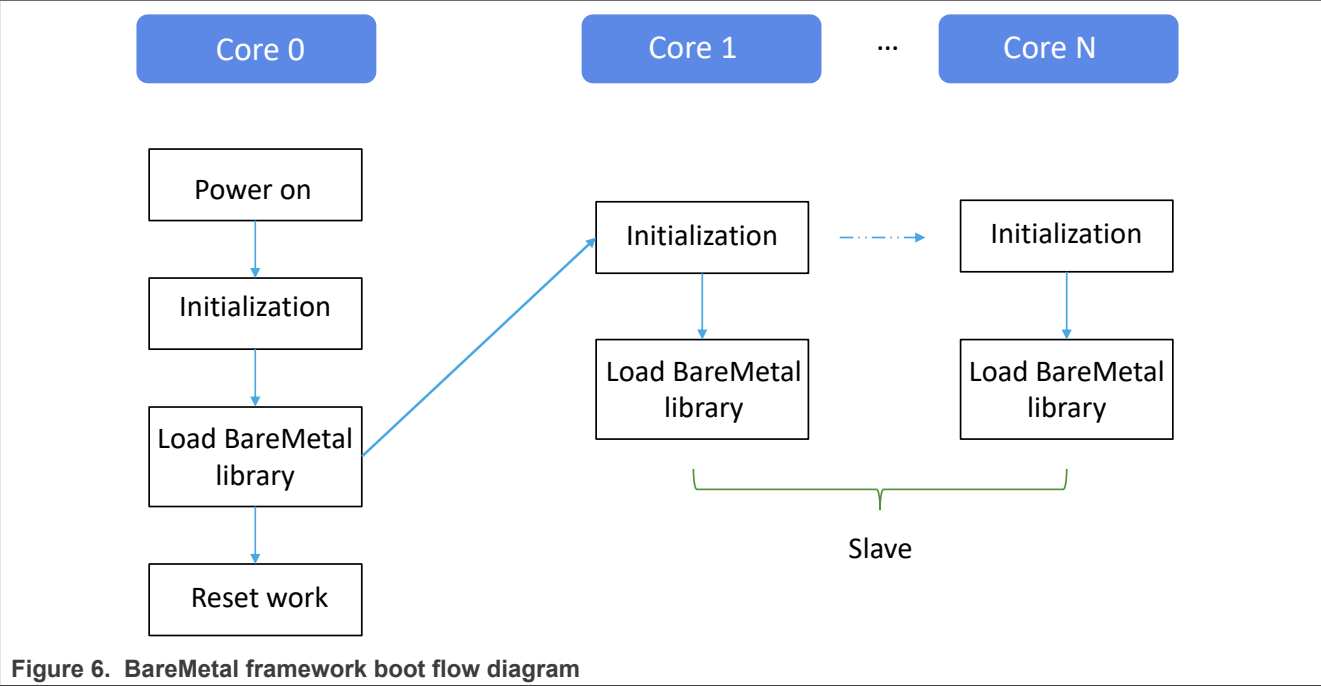


Figure 6. BareMetal framework boot flow diagram

The [Table 5](#) lists the industrial IoT features supported by various NXP processors and boards.

Table 5. BareMetal features supported by NXP processors

Processor	Board	Main features supported
i.MX 8M Mini	i.MX 8M Mini LPDDR4 EVK	UART, IPI, data transfer, Ethernet, GPIO
i.MX 8M Plus	i.MX 8M Plus LPDDR4 EVK	UART, IPI, data transfer, Ethernet, GPIO
i.MX 93	i.MX 93 EVK	UART, IPI, data transfer, Ethernet
	i.MX 93 9x9 LPDDR4 QSB	UART, IPI, data transfer
LS1028A	LS1028ARDB	I2C, UART, ENETC, IPI, data transfer, SAI
LS1043A	LS1043ARDB	IRQ, IPI, data transfer, Ethernet, IFC, I2C, UART, FMan, USB, PCIe
LS1046A	LS1046ARDB	IRQ, IPI, data transfer, Ethernet, IFC, I2C, UART, FMan, QSPI, USB, PCIe, GPIO
LX2160A/Rev2	LX2160ARDB	UART, IPI, data transfer

Typical use cases are as follows:

1. Core0 as a master core runs Linux to manage slave cores and communicate with server. Slave cores run Baremetal application for real-time processing. Refer [Figure 7](#).




Figure 7. BareMetal Use Case 1

2. All cores run BareMetal application for real-time processing. Refer [Figure 8](#).




Figure 8. BareMetal Use Case 2

2.5.2 Getting started

This section describes how to set up the environment and run the Baremetal examples on slave cores (assuming that the core0 is the master core and the other cores are the slave cores).

2.5.2.1 Hardware and software requirements

In order to run baremetal framework scenarios, the following are required:

- **Hardware:** i.MX 8M Mini LPDDR4 EVK, i.MX 8M Plus LPDDR4 EVK, i.MX 93 EVK, i.MX 93 9x9 LPDDR4 QSB, LS1028ARDB, LS1043ARDB, LS1046ARDB, LX2160ARDB, and serial cables.
- **Software:** Real-time Edge Software v2.7 release or later.

2.5.2.2 Hardware setup

This section describes the hardware setup required for the NXP boards for running the Baremetal framework examples.

2.5.2.2.1 i.MX 8M Mini LPDDR4 EVK and i.MX 8M Plus LPDDR4 EVK board

Follow the steps below.

1. **i.MX 8M Plus LPDDR4 EVK:** There is one USB MicroB Debug port on board. Four UART ports can be found when the MicroB cable connects to PC.

```
/dev/ttyUSB0  
/dev/ttyUSB1  
/dev/ttyUSB2  
/dev/ttyUSB3
```

Use `/dev/ttyUSB2` for core0 (master core) and `/dev/ttyUSB3` for core1, core2, and core3 (slave cores).

2. **i.MX 8M Mini LPDDR4 EVK:** There is one USB MicroB Debug port on board. Two UART ports can be found when the MicroB cable connects to PC.

```
/dev/ttyUSB0  
/dev/ttyUSB1
```

Use `/dev/ttyUSB1` for core0 (master core) and `/dev/ttyUSB0` for core1, core2, and core3 (slave cores).

3. **GPIO setup**

For GPIO test on i.MX 8M Plus LPDDR4 EVK, connect pin 7 and pin 8 of J21 by a jumper, as shown in [Figure 9](#).




Figure 9. Connections for GPIO test on i.MX 8M Plus LPDDR4 EVK board

4. For GPIO test on i.MX 8M Mini LPDDR4 EVK, connect pin7 and pin8 of J1003 by a jumper as shown in [Figure 10](#).




Figure 10. Connections for GPIO test on i.MX 8M Mini LPDDR4 EVK board

2.5.2.2.2 LS1028ARDB, LX2160ARDB, LS1043ARDB, or LS1046ARDB

To develop the Real-time Edge BareMetal framework using one of the boards (LS1028ARDB, LX2160ARDB, LS1043ARDB, or LS1046ARDB), two serial cables are needed for connection. One serial cable is used for core0, to connect to UART1 port. The other cable is used for slave cores and connects to the UART2 port.

To support the SAI feature on LS1028ARDB, set the switch SW5_8 to "ON".

2.5.2.2.3 i.MX 93 EVK

On the i.MX 93 EVK board, the USB Type-C connector (J1401) provides four UART ports when connected to a PC using USB cable. The third port (LPUART1) is used for core0 (master core) and the fourth port (LPUART2) is used for core1 (slave core).

2.5.2.2.4 i.MX93 9x9 LPDDR4 QSB

On the i.MX 93 9x9 LPDDR4 QSB board, the USB Type-C connector (J1708) provides four UART ports when connected to a PC using USB cable. The third port (LPUART2) is used for core0 (master core) and the fourth port (LPUART3) is used for core1 (slave core).

2.5.2.3 Building the software

There are two methods to build the BareMetal images:

- The first method is to compile the images in a standalone way, and is described in the following section.
- The second method is to build the BareMetal images using the Real-time Edge framework. This method is described in the document, Real-time Edge Yocto Project User Guide in the section "*Building the image through Yocto*".

2.5.2.3.1 Building the BareMetal binary for slave cores

Perform the steps mentioned below:

1. Download the project source from the following path:
<https://github.com/nxp-real-time-edge-sw/real-time-edge-uboot.git>
2. Check it out to the tag:
 - Real-Time-Edge-v3.1-baremetal-202503
3. Configure the cross-toolchain on your host environment.
4. Then, run the following commands:

```
/* build Baremetal image for i.MX 8M Mini LPDDR4 EVK Rev.C board */
$ make imx8mm_evk_baremetal_slave_defconfig
$ make
/* build Baremetal image for i.MX 8M Plus LPDDR4 EVK board */
$ make imx8mp_evk_baremetal_slave_defconfig
$ make
/* build Baremetal image for i.MX 93 EVK board */
$ make imx93_11x11_evk_baremetal_slave_defconfig
$ make
/* build Baremetal image for i.MX 93 9x9 LPDDR4 QSB board */
$ make imx93_9x9_qsb_baremetal_slave_defconfig
$ make
/* build Baremetal image for LS1028ARDB board */
$ make ls1028ardb_baremetal_slave_defconfig
$ make
/* build Baremetal image with SAI for LS1028ARDB board */
$ make ls1028ardb_baremetal_slave_sai_defconfig
$ make
/* build Baremetal image for LS1043ARDB board */
$ make ls1043ardb_baremetal_slave_defconfig
$ make
/* build Baremetal image for LS1046ARDB board */
$ make ls1046ardb_baremetal_slave_defconfig
$ make
/* build Baremetal image for LX2160ARDB board */
$ make lx2160ardb_baremetal_slave_defconfig
$ make
```

5. Finally, the file `u-boot-dtb.bin` used for BareMetal is generated.

Follow the Real-time Edge Software Yocto Project to get the code and build images for these platforms.

2.5.2.3.2 Building the image through Yocto

There are two methods to build the Baremetal images. The first method, which is used to compile the images in a standalone way, is described in [Section 2.5.2.3](#). The other method is to build the Baremetal images using Real-time Edge software framework, and is described in this section.

The Real-time Edge software is designed for embedded industrial usage. It is an integrated Linux distribution for industry. With the current version, the Baremetal can be built and implemented conveniently.

2.5.2.3.2.1 Getting Real-time Edge software

The latest release is available at the following URL:

<https://github.com/nxp-real-time-edge-sw/yocto-real-time-edge.git>

Follow Yocto documentation "Real-time Edge Yocto Project User Guide" to get the code and build the image. Refer to [Section 1.5](#).

2.5.2.3.2.2 Building the Baremetal images

This section describes the steps for building the Baremetal images for various boards. The steps described are applicable to the boards such as LS1043ARDB, LS1046ARDB, LX2160ARDB, i.MX 8M Plus LPDDR4 EVK, and i.MX 8M Mini LPDDR4 EVK board.

Building the Baremetal images for various boards

Run the following commands to build the final Baremetal image for Layerscape and i.MX platforms.

```
$ cd yocto-real-time-edge
```

For i.MX 93 EVK Baremetal image:

```
$ DISTRO=nxp-real-time-edge-baremetal MACHINE=imx93evk source real-time-edge-setup-env.sh -b build-imx93evk-bm
```

For i.MX 93 9x9 LPDDR4 QSB Baremetal image:

```
$ DISTRO=nxp-real-time-edge-baremetal MACHINE=imx93-9x9-lpddr4-qsb source real-time-edge-setup-env.sh -b build-imx93qsb-bm
```

For LS1028ARDB Baremetal image:

```
$ DISTRO=nxp-real-time-edge-baremetal MACHINE=ls1028ardb source real-time-edge-setup-env.sh -b build-ls1028ardb-bm
```

For LS1043ARDB Baremetal image:

```
$ DISTRO=nxp-real-time-edge-baremetal MACHINE=ls1043ardb source real-time-edge-setup-env.sh -b build-ls1043ardb-bm
```

For LS1046ARDB Baremetal image:

```
$ DISTRO=nxp-real-time-edge-baremetal MACHINE=ls1046ardb source real-time-edge-setup-env.sh -b build-ls1046ardb-bm
```

For LX2160ARDB Baremetal image:

```
$ DISTRO=nxp-real-time-edge-baremetal MACHINE=lx2160ardb-rev2 source real-time-edge-setup-env.sh -b build-lx2160ardb-bm
```

For i.MX 8M Plus LPDDR4 EVK Baremetal image:

```
$ DISTRO=nxp-real-time-edge-baremetal MACHINE=imx8mp-lpddr4-evk source real-time-edge-setup-env.sh -b build-imx8mpevk-bm
```

For i.MX 8M Mini LPDDR4 EVK Baremetal image:

```
$ DISTRO=nxp-real-time-edge-baremetal MACHINE=imx8mm-lpddr4-evk source real-time-edge-setup-env.sh -b build-ix8mmevk-bm
```

Then, use:

```
$ bitbake nxp-image-real-time-edge
```

2.5.2.4 Booting up the Linux with Baremetal

Use the following steps to boot up the system with the images built from Real-time Edge software.

For platforms that can be booted up from the SD card, there are just two steps required to program the image into SD card.

1. Insert an SD card (at least 4 GB size) into any Linux host machine.
2. Find the image file in building directory (for example: ls1028ardb):

```
tmp/deploy/images/ls1028ardb/nxp-image-real-time-edge-ls1028ardb.wic.zst
```

3. Then, run the following commands:

```
$ zstd -d nxp-image-real-time-edge-ls1028ardb.wic.zst
$ sudo dd if=./nxp-image-real-time-edge-ls1028ardb.wic of=/dev/sdx bs=1M
conv=fsync
# or in some other host machine:
$ sudo dd if=./nxp-image-real-time-edge-ls1028ardb.wic of=/dev/mmcblkx bs=1M
conv=fsync
# find the right SD Card device name in your host machine and replace the
"sdx" or "mmcblkx".
```

4. Then, insert the SD card into the target board (for example ls1028ardb) and power on.

After completion of the above mentioned steps, the Linux system boots up on the master core (core 0), and the Baremetal system boots up on slave core (core 1) automatically.

2.5.3 Running examples

The following sections describe how to run the Baremetal examples on the host environment for LS1028ARDB board. Similar steps can be followed for LS1043ARDB, LS1046ARDB, i.MX 8M Mini LPDDR4 EVK, i.MX 8M Plus LPDDR4 EVK, i.MX 93 EVK and i.MX 93 9x9 LPDDR4 QSB board.

2.5.3.1 Preparing the console

In current Baremetal framework design, two UART ports are used as console. One UART is used for master core and the other UART is used for slave cores. Refer to Section 3.2.2.2 for preparing the console.

2.5.3.2 Running the Baremetal binary

As described earlier, there are two methods to compile the Baremetal framework. One is a standalone method and the other method uses the Real-time Edge software. These methods are described in [Section 2.5.2.3](#) and [Section 2.5.2.3.2](#) respectively.

- If the Real-time Edge software is used to compile the Baremetal image, the Baremetal image is included in the `nxp-image-real-time-edge-xxxx.wic.zst`. In this case, the master core starts the Baremetal image on slave cores automatically.
- If standalone compilation method is used, perform the steps below to run the Baremetal binary from U-Boot prompt of master core. See the below example run on Layerscape platform:
 1. After starting U-Boot on the master, download the bare metal image: `u-boot-dtb.bin` on `0x84000000` using the command below:

```
=> tftp 0x84000000 xxxx/u-boot-dtb.bin
```

Where

- `xxxx` is your tftp server directory.
- `0x84000000` is the address of `CONFIG_TEXT_BASE` on bare metal for Layerscape platforms.

Note:

- a. The address is `0x50200000` for i.MX 8M Plus LPDDR4 EVK and i.MX 8M Mini LPDDR4 EVK boards.
- b. The address is `0x90200000` for i.MX 93 EVK board.

2. Then, start the Baremetal cores using the command below:

```
=> dcache flush; cpu 1 release 0x84000000
```

Note: In the command `cpu <num> release 0x84000000`, the 'num' can be 1, 2, 3, ... to the maximum CPU number.

For i.MX 8M Plus LPDDR4 EVK and i.MX 8M Mini LPDDR4 EVK boards, use the below command:

```
=> dcache flush;cpu 1 release 50200000;sleep 6;cpu 2 release 50200000;sleep 2;cpu 3 release 50200000;
```

3. Last, the UART2 port displays the logs, and the bare metal application runs on slave cores successfully.

2.5.4 Development based on Baremetal framework

This chapter describes how to develop customer-specific application based on Baremetal framework.

2.5.4.1 Developing the Baremetal application

The “app” directory in the U-boot repository includes the test cases for testing the I2C, GPIO, and IRQ init features. Users can write their custom applications and store them in this directory.

2.5.4.2 Main file `app.c`

The file `<U-boot path>/app/app.c`, is the main file to add all applications. Users can modify the `app.c` file to add their applications.

- When the standalone method is used to build the Baremetal image as described in [Section 2.5.2.3](#), change the directory to `U-boot path` to read or edit the `app.c` file.
- When the Real-time Edge software is used to compile the Baremetal binary, change to the building directory to view or edit the `app.c` file.

The following is a sample code of the file `app.c` that shows how to add the example test cases of I2C, IRQ, and GPIO.

```
void core1_main(void)
{
    test_i2c();
    test_irq_init();
    test_gpio();
    return;
}
```

2.5.4.3 Common header files

There are some common APIs provided by Baremetal. The table below describes the header files that include the APIs.

Table 6. Common header file description

Header file	Description
asm/io.h	Read/Write IO APIs. For example, <code>_raw_readb</code> , <code>_raw_writeb</code> , <code>out_be32</code> , and <code>in_be32</code> .
linux/string.h	APIs for manipulating strings. For example, <code>strlen</code> , <code>strcpy</code> , and <code>strcmp</code> .
linux/delay.h	APIs used for small pauses. For example, <code>udelay</code> and <code>mdelay</code> .
linux/types.h	APIs specifying common types. For example, <code>_u32</code> and <code>_u64</code> .
common.h	Common APIs For example, <code>printf</code> and <code>puts</code> .

2.5.4.4 GPIO example

The file `uboot/app/test_gpio.c` is an example to test the GPIO feature, and shows how to write a GPIO application.

Here is an example for the i.MX 8M Mini board:

1. First, you need the GPIO header file, `asm-generic/gpio.h` and `dm.h`, which include all interfaces for the GPIO.
2. Then, find the corresponding GPIO description according to the name of the GPIO (such as `GPIO5_7`), configure `GPIO5_7` to OUT direction, configure `GPIO5_8` to IN direction and request it.
3. Now, by writing the value 1 or 0 to `GPIO5_7`, you can receive the same value from `GPIO5_8`.

The [Table 7](#) shows the APIs used in the file `test_gpio.c` application example.

Table 7. GPIO APIs and their description

Function declaration	Description
<code>int dm_gpio_lookup_name(const char *name, struct gpio_desc *desc)</code>	Look up a named GPIO and return its description name- Name to look up, such as <code>GPIO5_7</code> desc - GPIO description Returns: 0 if OK, -ve on error

Table 7. GPIO APIs and their description...continued

Function declaration	Description
<code>int dm_gpio_request(struct gpio_desc *desc, const char *label)</code>	Manually request a GPIO desc - GPIO description of GPIO to request label- Label to attach to the GPIO while claimed, such as "output1" Returns: 0 if OK, -ve on error
<code>int dm_gpio_set_value(const struct gpio_desc *desc, int value)</code>	Configures the direction of GPIO to OUT and writes the value to it. desc - GPIO description value- the value written to this GPIO Returns: 0 if OK, -ve on error
<code>int dm_gpio_set_dir_flags(struct gpio_desc *desc, ulong flags)</code>	Set direction using description and added flags desc - GPIO description flags - New flags to use Returns: 0 if OK, -ve on error
<code>int dm_gpio_free(struct udevice *dev, struct gpio_desc *desc)</code>	Free a single GPIO dev: Device that requested the GPIO desc - GPIO description Returns: 0 if OK, -ve on error

2.5.4.5 I2C example

The file `uboot/app/test_i2c.c` can be used as an example to test the I2C feature and shows how to write an I2C application.

On LS1043ARDB board, read a fixed data from offset 0 of INA220 device(0x40). If the data is 0x39, a message, [ok]I2C test ok is displayed on the console.

The table below shows the APIs used in the sample file, `test_i2c.c`.

Table 8. I2C APIs and their description

Function declaration	Description
<code>int i2c_set_bus_num (unsigned int bus)</code>	Sets the I2C bus. bus- bus index, zero based Returns 0 if OK, -1 on error.
<code>int i2c_read (uint8_t chip, unsigned int addr, int alen, uint8_t *buffer, int len)</code>	Read data from I2C device chip. <ul style="list-style-type: none"> chip - I2C chip address, range 0..127 addr - Memory (register) address within the chip alen - Number of bytes to use for address (typically 1, 2 for larger memories, 0 for register type devices with only one register) buffer - Where to read/write the data len - How many bytes to read/write Returns 0 if OK, not 0 on error.
<code>int i2c_write (uint8_t chip, unsigned int addr, int alen, uint8_t *buffer, int len)</code>	Writes data to I2C device chip. <ul style="list-style-type: none"> chip - I2C chip address, range 0..127 addr - Memory (register) address within the chip alen - Number of bytes to use for address (typically 1, 2 for larger memories, 0 for register type devices with only one register) buffer - Where to read/write the data len - How many bytes to read/write Returns 0 if OK, not 0 on error.

2.5.4.6 IRQ example

The file, `uboot/app/test_irq_init.c` is an example to test the IRQ and IPI (Inter-Processor Interrupts) feature, and shows how to write an IRQ application. The process is described in brief below.

The file `asm/interrupt-gic.h`, is the header file of IRQ, and includes all its interfaces. Then, register an IRQ function for SGI 0. After setting an SGI signal, the CPU gets this IRQ and runs the IRQ function. Then, register a hardware interrupt function to show how to use the external hardware interrupt.

SGI IRQ is used for inter-processor interrupts, and it can only be used between bare metal cores. In case you want to communicate between Baremetal core and Linux core, refer to [Section 2.5.4.16](#). SGI IRQ id is 0-15. The SGI IRQ id '8' is reserved for ICC.

Note: For *i.MX 8M Mini LPDDR4 EVK*, *i.MX 8M Plus LPDDR4 EVK*, *i.MX 93 EVK* and *i.MX 93 9x9 LPDDR4 QSB boards*, SGI IRQ id is 9.

The [Table 9](#) shows the APIs used in the sample file, `test_irq_init.c`.

Table 9. IRQ APIs and their description

Return type API name (parameter list)	Description
<code>int irq_desc_register(struct irq *irq_data, void (*irq_handle)(int, int, void *), void *data)</code>	Registers an IRQ function. <ul style="list-style-type: none"> <code>irq_data</code> - Include IRQ id(0-15 for SGI, 16-31 for PPI, 32-1019 for SPI), and IRQ dev <code>irq_handle</code> - IRQ function <code>data</code> - IRQ data
<code>void gic_send_sgi(u32 id, int core_mask)</code>	Sets a SGI IRQ signal. <ul style="list-style-type: none"> <code>core_mask</code> - target core mask <code>id</code> - IRQ id
<code>int irq_set_affinity(struct irq *irq, int core_mask)</code>	Sends the target core for hw IRQ. <ul style="list-style-type: none"> <code>core_mask</code> - target core mask <code>irq</code> - IRQ data, include IRQ id
<code>int irq_set_polarity(struct udevice *dev, uint id, bool active_low)</code>	Sets the type for hardware IRQ to identify whether the corresponding interrupt is edge-triggered or level-sensitive. <ul style="list-style-type: none"> <code>dev</code> - IRQ dev <code>id</code> - IRQ id <code>active_low</code> - true if active low, false for active high

2.5.4.7 QSPI example

The file `uboot/app/test_qspi.c` provides an example that can be used to test the QSPI feature. The below steps show how to write a QSPI application:

1. First, locate the QSPI header files `spi_flash.h` and `spi.h`, which include all interfaces for QSPI.
2. Then, initialize the QSPI flash. Subsequently, erase the corresponding flash area and confirm that the erase operation is successful.
3. Now, read or write to the flash with an offset of 0x3f00000 and size of 0x40000.

The [Table 10](#) shows the APIs used in the file `test_qsip.c` example.

Table 10. QSPI APIs

API name (type)	Description
<code>spi_find_bus_and_cs(bus, cs, &bus_dev, &new)</code>	The API finds if a SPI device already exists. <ul style="list-style-type: none"> "bus" - bus index, zero based. "cs" - the value to chip select mode.

Table 10. QSPI APIs...continued

API name (type)	Description
	<ul style="list-style-type: none"> • "bus_dev" - If the bus is found. • "new" - If the device is found. Returns 0 if OK, -ENODEV on error.
<code>spi_flash_probe_bus_cs(bus, cs, speed, mode, &new)</code>	Initializes the SPI flash device. <ul style="list-style-type: none"> • "bus" - bus index, zero based. • "cs" - the value to Chip Select mode. • "speed" - SPI flash speed, can use 0 or CONFIG_SF_DEFAULT_SPEED. • "mode" - SPI flash mode, can use 0 or CONFIG_SF_DEFAULT_MODE. • "new" - If the device is initialized. Returns 0 if OK, -ENODEV on error.
<code>dev_get_uclass_priv(new)</code>	Gets the SPI flash. <ul style="list-style-type: none"> • "new" - The device being initialized. Returns flash if OK, NULL on error.
<code>spi_flash_erase(flash, offset, size)</code>	Erases the specified location and length of the flash content, erases the content of all. <ul style="list-style-type: none"> • "flash" - Flash is being initialized. • "offset" - Flash offset address. • "size" - Erase the length of the data. Returns 0 if OK, !0 on error.
<code>spi_flash_read(flash, offset, len, vbuf)</code>	Reads flash data to memory. <ul style="list-style-type: none"> • "flash" - The flash being initialized. • "offset" - Flash offset address. • "len" - Read the length of the data. • "vbuf" - the buffer to store the data read Returns 0 if OK, !0 on error.
<code>spi_flash_write(flash, offset, len, buf)</code>	Writes memory data to flash. <ul style="list-style-type: none"> • "flash" - The flash being initialized. • "offset" - Flash offset address. • "len" - Write the length of the data. • "buf" - the buffer to store the data write Returns 0 if OK, !0 on error.

2.5.4.8 IFC example

Both LS1043ARDB and LS1046ARDB have IFC controller. However, LS1043ARDB supports both NOR flash and NAND flash, whereas LS1046ARDB supports only NAND flash.

NOR and NAND flash messages are displayed while booting Baremetal cores, as shown below:

```
1:NAND: 512 MiB
1:Flash: 128 MiB
```

or (LS1046ARDB)

```
1:NAND: 512 MiB
```

There is no example code to test it, but we can use a few commands to verify these features.

For LS1043ARDB NOR Flash (the map memory address is 0x60000000), the below command can be used to verify it:

```
=> md 0x60000000
1:60000000: 55aa55aa 0001ee01 10001008 0000000a .U.U.....
1:60000010: 00000000 00000000 02005514 12400080 .....U....@.
1:60000020: 005002e0 002000c1 00000000 00000000 ..P... ..
1:60000030: 00000000 00880300 00000000 01110000 .....
1:60000040: 96000000 01000000 78015709 10e00000 .....W.x....
1:60000050: 00001809 08000000 18045709 9e000000 .....W.....
1:60000060: 1c045709 9e000000 20045709 9e000000 .W.....W. ....
1:60000070: 00065709 00000000 04065709 00001060 .W.....W..`...
1:60000080: c000ee09 00440000 58015709 00220000 .....D..W.X..".
1:60000090: 40800089 01000000 40006108 f56b710a ...@.....a.@.qk.
1:600000a0: ffffffff ffffffff ffffffff ffffffff .....
```

For NAND flash on LS1043ARDB and LS1046ARDB, "nand" command can be used to verify it (nand erase, nand read, nand write, and so on.):

```
=> nand info
1:Device 0: nand0, sector size 128 KiB
1: Page size      2048 b
1: OOB size       64 b
1: Erase size     131072 b
1: subpagesize    2048 b
1: options        0x00004200
1: bbt options    0x00028000
```

```
=> nand
1:nand - NAND sub-system
1:Usage:
nand info - show available NAND devices
nand device [dev] - show or set current device
nand read - addr off|partition size
nand write - addr off|partition size
             read/write 'size' bytes starting at offset 'off'
             to/from memory address 'addr', skipping bad blocks.
nand read.raw - addr off|partition [count]
nand write.raw [.noverify] - addr off|partition [count]
             Use read.raw/write.raw to avoid ECC and access the flash as-is.
nand erase[.spread] [clean] off size - erase 'size' bytes from offset 'off'
             With '.spread', erase enough for given file size, otherwise,
             'size' includes skipped bad blocks.
nand erase.part [clean] partition - erase entire mtd partition'
nand erase.chip [clean] - erase entire chip'
nand bad - show bad blocks
nand dump[.oob] off - dump page
nand scrub [-y] off size | scrub.part partition | scrub.chip
             really clean NAND erasing bad blocks (UNSAFE)
nand markbad off [...] - mark bad block(s) at offset (UNSAFE)
nand biterr off - make a bit error at offset
```

2.5.4.9 Ethernet example

The file `uboot/app/test_net.c` provides an example to test the Ethernet feature and shows how to write a net application for using this feature.

Here is an example for the LS1043ARDB (or LS1046ARDB) board.

Note: For LS1046ARDB board, network could be assigned by setting `CONFIG_FMAN_FMAN1_COREID=<core num>` (core num could be 0 - 3). If you want to verify it on core 1, please set `CONFIG_FMAN_FMAN1_COREID=1` and compile baremetal image in standalone way to enable the network for the target core.

1. Connect one Ethernet port of LS1043ARDB board to one host machine using Ethernet cable.
 - For LS1046ARDB, the default `ethact` is FM1@DTSEC5. Network cable should be connected to SGMII1 port.
 - For LS1043ARDB, the default `ethact` is FM1@DTSEC3. Network cable should be connected to RGMII1 port.
2. Configure the IP address of the host machine as 192.168.1.2.
3. Power up the LS1043ARDB board. If the network is connected, the message `host 192.168.1.2 is alive` is displayed on the console.
4. The IP addresses of the board and host machine are defined in the file `test_net.c`. In this file, modify the IP address of LS1043ARDB board using variable `ipaddr` and change the IP address of host machine using variable `ping_ip`.

The table below lists the Net APIs and their description.

Table 11. Net APIs and their description

API name (type)	Description
<code>void net_init (void)</code>	Initializes the network
<code>int net_loop (enum proto_t protocol)</code>	Main network processing loop. <ul style="list-style-type: none"> • <code>enum proto_t protocol</code> - protocol type
<code>int eth_receive (void *packet, int length)</code>	Reads data from NIC device chip. <ul style="list-style-type: none"> • <code>void *packet</code> • <code>length</code> - Network packet length Returns length
<code>int eth_send (void *packet, int length)</code>	Writes data to NIC device chip. <ul style="list-style-type: none"> • <code>packet</code> - pointer to the packet is sent • <code>length</code> - Network packet length Returns length.

2.5.4.10 USB example

The file `uboot/app/test_usb.c` provides an example that can be used to test the USB features. The steps below show how to write a USB application:

1. Connect a USB disk to the USB port.
2. Include the header file, `usb.h`, which includes all APIs for USB.
3. Initialize the USB device using the `usb_init` API.
4. Scan the USB storage device on the USB bus using the `usb_stor_scan` API.
5. Get the device number using the `blk_get_devnum_by_type` API.
6. Read data from the USB disk using the `blk_dread` API.
7. Write data to the USB disk using the `blk_dwrite` API.

The table below shows the APIs used in the file `test_usb.c` example:

Table 12. USB APIs and their description

API name (type)	Description
<code>int usb_init(void)</code>	Initializes the USB controller.
<code>int usb_stop(void)</code>	Stops the USB controller.
<code>int usb_stor_scan(int mode)</code>	Scans the USB and reports device information to the user if <code>mode = 1</code> <ul style="list-style-type: none"> • <code>Mode</code> – if <code>mode = 1</code>, the information is returned to user. Returns <ul style="list-style-type: none"> • the current device, or • -1 (if device not found).
<code>struct blk_desc *blk_get_devnum_by_type(enum if_type if_type, int devnum)</code>	Get a block device by type and number. <ul style="list-style-type: none"> • <code>If_type</code> – Block device type • <code>devnum</code> - device number Returns <ul style="list-style-type: none"> • Points to block device descriptor, or • NULL (if not found).
<code>unsigned long blk_dread(struct blk_desc *block_dev, lbaint_t start, lbaint_t blkcnt, void *buffer);</code>	Reads data from USB device. <ul style="list-style-type: none"> • <code>block_dev</code> – block device descriptor • <code>start</code> – start block • <code>blkcnt</code> – block number • <code>buffer</code> – buffer to store the data Returns the block number from which, data is read.
<code>unsigned long blk_dwrite(struct blk_desc *block_dev, lbaint_t start, lbaint_t blkcnt, const void *buffer);</code>	Writes data to USB device. <ul style="list-style-type: none"> • <code>block_dev</code> – block device descriptor • <code>start</code> – start block • <code>blkcnt</code> – block number • <code>buffer</code> – buffer to store the data Returns the block number to which data is written.

2.5.4.11 PCIe example

The file `app/test_pcie.c` provides a sample code to test PCIe and network card (such as e1000) features. The steps below show how to write a PCIe and net application:

1. Insert a PCIe network card (such as e1000) into PCIe2, or PCIe3 slot (if it exists).
2. Configure the IP address of the host machine to 192.168.1.2.
3. Include the files `include/pci.h` and `include/netdev.h`.
4. Initialize the PCIe controller using the `pci_init` API.
5. Get `uclass` device by its name using the `uclass_get_device_by_seq` API.
6. Initialize the PCIe network device using the `pci_eth_init` API.
7. Begin pinging the host machine using the `net_loop` API.

The table below shows the APIs used in the file `test_pcie.c` example.

Table 13. PCIe APIs and their description

API name (type)	Description
<code>void pci_init(void)</code>	Initializes the PCIe controller. Does not return a value.

Table 13. PCIe APIs and their description...continued

API name (type)	Description
<code>int uclass_get_device_by_seq(enum uclass_id id, int seq, struct udevice **devp)</code>	Gets the uclass device based on an ID and sequence: <ul style="list-style-type: none"> • <code>id</code> – uclass ID • <code>seq</code> – sequence • <code>devp</code> – Pointer to device Returns: <ul style="list-style-type: none"> • 0 if Ok. • Negative value on error.
<code>static inline int pci_eth_init(bd_t *bis)</code>	Initializes network card on the PCIe bus. <ul style="list-style-type: none"> • <code>Bis</code> – struct containing variables accessed by shared code Returns the number of network cards.
<code>int net_loop (enum proto_t protocol)</code>	Main network processing loop. <ul style="list-style-type: none"> • <code>enum proto_t protocol</code> - protocol type Returns: <ul style="list-style-type: none"> • 0 if Ok. • Negative value on error.

2.5.4.12 ENETC example

The file `app/test_net.c` provides an example to test ENETC Ethernet feature and shows how to write a net application for using this feature. This example is a special case of using Net APIs.

The file `test_net` for ENETC is only an example for the LS1028ARDB board with (`CONFIG_ENETC_COREID_SET` enabled).

1. Connect ENETC port of LS1028ARDB board to one host machine using Ethernet cable.
2. Configure the IP address of the host machine as 192.168.1.2.
3. Power up the LS1028ARDB board. If the network is connected, the message `host 192.168.1.2 is alive` is displayed on the console.
4. The IP addresses of the board and host machines are defined in the file `test_net.c`. In this file, modify the IP address of LS1028ARDB board using variable `ipaddr` and change the IP address of host machine using variable `ping_ip`.

The table below lists the Net APIs for ENETC and their description, refer to [Section 2.5.4](#) for other Net APIs.

Table 14. ENETC APIs and their description

API name (type)	Description
<code>void pci_init(void)</code>	Initializes the PCIe controller. Does not return a value.
<code>void eth_initialize(void)</code>	Initializes the Ethernet.

2.5.4.13 SAI example

The audio feature needs SAI module and codec drivers. The following sections provide an introduction to SAI module and the audio codec (SGTL5000). These sections also describe the steps for integrating audio with Baremetal and running an audio application on Baremetal.

2.5.4.13.1 Synchronous Audio Interface (SAI)

The LS1028A integrates six SAI modules, but only SAI4 is used by LS1028ARDB board. The synchronous audio interface (SAI) supports full duplex serial interfaces with frame synchronization. The bit clock and frame sync of SAI are both generated externally (SGTL5000).

- Transmitter with independent bit clock and frame sync supporting 1 data line
- Receiver with independent bit clock and frame sync supporting 1 data line
- Maximum Frame Size of 32 words
- Word size of between 8-bits and 32-bits
- Word size configured separately for first word and remaining words in frame
- Asynchronous 32 × 32-bit FIFO for each transmit and receive channel
- Supports graceful restart after FIFO error




Figure 11. SAI block diagram

2.5.4.13.2 Audio codec (SGTL5000)

The SGTL5000 is a low-power stereo codec with headphone amplifier from NXP. It is designed to provide a complete audio solution for products requiring LINEIN, MIC_IN, LINEOUT, headphone-out, and digital I/Os. It allows an 8.0 MHz to 27 MHz system clock as input. The codec supports 8.0 kHz, 11.025 kHz, 12 kHz, 16 kHz, 22.05 kHz, 24 kHz, 32 kHz, 44.1 kHz, 48 kHz, and 96 kHz sampling frequencies. The LS1028ARDB board provides a 25 MHz crystal oscillator to the SGTL5000.

The SGTL5000 provides two interfaces (I2C and SPI) to setup registers. The LS1028ARDB board uses I2C interface.

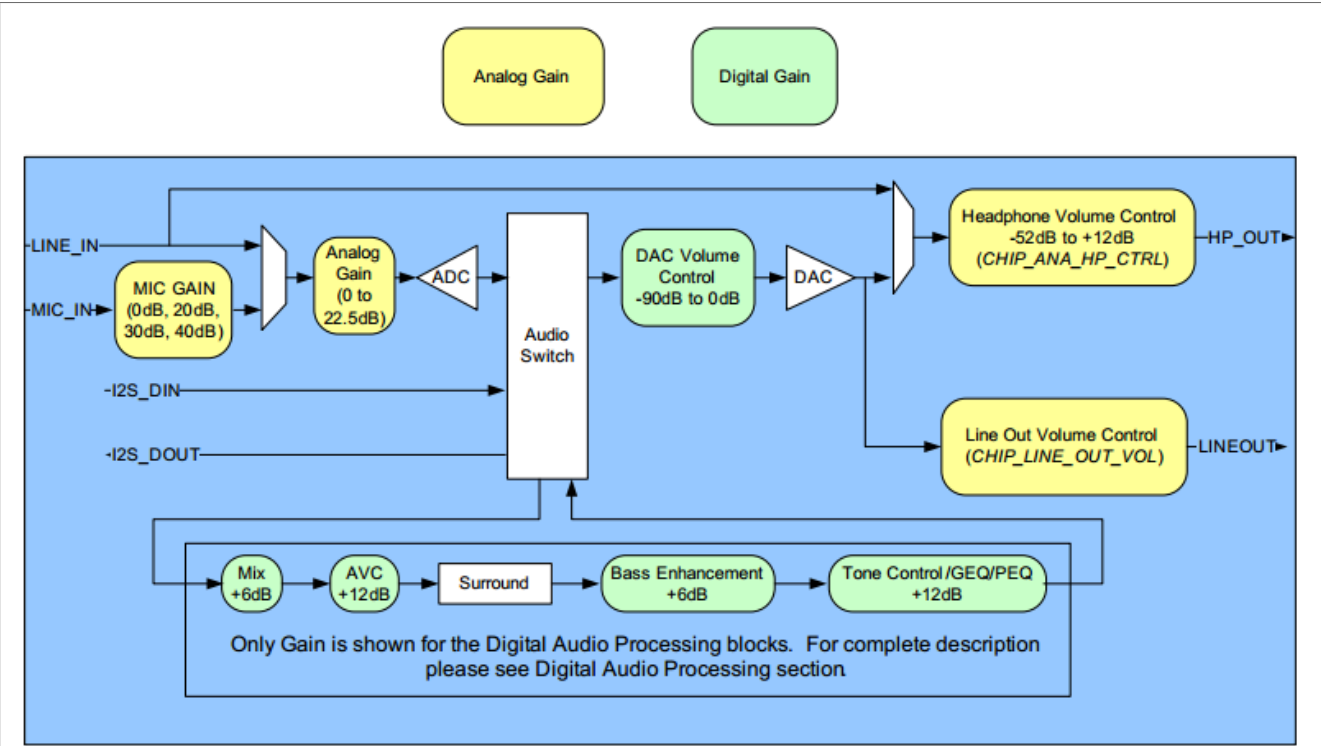


Figure 12. System block diagram, signal flow, and gain

2.5.4.13.3 Digital interface formats

The SGTL5000 provides five common digital interface formats. The SAI and SGTL5000 digital interface formats must be the same.

• I2S Format (n = bit length)

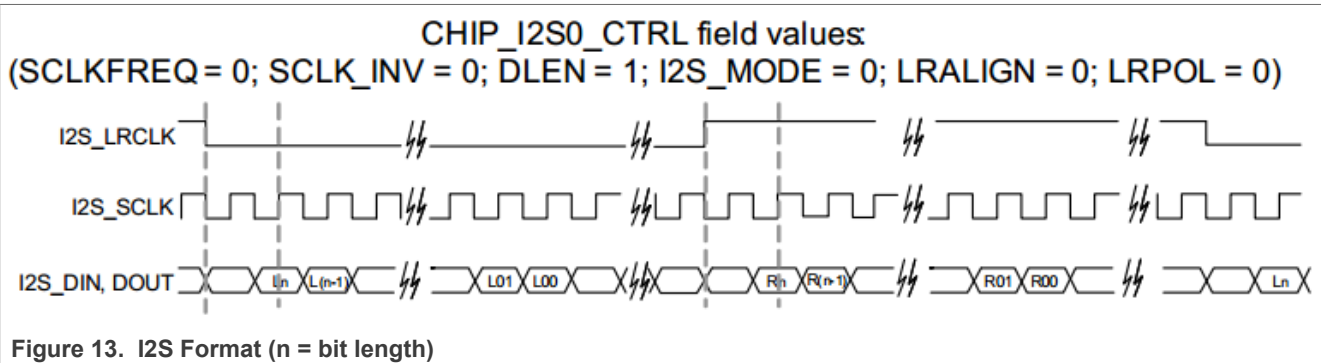
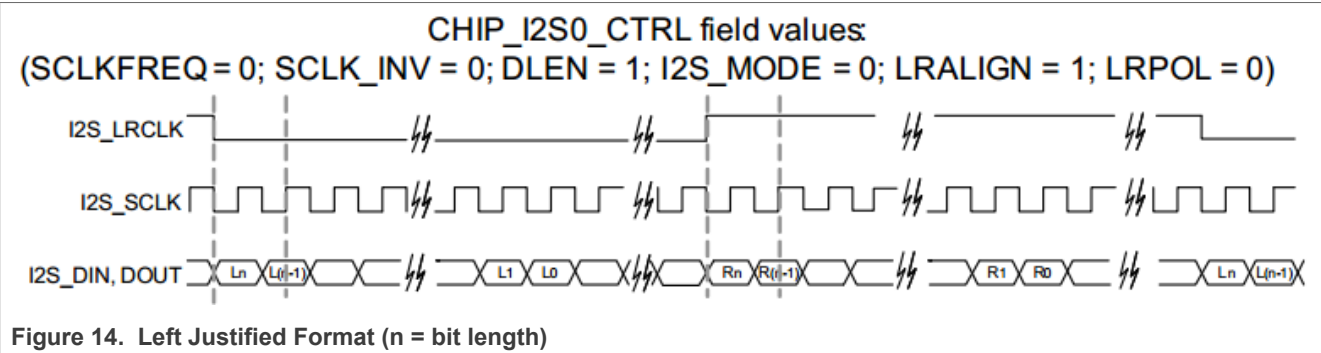
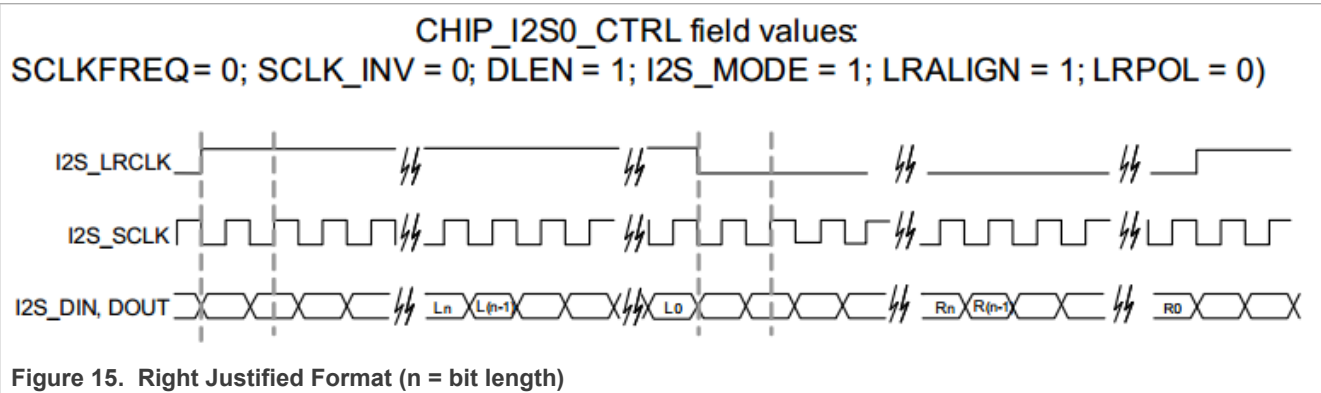


Figure 13. I2S Format (n = bit length)

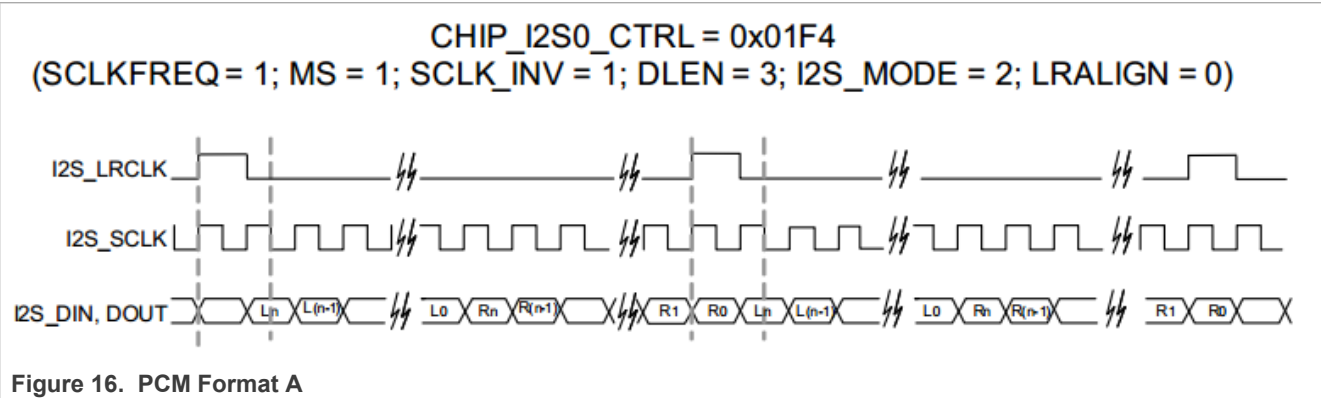
• Left Justified Format (n = bit length)



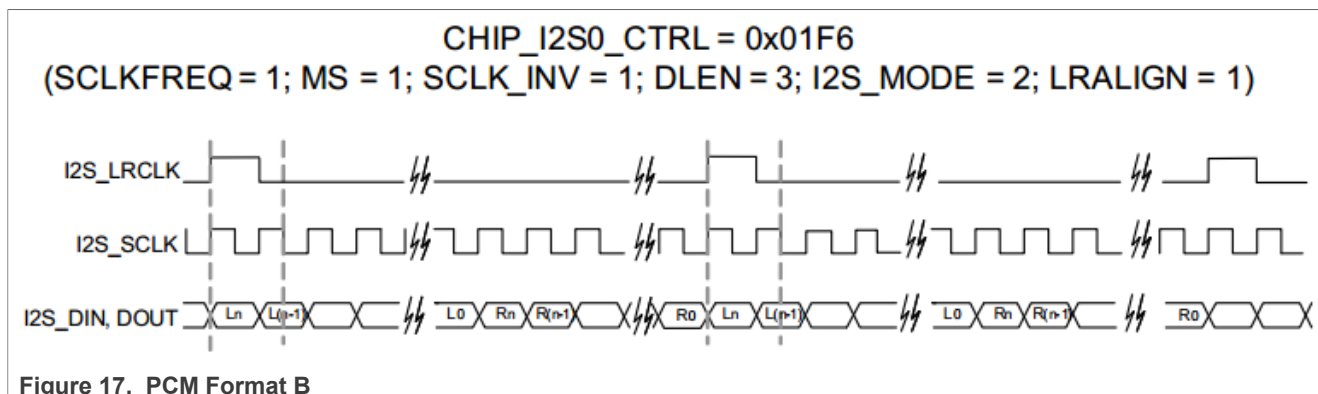
• Right Justified Format (n = bit length)



• PCM Format A



• PCM Format B



2.5.4.13.4 Running the SAI application

In order to run SAI application, Baremetal images should be rebuilt with SAI support.

1. Enable SAI support in Real-time Edge software

```
$ cd yocto-real-time-edge/sources/meta-real-time-edge
# Open file "conf/distro/include/real-time-edge-base.inc", add "sai" to
"DISTRO_FEATURES:append:ls1028ardb" like this:
DISTRO_FEATURES:append:ls1028ardb = " jailhouse real-time-edge-libbee real-time-
edge-libbblep libnfc-nci \
    wayland-protocols weston imx-gpu-viv libdrm kmscube \
    real-time-edge-sysrepo tsn-scripts wayland alsa sai"
```

2. Build the image

```
$ cd yocto-real-time-edge
$ DISTRO=nxp-real-time-edge-baremetal MACHINE=ls1028ardb source real-time-edge-
setup-env.sh -b build-ls1028ardb-bm
$ bitbake nxp-image-real-time-edge
```

3. Play a demo audio file in slave core after booting the board:

```
=> wavplayer
*****
audioformat: PCM nchannels: 1 samplerate: 16000 bitrate: 256000 blockalign: 2
bps: 16 datasize: 67968 datastart: 44
*****
sgtl5000 revision 0x11 fsl_sai_ofdata_to_platdata Probed
sound 'sound' with codec 'codec@a' and i2s 'sai@f130000'
i2s transfer tx data The music waits for the end! The music is finished!
*****
```

2.5.4.14 FlexTimer example

The FlexTimer module (FTM) works on BareMetal core as the wakeup source for LS1046ARDB. It can support nanosecond (ns) level alarm setting. There is no example code to test the module. However, a few commands can be used to verify these features.

Use the below commands to verify FTM feature:

- Use the “ftm” command to get help information:

```
=> ftm
1:ftm - ftm alarm test

1:Usage:
ftm test ftm alarm
show      - show FTM test result
start [count] us      - start FTM test
stop      - stop FTM test
```

- Use the “ftm start [count]” command to start ftm test:

```
=> ftm start
1:Use default alarm time - 5 us
1:FTM test start.

=> ftm start 100
1:FTM test start.
```

- Using “ftm stop” command to stop ftm test and show the test result:

```
=> ftm stop
1:FTM test stop.
1:irq count | total (us) | average (us) | max (us) | min (us) |
1:3087560   | 309579251 | 100         | 102      | 100       |
```

- Use “ftm show” command to show the test result:

```
=> ftm show
1:irq count | total (us) | average (us) | max (us) | min (us) |
1:317803    | 31854521  | 100          | 102      | 100       |
```

The table below lists the attributes for “ftm show” and “ftm stop” result:

Table 15. FlexTimer module attributes and their description

Attribute Name	Description
irq count	Generated interrupt single count since “ftm start” command
total (us)	The time since “ftm start” command
average (us)	The average time between two interrupt signals
max (us)	The maximum time between two interrupt signals
min (us)	The minimum time between two interrupts signals

The table below lists the FTM APIs and their description.

Table 16. FlexTimer module APIs

API Name	Description
int ftm_rtc_set_alarm_by_us (struct udevice *dev, unsigned long us, void (* func) (void *))	Setting alarm by us count <ul style="list-style-type: none"> struct udevice *dev – device struct of ftm unsigned long us – the time for ftm alarm void (* func) (void *) – the handle function when timeup

Table 16. FlexTimer module APIs...continued

API Name	Description
<code>void ftm_rtc_set_alarm (struct udevice *dev, u16 ticks, void (* func)(void *));</code>	Setting alarm by ftm timer count <ul style="list-style-type: none">• <code>struct udevice *dev</code> – device struct of ftm• <code>u16 ticks</code> – the timer counter for ftm alarm• <code>void (* func)(void *)</code> – the handle function when timeup
<code>void ftm_rtc_alarm_stop(struct udevice *dev)</code>	Stop and reset ftm alarm <ul style="list-style-type: none">• <code>struct udevice *dev</code> – device struct of ftm
<code>unsigned long ftm_rtc_get_max_alarm_us (struct udevice *dev)</code>	Get the max alarm time value for ftm alarm <ul style="list-style-type: none">• <code>struct udevice *dev</code> – device struct of ftm

2.5.4.15 Newlib’s math library

In order to control IO devices such as changing the speed or angle, mathematical calculations are required. Newlib’s math library is added to support such calculations. Newlib is a C library intended for use on embedded systems.

All math related files are under `math` folder. The file directory structure is as follows:

math

```
├── COPYING
├── include
│   └── math.h
├── lib
│   └── libm.a
└── README
```

To use math library, the below code should be in the header of the file, and then we can directly call all kinds of math APIs.

```
#include <math.h>
#undef __always_inline
#undef __section
#include <stdlib.h>
#include <common.h>
#include <command.h>
#undef log
```

For the detailed usage, refer to the example file which is `math.c` under `cmd` folder, The example shows how to call the API of math library including `acos/asin/atan/cos/sin/tan` and `log/pow/sqrt`. We can use the `math` command to verify these APIs under U-Boot command.

For example:

```
=> math
```

math - Test Math Functions

Usage:

math - Only test some simple math functions:

```
math acos x(double)
math asin x(double)
math atan x(double)
```

```
math atan2 y x(double)
math cos x(double)
math cosh x(double)
math sin x(double)
math sinh x(double)
math tanh x(double)
math exp x(double)
math ldexp x(double) exp(int)
math log x(double)
math log10 x(double)
math pow x(double) y(double)
math sqrt x(double)
math ceil x(double)
math fabs x(double)
math floor x(double)
math fmod x(double) y(double)

=> math asin 0.8
= 0.927
1:=> math sin 1.0
= 0.841
1:=> math cos 1.0
= 0.540
=> math log 10
= 2.302
1:=> math log10 10
= 1.000
```

2.5.4.16 ICC module

Inter-core communication (ICC) module works on Linux core (master) and Baremetal core (slave). It provides the data transfer between cores via SGI inter-core interrupt and shared memory blocks. It can support multicore silicon platform and transfer the data concurrently and efficiently.

ICC module structure is based on two basics:

- SGI: Software-generated Interrupts in Arm GIC, used to generate inter-core interrupts. The ICC module uses the number 8 SGI interrupt for all Linux and Baremetal cores.
- Shared memory: A memory space shared by all platform cores. The base address and size of the share memory should be defined in header files before compilation.

ICC modules can work concurrently, lock-free among multicore platform, and support broadcast case with Buffer Descriptor Ring mechanism.

The figure below shows the basic operating principle for data transfer from Core 0 to Core 1. After the data writing and head point moving to next, Core 0 triggers a SGI (8) to Core 1. After this step, the Core 1 gets the BD ring updated status and reads the new data, then moves the tail point to next.

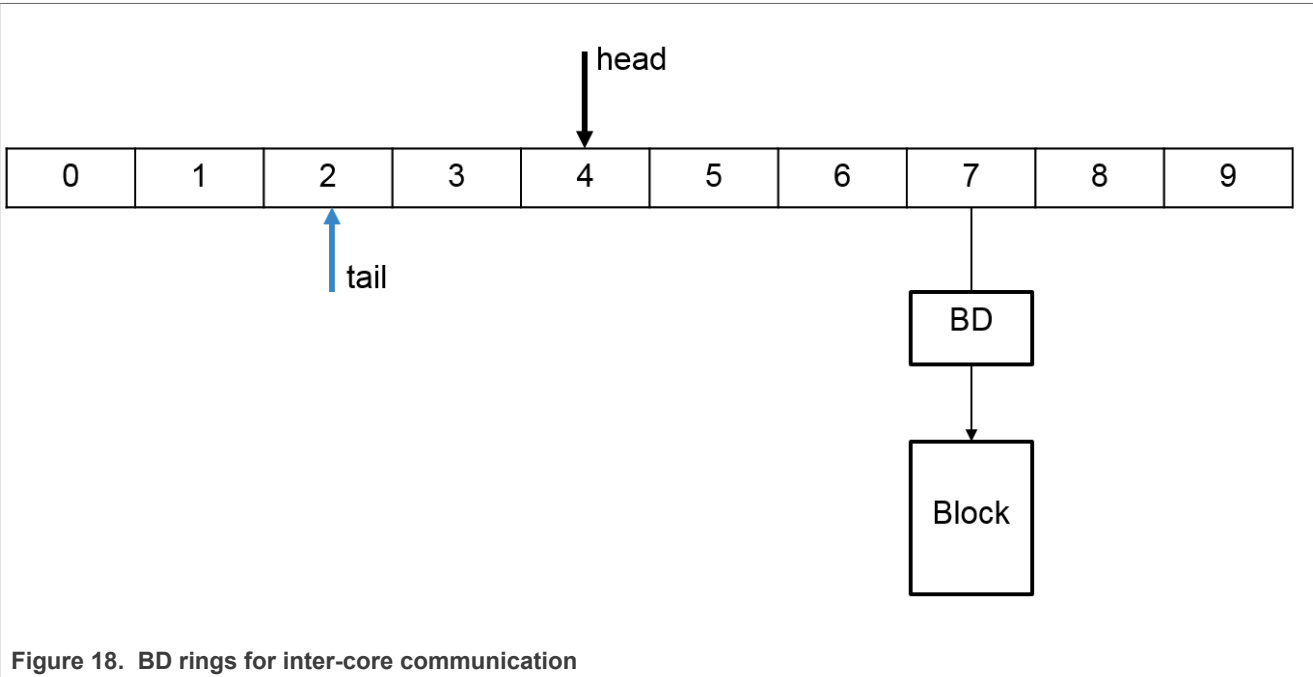


Figure 18. BD rings for inter-core communication

For a multicore platform (that is, four cores), the total BD rings are arranged as shown in the following figure. (See the BD rings on Core 0 and Core 1.)




Figure 19. BD rings for multicore platform

All the ICC ring structures, BD structures, and blocks for data are in the shared memory. A four-core platform ICC module would map the shared memory as shown in the figure below.

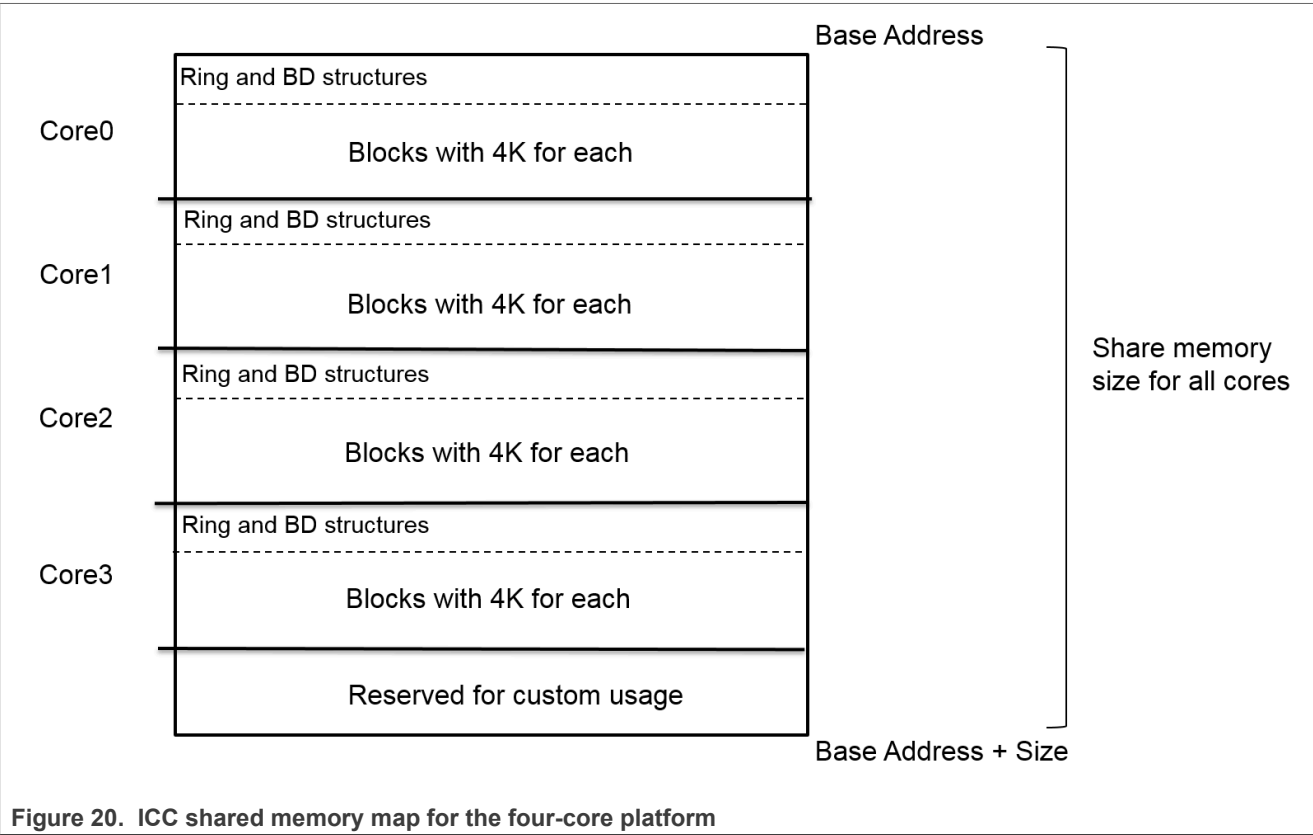


Figure 20. ICC shared memory map for the four-core platform

Generally, Core 0 runs Linux as master core while other cores run Baremetal as slaves. They obtain the same size of share memory to structure the rings and BDs, and split the blocks space with 4k unit for each block. The reserved space at the top of the share memory is out of the ICC module and for the custom usage.

For LS1028ARDB platform with two cores, the shared memory map is defined as:

- The total shared memory size is 256 MB.
- The reserved space for custom usage is 16 MB at the top of the share memory space.
- Core 0 runs Linux as master core, the share memory size for ICC is 120 MB, in which the ring and BD structure space is 2 M, and the block space for data is 118 MB with 4K for each block.
- Core 1 runs Baremetal as slave core, the share memory size for ICC is 120 MB, in which the ring and BD structure space is 2M, and the block space for data is 118 MB with 4K for each block.

The ICC module includes two parts of the code:

- ICC code for Linux user space, works for data transfer between master core and slave cores. The code is integrated into the Real-time Edge software and named `real-time-edge-icc`. After the compilation, the `icc` binary is put into the Linux file system.
- ICC code for Baremetal runs on every slave core, works for data transfer between Baremetal cores and master core.

The ICC code for Linux user space in the repository: <https://github.com/nxp-real-time-edge-sw/real-time-edge-icc.git>.

- `icc-main.c` --- the example case commands
- `inter-core-comm.c`
- `inter-core-comm.h` --- include the header file to use ICC module

└─ Makefile

The ICC code for Baremetal in `baremetal` directory:

`baremetal/`

└─ `arch/arm/lib/inter-core-comm.c`

└─ `arch/arm/include/asm/inter-core-comm.h` --- includes the header file to use ICC module

└─ `cmd/icc.c` --- the example case commands

The ICC modules of the APIs are exported out for usage in both Linux user space and Baremetal code.

Table 17. ICC APIs

APIs	Description
<code>unsigned long icc_ring_state(int coreid)</code>	Checks the ring and block state. Returns: <ul style="list-style-type: none"> 0 - if empty. !0 - the working block address currently.
<code>Unsigned long icc_block_request(void)</code>	Requests a block, which is <code>ICC_BLOCK_UNIT_SIZE</code> size. Returns: <ul style="list-style-type: none"> 0 - failed. !0 - block address can be used.
<code>void icc_block_free(unsigned long block)</code>	Frees a block requested. Be careful if the destination cores are working on this block.
<code>int icc_irq_register(int src_coreid, void (*irq_handle)(int, unsigned long, unsigned int))</code>	Registers ICC callback handler for received data. Returns: <ul style="list-style-type: none"> 0 - on success -1 - if failed.
<code>int icc_set_block(int core_mask, unsigned int byte_count, unsigned long block)</code>	Sends the data in the block to a core or multicore. This triggers the SGI interrupt. Returns: <ul style="list-style-type: none"> 0 - on success -1 - if failed.
<code>void icc_show(void)</code>	Shows the ICC basic information.
<code>int icc_init(void)</code>	Initializes the ICC module.

2.5.4.16.1 ICC examples

This section provides example commands for use cases in both Linux user space and Baremetal code. They can be used to check and verify the ICC module conveniently.

1. In Linux user space, use the command `icc` to display the supported cases.

```
[root@LS1046ARDB ~] # icc
icc show - Shows all icc rings status at this core
icc perf <core_mask> <counts> - ICC performance to cores <core_mask> with
<counts> bytes
icc send <core_mask> <data> <counts> - Sends <counts> <data> to cores
<core_mask>
icc irq <core_mask> <irq> - Sends SGI <irq> ID[0 - 15] to <core_mask>
icc read <addr> <counts> - Reads <counts> 32bit register from <addr>
icc write <addr> <data> - Writes <data> to a register <addr>
```


Likewise, in Baremetal system, use the command `icc` to view the supported cases.

```
=> icc
1:icc - Inter-core communication via SGI interrupt
1:Usage:
icc show                      - Show all icc rings status at this
    core
icc perf <core_mask> <counts> - ICC performance to cores
    <core_mask> with <counts> bytes
icc send <core_mask> <data> <counts> - Send <counts> <data> to cores
    <core_mask>
icc irq <core_mask> <irq>      - Send SGI <irq> ID[0 - 15] to
    <core_mask>
```

2. The ICC module command examples on LS1046ARDB with Linux (Core 0) + Baremetal (Core 1, 2, 3) system:

Run `icc send 0x2 0x55 128` to send 128 bytes data 0x55 to core 1.

```
[root@LS1046ARDB ~] # icc send 0x2 0x55 128
gic_base: 0xfffffa033f000, share_base: 0xffff9133f000, share_phy: 0xd0000000,
  block_phy: 0xd0200000
ICC send testing ...
Target cores: 0x2, bytes: 128
ICC send: 128 bytes to 0x2 cores success
all cores: reserved share_memory_base: 0xdf000000; size: 16777216
mycoreid: 0; ICC_SGI: 8; share_memory_size: 62914560
block_unit_size: 4096; block_number: 14848; block_idx: 0
#ring_0 base: 0xffff9133f000; dest_core: 0; SGI: 8
desc_num: 128; desc_base: 0xd00000c0; head: 0; tail: 0
busy_counts: 0; interrupt_counts: 0
#ring_1 base: 0xffff9133f030; dest_core: 1; SGI: 8
desc_num: 128; desc_base: 0xd00008c0; head: 1; tail: 1
busy_counts: 0; interrupt_counts: 1
#ring_2 base: 0xffff9133f060; dest_core: 2; SGI: 8
desc_num: 128; desc_base: 0xd00010c0; head: 0; tail: 0
busy_counts: 0; interrupt_counts: 0
#ring_3 base: 0xffff9133f090; dest_core: 3; SGI: 8
desc_num: 128; desc_base: 0xd00018c0; head: 0; tail: 0
busy_counts: 0; interrupt_counts: 0
```

At the same time, Core 1 displays the received information.

```
=> 1:Get the ICC from core 0; block: 0xd0200000, bytes: 128, value: 0x55
```

3. ICC command run on Baremetal side

```
=> icc send 0x1 0xaa 128
1:ICC send testing ...
1:Target cores: 0x1, bytes: 128
1:ICC send: 128 bytes to 0x1 cores success
1:all cores: reserved share_memory_base: 0xdf000000; size: 16777216
1:mycoreid: 1; ICC_SGI: 8; share_memory_size: 62914560
1:block_unit_size: 4096; block_number: 14848; block_idx: 0
1:#ring_0 base: 00000000d3c00000; dest_core: 0; SGI: 8
1:desc_num: 128; desc_base: 00000000d3c000c0; head: 1; tail: 1
1:busy_counts: 0; interrupt_counts: 1
1:#ring_1 base: 00000000d3c00030; dest_core: 1; SGI: 8
1:desc_num: 128; desc_base: 00000000d3c008c0; head: 0; tail: 0
1:busy_counts: 0; interrupt_counts: 0
1:#ring_2 base: 00000000d3c00060; dest_core: 2; SGI: 8
1:desc_num: 128; desc_base: 00000000d3c010c0; head: 0; tail: 0
1:busy_counts: 0; interrupt_counts: 0
```

```
1:#ring 3 base: 00000000d3c00090; dest_core: 3; SGI: 8
1:desc_num: 128; desc_base: 00000000d3c018c0; head: 0; tail: 0
1:busy_counts: 0; interrupt_counts: 0
```

Then, Core 0 side (Linux) receives this data:

```
[root@LS1046ARDB ~] # [ 4247.733753] 000: Get the ICC from core 1; block:
0xd3e00000, bytes: 128, value: 0xaa
```

2.5.4.17 Single hardware interrupt routed to multiple cores

This section describes how to use GPIO to simulate external interrupt to notify all slave cores. With this feature, all the slave cores can be triggered to perform operations almost at the same time via a single hardware interrupt.

Two GPIO pins are selected. One pin is used to output 0 and 1 to simulate an external hardware. The other one is used as an interrupt pin to trigger an interrupt to a core under pull-down mode. When the core receives the interrupt, it triggers other slave cores via ICC SGI interrupt.

This feature is supported on LS1046ARDB and i.MX 8M Mini. The GPIO interrupt number and two GPIO pins required for the interrupt test can be obtained from the corresponding `dtb` file.

In `fsl-ls1046a-rdb.dts`

```
gpio_int {
    compatible = "fsl,gpio-int";
    gpios = <&gpio1 1 0>,
           <&gpio1 2 0>;
    interrupts = <0 99 0>;
};
```

In `imx8mm-evk.dts`

```
gpio_int {
    compatible = "fsl,gpio-int";
    gpios = <&gpio5 7 0>,
           <&gpio5 8 0>;
    interrupts = <0 104 0>;
};
```

On LS1046ARDB, GPIO2_01 and GPIO2_02 are selected. GPIO2_01 is used to simulate an external hardware whereas GPIO2_02 triggers an interrupt. Connect the GPIO2_01 and GPIO2_02 pins on the board. TP14 and TP13 are connected to GPIO2_01 and GPIO2_02 separately. The figure below shows how to connect TP14 and TP13.




Figure 21. LS1046ARDB hardware interrupt routing to multiple cores

On i.MX 8M Mini, connect the pins GPIO5_07 and GPIO5_08 on the board. GPIO5_07 is used to simulate an external hardware whereas GPIO5_08 triggers an interrupt. The figure below shows how to connect GPIO5_07 and GPIO5_08.




Figure 22. i.MX 8M Mini hardware interrupt routing to multiple cores

On LS1046ARDB, GPIO2_01 and GPIO2_02 are multiplexed with SPI_CS_B[0] and SDHC_DAT[4] signals. User must configure RCW[382 ~ 383] to 0b'10 to enable GPIO2[0] signal.

Since GPIO2 is assigned to Baremetal core, Linux should not use it again. We can disable GPIO2 under Linux via dts file. The below code should be added in Linux kernel file fsl-ls1046a-rdb-sdk-bm.dts to disable GPIO2.

```
&gpio1 {
    status = "disabled";
};
```

Table 18. GPIO_INT driver APIs and their description

Function declaration	Description
<pre>int gpio_request_by_name(struct udevice *dev, const char *list_name, int index, struct gpio_desc *desc, int flags)</pre>	<p>Locate and request a GPIO by name</p> <p>dev- Device requesting the GPIO</p> <p>index - Index number of the GPIO in that list use request (0 = first)</p> <p>desc - Returns GPIO description information</p> <p>flags - Indicates the GPIO input/output settings</p> <p>Returns: 0 if OK, -ve on error</p>

Table 18. GPIO_INT driver APIs and their description...continued

Function declaration	Description
<code>int dm_gpio_set_value(const struct gpio_desc *desc, int value)</code>	Configures the direction of GPIO to OUT and writes the value to it. desc - GPIO description value- the value written to this GPIO Returns: 0 if OK, -ve on error
<code>int dm_gpio_set_interrupt(const struct gpio_desc *desc)</code>	Enable GPIO interrupt desc - GPIO description Returns: 0 if OK, -ve on error
<code>int dm_gpio_clr_interrupt(const struct gpio_desc *desc)</code>	Disable GPIO interrupt desc - GPIO description Returns: 0 if OK, -ve on error

Under Baremetal, `gpio_interrupt` command provides “*enable*”, “*start*” and “*stop*” commands to control these two pins.

- **gpio_interrupt enable** - Initializes the `gpio_int` driver and enables the interrupt.
- **gpio_interrupt start** - Sets GPIO2_01 to high.
- **gpio_interrupt stop** - Sets GPIO2_01 to low.

`gpio_interrupt` can run under BareMetal console.

The `gpio_interrupt enable` command should be run first to initialize two pins. Then, use the command pair `gpio_interrupt start` and `gpio_interrupt stop` to pull high and pull down GPIO2_01/GPIO5_07. After the command pair, GPIO2_02/GPIO5_08 triggers an interrupt when getting pull-down signal. The core1 sends the SGI interrupt to other slave cores. The time of GPIO interrupt and SGI interrupt is dumped by each slave core. The latency is the time difference between GPIO interrupt and SGI interrupt. The below example shows the latency is about 1 μ s. It means all slave cores could be triggered within 1 μ s.

```
=> gpio_interrupt enable
=> gpio_interrupt start
=> gpio_interrupt stop
1:Time(us): 0x33cc582
3:Time(us): 0x33cc585, Get the SGI from CoreID: 1
2:Time(us): 0x33cc584, Get the SGI from CoreID: 1
=>
```

2.5.4.18 Hardware resource allocation

This section describes how to modify the hardware resource allocation depending on the application and used reference design board.

2.5.4.18.1 LS1028ARDB board

This section describes the hardware resource allocation for LS1028A reference design board for ENETC, I2C, and SAI configuration settings.

2.5.4.18.1.1 ENETC

LS1028ARDB has only one ENETC controller in use, which is assigned to core1 as the default setting. The controller can be reconfigured by using the command, `make menuconfig`.

See the following:

```
ARM architecture --->
[*] Enable baremetal
[*] Enable ENETC for baremetal
    (1) Enetc1 is assigned to core1
    (1) ENETC Controller numbers
```

2.5.4.18.1.2 I2C

This section describes how to configure the I2C bus on LS1028A reference design boards.

LS1028ARDB has eight I2C controllers, but only controller 0 is used for I2C devices. For example, RTC, Thermal Monitor, and Linux (core 0) use this controller for some features (for example, RTC). Therefore, the code below just shows how to enable I2C on Baremetal side.

Note:

Operate the I2C devices in Baremetal side CAREFULLY.

```
#define CONFIG_SYS_I2C_MXC_I2C1 /* enable I2C bus 0 */
#define CONFIG_SYS_I2C_MXC_I2C2 /* enable I2C bus 1 */
#define CONFIG_SYS_I2C_MXC_I2C3 /* enable I2C bus 2 */
#define CONFIG_SYS_I2C_MXC_I2C4 /* enable I2C bus 3 */
#define CONFIG_I2C_BUS_CORE_ID SET
#define CONFIG_SYS_I2C_MXC_I2C0_COREID 1
```

The `CONFIG_SYS_I2C_MXC_I2C0_COREID` defines the slave core that runs the I2C bus.

Since I2C is enabled in DM mode on Baremetal side, there is no automatic code to test it. Follow the below steps to read RTC (0x51 address, is on bus 2) on Baremetal side:

```
=> i2c bus
Bus 0: i2c@20000000 (active 0)
    77: i2c-mux@77, offset len 1, flags 0
    57: generic 57, offset len 1, flags 0
Bus 1: i2c@20000000->i2c-mux@77->i2c@1
Bus 2: i2c@20000000->i2c-mux@77->i2c@3
    51: rtc@51, offset len 1, flags 0
Bus 3: i2c@20100000
Bus 4: i2c@20200000
Bus 5: i2c@20300000
Bus 6: i2c@20400000
```

```
Bus 7: i2c@2050000
Bus 8: i2c@2060000
Bus 9: i2c@2070000
=> i2c md 0x51 0
Error reading the chip: -121
=> i2c dev 2
Setting bus to 2
=> i2c md 0x51 0
0000: 04 00 36 03 12 15 02 12 20 80 80 80 80 80 00 c2    ..6..... .....
```

2.5.4.18.1.3 SAI

LS1028ARDB has only one SAI module in use, which is assigned to core1 in the default setting. The SAI module can be reconfigured by using the command, `make menuconfig`.

See the following:

```
Command line interface --->
  Misc commands --->
    [*] wavplayer
  Device Drivers --->
    Sound support --->
      [*] Enable sound support
      [*]   Enable I2S support
      [*] Freescale sound
      [*] Freescale sgtl5000 audio codec
      [*] Freescale SAI module
```

2.5.4.18.2 LS1043ARDB or LS1046ARDB board

The following sections describe the hardware resource allocation for the LS1043ARDB or LS1046ARDB boards for implementing the supported features.

2.5.4.18.2.1 Linux DTS

Remove cpu1, cpu2, cpu3 nodes on DTS, and remove all the devices that bare metal has used.

2.5.4.18.2.2 Memory configuration

This section describes the memory configuration for LS1043ARDB or LS1046ARDB boards.

The LS1043ARDB or LS1046ARDB boards have a DDR of size 2 GB. To use the bare metal framework, configure DDR into three partitions:

- 512M for core0 (Linux)
- 256M for core1 (bare metal)
- 256M for core2 (bare metal)
- 256M for core3 (bare metal), and 256M for shared memory.

The configuration can be defined in the file `include/configs/ls1043a_baremetal.h`.

```
#define CFG_BAREMETAL_SYS_SDRAM_MASTER_SIZE (512 * 1024 * 1024UL)
#define CFG_BAREMETAL_SYS_SDRAM_SLAVE_SIZE (256 * 1024 * 1024UL)
#define CFG_BAREMETAL_SYS_SDRAM_RESERVE_SIZE (16 * 1024 * 1024UL)
#define CFG_BAREMETAL_SYS_SDRAM_SHARE_SIZE \
    ((256 * 1024 * 1024UL) - CFG_BAREMETAL_SYS_SDRAM_RESERVE_SIZE)
```

Note: The memory configuration must be consistent with the U-Boot configuration of core0.

The memory configuration for bare metal is shown in the figure below.

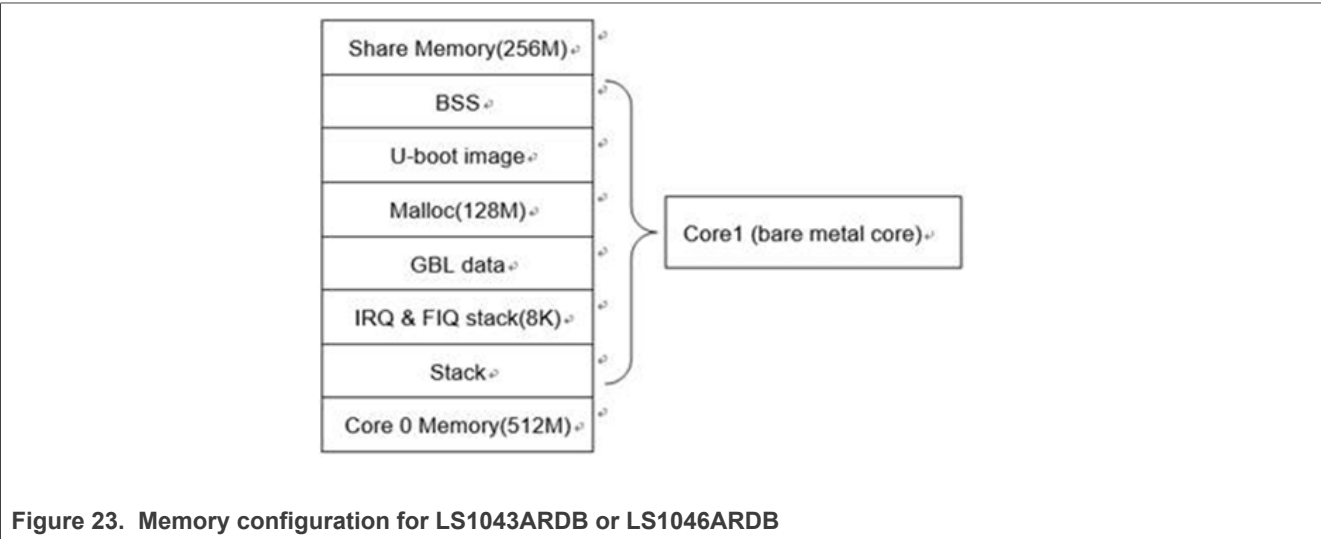


Figure 23. Memory configuration for LS1043ARDB or LS1046ARDB

The functions included in `malloc.h` in the table below can be used to allocate or free memory in program. Modify `CONFIG_SYS_MALLOC_LEN` in `defconfig` of the board to change the maximum size of `malloc`.

Table 19. Memory APIs description

API name (type)	Description
<code>void* malloc (size_t n)</code>	Allocates memory <ul style="list-style-type: none"> • “n” – length of allocated chunk • Returns a pointer to the newly allocated chunk
<code>void free (void *ptr)</code>	Releases the chunk of memory pointed to by ptr (where “ptr” is a pointer to the chunk of memory)

The GPIO for LS1043ARDB (or LS1046ARDB) has four GPIO controllers. You need to add a GPIO node in the file `ls1043/6a-rdb.dts` to assign a GPIO resource to different cores. The configuration can be done in the file `arch/arm/dts/fsl-ls1043/6a-rdb.dts`.

2.5.4.18.2.3 GPIO

LS1043 and LS1046A have four GPIO controllers. You can add a GPIO node in the file `ls1043-rdb.dts` or `ls1046a-rdb.dts` to assign a GPIO resource to different cores. The configuration is in `arch/arm/dts/fsl-ls1043a-rdb.dts` / `arch/arm/dts/fsl-ls1046a-rdb.dts`. Use the command below to add a GPIO node:

```
&gpio2 {
    status = "okay";
};
```

2.5.4.18.2.4 I2C

This section describes how to configure the I2C bus on LS1028A, LS1043A, or LS1046A reference design boards.

The LS1043ARDB (or LS1028ARDB / LS1046ARDB) has four I2C controllers. You can configure the I2C bus using the `ls1043ardb_bm_defconfig` file using the commands below:

```
CONFIG_SYS_I2C_MXC_I2C1=y
CONFIG_SYS_I2C_MXC_I2C2=y
CONFIG_SYS_I2C_MXC_I2C3=y
CONFIG_SYS_I2C_MXC_I2C4=y
CONFIG_I2C_COREID_SET=y
CONFIG_SYS_I2C_MXC_I2C0_COREID=1
CONFIG_SYS_I2C_MXC_I2C1_COREID=2
CONFIG_SYS_I2C_MXC_I2C2_COREID=3
CONFIG_SYS_I2C_MXC_I2C3_COREID=1
```

The `CONFIG_SYS_I2C_MXC_I2C0_COREID` defines the slave core that runs the I2C bus.

2.5.4.18.2.5 Hardware interrupts

LS1043A has twelve IRQs as external IO signals connected to interrupt the controller. These twelve IRQs can be used on baremetal cores. The ids for these signals, IRQ0-IRQ11 are: 163, 164, 165, 167, 168, 169, 177, 178, 179, 181, 182, and 183. GIC interrupt APIs are defined in `asm/interrupt-gic.h`. The following example shows how to register a hardware interrupt:

```
//register HW interrupt
int irq_desc_register(struct irq *irq_data, void (*irq_handle)(int, int, void
*), void *data);
int irq_set_polarity(struct udevice *dev, uint irq, bool active_low);
```

```
int irq_set_affinity(struct irq *irq, int core_mask);
```

2.5.4.18.2.6 QSPI

LS1046ARDB has a QSPI flash device. To configure the QSPI on ls1046ardb_config.h, use the command below:

```
#define CONFIG_FSL_QSPI_COREID 1
```

Here, the CONFIG_FSL_QSPI_COREID defines the slave core that runs this QSPI.

2.5.4.18.2.7 IFC

LS1043A and LS1046A have IFC controller. LS1043RDB supports both NOR flash and NAND flash, whereas LS1046RDB supports only NAND flash.

1. IFC is disabled in Linux kernel via disabling "ifc" node:

```
&ifc {
    status = "disabled";
};
```

2. Enter the Baremetal-Framework directory path and then execute the commands below: (IFC is enabled by default)

```
make menuconfig ARM architecture ---> [*] Enable baremetal [*] Enable IFC for
baremetal (1) IFC is assigned to that core
```

2.5.4.18.2.8 Ethernet

This section describes the Ethernet configuration settings for LS1043A or LS1046A reference design boards.

LS1043A or LS1046A has only one FMan, so you should remove the DPAA driver in Linux.

1. Disable the DPAA driver in Linux kernel:

```
Device Drivers --->
  Staging drivers--->
    < > Freescale Datapath Queue and Buffer management
```

2. Enter the Baremetal-Framework directory and then execute the commands below:

```
make menuconfig ARM architecture ---> [*] Enable baremetal [*] Enable fman
for baremetal (1) FMAN1 is assigned to that core
```

Configure FMan to the specified core by modifying the FMan1 is assigned to that core value, which is the default configuration, to core1.

2.5.4.18.2.9 USB

This section describes the USB configuration setting for LS1043A and LS1046A reference design boards.

Both LS1043A and LS1046A have three DW3 USB controllers. By default, these are assigned as core1, core2, and core3. Users can reconfigure the controllers by using the 'make menuconfig' command as shown below.

```
ARM architecture --->
[*] Enable baremetal
[*] Enable USB for baremetal
```

```
(1)    USB0 is assigned to core1
(2)    USB1 is assigned to core2
(3)    USB2 is assigned to core3
(3)    USB Controller numbers
```

2.5.4.18.2.10 PCI Express (PCIe)

This section describes the PCIe configuration setting for LS1043A and LS1046A reference design boards.

Both LS1043A and LS1046A have three PCIe controllers. By default, these are assigned as core0, core1, and core2. To reconfigure them, use the command 'make menuconfig', as shown below:

```
ARM architecture  --->
[*] Enable baremetal
(0)    PCIe1 is assigned to core0
(1)    PCIe2 is assigned to core1
(2)    PCIe3 is assigned to core2
(3)    PCIe Controller numbers
```

2.5.4.18.3 LX2160ARDB board

The following sections describe the hardware resource allocation for the LX2160ARDB boards for implementing the supported features.

2.5.4.18.3.1 Memory configuration

This section describes the memory configuration for LX2160ARDB boards.

The LX2160ARDB boards have a 16 GB size DDR. To use the Baremetal framework, configure DDR into three partitions:

- 15G for core0 (Linux)
- 64M per core from core1 to core15 (baremetal), and 64M for shared memory.

The configuration can be defined in the file `include/configs/lx2160ardb_config.h`.

```
#define CFG_BAREMETAL_SYS_SDRAM_MASTER_SIZE      (512 * 1024 * 1024UL)
#define CFG_BAREMETAL_SYS_SDRAM_SLAVE_SIZE      (64 * 1024 * 1024UL)
#define CFG_BAREMETAL_SYS_SDRAM_RESERVE_SIZE    (16 * 1024 * 1024UL)
#define CFG_BAREMETAL_SYS_SDRAM_SHARE_SIZE \
    ((64 * 1024 * 1024UL) - CFG_BAREMETAL_SYS_SDRAM_RESERVE_SIZE)
```

The functions included in `malloc.h` in the table below can be used to allocate or free memory in program. Modify `CONFIG_SYS_MALLOC_LEN` in `include/configs/lx2160ardb.h` to change the maximum size of `malloc`.

Table 20. Memory API description

API name (type)	Description
<code>void_t* malloc (size_t n)</code>	Allocates memory <ul style="list-style-type: none">• “n” – length of allocated chunk• Returns a pointer to the newly allocated chunk
<code>void free (void *ptr)</code>	Releases the chunk of memory pointed to by ptr (where “ptr” is a pointer to the chunk of memory)

2.5.4.18.4 i.MX 8M Mini LPDDR4 EVK and i.MX 8M Plus LPDDR4 EVK boards

The following sections describe the hardware resource allocation for the i.MX 8M Mini LPDDR4 EVK and i.MX 8M Plus LPDDR4 EVK boards for implementing the supported features.

2.5.4.18.4.1 Linux DTS

When using Baremetal, users should remove all the devices from kernel that Baremetal has used, for example:

```
&fec1 {
    status = "disabled";
};
&gpio5 {
    status = "disabled";
};
&uart3 {
    status = "disabled";
};
```

2.5.4.18.4.2 Memory configuration

This section describes the memory configuration for i.MX 8M Mini LPDDR4 EVK or i.MX 8M Plus LPDDR4 EVK boards.

1. The boards have a 6 GB DDR memory. To use the Baremetal framework, configure DDR into five partitions:

- 6016M for core0 (Linux)
- 32M for core1 (bare metal)
- 32M for core2 (bare metal)
- 32M for core3 (bare metal)
- 32M for shared memory.

The configuration can be defined in the file `include/configs/imx8mm_baremetal.h`. or `include/configs/imx8mp_baremetal.h`.

```
#define CFG_BAREMETAL_SYS_SDRAM_SLAVE_SIZE    (32 * 1024 * 1024UL)
#define CFG_BAREMETAL_SYS_SDRAM_RESERVE_SIZE  (4 * 1024 * 1024UL)
```

2. Memory Reserve

For IPI data transfer, Baremetal needs to share memory between master core and slave core. Hence, users should reserve some memory from the Linux kernel, as shown in the following `dtb` file:

```
reserved-memory { #address-cells = <2>; #size-cells = <2>; ranges; bm_reserved:
    baremetal@0x600000000 { no-map; reg = <0 0x600000000 0 0x100000000>; }; };
```

2.5.4.18.4.3 GPIO

1. Connect pin7 and pin8 of J1003. The `test_gpio` case in Baremetal uses pin7 and pin8 of J1003, so connect these two pins.

2. Boot the Baremetal on slave core. If the GPIO is working fine, the message below is displayed:

```
[ok]GPIO test ok
```

3. Disable the devices from kernel.

For the test_gpio case, use GPIO5_7 (pin8 of J1003) and GPIO5_8 (pin7 of J1003). These two pins are muxed as UART3_TXD and UART3_CTS, so should disable GPIO5 and UART3 from kernel.

```
&gpio5 { status = "disabled"; }; &uart3 { status = "disabled"; };
```

2.5.4.18.4.4 Ethernet

This section describes the Ethernet configuration settings for i.MX 8M Mini LPDDR4 EVK or i.MX 8M Plus LPDDR4 EVK boards.

1. Disable the Ethernet card from dts files:

```
&fec1 {  
    status = "disabled";  
};
```

Note:

1. i.MX 8M Mini LPDDR4 EVK has only one NIC, default status of eth0(fec1) is disabled. if user does not use eth0 in Baremetal, can enable fec1 in kernel dts file.
2. i.MX 8M Plus LPDDR4 EVK has two NICs, default setting is eth0 for Baremetal, eth1 for Linux.

2. Confirm Baremetal configuration using the command below:

```
make menuconfig ARM architecture ---> [*] Enable baremetal [*] Enable NIC for  
baremetal (1) which core that NIC is assigned to
```

Configure NIC to the specified core by modifying the NIC to assign that core value, which is the default configuration, to core1.

2.6 Native RTOS on Cortex-A core

Native RTOS refer to the RTOS running without using Hypervisor (Jailhouse), it is kicked to specified Cortex-A Core by U-Boot commands or remoteproc under Linux.

2.6.1 Overview

Real-time Edge system supports Native RTOS including FreeRTOS and Zephyr running on Cortex-A cores.

Currently there are two modes in which RTOS can run on Cortex-A Cores:

- Jailhouse RTOS: leverage Jailhouse Hypervisor to run RTOS in Jailhouse inmate cells
- Native RTOS: running on Cortex-A Core without any Hypervisor, similar with BareMetal mode.

Jailhouse Hypervisor provides a mechanism to isolate hardware resources, such as memory and peripherals. However, hypervisor implementation causes Cortex-A core's privileged execution level switching between EL1 and EL2 at runtime, so it introduces extra real-time latency.

Native RTOS runs on Cortex-A Core directly similar to how the Linux kernel runs, without any Hypervisor. Therefore, native RTOS it has a better real-time performance in comparison with RTOS in Jailhouse inmate. Native RTOS is targeted for high real-time performance use cases with less real-time latency.

In order to run Native RTOS on Cortex-A Core, hardware resources used by each OS must be defined to ensure that there is no resource conflict between different operating systems. If Linux is used with RTOS simultaneously, the device nodes for peripherals used by RTOS should be disabled or removed from the Linux device tree.

2.6.2 Building Native RTOS on Cortex-A core

There are two methods to build Native RTOS running on Cortex-A Core, one method is to leverage Yocto, another method is to build the image by using ARM gcc directly.

Some native RTOS examples are available in the [Heterogeneous Multicore repo](#), refer to [Building Heterogeneous Multicore RTOS Application](#) for how to build Native RTOS.

2.6.3 Booting Native RTOS on Cortex-A core

Cortex-A Core Native RTOS can be booted up or stopped on specified Cortex-A Core by using the U-Boot commands or remoteproc under Linux

Refer to [Unified Life Cycle Management](#) for details.

2.6.4 Debugging Native RTOS on Cortex-A Core

Some commercial debugging tools can be used to debug Native RTOS running on Cortex-A Core, such as TRACE32 Development System which is a product from Lauterbach Datentechnik GmbH, and JLink which is product from SEGGER.


For using JLink debugger, refer to application note [Debugging Cortex-A U-Boot and Native RTOS on i.MX 8M Plus and i.MX 93 EVKs](#) which focuses on the NXP MPU Cortex-A Core debugging method and flow. It provides two methods, one is to use the command-line debugging tool (GDB), the other one is to use the GUI debugging tool (Eclipse in this AN), and takes U-Boot and RTOS debugging as examples.

2.6.5 Native RTOS real-time latency benchmark

2.6.5.1 rt_latency introduction

`rt_latency` is a real-time latency test application developed by NXP to measure the interrupt and scheduling latency of Jailhouse + RTOS under different workloads. It measures the latency (time delta, in nanoseconds) between hardware IRQ events and software actions.

A sample illustration is provided in the [Figure 24](#).



irq delay:

The time to enter in the software IRQ handler after a hardware IRQ occurs (hardware + hypervisor + IRQ vector latency). It is equal to $T2 - T1$ in the above graph.

irq to sched:

The time to enter in an RTOS task, scheduled by the IRQ handler (irq delay + RTOS scheduler). It is equal to $T4 - T1$ in the above graph.

All measurements are done using hardware timer and relative to the hardware IRQ event time, with submicrosecond precision. The benchmark task sets the timeout value to 20 ms for hardware timer interrupt. When running, the `rt_latency` application displays latency statistics every 10 seconds, based on the measurements taken, to help characterize the system real-time latency.

`rt_latency` application provides different latency test cases. For example, it is possible to engage some CPU load and/or IRQ load to measure their impact on the latency. The [Table 21](#) shows the six cases that are currently implemented in the `rt_latency` application. Each case has a corresponding case ID and can be selected at running time.

Table 21. `rt_latency` application use cases

Test Case ID	Case Description
1	No extra load
2	Extra CPU load (low-priority task, executing busy loop and consuming all available CPU time)
3	Extra IRQ load
4	Extra CPU load + semaphore load
5	Extra CPU load + Linux load (not provided by the <code>rt_latency</code> RTOS application)
6	Extra CPU load + cache flush (instruction cache only for this release)

FreeRTOS related information of `rt_latency`

FreeRTOS is customized using a configuration file named `FreeRTOSConfig.h`. In the **rt_latency FreeRTOS** project, the macro `configMAX_PRIORITIES` is defined to 5 in `FreeRTOSConfig.h`. It is the number of priorities available to the application tasks.

When **rt_latency** is started, it creates a task called "main_task" of priority 1. The task runs an endless loop to process the input of UART console for native RTOS or control application `harpoon_ctrl` for Jailhouse RTOS and then delay 100 ms using the `xTaskDelay()` FreeRTOS API function. This task always runs and never ends.

When a specified benchmark use case is selected to run, the **rt_latency** application creates a task called "benchmark_task" of priority 3 by the `xTaskCreate()` FreeRTOS API function to measure the latency. It also creates a task called "log_task" of priority 1 to display the latency statistics every 10 seconds.

Running benchmark cases that include extra load uses more system resources and creates more tasks. The details are shown in the [Table 22](#).

Table 22. Running benchmark cases that include extra load

Case ID	Additional FreeRTOS tasks (except "benchmark_task" and "log_task")
1	None
2	A task called "cpu_load_task" of priority 0, executing busy loop and consuming all available CPU time.
3	No additional task is created but additional IRQ load is added in the "benchmark_task". Note that the IRQ latency testing in "benchmark_task" uses GPT1 while the IRQ load uses GPT2.
4	A task called "cpu_load_task" of priority 0, executing busy loop and consuming all available CPU time plus obtaining/releasing a semaphore using the <code>xSemaphoreTake()/xSemaphoreGive()</code> FreeRTOS API functions in the loop.
5	A task called "cpu_load_task" of priority 0, executing busy loop and consuming all available CPU time. Linux load (that is, <code>hackbench</code>) has to be run separately by the user.
6	A task called "cpu_load_task" of priority 0, executing busy loop and consuming all available CPU time. A task called "cache_inval_task" of priority 1, executing a loop to execute cache invalidation instructions and then delay 100 ms using the <code>xTaskDelay()</code> FreeRTOS API function.

2.6.5.2 Building rt_latency benchmark application

RTOS latency benchmark application "rt_latency" runs on FreeRTOS or Zephyr, it is an application in the repo: [heterogeneous-multicore](#).

Refer to "[Building Heterogeneous Multicore RTOS Application](#)" for how to build RPSMSG performance evaluation application.

2.6.5.3 Running Native RTOS rt_latency application

Native RTOS **rt_latency** application can run on the following boards:

- i.MX 8M Plus EVK
- i.MX 8M Mini EVK
- i.MX 91 EVK and QSB
- i.MX 93 11x11 and 14x14 EVK
- i.MX 95 15x15 and 19x19 EVK

Test Environment

Requirements:

- i.MX 8M Plus EVK, i.MX 8M Mini EVK, i.MX 91 EVK and QSB, i.MX 93 11x11 and 14x14 EVK, or the i.MX 95 15x15 and 19x19 EVK board
- Real-time Edge Software v2.8 or above
- Host PC
- Console cables

Test Procedure

1. Set up the UART Console for Native RTOS

Refer to the steps listed in *Setup UART console* in the section [Section 3.7.2](#).

The first UART console is used for Linux boots up, whereas the other one is used for RTOS running on Cortex-A Core.

2. Booting Native RTOS rt_latency Application

After powering up the board and entering U-Boot command line, execute the following U-Boot commands to run the `rt_latency` application.

- Boot the FreeRTOS `rt_latency` application on i.MX 8M Mini EVK board

```
u-boot=> ext4load mmc 1:2 0x93C00000 /examples/heterogeneous-multicore/rt-
latency-ca/ddr_release/rt_latency_ca53.bin
u-boot=> dcache flush; icache flush; cpu 3 release 0x93C00000;
```

- Boot the Zephyr `rt_latency` application on i.MX 8M Mini EVK board

```
u-boot=> ext4load mmc 1:2 0x93C00000 /examples/heterogeneous-multicore/rt-
latency-ca-zephyr/rt_latency_ca53.bin
u-boot=> dcache flush; icache flush; cpu 3 release 0x93C00000;
```

- Boot the FreeRTOS `rt_latency` application on i.MX 8M Plus EVK board:

```
u-boot=> ext4load mmc 1:2 0xC0000000 /examples/heterogeneous-multicore/rt-
latency-ca/ddr_release/rt_latency_ca53.bin
u-boot=> dcache flush; icache flush; cpu 3 release 0xC0000000;
```

- Boot Zephyr `rt_latency` application on i.MX 8M Plus EVK board:

```
u-boot=> ext4load mmc 1:2 0xC0000000 /examples/heterogeneous-multicore/rt-
latency-ca-zephyr/rt_latency_ca53.bin
u-boot=> dcache flush; icache flush; cpu 3 release 0xC0000000;
```

- Boot FreeRTOS `rt_latency` application on i.MX 91 EVK and QSB board:

```
u-boot=> ext4load mmc 1:2 0x80000000 /examples/heterogeneous-multicore/rt-
latency-ca/ddr_release/rt_latency_ca55.bin
u-boot=> dcache flush; icache flush; go 0x80000000;
```

- Boot Zephyr `rt_latency` application on i.MX 91 EVK and QSB board:

```
u-boot=> ext4load mmc 1:2 0x80000000 /examples/heterogeneous-multicore/rt-
latency-ca-zephyr/rt_latency_ca55.bin
u-boot=> dcache flush; icache flush; go 0x80000000;
```

- Boot FreeRTOS `rt_latency` application on i.MX 93 11x11 and 14x14 EVK board:

```
u-boot=> ext4load mmc 1:2 0xD0000000 /examples/heterogeneous-multicore/rt-
latency-ca/ddr_release/rt_latency_ca55.bin
u-boot=> dcache flush; icache flush; cpu 1 release 0xD0000000;
```

- Boot Zephyr `rt_latency` application on i.MX 93 11x11 and 14x14 EVK board:

```
u-boot=> ext4load mmc 1:2 0xD0000000 /examples/heterogeneous-multicore/rt-
latency-ca-zephyr/rt_latency_ca55.bin
u-boot=> dcache flush; icache flush; cpu 1 release 0xD0000000;
```

- Boot FreeRTOS `rt_latency` application on i.MX 95 15x15 and 19x19 EVK board:

```
u-boot=> ext4load mmc 1:2 0xD0000000 /examples/heterogeneous-multicore/rt-
latency-ca/ddr_release/rt_latency_ca55.bin
u-boot=> dcache flush; icache flush; cpu 5 release 0xD0000000;
```

- Boot Zephyr `rt_latency` application on i.MX 95 15x15 and 19x19 EVK board:

```
u-boot=> ext4load mmc 1:2 0xD0000000 /examples/heterogeneous-multicore/rt-
latency-ca-zephyr/rt_latency_ca55.bin
u-boot=> dcache flush; icache flush; cpu 5 release 0xD0000000;
```

3. On the RTOS serial console connected, the following output is displayed.

For FreeRTOS:

```
*****
FreeRTOS Latency Benchmark on Cortex-A
*****

Test Case ID (1-6):
1: no extra load
2: extra CPU load
3: extra IRQ load
4: extra CPU load + semaphore load
5: extra CPU load + Linux load (not provided by the test case)
6: extra CPU load + cache flush

Please input the test case ID (1-6):
```

For Zephyr:

```
*** Booting Zephyr OS build v2.7.0-rc1-14566-g3b58bc3af6bb ***

*****
Zephyr Latency Benchmark on Cortex-A
*****

Test Case ID (1-6):
1: no extra load
2: extra CPU load
3: extra IRQ load
4: extra CPU load + semaphore load
5: extra CPU load + Linux load (not provided by the test case)
6: extra CPU load + cache flush

Please input the test case ID (1-6):
```

4. Add Extra Linux CPU Load (Only feasible for test case 5 and 6)

For test case 5 and 6, need to add extra load in Linux, so need to boot into Linux by running the following command in U-Boot command line:

Note: *i.MX 91 implements only one Cortex-A55 core, so the extra load is not applicable.*

- For i.MX 8M Mini EVK:

```
u-boot=> setenv fdtfile imx8mm-evk-multicore-rtos.dtb
u-boot=> setenv mmcargs $mmcargs clk_ignore_unused
u-boot=> boot
```

- For i.MX 8M Plus EVK:

```
u-boot=> setenv fdtfile imx8mp-evk-multicore-rtos.dtb
u-boot=> setenv mmcargs $mmcargs clk_ignore_unused
u-boot=> boot
```

- For i.MX 93 11x11 and 14x14 EVK:

```
# For i.MX 93 11x11 EVK
u-boot=> setenv fdtfile imx93-11x11-evk-multicore-rtos.dtb
# For i.MX 93 14x14 EVK
u-boot=> setenv fdtfile imx93-14x14-evk-multicore-rtos.dtb
u-boot=> setenv mmcargs $mmcargs clk_ignore_unused
u-boot=> boot
```

- For i.MX 95 15x15 and 19x19 EVK:

```
# For i.MX 95 15x15 EVK
u-boot=> setenv fdtfile imx95-15x15-evk-multicore-rtos.dtb
# For i.MX 95 19x19 EVK
u-boot=> setenv fdtfile imx95-19x19-evk-multicore-rtos.dtb
u-boot=> setenv mmcargs $mmcargs clk_ignore_unused
u-boot=> boot
```

Once the Linux system has booted up, optimize the system configuration to achieve maximum performance:

```
$ real_time_setup.sh
```

Then, run the following command to add extra Linux load:

```
$ dd if=/dev/urandom of=/dev/null &>/dev/null & disown;
$ dd if=/dev/urandom of=/dev/null &>/dev/null & disown;
$ dd if=/dev/urandom of=/dev/null &>/dev/null & disown;
```

5. Select test case ID

When user inputs a test case ID (1 ~ 6) on the RTOS console, then benchmarking process starts. Then, the latency statistics is displayed on the RTOS console.

6. Capture the test log

Save the logs displayed by the `rt_latency` application in the RTOS console to a file (that is, `rt_latency.log`) on host PC for further analysis. Below is a snippet of an example log:

```
INFO: start_test_case      : Running test case 1:
INFO: benchmark_task      : running
INFO: stats_print          : stats(D0601B30) irq delay (ns) min 250 mean 306
max 416 rms^2 94616 stddev^2 747 absmin 250 absmax 416
INFO: hist_print           : n_slot 21 slot_size 1000
INFO: hist_print           : 100000 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0
INFO: stats_print          : stats(D0601F90) irq to sched (ns) min 1500 mean
1596 max 2583 rms^2 2550950 stddev^2 798 absmin 1500 absmax 2583
INFO: hist_print           : n_slot 21 slot_size 1000
INFO: hist_print           : 0 99999 1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
INFO: print_stats          : late alarm scheduling: 0
INFO: print_stats          :
INFO: stats_print          : stats(D0601B30) irq delay (ns) min 250 mean 310
max 333 rms^2 97174 stddev^2 681 absmin 250 absmax 416
INFO: hist_print           : n_slot 21 slot_size 1000
INFO: hist_print           : 200000 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0
INFO: stats_print          : stats(D0601F90) irq to sched (ns) min 1500 mean
1601 max 1625 rms^2 2565317 stddev^2 695 absmin 1500 absmax 2583
INFO: hist_print           : n_slot 21 slot_size 1000
INFO: hist_print           : 0 199999 1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0
INFO: print_stats          : late alarm scheduling: 0
INFO: print_stats          :
```

```

INFO: stats_print      : stats(D0601B30) irq delay (ns) min 250 mean 310
max 333 rms^2 97001 stddev^2 689 absmin 250 absmax 416
INFO: hist_print      : n_slot 21 slot_size 1000
INFO: hist_print      : 300000 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0
INFO: stats_print      : stats(D0601F90) irq to sched (ns) min 1541 mean
1601 max 1625 rms^2 2564434 stddev^2 704 absmin 1500 absmax 2583
INFO: hist_print      : n_slot 21 slot_size 1000
INFO: hist_print      : 0 299999 1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0
INFO: print_stats      : late alarm scheduling: 0
INFO: print_stats      :
..... (omitted)

```

2.6.5.4 Output of rt_latency

The following latency statistics are displayed after every 10 seconds. The first three lines are for "irq delay" and the other three lines are for "irq to sched".

```

INFO: stats_print      : stats(C0601260) irq delay (ns) min 708 mean 711
max 833 rms^2 506896 stddev^2 174 absmin 708 absmax 3083
INFO: hist_print      : n_slot 21 slot_size 200
INFO: hist_print      : 0 0 0 999 2 0 0 0 0 0 0 0 0 0 1 1 0 0
0 0 0 0
INFO: stats_print      : stats(C06016C0) irq to sched (ns) min
2416 mean 2499 max 5333 rms^2 6341435 stddev^2 92571 absmin 2416 absmax 5625
INFO: hist_print      : n_slot 21 slot_size 1000
INFO: hist_print      : 0 0 986 0 11 7 0 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0
INFO: print_stats      :

```

As stated before, the benchmark task sets the timeout value to 20 ms for GPT interrupt. This means that one sample of "irq delay" latency data and one sample of "irq to sched" latency data are measured about every 20 ms. So a 10 s time would have approximately 500 samples.

1. **min/mean/max**: minimum, average and maximum latency value measured within the last 10 s period of time. These statistics are collected and calculated for approximately 500 samples.
2. **rms^2**: square of the RMS (Root Mean Square) measured within the last 10 s period of time.
3. **stddev^2**: square of the standard deviation measured within the last 10 s period of time. The square of the standard deviation is also called "variance".
4. **absmin/absmax**: minimum and maximum latency value measured since the beginning of the test. Therefore, we take the value of `absmin` and `absmax` in the last set of latency statistics as the benchmarking results.
5. A histogram is also shown to give an idea of repartition of the measured latency values. Note that the histogram data is accumulated so the sum of one line of histogram data increases by about 500 compared to the previous one. Therefore, we only care about the histogram data in the last set of latency statistics. That is the repetition of the measured latency values over the whole test period (that is, 12h).

2.7 RTOS on Cortex-A core with Jailhouse

2.7.1 Jailhouse

2.7.1.1 Overview

Jailhouse is a partitioning Hypervisor based on Linux. It is able to run baremetal applications or (adapted) operating systems besides Linux. For this purpose, it configures CPU and device virtualization features of the hardware platform in a way that none of these domains, called "cells" here, can interfere with each other in an unacceptable way.

Jailhouse is optimized for simplicity rather than feature richness. Jailhouse does not support overcommitment of resources such as CPUs, RAM, or devices. This feature makes it different from full-featured Linux-based hypervisors such as KVM or Xen. It performs no scheduling and only virtualizes those resources in software, which are essential for a platform and cannot be partitioned in hardware.

Once Jailhouse is activated, it runs Baremetal. This implies that it takes full control over the hardware and needs no external support. However, in contrast to other baremetal hypervisors, it requires a normal Linux system to be loaded and configured. Its management interface is based on Linux infrastructure. So, you boot Linux first, then, enable Jailhouse and finally split off parts of the system's resources and assign them to additional cells.

2.7.1.2 Running PREEMPT_RT Linux in Inmate

2.7.1.2.1 i.MX 8M Plus LPDDR4 EVK

Perform the following steps on i.MX 8M Plus LPDDR4 EVK board:

1. Execute `run jh_mmcboot` at U-Boot prompt on the terminal of UART2.
2. Wait for Linux OS to boot up and login.
3. Execute non-root Linux demo (Assuming rootfs have been deployed in `/dev/mmcblk2p2`):

```
# cd /usr/share/jailhouse/scripts
# ./linux-demo-imx8mp.sh
```

4. Check the output on the terminal of UART4:

```
[ 0.717545] printk: console [ttymxc3] enabled
[ 0.721628] printk: bootconsole [ec_imx6q0] disabled
[ 0.732428] loop: module loaded
[ 0.732902] of_reserved_mem_lookup() returned NULL
[ 0.732952] megasas: 07.714.04.00-rc1
[ 0.733632] imx ahci driver is registered.
[ 0.735615] tun: Universal TUN/TAP device driver, 1.6
[ 0.735835] thunder_xcv, ver 1.0
[ 0.735863] thunder_bgx, ver 1.0
[ 0.735889] nicpf, ver 1.0
[ 0.736340] hclge is initializing
[ 0.736351] hns3: Hisilicon Ethernet Network Driver for Hip08 Family -
version
[ 0.736354] hns3: Copyright (c) 2017 Huawei Corporation.
[ 0.736382] e1000: Intel(R) PRO/1000 Network Driver
[ 0.736384] e1000: Copyright (c) 1999-2006 Intel Corporation.
[ 0.736416] e1000e: Intel(R) PRO/1000 Network Driver
[ 0.736418] e1000e: Copyright(c) 1999 - 2015 Intel Corporation.
[ 0.736447] igb: Intel(R) Gigabit Ethernet Network Driver
[ 0.736450] igb: Copyright (c) 2007-2014 Intel Corporation.
```

```
[ 0.736473] igbvf: Intel(R) Gigabit Virtual Function Network Driver
[ 0.736475] igbvf: Copyright (c) 2009 - 2012 Intel Corporation.
...
NXP Real-time Edge Distro 2.2 imx8mp-lpddr4-evk ttyMXC3
imx8mp-lpddr4-evk login: root
root@imx8mp-lpddr4-evk:~#
root@imx8mp-lpddr4-evk:~# uname -a
Linux imx8mpevk 5.10.72-rt53-lts-5.10.y+g5304e5555731 #1 SMP PREEMPT_RT Tue
Mar 1 06:03:05 UTC 2022 aarch64 aarch64 aarch64 GNU/Linux
root@imx8mp-lpddr4-evk:~#
```

Note: If the case fails because of rootfs error, update rootfs using the following command:

```
# rm -fr /run/media/mmcblk2p2/*
# cp -frd /usr /bin /etc /home /fat /lib /linuxrc /lost+found/ /media/ /mnt /
opt /root /sbin /run/media/mmcblk2p2/
```

5. Exit Jailhouse.

2.7.1.2.2 LS1028ARDB

Perform the steps listed in the following section to run PREEMPT_RT Linux on LS1028ARDB board.

2.7.1.2.2.1 Linux in non-root cell

Perform the following steps to run PREEMPT_RT Linux in Inmate on LS1028ARDB platform:

1. Execute `run jh_mmcboot` from U-Boot prompt.
2. Wait for Linux OS to boot up and log in it.
3. Execute non-root Linux demo:

```
# cd /usr/share/jailhouse/scripts
# ./linux-demo-ls1028ardb.sh
```

4. Exit Jailhouse.

```
# ../tools/jailhouse disable
```

2.7.1.2.2.2 ENETC in non-root cell

Follow the below steps for ENETC that is assigned to non-root cell:

1. Under U-Boot prompt, run the below commands to set the device tree blob, which has ENETC nodes removed and then boot up Linux:

```
=> setenv jh_mmcboot `setenv dtb fsl-ls1028a-rdb-jailhouse-without-
enetc.dtb;run bootcmd`
=> run jh_mmcboot
```

2. Wait for Linux OS to boot up and then log in.
3. Execute non-root Linux demo:

```
# cd /usr/share/jailhouse/scripts
# ./linux-demo-ls1028ardb-enetc.sh
```

Then, network can be available in non-root cell Linux.

4. Exit Jailhouse.

```
# ../tools/jailhouse disable
```

Note:

In this case, the GICv3 ITS node is also removed from the root cell Linux device tree. The node is assigned to non-root cell to service ENETC MSI-X interrupts, so the root cell Linux does not support the MSI/MSI-X service anymore.

2.7.1.2.2.3 GPIO in non-root cell

GPIO3 controller is assigned to non-root cell, below steps is for GPIO that is assigned to non-root cell:

1. Hardware setup

Connect J11 Pin5 (1588_ALARM_OUT1/GPIO3_DAT11) to Pin 8 (1588_CLK_OUT/GPIO3_DAT10)

2. RCW setting

In dash-rcw/ls1028ardb/R_SQPP_0x85bb/rcw_1500_sdboot.rcw, change as below:

```
EC1_SAI4_5_POMUX=1
EC1_SAI3_6_POMUX=1
```

EC1_SAI4_5_POMUX is set to 0b001, EC1_SAI3_6_POMUX is set to 0b001 to select GPIO.

3. Software configuration required:

- a. Configure CPLD register BRDCFG3 (offset 053h) bit 2 to 0 (IEEE signals connect to the IEEE header) in U-Boot prompt:

```
=> i2c mw 66 53 00
```

- b. Boot up Linux using Jailhouse DTB and bring up non-root Linux:

```
=> run jh_mmcboot
```

- c. Wait for Linux OS to boot up and login.

- d. Execute non-root Linux demo.

```
# cd /usr/share/jailhouse/scripts
# ./linux-demo-ls1028ardb.sh
```

4. Test GPIO function in non-root Linux.

- a. Export GPIO pin

```
# ls /sys/class/gpio
# echo 490 > /sys/class/gpio/export
# echo 491 > /sys/class/gpio/export
```

- b. Configure GPIO output and input.

```
# echo out > /sys/class/gpio/gpio490/direction
# cat /sys/class/gpio/gpio490/direction
# cat /sys/class/gpio/gpio491/direction
```

- c. Verify write 1 to GPIO output.

```
# echo 1 > /sys/class/gpio/gpio490/value
# cat /sys/class/gpio/gpio490/value
# cat /sys/class/gpio/gpio491/value
```

- d. Verify write 0 to GPIO output.

```
# echo 0 > /sys/class/gpio/gpio490/value
# cat /sys/class/gpio/gpio490/value
# cat /sys/class/gpio/gpio491/value
```


5. Exit Jailhouse

```
# ../tools/jailhouse disable
```

2.7.1.2.3 LS1046ARDB

Perform the following steps:

1. Execute `run jh_mmcboot` in U-Boot stage.
2. Wait for Linux OS to boot up and login in it.
3. Execute non-root Linux demo:

```
# cd /usr/share/jailhouse/scripts
# ./linux-demo-ls1046ardb.sh
```

4. Exit Jailhouse:

```
# ../tools/jailhouse disable
```

2.7.1.3 Running Jailhouse examples In Inmate

2.7.1.3.1 i.MX 8M Plus LPDDR4 EVK

1. Execute `run jh_mmcboot` in U-Boot stage
2. Wait for Linux OS to boot up and log in.
3. Execute the GIC demo using the commands below.

```
# cd /usr/share/jailhouse/scripts
# ./gic-demo-imx8mp.sh
```

4. Check the result on the serial port:

```
Initializing the GIC...
Initializing the timer...
Timer fired, jitter: 2039 ns, min: 2039 ns, max: 2039 ns
Timer fired, jitter: 1039 ns, min: 1039 ns, max: 2039 ns
Timer fired, jitter: 879 ns, min: 879 ns, max: 2039 ns
Timer fired, jitter: 959 ns, min: 879 ns, max: 2039 ns
Timer fired, jitter: 1039 ns, min: 879 ns, max: 2039 ns
Timer fired, jitter: 919 ns, min: 879 ns, max: 2039 ns
Timer fired, jitter: 919 ns, min: 879 ns, max: 2039 ns
Timer fired, jitter: 919 ns, min: 879 ns, max: 2039 ns
Timer fired, jitter: 1079 ns, min: 879 ns, max: 2039 ns
Timer fired, jitter: 919 ns, min: 879 ns, max: 2039 ns
Timer fired, jitter: 919 ns, min: 879 ns, max: 2039 ns
Timer fired, jitter: 959 ns, min: 879 ns, max: 2039 ns
```

Figure 25. Console log after running the GIC demo

5. Execute the UART demo:

```
# ./uart-demo-imx8mp.sh
```

6. Check the result on serial port:

```
Hello 1 from cell!  
Hello 2 from cell!  
Hello 3 from cell!  
Hello 4 from cell!  
Hello 5 from cell!  
Hello 6 from cell!  
Hello 7 from cell!  
Hello 8 from cell!  
Hello 9 from cell!  
Hello 10 from cell!  
Hello 11 from cell!  
Hello 12 from cell!  
Hello 13 from cell!  
Hello 14 from cell!  
Hello 15 from cell!  
Hello 16 from cell!  
Hello 17 from cell!  
Hello 18 from cell!  
Hello 19 from cell!  
Hello 20 from cell!
```

Figure 26.

7. Exit Jailhouse.

```
# ../tools/jailhouse disable
```

2.7.1.3.2 LS1028ARDB Jailhouse example in Inmate

Perform the following steps for running LS1028ARDB Jailhouse example In Inmate:

1. Execute `run jh_mmcboot` in U-Boot stage.
2. Wait for Linux OS to boot up and then log in.
3. Execute GIC demo using the command below:

```
# cd /usr/share/jailhouse/scripts  
# ./gic-demo-ls1028ardb.sh
```

4. Execute UART demo using the command below:

```
# ./uart-demo-ls1028ardb.sh
```

5. Execute ivshmem demo using the command below:

```
# ./ivshmem-demo-ls1028ardb.sh
```

Note: If *ivshmem* case fails, then, reboot the board and test the case again.
Check the result on the second serial port:

```
IVSHMEM: Found device at 00:00.0
IVSHMEM: bar0 is at 0x00000000ff000000
IVSHMEM: bar1 is at 0x00000000ff001000
IVSHMEM: ID is 1
IVSHMEM: max. peers is 1
IVSHMEM: state table is at 0x00000000c0500000
IVSHMEM: R/W section is at 0x00000000c0501000
IVSHMEM: input sections start at 0x00000000c050a000
IVSHMEM: output section is at 0x00000000c050c000
IVSHMEM: initialized device
state[0] = 1
state[1] = 2
state[2] = 0
rw[0] = 1
rw[1] = 0
rw[2] = -1001599800
in@0x0000 = 10
in@0x2000 = 0
in@0x4000 = 1758252876

IVSHMEM: got interrupt 0 (#1)
state[0] = 1
state[1] = 2
state[2] = 0
rw[0] = 1
rw[1] = 1
rw[2] = -1001599800
in@0x0000 = 10
in@0x2000 = 10
in@0x4000 = 1758252876
```

6. Exit Jailhouse.

2.7.1.3.3 LS1046ARDB Jailhouse example

Perform the below steps for running Jailhouse examples in Inmate on LS1046ARDB:

1. Execute `run jh_mmcboot` in U-Boot stage.
2. Wait for Linux OS to boot up and login it.
3. Execute GIC demo:

```
# cd /usr/share/jailhouse/scripts
# ./gic-demo-ls1046ardb.sh
```

4. Execute UART demo:

```
# ./uart-demo-ls1046ardb.sh
```

5. Execute ivshmem demo:

```
# ./ivshmem-demo-ls1046ardb.sh
```

6. Exit Jailhouse.

```
# ../tools/jailhouse disable
```

2.7.2 Harpoon (RTOS on Cortex-A)

2.7.2.1 Overview

Harpoon RTOS provides an environment for developing real-time demanding applications on an RTOS running on one (or several) Cortex-A core(s) in parallel of a Linux distribution.

Harpoon leverages Jailhouse to partition the hardware and run the RTOS as a Linux guest.

The Harpoon RTOS is based on either FreeRTOS or Zephyr plus MCUXpresso drivers and provides several example applications:

- Audio application
- Industrial application
- Real-time latency test application

For details about Harpoon OS, refer to its user guide available at the following location:

https://www.nxp.com/design/software/development-software/real-time-edge-software:REALTIME-EDGE-SOFTWARE?tab=Documentation_Tab

2.8 RTOS and Baremetal on Cortex-M core

Regarding RTOS and Baremetal building, please refer to [Section 2.3](#)

Real-time Edge images has some demo for testing.

For more example, please refer to <https://mcuxpresso.nxp.com/en/welcome>

2.8.1 Booting Cortex-M Core RTOS Image

There are two ways to boot ARM Cortex-M Core: booting from U-Boot, or using RemoteProc to boot from Linux.

2.8.1.1 Using U-Boot Commands to Manage Life Cycle of RTOS on Cortex-M Core

U-boot command "bootaux" is used to boot Cortex-M Core RTOS Image from U-Boot, for example, after the board is booted into the U-Boot console.

- Use the following command to boot Arm Cortex-M core on i.MX 8M Mini LPDDR4 EVK board:

```
=> ext4load mmc 1:2 0x48000000 /examples/heterogeneous-multicore/hello-world-cm/
release/hello_world_cm4_UART4.bin; cp.b 0x48000000 0x7e0000 20000;
=> bootaux 0x7e0000
```

- Use the following command on i.MX 8M Plus LPDDR4 EVK board:

```
=> ext4load mmc 1:2 0x48000000 /examples/heterogeneous-multicore/hello-world-cm/
release/hello_world_cm7_UART4.bin; cp.b 0x48000000 0x7e0000 20000;
=> bootaux 0x7e0000
```

- Use the following command on i.MX 93 EVK board:

```
=> ext4load mmc 1:2 0xd0000000 /examples/heterogeneous-multicore/hello-world-cm/
release/hello_world_cm33_UART2.bin; cp.b 0xd0000000 0x201e0000 20000;
=> bootaux 0x1ffe0000
```

- Then, the below RTOS log is displayed on the UART console (taking i.MX 8M Mini EVK as an example)

```
Cortex-M4: RTOS0: Hello world! Real-time Edge on MIMX8MM-EVK
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac
```

2.8.1.2 Using RemoteProc under Linux to Manage Life Cycle of RTOS on Cortex-M Core

Use Linux remoteproc to start and stop RTOS on Cortex-M Core, taking the Heterogeneous Multicore hello_world application as an example:

- Booting Arm Cortex-M core on i.MX 8M Mini EVK, i.MX 8M Plus EVK and i.MX 93 EVK boards:

1. Starting Linux kernel:

```
=> run prepare_mcore
=> setenv fdtfile imx8mm-evk-multicore-rtos.dtb # for i.MX 8M Mini EVK
=> setenv fdtfile imx8mp-evk-multicore-rtos.dtb # for i.MX 8M Plus EVK
=> setenv fdtfile imx93-11x11-evk-multicore-rtos.dtb # for i.MX 93 EVK
=> boot
```

2. After Linux booting up, boot up Cortex-M core RTOS application:

```
# i.MX 8M Mini EVK
root@imx8mm-lpddr4-evk:~# echo -n /examples/heterogeneous-multicore/hello-
world-cm/release/hello_world_cm4_UART4.elf > /sys/devices/platform/imx8mm-
cm4/remoteproc/remoteproc4/firmware
root@imx8mm-lpddr4-evk:~# echo start > /sys/devices/platform/imx8mm-cm4/
remoteproc/remoteproc4/state

# i.MX 8M Plus EVK
root@imx8mp-lpddr4-evk:~# echo -n /examples/heterogeneous-multicore/hello-
world-cm/release/hello_world_cm7_UART4.elf > /sys/devices/platform/imx8mp-
cm7/remoteproc/remoteproc4/firmware
root@imx8mp-lpddr4-evk:~# echo start > /sys/devices/platform/imx8mp-cm7/
remoteproc/remoteproc4/state

# i.MX 93 EVK
root@imx93evk:~# echo -n /examples/heterogeneous-multicore/hello-world-cm/
release/hello_world_cm33_UART2.elf > /sys/devices/platform/remoteproc-cm33/
remoteproc/remoteproc1/firmware
root@imx93evk:~# echo start > /sys/devices/platform/remoteproc-cm33/
remoteproc/remoteproc1/state
```

Then, the below RTOS log is displayed on the UART console (taking i.MX 8M Mini EVK as an example):

```
Cortex-M4: RTOS0: Hello world! Real-time Edge on MIMX8MM-EVK
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac
tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac
```

3. Then the Cortex-M core RTOS application can be shut down:

```
# i.MX 8M Mini EVK
root@imx8mm-lpddr4-evk:~# echo stop > /sys/devices/platform/imx8mm-cm4/
remoteproc/remoteproc4/state

# i.MX 8M Plus EVK
root@imx8mp-lpddr4-evk:~# echo stop > /sys/devices/platform/imx8mp-cm7/
remoteproc/remoteproc4/state

# i.MX 93 EVK
```

```
root@imx93evk:~# echo stop > /sys/devices/platform/remoteproc-cm33/  
remoteproc/remoteproc1/state
```

Note: *The remoteproc portal instance ID /sys/devices/platform/remoteproc-cm7/remoteproc/remoteproc4 might be different, depending on the device tree configuration. Therefore, use the name under each remoteproc instance portal to check which instance must be used.*

3 Heterogeneous Multicore Framework

3.1 Overview

Heterogeneous Multicore Framework provides a general software framework to support Heterogeneous AMP. It enables AMP to be interconnected and provides a unified resource management and life-cycle management as shown in [Figure 27](#).

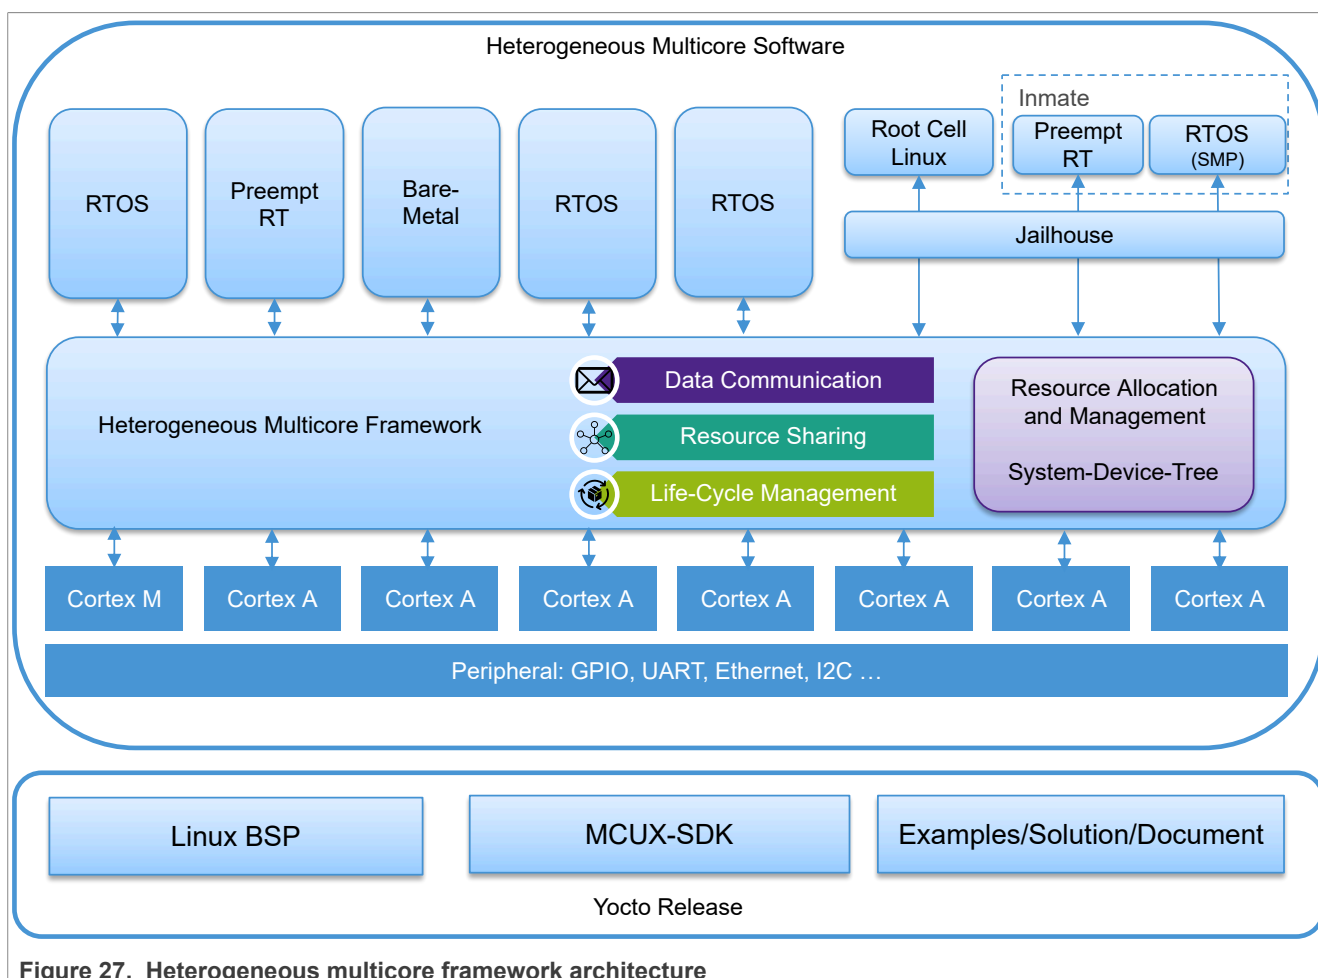


Figure 27. Heterogeneous multicore framework architecture

It provides the below key functions to help users to accelerate solution development based on multicore platforms:

1. Data communication between different operating systems

The following technical implementation can be used to pass common data between different operating systems. The data transfer can be between Cortex-M Core and Cortex-A Core, or between different Cortex-A cores, or between multiple CPU cores simultaneously.

- **RPMMsg**
RPMMsg is a standard intercore communication protocol supported on Linux and RTOS.
- **Heterogeneous Multicore VirtIO**
Heterogeneous Multicore VirtIO applies para-virtualization VirtIO technology to build high-performance intercore data path. A customized data path is defined according to different use cases.

2. Resource sharing between different operating systems

Resource sharing enables sharing physical resources between different OSes. In general, one OS owns and controls hardware resources while the other OS uses a virtual device. The following mechanism is followed to build the control path and data path to access physical resource.

- **RPMMsg**

Use RPMMsg to build control and data path crossing OS, physical resource is shared with another OS in terms of virtual device. The simplified Real-Time Messaging (SRTM) protocol provided in Real-time Edge is an implementation based on RPMMsg. It is used to share the physical resources of the Cortex-M core with the Cortex-A core in terms of virtual devices in Linux.

- **Heterogeneous Multicore VirtIO**

Heterogeneous Multicore VirtIO have a better performance than RPMSG, and it can also be used for resource sharing. POSIX compatible API can be used to access virtual device, and some existing VirtIO device drivers in Linux can be reused. Networking sharing is provided in Real-time Edge to share the same networking interface between multiple OSes.

3. **Unified Life-Cycle Management**

Heterogeneous Multicore Framework provides unified Life Cycle Management both for Cortex-A Core and Cortex-M Core. It supports using U-Boot commands or Linux remoteproc to start or stop RTOS.

4. **Industrial Applications**

The Heterogeneous Multicore framework provides a flexible mechanism to run Preempt-RT Linux or RTOS on different cores with Industrial applications.

There are some sample applications provided in Heterogeneous Multicore Framework. These applications can be used to demo and evaluate the features in Heterogeneous Multicore Framework:

1. `hello_world`

The `hello_world` application demonstrates a sample flexible Real-time System on MPU platforms. The multiple images provided can be used to run single or multiple RTOS on Cortex-A Core or Cortex-M with or without running Linux simultaneously.

2. **RPMSG Applications**

Heterogeneous Multicore Framework supports RPMSG communication between any Real-time Systems on MPU Platforms, such as:

- RPMSG between RTOS on Cortex-M Core and Linux on Cortex-A core
- RPMSG between RTOS on Cortex-A Core and Linux on Cortex-A Core
- RPMSG between RTOS on Cortex-A Core and RTOS on Cortex-A Core
- RPMSG between RTOS on Cortex-M Core and RTOS on Cortex-A Core

The following applications provide field trail for RPMSG-related features.

- `rpmsg_str_echo`

This demo demonstrates building up multiple RPMSG endpoints between RTOS and Linux. For example, on the i.MX 8M Plus EVK board, images provided in Real-time Edge can be used to run three RTOS `rpmsg_str_echo` applications on two Cortex-A Core and one Cortex-M Core. The other two Cortex-A Cores run SMP Linux, then each RTOS establishes three RPMSG channels with Linux.

- `rpmsg_pingpong`

This demo demonstrates RPMSG communication between RTOS and RTOS, one is RPMSG master and the other is RPMSG slave.

- `rpmsg_perf`

`rpmsg_perf` is a tool to evaluate RPMSG bandwidth performance between RTOS and Linux Kernel.

- **RPMSG enhanced 8MB buffer**

The application demonstrates how to change the default RPMSG buffer size and count.

- **UART sharing based on RPMSG**

This application demonstrates how to use RPSMG to share physical peripherals or other resource between different CPU Core or OS. This demo share physical UART controlled by Cortex-M Core with Cortex-A Core on which virtual UART device driver is provided in Linux.

3. Heterogeneous Multicore VirtIO Applications

- `virtio_perf`
`virtio_perf` is a tool that evaluates Heterogeneous Multicore VirtIO bandwidth performance between RTOS and Linux.
- VirtIO Networking Sharing
This application demonstrates how to use a Heterogeneous Multicore VirtIO to share physical peripherals or other resources between different CPU Cores or the OSes. The applications provide networking sharing, where the physical networking interface is controlled by Cortex-M Core or Cortex-A Core. Then, the information about which virtual NIC device driver is provided in Linux is shared with the Cortex-A Core. Heterogeneous Multicore VirtIO is used to establish a high performance data path between two sides.

4. SOEM Applications

- `soem_digital_io`
SOEM `digital_io` example for IO control using BECKHOFF EK1100 as EtherCAT slave.
- `soem_servo`
SOEM `servo_motor` example for motor control using Inovance SV680 servo as EtherCAT slave.
- `soem_servo_rt1180`
SOEM `servo_motor_rt1180` example for motor control using i.MX RT1180 EVK board as EtherCAT slave.

Table 23 shows the support matrix on NXP platforms:

Table 23. Heterogeneous Multicore Application Support Matrix

Heterogeneous Multicore Framework				i.MX 8M Mini LPDDR4 EVK		i.MX 8M Plus LPDDR4 EVK		i.MX 91 QSB	i.MX 91 EVK	i.MX 93 9x9 LPDDR4 QSB		i.MX 93 EVK		i.MX 93 14x14 EVK		i.MX 95 15x15 EVK		i.MX 95 19x19 EVK	
Feature	Sub-feature	RTOS	Application	4 X A53	1 X M4	4 X A53	1 X M7	1 X A55	1 X A55	2 X A55	1 X M33	2 X A55	1 X M33	2 X A55	1 X M33	6 X A55	1 X M7	6 X A55	1 X M7
Flexible Real-time System	Flexible Real-time System	Free RTOS	hello_world	Y	Y	Y	Y					Y	Y	Y	Y	Y		Y	
		Zephyr	hello_world	Y		Y						Y		Y		Y		Y	
	RAM Console	Free RTOS	hello_world	Y		Y						Y		Y		Y		Y	
		Zephyr	hello_world	Y		Y						Y		Y		Y		Y	
	networking stack on A-Core RTOS	Free RTOS	lwip_ping	Y		Y						Y		Y					
Unified Life Cycle Management	U-Boot commands booting/ stopping A-Core Native RTOS	Free RTOS	all	Y		Y		Y	Y			Y		Y		Y		Y	
		Zephyr	all	Y		Y		Y	Y			Y		Y		Y		Y	
	Linux remoteproc booting/ stopping A-Core Native RTOS	Free RTOS	all	Y		Y						Y		Y		Y		Y	
		Zephyr	all	Y		Y						Y		Y		Y		Y	
	U-Boot commands booting M-Core Native RTOS	Free RTOS	all		Y		Y				Y		Y		Y				
	Linux remoteproc booting M-Core Native RTOS	Free RTOS	all		Y		Y				Y		Y		Y				
RPMMSG	RPMMSG between Linux and M-Core RTOS	Free RTOS	rpmsg_str_echo		Y		Y				Y ^[1]		Y		Y				
	RPMMSG between Linux and A-Core RTOS	Free RTOS	rpmsg_str_echo	Y		Y						Y		Y					
	RPMMSG between 2 A-Core RTOS	Free RTOS	rpmsg_pingpong			Y													
	RPMMSG between Linux and M-Core RTOS with enhanced 8MB buffer	Free RTOS	rpmsg_str_echo		Y														
	RPMMsg Performance Evaluation	Free RTOS	rpmsg_perf			Y	Y												
	UART Sharing based on RPMMsg	Free RTOS	rpmsg_uart_sharing		Y						Y		Y		Y				
Heterogeneous Multicore VirtIO	Heterogeneous Multicore VirtIO	Free RTOS	virtio_perf	Y	Y	Y	Y												

Table 23. Heterogeneous Multicore Application Support Matrix...continued

Heterogeneous Multicore Framework				i.MX 8M Mini LPDDR4 EVK		i.MX 8M Plus LPDDR4 EVK		i.MX 91 QSB		i.MX 91 EVK		i.MX 93 9x9 LPDDR4 QSB		i.MX 93 EVK		i.MX 93 14x14 EVK		i.MX 95 15x15 EVK		i.MX 95 19x19 EVK	
Feature	Sub-feature	RTOS	Application	4 X A53	1 X M4	4 X A53	1 X M7	1 X A55	1 X A55	2 X A55	1 X M33	2 X A55	1 X M33	2 X A55	1 X M33	2 X A55	1 X M33	6 X A55	1 X M7	6 X A55	1 X M7
	Performance Evaluation																				
	Heterogeneous Multicore VirtIO Network Sharing	Free RTOS	virtio_net_backend	Y	Y	Y	Y						Y	Y	Y	Y					
Industrial Applications	SOEM Digital IO	Free RTOS	soem_digital_io	Y	Y	Y	Y						Y								
	SOEM Servo Motor	Free RTOS	soem_servo	Y	Y	Y	Y						Y								
	SOEM Servo Motor RT1180	Free RTOS	soem_servo_rt1180	Y	Y	Y	Y						Y								

[1] Using the rpmsg_uart_sharing

3.2 Building Heterogeneous Multicore RTOS Application

Heterogeneous Multicore Framework provides some RTOS applications in the repo: [heterogeneous-multicore](#).

There are two methods to build Heterogeneous Multicore RTOS Applications, one method is to leverage Yocto, another method is standalone method by using ARM gcc directly.

3.2.1 Build with Yocto

Real-time Edge supports to build all images by using Yocto. Refer to [Section 2.3](#) for how to leverage Yocto to build RTOS Applications on Cortex-A or Cortex-M cores.

The followings are some Yocto quick commands:

Build all RTOS application running both on Cortex-A Core and Cortex-M Core:

```
bitbake packagegroup-real-time-edge-rtos
```

Build single Heterogeneous Multicore RTOS application separately:

```
bitbake APP-NAME-CORE
```

The "**APP-NAME-CORE**" could be the following applications with suffix "-ca" for Core-A application or "-cm" for Core-M application:

- hello-world-ca
- lwip-ping-ca
- rpmsg-str-echo-ca
- rpmsg-str-echo-cm
- rpmsg-pingpong-ca
- rpmsg-lite-uart-sharing-rtos
- virtio-net-backend-ca
- virtio-net-backend-cm
- virtio-perf-ca

- virtio-perf-cm

3.2.2 Build with Standalone Mode

Some RTOS applications are in the repository: [heterogeneous-multicore](#), which is managed by the "west" tool. Therefore, you can use "west" to download all the software components and then use the ARM gcc toolchain to build the application directly.

1. Download and install the toolchain

The toolchain for building RTOS on Cortex-A Core:

```
mkdir ~/toolchains/;cd ~/toolchains/
wget https://developer.arm.com/-/media/Files/downloads/gnu/13.2.rel1/binrel/
arm-gnu-toolchain-13.2.rel1-x86_64-aarch64-none-elf.tar.xz
tar xf arm-gnu-toolchain-13.2.rel1-x86_64-aarch64-none-elf.tar.xz
```

The toolchain for building RTOS on Cortex-M Core:

```
cd ~/toolchains/
wget https://developer.arm.com/-/media/Files/downloads/gnu/13.2.rel1/binrel/
arm-gnu-toolchain-13.2.rel1-x86_64-arm-none-eabi.tar.xz
tar xf arm-gnu-toolchain-13.2.rel1-x86_64-arm-none-eabi.tar.xz
```

2. Download source code

This Heterogeneous Multicore project uses west to manage all related repos, west.yml provides the description and revision for other projects used by Heterogeneous Multicore.

Use the following command to download all the source code:

```
export revision=Real-Time-Edge-v3.1-202503
west init -m https://github.com/nxp-real-time-edge-sw/heterogeneous-
multicore.git --mr ${revision} workspace
cd workspace
west update
```

Note: Replace `${revision}` with any Real-Time Edge release you wish to use, it can also be "main" if you want to use the latest release.

3. Build the RTOS application

Each application provides a building script that you can use to build the application separately. You can even use a common building helper script "build_apps.sh" in the root directory of "heterogeneous-multicore" to build a single or an application for all boards.

- Build FreeRTOS Application

Take 'virtio_net_backend' application as an example. Use the following command to build the application for Cortex-A Core on i.MX 8M Plus EVK:

```
export ARMGCC_DIR=~/toolchains/arm-gnu-toolchain-13.2.Rel1-x86_64-aarch64-
none-elf
cd ~/workspace/heterogeneous-multicore/apps/virtio_net_backend/freertos/
boards/evkmimx8mm_ca53/armgcc_aarch64
./build_dds_release.sh
```

Then the RTOS images "virtio_net_backend_ca53.bin virtio_net_backend_ca53.elf" can be found in directory "dds_release".

And use the following commands to build the application for Cortex-M Core on i.MX 8M Plus EVK:

```
export ARMGCC_DIR=~/toolchains/arm-gnu-toolchain-13.2.Rel1-x86_64-arm-none-
eabi
cd ~/workspace/heterogeneous-multicore/apps/virtio_net_backend/freertos/
boards/evkmimx8mm_cm4/armgcc
./build_release.sh
```

Then the RTOS image "virtio_net_backend_cm4.bin virtio_net_backend_cm4.elf" can be found in directory "release".

- **Build the Zephyr application**

Building application on Cortex-A Core, for example, building Zephyr hello_world running on Cortex-A Core:

```
export ARMGCC_DIR=~/.toolchains/arm-gnu-toolchain-13.2.Rel1-x86_64-aarch64-
none-elf
export Zephyr_DIR=~/.workspace/zephyr
cd ~/.workspace/heterogeneous-multicore/apps/hello_world/zephyr/boards/
evkmimx8mm_ca53/armgcc_aarch64
./build_release.sh
```

Then, the following binary Zephyr images are built out:

```
./build_RTOS0_UART4/zephyr/hello_world_ca53_RTOS0_UART4.bin
./build_RTOS0_RAM_CONSOLE/zephyr/
hello_world_ca53_RTOS0_RAM_CONSOLE-0x93d00000.bin
./build_RTOS1_RAM_CONSOLE/zephyr/
hello_world_ca53_RTOS1_RAM_CONSOLE-0x94d00000.bin
./build_RTOS2_RAM_CONSOLE/zephyr/
hello_world_ca53_RTOS2_RAM_CONSOLE-0x95d00000.bin
./build_RTOS3_UART2/zephyr/hello_world_ca53_RTOS3_UART2.bin
./build_RTOS3_RAM_CONSOLE/zephyr/
hello_world_ca53_RTOS3_RAM_CONSOLE-0x96d00000.bin
```

- **Build with Helper script**

The "build_apps.sh" file in the root directory of "heterogeneous-multicore" can be used to build a single or all applications for all boards.

The following is help information for "build_apps.sh" tool:

```
./build_apps.sh [clean] - build or
clean all applications
./build_apps.sh [clean] [os] [board-list] [app-list] [core] - build or
clean specified applications
- os: specify freertos or zephyr or both if no specified.
- core: a-core or m-core
- board-list: specify one or some or all boards if no specified:
evkmimx8mm_ca53 evkmimx8mp_ca53 mcimx93evk_ca55 mcimx95evk_ca55
evkmimx8mm_cm4 evkmimx8mp_cm7 mcimx93evk_cm33
- app-list: specify one or some or all applications if no
specified: hello_world lwip_ping rpmsg_perf rpmsg_pingpong rpmsg_str_echo
rpmsg_uart_sharing soem_digital_io soem_servo virtio_net_backend
virtio_perf
```

Need to set toolchain environment variables "ARMGCC_DIR" firstly before using the tool, and set environment variables "Zephyr_DIR" for Zephyr building.

For example, use the tool to build all hello_world application on Cortex-M Core for all supported boards:

```
export ARMGCC_DIR=~/.toolchains/arm-gnu-toolchain-13.2.Rel1-x86_64-arm-none-
eabi
cd ~/.workspace/heterogeneous-multicore/
./build_apps.sh m-core hello_world
```

Use the tool to build all Zephyr application on Cortex-A Core for all supported boards:

```
export ARMGCC_DIR=~/.toolchains/arm-gnu-toolchain-13.2.Rel1-x86_64-aarch64-
none-elf
export Zephyr_DIR=~/.workspace/zephyr
cd ~/.workspace/heterogeneous-multicore/
```

```
./build_apps.sh a-core zephyr
```

After executing the tool, binary images built out can be found in the directory: "deploy/images", for example, the following are all the images both for A-Core and M-Core on both Zephyr and FreeRTOS.

```

deploy/
├── images
│   ├── evkmimx8mm_ca53
│   │   ├── freertos
│   │   │   ├── hello_world_ca53_RTOS0_RAM_CONSOLE-0x94bfff000.bin
│   │   │   ├── hello_world_ca53_RTOS0_RAM_CONSOLE-0x94bfff000.elf
│   │   │   ├── hello_world_ca53_RTOS0_UART4.bin
│   │   │   ├── hello_world_ca53_RTOS0_UART4.elf
│   │   │   ├── hello_world_ca53_RTOS1_RAM_CONSOLE-0x95bfff000.bin
│   │   │   ├── hello_world_ca53_RTOS1_RAM_CONSOLE-0x95bfff000.elf
│   │   │   ├── hello_world_ca53_RTOS2_RAM_CONSOLE-0x96bfff000.bin
│   │   │   ├── hello_world_ca53_RTOS2_RAM_CONSOLE-0x96bfff000.elf
│   │   │   ├── hello_world_ca53_RTOS3_RAM_CONSOLE-0x97bfff000.bin
│   │   │   ├── hello_world_ca53_RTOS3_RAM_CONSOLE-0x97bfff000.elf
│   │   │   ├── hello_world_ca53_RTOS3_UART2.bin
│   │   │   ├── hello_world_ca53_RTOS3_UART2.elf
│   │   │   ├── lwip_ping_ca53.bin
│   │   │   ├── lwip_ping_ca53.elf
│   │   │   ├── rpmsg_str_echo_ca53_RTOS0_RAM_CONSOLE-0x94bfff000.bin
│   │   │   ├── rpmsg_str_echo_ca53_RTOS0_RAM_CONSOLE-0x94bfff000.elf
│   │   │   ├── rpmsg_str_echo_ca53_RTOS0_UART4.bin
│   │   │   ├── rpmsg_str_echo_ca53_RTOS0_UART4.elf
│   │   │   ├── rpmsg_str_echo_ca53_RTOS1_RAM_CONSOLE-0x95bfff000.bin
│   │   │   ├── rpmsg_str_echo_ca53_RTOS1_RAM_CONSOLE-0x95bfff000.elf
│   │   │   ├── rt_latency_ca53.bin
│   │   │   ├── rt_latency_ca53.elf
│   │   │   ├── soem_digital_io_ca53.bin
│   │   │   ├── soem_digital_io_ca53.elf
│   │   │   ├── soem_servo_ca53.bin
│   │   │   ├── soem_servo_ca53.elf
│   │   │   ├── virtio_net_backend_ca53.bin
│   │   │   ├── virtio_net_backend_ca53.elf
│   │   │   ├── virtio_perf_ca53.bin
│   │   │   └── virtio_perf_ca53.elf
│   │   └── zephyr
│   │       ├── hello_world_ca53_RTOS0_RAM_CONSOLE-0x93d000000.bin
│   │       ├── hello_world_ca53_RTOS0_RAM_CONSOLE-0x93d000000.elf
│   │       ├── hello_world_ca53_RTOS0_UART4.bin
│   │       ├── hello_world_ca53_RTOS0_UART4.elf
│   │       ├── hello_world_ca53_RTOS1_RAM_CONSOLE-0x94d000000.bin
│   │       ├── hello_world_ca53_RTOS1_RAM_CONSOLE-0x94d000000.elf
│   │       ├── hello_world_ca53_RTOS2_RAM_CONSOLE-0x95d000000.bin
│   │       ├── hello_world_ca53_RTOS2_RAM_CONSOLE-0x95d000000.elf
│   │       ├── hello_world_ca53_RTOS3_RAM_CONSOLE-0x96d000000.bin
│   │       ├── hello_world_ca53_RTOS3_RAM_CONSOLE-0x96d000000.elf
│   │       ├── hello_world_ca53_RTOS3_UART2.bin
│   │       ├── hello_world_ca53_RTOS3_UART2.elf
│   │       ├── rt_latency_ca53.bin
│   │       └── rt_latency_ca53.elf
│   └── evkmimx8mm_cm4
│       ├── freertos
│       │   ├── hello_world_cm4_UART4.bin
│       │   ├── hello_world_cm4_UART4.elf
│       │   ├── rpmsg_str_echo_cm4.bin
│       │   ├── rpmsg_str_echo_cm4.elf
│       │   └── rpmsg_uart_sharing.bin

```

```

— rpmsg_uart_sharing.elf
— virtio_net_backend_cm4.bin
— virtio_net_backend_cm4.elf
— virtio_perf_cm4.bin
— virtio_perf_cm4.elf
— evkmimx8mp_ca53
— freertos
— hello_world_ca53_RTOS0_RAM_CONSOLE-0xc0fff000.bin
— hello_world_ca53_RTOS0_RAM_CONSOLE-0xc0fff000.elf
— hello_world_ca53_RTOS0_UART4.bin
— hello_world_ca53_RTOS0_UART4.elf
— hello_world_ca53_RTOS1_RAM_CONSOLE-0xc1fff000.bin
— hello_world_ca53_RTOS1_RAM_CONSOLE-0xc1fff000.elf
— hello_world_ca53_RTOS2_RAM_CONSOLE-0xc2fff000.bin
— hello_world_ca53_RTOS2_RAM_CONSOLE-0xc2fff000.elf
— hello_world_ca53_RTOS3_RAM_CONSOLE-0xc3fff000.bin
— hello_world_ca53_RTOS3_RAM_CONSOLE-0xc3fff000.elf
— hello_world_ca53_RTOS3_UART2.bin
— hello_world_ca53_RTOS3_UART2.elf
— lwip_ping_ca53.bin
— lwip_ping_ca53.elf
— rpmsg_perf_ca53.bin
— rpmsg_perf_ca53.elf
— rpmsg_pingpong_master_ca53_RAM_CONSOLE-0xc1fff000.bin
— rpmsg_pingpong_master_ca53_RAM_CONSOLE-0xc1fff000.elf
— rpmsg_pingpong_master_ca53_UART4.bin
— rpmsg_pingpong_master_ca53_UART4.elf
— rpmsg_pingpong_remote_ca53_RAM_CONSOLE-0xc0fff000.bin
— rpmsg_pingpong_remote_ca53_RAM_CONSOLE-0xc0fff000.elf
— rpmsg_pingpong_remote_ca53_UART4.bin
— rpmsg_pingpong_remote_ca53_UART4.elf
— rpmsg_str_echo_ca53_RTOS0_RAM_CONSOLE-0xc0fff000.bin
— rpmsg_str_echo_ca53_RTOS0_RAM_CONSOLE-0xc0fff000.elf
— rpmsg_str_echo_ca53_RTOS0_UART4.bin
— rpmsg_str_echo_ca53_RTOS0_UART4.elf
— rpmsg_str_echo_ca53_RTOS1_RAM_CONSOLE-0xc1fff000.bin
— rpmsg_str_echo_ca53_RTOS1_RAM_CONSOLE-0xc1fff000.elf
— rt_latency_ca53.bin
— rt_latency_ca53.elf
— soem_digital_io_ca53.bin
— soem_digital_io_ca53.elf
— soem_servo_ca53.bin
— soem_servo_ca53.elf
— virtio_net_backend_ca53.bin
— virtio_net_backend_ca53.elf
— virtio_perf_ca53.bin
— virtio_perf_ca53.elf
— zephyr
— hello_world_ca53_RTOS0_RAM_CONSOLE-0xc0100000.bin
— hello_world_ca53_RTOS0_RAM_CONSOLE-0xc0100000.elf
— hello_world_ca53_RTOS0_UART4.bin
— hello_world_ca53_RTOS0_UART4.elf
— hello_world_ca53_RTOS1_RAM_CONSOLE-0xc1100000.bin
— hello_world_ca53_RTOS1_RAM_CONSOLE-0xc1100000.elf
— hello_world_ca53_RTOS2_RAM_CONSOLE-0xc2100000.bin
— hello_world_ca53_RTOS2_RAM_CONSOLE-0xc2100000.elf
— hello_world_ca53_RTOS3_RAM_CONSOLE-0xc3100000.bin
— hello_world_ca53_RTOS3_RAM_CONSOLE-0xc3100000.elf
— hello_world_ca53_RTOS3_UART2.bin
— hello_world_ca53_RTOS3_UART2.elf

```



```

├── rt_latency_ca53.bin
├── rt_latency_ca53.elf
├── evkmimx8mp_cm7
│   └── freertos
│       ├── hello_world_cm7_UART4.bin
│       ├── hello_world_cm7_UART4.elf
│       ├── rpmsg_perf_cm7.bin
│       ├── rpmsg_perf_cm7.elf
│       ├── rpmsg_str_echo_cm7.bin
│       ├── rpmsg_str_echo_cm7.elf
│       ├── virtio_net_backend_cm7.bin
│       ├── virtio_net_backend_cm7.elf
│       ├── virtio_perf_cm7.bin
│       └── virtio_perf_cm7.elf
├── mcimx93evk_ca55
│   └── freertos
│       ├── hello_world_ca55_RTOS0_RAM_CONSOLE-0xd0fff000.bin
│       ├── hello_world_ca55_RTOS0_RAM_CONSOLE-0xd0fff000.elf
│       ├── hello_world_ca55_RTOS0_UART2.bin
│       ├── hello_world_ca55_RTOS0_UART2.elf
│       ├── hello_world_ca55_RTOS1_RAM_CONSOLE-0xd1fff000.bin
│       ├── hello_world_ca55_RTOS1_RAM_CONSOLE-0xd1fff000.elf
│       ├── hello_world_ca55_RTOS1_UART1.bin
│       ├── hello_world_ca55_RTOS1_UART1.elf
│       ├── lwip_ping_ca55.bin
│       ├── lwip_ping_ca55.elf
│       ├── rt_latency_ca55.bin
│       ├── rt_latency_ca55.elf
│       ├── soem_digital_io_ca55.bin
│       ├── soem_digital_io_ca55.elf
│       ├── soem_servo_ca55.bin
│       ├── soem_servo_ca55.elf
│       ├── virtio_net_backend_ca55.bin
│       └── virtio_net_backend_ca55.elf
│   └── zephyr
│       ├── hello_world_ca55_RTOS0_RAM_CONSOLE-0xd0100000.bin
│       ├── hello_world_ca55_RTOS0_RAM_CONSOLE-0xd0100000.elf
│       ├── hello_world_ca55_RTOS0_UART2.bin
│       ├── hello_world_ca55_RTOS0_UART2.elf
│       ├── hello_world_ca55_RTOS1_RAM_CONSOLE-0xd1100000.bin
│       ├── hello_world_ca55_RTOS1_RAM_CONSOLE-0xd1100000.elf
│       ├── hello_world_ca55_RTOS1_UART1.bin
│       ├── hello_world_ca55_RTOS1_UART1.elf
│       ├── rt_latency_ca55.bin
│       └── rt_latency_ca55.elf
├── mcimx93evk_cm33
│   └── freertos
│       ├── hello_world_cm33_UART2.bin
│       ├── hello_world_cm33_UART2.elf
│       ├── rpmsg_uart_sharing.bin
│       ├── rpmsg_uart_sharing.elf
│       ├── virtio_net_backend_cm33.bin
│       └── virtio_net_backend_cm33.elf
├── mcimx95evk_ca55
│   └── freertos
│       ├── hello_world_ca55_RTOS0_RAM_CONSOLE-0xd0fff000.bin
│       ├── hello_world_ca55_RTOS0_RAM_CONSOLE-0xd0fff000.elf
│       ├── hello_world_ca55_RTOS0_UART3.bin
│       ├── hello_world_ca55_RTOS0_UART3.elf
│       └── hello_world_ca55_RTOS1_RAM_CONSOLE-0xd1fff000.bin

```

```
├── hello_world_ca55_RTOS1_RAM_CONSOLE-0xd1fff000.elf
├── hello_world_ca55_RTOS2_RAM_CONSOLE-0xd2fff000.bin
├── hello_world_ca55_RTOS2_RAM_CONSOLE-0xd2fff000.elf
├── hello_world_ca55_RTOS2_UART1.bin
├── hello_world_ca55_RTOS2_UART1.elf
├── rt_latency_ca55.bin
├── rt_latency_ca55.elf
└── zephyr
    ├── hello_world_ca55_RTOS0_RAM_CONSOLE-0xd0100000.bin
    ├── hello_world_ca55_RTOS0_RAM_CONSOLE-0xd0100000.elf
    ├── hello_world_ca55_RTOS0_UART3.bin
    ├── hello_world_ca55_RTOS0_UART3.elf
    ├── hello_world_ca55_RTOS1_RAM_CONSOLE-0xd1100000.bin
    ├── hello_world_ca55_RTOS1_RAM_CONSOLE-0xd1100000.elf
    ├── hello_world_ca55_RTOS2_RAM_CONSOLE-0xd2100000.bin
    ├── hello_world_ca55_RTOS2_RAM_CONSOLE-0xd2100000.elf
    ├── hello_world_ca55_RTOS2_UART1.bin
    ├── hello_world_ca55_RTOS2_UART1.elf
    ├── rt_latency_ca55.bin
    └── rt_latency_ca55.elf
```

3.3 Flexible Real-time System

NXP MPU platforms support the Flexible Real-time System. This system can run single or multiple RTOS on Cortex-M core and Cortex-A core with or without running Linux on Cortex-A Core simultaneously.

The system provides RAM Console to make it easy to debug multiple OSES in case of not enough physical UART Consoles. The system also provides some common software stacks such as lwIP networking stack on Cortex-A core or Cortex-M core.

3.3.1 Heterogeneous Multicore RAM Console

RAM Console is a virtual debug console that can be used by the RTOS. It provides common console APIs to print and save console log to a reserved memory region.

If multiple OSES run on an MPU platform, some OSES can use physical UART as the debug console. In such a case, the other OSES can use the RAM Console. This requires to dump the other OS's RAM Console from the OS which is using the physical UART Console.

There are two methods to dump the RAM Console log: by using U-Boot commands or by using the Linux Userspace tool.

3.3.1.1 Using RAM Console in FreeRTOS

This chapter describes how to develop a FreeRTOS application based on the RAM Console.

3.3.1.1.1 RAM Console technical details

The RAM Console driver is located at the debug console (utilities/debug_console) in the repository [mcux-sdk](#). There is a 64-byte Console Header at the start of the RAM Console log memory. [Figure 28](#) shows the memory layout.




Figure 28. RAM console memory layout

The RAM Console Header includes the following parts:

- 16 bytes fixed string flag: "RAM_CONSOLE".
- console_start_address: the memory physical address for the console_start of Console buffer.
- console_buffer_length: the length in bytes of Console Buffer from console_start to console_end.
- console_cursor_position: The current console cursor increased position from console starting.

3.3.1.1.2 Develop with RAM Console

Take the "hello_world" program for the "evkmimx8mm_ca53" board in the repository [heterogeneous-multicore](#) as an example. Then follow the steps listed below to use RAM Console as debug console in RTOS application.

1. Reserve a memory block for RAM Console, and add mmu mapping entry in `app_mmu.h`.
2. In application's `CMakeLists.txt`

```
...  
SET(DEBUG_CONSOLE_CONFIG "-DSDK_DEBUGCONSOLE=2 -DCONFIG_RAM_CONSOLE")  
...  
include(utility_assert_ram_console)  
include(utility_ram_console)  
...
```

3. Call `RamConsole_Init()` with RAM Console memory block's address and size to initialize RAM Console

```
#ifdef CONFIG_RAM_CONSOLE  
    RamConsole_Init(RAM_CONSOLE_ADDR, RAM_CONSOLE_SIZE);  
#else  
    BOARD_InitDebugConsole();  
#endif
```

4. In general, while running multiple RTOS instances, all the RAM Console memory addresses can be printed out to physical debug console from the RTOS instances using physical debug console.

3.3.1.2 Use RAM Console in Zephyr

The following sections describe how to develop a Zephyr application based on the RAM Console.

3.3.1.2.1 RAM Console Technical Details

In Zephyr RAM console, there is the same 64 bytes Console Header as the one in FreeRTOS, refer to [Section 3.3.1.1.1](#).

The RAM Console driver in Zephyr is `drivers/console/ram_console.c`.

3.3.1.2.2 Develop with the RAM Console

Take the "hello_world" program for the "evkmimx8mm_ca53" board in the repository [heterogeneous-multicore](#) as an example. Then follow the steps listed below to use RAM Console as debug console in Zephyr application.

1. Enable the RAM console driver in the `prj.conf` file, and disable the UART console driver.

```
CONFIG_RAM_CONSOLE=y
CONFIG_UART_CONSOLE=n
```

2. Add a device tree node to add a memory region and section for the RAM console buffer in the board `dtb` file.

```
chosen {
    /delete-property/ zephyr,console;
    /delete-property/ zephyr,shell-uart;
    zephyr,ram-console = &ram_console;
};
ram_console: memory@93d00000 {
    compatible = "zephyr,memory-region";
    reg = <0x93d00000 DT_SIZE_K(4)>;
    zephyr,memory-region = "RAM_CONSOLE";
};
```

3. In general, while running multiple RTOS instances, all the RAM Console memory addresses can be printed out to physical debug console from the RTOS0 instance using the physical debug console.

3.3.1.3 Dump RAM Console Log

RAM Console Log can be dumped from the U-Boot command line or from Linux userspace. While running multiple OSes on the MPU platform, the OS instances must be started using RAM Console first. Then start the OS using the physical UART Console. If the last OS using physical UART Console is RTOS, it can dump the other OS's RAM Console log before U-Boot command line starts this last RTOS. Otherwise it has to dump the console memory by using JTAG tools. But if the last OS using physical UART Console is Linux, we can still dump the other OS's RAM Console log by using Linux userspace tool after Linux boots up.

- Dump from U-Boot

In U-Boot command line, use the "md" command to dump the whole RAM Console memory including RAM Console Header

```
u-boot=> dcache flush; md C1FFF000
c1fff000: 5f4d4152 534e4f43 00454c4f 00000000  RAM_CONSOLE.....
c1fff010: c1fff040 00000000 00000fc0 00000000  @.....
c1fff020: 0000024d 00000000 00000000 00000000  M.....
c1fff030: 00000000 00000000 00000000 00000000  .....
c1fff040: 6f430a0d 78657472 3335412d 5452203a  ..Cortex-A53: RT
```

```

c1fff050: 3a31534f 6c654820 77206f6c 646c726f OS1: Hello world
c1fff060: 65522021 742d6c61 20656d69 65676445 ! Real-time Edge
c1fff070: 206e6f20 584d494d 2d504d38 0d4b5645 on MIMX8MP-EVK.
c1fff080: 6369740a 63617420 63697420 63617420 .tic tac tic tac
c1fff090: 63697420 63617420 63697420 63617420 tic tac tic tac
c1fff0a0: 63697420 63617420 63697420 63617420 tic tac tic tac
c1fff0b0: 63697420 63617420 63697420 63617420 tic tac tic tac
c1fff0c0: 63697420 63617420 63697420 63617420 tic tac tic tac

```

- Dump from Linux

There is a Linux userspace tool provided in the repository [heterogeneous-multicore](#) "tool" directory. Use this tool to dump RAM Console log.

The following shows the help information of the tool:

```

root@imx8mp-lpddr4-evk:~# ram_console_dump

RAM Console tool to dump console log.
Usage:ram_console_dump -a buffer_address [-r refresh_period_in_seconds]
        -a buffer_address:      RAM Console buffer physical address
        -r refresh_period:      Refresh period time in seconds for
        continuous dump, ctrl+c to exit

```

The following log shows the example to use this tool. It dumps the RAM Console from address 0xC1FFF000 in the loop period of one second:

```

root@imx8mp-lpddr4-evk:~# ram_console_dump -a 0xC1FFF000 -r 1
RAM Console@0xc1fff000:

Cortex-A53: RTOS1: Hello world! Real-time Edge on MIMX8MP-EVK
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac
tic tac tic tac tic >

```

- Dump by using debugger tools:

In case there is no U-Boot or Linux running on the platform, users can use hardware debugger tools such as JLink or TRACE32 to dump the console log from a specified buffer memory.

3.3.2 Heterogeneous Multicore hello_world

This section describes details of the Heterogeneous Multicore `hello_world` application.

3.3.2.1 Overview

Heterogeneous Multicore `hello_world` application demonstrates flexible Real-time System on MPU platforms. It can run single or multiple FreeRTOS or Zephyr RTOS instances on Cortex-A Core and single FreeRTOS instance on Cortex-M with or without running Linux simultaneously.

Note: In this release, it only supports the Zephyr `hello_world` application on Cortex-A cores.

Take i.MX 8M Plus Applications Processor as an example. It has four Cortex-A53 cores and one Cortex-M7 Core. It could run the following use cases on this MPU platform:

Table 24. Flexible Real-time System on i.MX 8M Plus

ID	M7	A53	A53	A53	A53
0	RTOS	SMP Linux			
1	RTOS	SMP Linux			RTOS

Table 24. Flexible Real-time System on i.MX 8M Plus...continued

ID	M7	A53	A53	A53	A53
2	RTOS	SMP Linux		RTOS	RTOS
3	RTOS	Linux	RTOS	RTOS	RTOS
4	RTOS	RTOS	RTOS	RTOS	RTOS

3.3.2.2 Technical Points

- Debug Console

Take i.MX 8M Plus Applications Processor as an example, in general, UART4 is used for Cortex-M Core RTOS or Cortex-A Core RTOS, UART2 is used for Linux, the other RTOS instance can use RAM Console. The hello_world application could build all RTOS instances with any possible UART Console or RAM Console, the following images are for i.MX 8M Plus EVK, different RTOS instance RTOS0, RTOS1, RTOS2 and RTOS3 use different memory space, for each RTOS instance also provides the images with different Debug Console(UART4, UART2 or RAM Console (the buffer address is appended in the image name), can run any of these images according to usecase setup in order to avoid memory and Debug Console confliction.

```
evkmimx8mp/
├─ hello_world_ca53_RTOS0_RAM_CONSOLE-0xc0fff000.bin      #RTOS0 with RAM
  Console at 0xc0fff000
├─ hello_world_ca53_RTOS0_RAM_CONSOLE-0xc0fff000.elf      #RTOS0 with RAM
  Console at 0xc0fff000
├─ hello_world_ca53_RTOS0_UART4.bin                        #RTOS0 with UART4
  Debug Console
├─ hello_world_ca53_RTOS0_UART4.elf                        #RTOS0 with UART4
  Debug Console
├─ hello_world_ca53_RTOS1_RAM_CONSOLE-0xc1fff000.bin      #RTOS1 with RAM
  Console at 0xc1fff000
├─ hello_world_ca53_RTOS1_RAM_CONSOLE-0xc1fff000.elf      #RTOS1 with RAM
  Console at 0xc1fff000
├─ hello_world_ca53_RTOS2_RAM_CONSOLE-0xc2fff000.bin      #RTOS2 with RAM
  Console at 0xc2fff000
├─ hello_world_ca53_RTOS2_RAM_CONSOLE-0xc2fff000.elf      #RTOS2 with RAM
  Console at 0xc2fff000
├─ hello_world_ca53_RTOS3_RAM_CONSOLE-0xc3fff000.bin      #RTOS3 with RAM
  Console at 0xc3fff000
├─ hello_world_ca53_RTOS3_RAM_CONSOLE-0xc3fff000.elf      #RTOS3 with RAM
  Console at 0xc3fff000
├─ hello_world_ca53_RTOS3_UART2.bin                        #RTOS3 with UART2
  Debug Console
└─ hello_world_ca53_RTOS3_UART2.elf                        #RTOS3 with UART2
  Debug Console
```

- Memory Usage

– For FreeRTOS applications in the repository: [heterogeneous-multicore](#), the memory resource used by FreeRTOS kernel is defined in "os/freertos/Core_AArch64/boards/<PLAT_NAME>/rtos_memory.h", for example, in the following rtos_memory.h for i.MX 8M Plus, RTOS0, RTOS1, RTOS2 and RTOS3, each RTOS uses 16M bytes memory from 0xC0000000, 0xC1000000, 0xC2000000 and 0xC3000000 respectively.

```
#ifndef _RTOS_MEMORY_H_
#define _RTOS_MEMORY_H_

/* Memory used by RTOS kernel */
#define M_INTERRUPTS_BASE (0xC0000000 + 0x1000000 * RTOSID)
#define M_INTERRUPTS_LEN 0x00002000 /* 8 kB */
```



```

#define M_TEXT_BASE                (M_INTERRUPTS_BASE + M_INTERRUPTS_LEN)
#define M_TEXT_LEN                 0x005FE000 /* ~6 MB */

#define M_DATA_BASE                (M_TEXT_BASE + M_TEXT_LEN)
#define M_DATA_LEN                 0x005FE000 /* ~6 MB */

#define M_STACKS_BASE              (M_DATA_BASE + M_DATA_LEN)
#define M_STACKS_LEN               0x00002000 /* 8 kB */

#define M_STACKS_NC_BASE           (M_STACKS_BASE + M_STACKS_LEN)
#define M_STACKS_NC_LEN            0x003FF000 /* ~4 MB */

/* Memory used by RAM Console */
#define RAM_CONSOLE_ADDR           (M_STACKS_NC_BASE + M_STACKS_NC_LEN)
#define RAM_CONSOLE_SIZE           0x00001000 /* 4KB */

#define RTOS_MEM_LEN                0x01000000 /* 16 MB */

#endif

```

- For Zephyr applications in the repository: [heterogeneous-multicore](#), the memory resource used by Zephyr kernel is defined in "apps/hello_world/zephyr/boards/<PLAT_NAME>/rtos_memory.h", for example, in the following rtos_memory.h for i.MX 8M Plus, RTOS0, RTOS1, RTOS2 and RTOS3, each RTOS uses 1M bytes memory for Zephyr kernel starting from 0xC0000000, 0xC1000000, 0xC2000000 and 0xC3000000 respectively, and 4K bytes memory for RAM console buffer appended.

```

#if (RTOSID == 0)
#define SRAM_BASE                  c0000000
#define RAM_CONSOLE_BASE           c0100000
#elif (RTOSID == 1)
#define SRAM_BASE                  c1000000
#define RAM_CONSOLE_BASE           c1100000
#elif (RTOSID == 2)
#define SRAM_BASE                  c2000000
#define RAM_CONSOLE_BASE           c2100000
#elif (RTOSID == 3)
#define SRAM_BASE                  c3000000
#define RAM_CONSOLE_BASE           c3100000
#else
#error "Unsupported RTOSID!"
#endif

```

- If run Linux kernel with RTOS simultaneously, need to reserve the memory space used by RTOS in device tree, for example, in imx8mp-evk-multicore-rtos.dts:

```

/*
 * Reserve up to 48MB (16MB x 3) for three FreeRTOS instances
 * running on three Cortex-A Cores when booting Linux on at least on Cortex-
 * A Core.
 */
ca53_reserved: ca53@c0000000 {
    no-map;
    reg = <0 0xc0000000 0x0 0x30000000>;
};

/* Reserve 16MB for RTOS running on CM7 */
m7_reserved: m7@80000000 {
    no-map;
    reg = <0 0x80000000 0 0x10000000>;
};

```

```
};
```

3.3.2.3 Running flexible multicore hello_world application

3.3.2.3.1 Running use cases on i.MX 8M Plus LPDDR4 EVK

3.3.2.3.1.1 Overview

The following use cases can run on i.MX 8M Plus LPDDR4 EVK:

Table 25. Flexible Real-time System on i.MX 8M Plus

ID	M7	A53 Core0	A53 Core1	A53 Core2	A53 Core3
0	RTOS	SMP Linux			
1	RTOS	SMP Linux			RTOS
2	RTOS	SMP Linux		RTOS	RTOS
3	RTOS	Linux	RTOS	RTOS	RTOS
4	RTOS	RTOS	RTOS	RTOS	RTOS

By default, the RTOS of the Cortex-M Core uses UART4 as the debug console. Refer to [Section 2.8.1](#) for how to boot Cortex-M Core's RTOS image.

RTOS of the Cortex-A Core could run with UART4 Console, RAM console, or UART2 console. The following RTOS images of the Cortex-A Core are provided:

```
/examples/heterogeneous-multicore/hello-world-ca/ddr_release/ #FreeRTOS A53
hello_world Images
— hello_world_ca53_RTOS0_RAM_CONSOLE-0xc0fff000.bin
— hello_world_ca53_RTOS0_RAM_CONSOLE-0xc0fff000.elf
— hello_world_ca53_RTOS0_UART4.bin
— hello_world_ca53_RTOS0_UART4.elf
— hello_world_ca53_RTOS1_RAM_CONSOLE-0xc1fff000.bin
— hello_world_ca53_RTOS1_RAM_CONSOLE-0xc1fff000.elf
— hello_world_ca53_RTOS2_RAM_CONSOLE-0xc2fff000.bin
— hello_world_ca53_RTOS2_RAM_CONSOLE-0xc2fff000.elf
— hello_world_ca53_RTOS3_RAM_CONSOLE-0xc3fff000.bin
— hello_world_ca53_RTOS3_RAM_CONSOLE-0xc3fff000.elf
— hello_world_ca53_RTOS3_UART2.bin
— hello_world_ca53_RTOS3_UART2.elf
/examples/heterogeneous-multicore/hello-world-ca-zephyr/ #Zephyr A53
hello_world Images
— hello_world_ca53_RTOS0_RAM_CONSOLE-0xc0100000.bin
— hello_world_ca53_RTOS0_RAM_CONSOLE-0xc0100000.elf
— hello_world_ca53_RTOS0_UART4.bin
— hello_world_ca53_RTOS0_UART4.elf
— hello_world_ca53_RTOS1_RAM_CONSOLE-0xc1100000.bin
— hello_world_ca53_RTOS1_RAM_CONSOLE-0xc1100000.elf
— hello_world_ca53_RTOS2_RAM_CONSOLE-0xc2100000.bin
— hello_world_ca53_RTOS2_RAM_CONSOLE-0xc2100000.elf
— hello_world_ca53_RTOS3_RAM_CONSOLE-0xc3100000.bin
— hello_world_ca53_RTOS3_RAM_CONSOLE-0xc3100000.elf
— hello_world_ca53_RTOS3_UART2.bin
— hello_world_ca53_RTOS3_UART2.elf
```

```
/examples/heterogeneous-multicore/hello-world-cm/release/      #FreeRTOS M7
hello_world Images
└─ hello_world_cm7_UART4.bin
└─ hello_world_cm7_UART4.el
```

RAM Console address is appended at the end of image name, for example `hello_world_ca53_RTOS3_RAM_CONSOLE-0xc3100000.bin`, it uses RAM Console which buffer address is `0xc3100000`.

Refer to [Section 2.6.3](#) for how to boot RTOS on Cortex-A Core, and refer to [Section 3.3.1.3](#) for how to dump RAM Console log.

The following chapters introduce how to use U-Boot commands or Linux remoteproc to boot FreeRTOS or Zephyr use cases.

3.3.2.3.1.2 Running flexible multicore use cases by using U-Boot Commands

- This section describes the steps to run FreeRTOS applications:

To run FreeRTOS on M4 core, follow the steps below. These steps are for running two FreeRTOS instances on A53 Core2 and Core3, and a SMP Linux kernel on A53 Core0 and Core1.

1. Boot RTOS of Cortex-M core.

```
u-boot=> ext4load mmc 1:2 0x48000000 /examples/heterogeneous-multicore/
hello-world-cm/release/hello_world_cm7_UART4.bin;
u-boot=> cp.b 0x48000000 0x7e0000 20000;
u-boot=> bootaux 0x7e0000
```

Then the following log is displayed from UART4:

```
Cortex-M7: RTOS0: Hello world! Real-time Edge on MIMX8MP-EVK
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac
```

2. Boot the first RTOS of Cortex-A core on A53 Core2.

```
u-boot=> ext4load mmc 1:2 0xC0000000 /examples/heterogeneous-multicore/
hello-world-ca/ddr_release/hello_world_ca53_RTOS0_RAM_CONSOLE-0xc0fff000.bin
u-boot=> dcache flush; icache flush; cpu 2 release 0xC0000000
```

Then, check the RAM Console log as follows:

```
u-boot=> dcache flush; md 0xC0FFF000
c0fff000: 5f4d4152 534e4f43 00454c4f 00000000 RAM_CONSOLE.....
c0fff010: c0fff040 00000000 00000fc0 00000000 @.....
c0fff020: 00000182 00000000 00000000 00000000 .....
c0fff030: 00000000 00000000 00000000 00000000 .....
c0fff040: 6f430a0d 78657472 3335412d 5452203a ..Cortex-A53: RT
c0fff050: 3a30534f 6c654820 77206f6c 646c326f OS0: Hello wo2ld
c0fff060: 65520021 742d6c61 20656d69 65676445 !.Real-time Edge
c0fff070: 206e4f20 584d494d 2d500d38 0d4b5645 On MIMX8.P-EVK.
c0fff080: 4f54560a 203a3153 204d4152 736e6f63 .VTOS1: RAM cons
c0fff090: 40656c6f 31637830 30666664 0a0d3030 ole@0xc1dff000..
c0fff0a0: 534f5452 52203a32 63204d40 6f73666f RTOS2: R@M cofso
c0fff0b0: 3040656c 66326130 30306666 520a0d30 le@00a2fff000..R
c0fff0c0: 33534f5c 4152203a 6f63204d 6c6f736e \OS3: RAM consol
c0fff0d0: 78304065 66663363 30303066 69740a0d e@0xc3fff000..ti
c0fff0e0: 61742063 69742023 61742063 69742063 c ta# tic tac ti
c0fff0f0: 61742063 69742063 61742043 69742063 c tac tiC tac ti
```

3. Boot the second Cortex-A Core's RTOS on A53 Core3

```
u-boot=> ext4load mmc 1:2 0xC1000000 /examples/heterogeneous-multicore/
hello-world-ca/ddr_release/hello_world_ca53_RTOS1_RAM_CONSOLE-0xc1fff000.bin
```

```
u-boot=> dcache flush; icache flush; cpu 3 release 0xC1000000
```

Then check the RAM Console's log as follows:

```

u-boot=> dcache flush; md 0xC1FFF000
c1fff000: 5f4d4152 534e4f43 00454c4f 00000000  RAM_CONSOLE.....
c1fff010: c1fff040 00000000 00000fc0 00000000  @_.....
c1fff020: 00000079 00000000 00000000 00000000  y.....
c1fff030: 00000000 00000000 00000000 00000000  .....
c1fff040: 6f430a0d 78657472 3335412d 5452203a  ..Cortex-A53: RT
c1fff050: 3a31534f 6c654820 77206f6c 646c726f  OS1: Hello world
c1fff060: 65522021 742d6c61 20656d69 65676445  ! Real-time Edge
c1fff070: 206e6f20 584d494d 2d504d38 0d4b5645  on MIMX8MP-EVK.
c1fff080: 6369740a 63617420 63697420 63617420  .tic tac tic tac
c1fff090: 63697420 63617420 63697420 63617420  tic tac tic tac
c1fff0a0: 63697420 63617420 63697420 63617420  tic tac tic tac
c1fff0b0: 63697420 63617420 00000020 00000000  tic tac .....
c1fff0c0: 00000000 00000000 00000000 00000000  .....
c1fff0d0: 00000000 00000000 00000000 00000000  .....
c1fff0e0: 00000000 00000000 00000000 00000000  .....
c1fff0f0: 00000000 00000000 00000000 00000000  .....

```

4. Boot SMP Linux from A53 Core0 and Core1:

If M-Core is booting up, use `imx8mp-evk-multicore-rpmsg.dtb`

```
u-boot=> setenv fdtfile imx8mp-evk-multicore-rpmsg.dtb
```

Otherwise, can also use `imx8mp-evk-multicore-rtos.dtb`:

```
u-boot=> setenv fdtfile imx8mp-evk-multicore-rtos.dtb
```

Then, boot up the kernel using the command below:

```
u-boot=> setenv mmcargs $mmcargs clk_ignore_unused
u-boot=> boot
```

After Linux kernel boots up, check Cortex-A Core's RTOS log from the RAM console:

[illegible]

- Run the Zephyr applications:

Follow the following steps to run two Zephyr instances on A53 Core2 and Core3, and run a SMP Linux kernel on A53 Core0 and Core1.

1. Boot the first Cortex-A Core's Zephyr instance on A53 Core2 with UART4:

```
u-boot=> ext4load mmc 1:2 0xC0000000 /examples/heterogeneous-multicore/
hello-world-ca-zephyr/hello_world_ca53_RTOS0_UART4.bin
u-boot=> dcache flush; icache flush; cpu 2 release 0xC0000000
```

Then the UART4 shows the logs as follows:

```
*** Booting Zephyr OS build zephyr-v3.5.0-27-g24a3599ab7a0 ***
Cortex-A53: RTOS0: Hello World! Real-time Edge on mimx8mp_evk_a53
RTOS1: RAM console@0xc1100000
RTOS2: RAM console@0xc2100000
RTOS3: RAM console@0xc3100000
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
```

2. Boot the second Cortex-A Core's Zephyr instance on A53 Core3

```
u-boot=> ext4load mmc 1:2 0xC1000000 /examples/heterogeneous-multicore/
hello-world-ca-zephyr/hello_world_ca53_RTOS1_RAM_CONSOLE-0xc1100000.bin
u-boot=> dcache flush; icache flush; cpu 3 release 0xC1000000
```

Then check the RAM Console's log as follows:

```
u-boot=> dcache flush; md 0xC1100000
c1100000: 202a2a2a 746f6f42 20676e69 6870655a *** Booting Zeph
c1100010: 4f207279 75622053 20646c69 6870657a yr OS build zeph
c1100020: 762d7279 2e352e33 37322d30 3432672d yr-v3.5.0-27-g24
c1100030: 39353361 37626139 2a203061 430a2a2a a3599ab7a0 ***.C
c1100040: 6574726f 35412d78 52203a33 31534f54 ortex-A53: RTOS1
c1100050: 6548203a 206f6c6c 6c726f57 52202164 : Hello World! R
c1100060: 2d6c6165 656d6974 67644520 6e6f2065 eal-time Edge on
c1100070: 6d696d20 706d3878 6b76655f 3335615f mimx8mp_evk_a53
c1100080: 6369740a 63617420 63697420 63617420 .tic tac tic tac
c1100090: 63697420 63617420 63697420 63617420 tic tac tic tac
c11000a0: 63697420 63617420 63697420 63617420 tic tac tic tac
c11000b0: 63697420 63617420 63697420 63617420 tic tac tic tac
c11000c0: 63697420 63617420 63697420 63617420 tic tac tic tac
c11000d0: 69740a20 61742063 69742063 61742063 .tic tac tic ta
c11000e0: 69742063 61742063 69742063 61742063 c tic tac tic ta
c11000f0: 69742063 61742063 69742063 61742063 c tic tac tic ta
```

3. Boot SMP Linux from A53 Core0 and Core1:

```
u-boot=> setenv fdtfile imx8mp-evk-multicore-rtos.dtb
```

Then, boot up the kernel using the command below:

```
u-boot=> setenv mmcargs $mmcargs clk_ignore_unused
u-boot=> boot
```

3.3.2.3.1.3 Running flexible multicore use cases by using remoteproc

In addition to using U-Boot commands, users can also use remoteproc under Linux to run RTOS. Take the following use case as example, using the flexible hello_world application images to set up the following use case on the i.MX 8M Plus board.

As can using remoteproc to start and stop RTOS dynamically without rebooting the board, so will demo switching from use case #1 to use case #2 dynamically by stopping FreeRTOS on Cortex-A Core2 and start a Zephyr on it.

Table 26. Dynamic Flexible Real-time System on i.MX 8M Plus

ID	M7	A53-Core0	A53-Core1	A53-Core2	A53-Core3
1	FreeRTOS	SMP Linux		FreeRTOS	FreeRTOS
2	FreeRTOS	SMP Linux		Zephyr	FreeRTOS

1. Boot SMP Linux from A53 Core0, Core1, Core2 and Core3:

```
u-boot=> run prepare_mcore
u-boot=> setenv fdtfile imx8mp-evk-multicore-rtos.dtb
u-boot=> setenv mmccargs $mmccargs clk_ignore_unused
u-boot=> run bsp_bootcmd
```

After Linux boots up, use the command `cat /proc/cpuinfo` to check that the four A53 cores (Core0, Core1, Core2 and Core3) are all used by SMP Linux.

2. Boot FreeRTOS on M4 Core

Use `remoteproc` to boot FreeRTOS image `hello_world_cm7_UART4.elf` on M7 Core, it uses UART4 as the debug console.

```
root@imx8mp-lpddr4-evk:~# echo /examples/heterogeneous-multicore/hello-world-
cm/release/hello_world_cm7_UART4.elf > /sys/devices/platform/imx8mp-cm7/
remoteproc/remoteproc4/firmware
root@imx8mp-lpddr4-evk:~# echo start > /sys/devices/platform/imx8mp-cm7/
remoteproc/remoteproc4/state
```

Then the following log is displayed on UART4:

```
Cortex-M7: RTOS0: Hello world! Real-time Edge on MIMX8MP-EVK
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac
```

3. Boot FreeRTOS on A53 Core2

Use the following commands to boot FreeRTOS image `hello_world_ca53_RTOS0_RAM_CONSOLE-0xc0fff000.elf` on the A53 Core2. Note that it uses the RAM Console, for which the buffer address is at `0xc0fff000`:

```
root@imx8mp-lpddr4-evk:~# echo /examples/heterogeneous-multicore/hello-world-
ca/ddr_release/hello_world_ca53_RTOS0_RAM_CONSOLE-0xc0fff000.elf > /sys/
devices/platform/remoteproc-ca53-2/remoteproc/remoteproc1/firmware
root@imx8mp-lpddr4-evk:~# echo start > /sys/devices/platform/remoteproc-
ca53-2/remoteproc/remoteproc1/state
```

Then, use the tool to dump RAM Console and log, which is as follows:

```
root@imx8mp-lpddr4-evk:~# ram_console_dump -a 0xc0fff000 -r 1
RAM Console@0xc0fff000:

Cortex-A53: RTOS0: Hello world! Real-time Edge on MIMX8MP-EVK
RTOS1: RAM console@0xc1fff000
RTOS2: RAM console@0xc2fff000
RTOS3: RAM console@0xc3fff000
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac
```

Using commands `cat /proc/cpuinfo` can check that A53 Core0, Core1 and Core3 are used by SMP Linux as Core2 is running FreeRTOS.

4. Boot FreeRTOS on A53 Core3

Using the following commands to boot FreeRTOS image `hello_world_ca53_RTOS1_RAM_CONSOLE-0xc1fff000.elf` on A53 Core3, it use RAM Console which buffer address is at `0xc1fff000`:

```
root@imx8mp-lpddr4-evk:~# echo /examples/heterogeneous-multicore/hello-world-ca/ddr_release/hello_world_ca53_RTOS1_RAM_CONSOLE-0xc1fff000.elf > /sys/devices/platform/remoteproc-ca53-3/remoteproc/remoteproc2/firmware
root@imx8mp-lpddr4-evk:~# echo start > /sys/devices/platform/remoteproc-ca53-3/remoteproc/remoteproc2/state
```

Then, use the tool to dump RAM Console and log is as follows:

```
root@imx8mp-lpddr4-evk:~# ram_console_dump -a 0xc1fff000 -r 1
RAM Console@0x95bff000:

Cortex-A53: RTOS1: Hello world! Real-time Edge on MIMX8MP-EVK
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac
```

Using commands `cat /proc/cpuinfo` can check that A53 Core0 and Core1 are used by SMP Linux as Core2 and Core3 are running FreeRTOS.

5. Stop FreeRTOS and then start Zephyr on A53 Core2:
Stop FreeRTOS on A53 Core2:

```
root@imx8mp-lpddr4-evk:~# echo stop > /sys/devices/platform/remoteproc-ca53-2/remoteproc/remoteproc1/state
```

Using commands `cat /proc/cpuinfo` can check that A53 Core0, Core1 and Core2 are used by SMP Linux as Core2 RTOS has been stopped and it is hot plugged back to Linux kernel.

Start Zephyr on A53 Core2:

```
root@imx8mp-lpddr4-evk:~# echo /examples/heterogeneous-multicore/hello-world-ca-zephyr/hello_world_ca53_RTOS0_RAM_CONSOLE-0xc0100000.elf > /sys/devices/platform/remoteproc-ca53-2/remoteproc/remoteproc1/firmware
root@imx8mp-lpddr4-evk:~# echo start > /sys/devices/platform/remoteproc-ca53-2/remoteproc/remoteproc1/state
```

Then, use the tool to dump RAM Console and log is as follows:

```
root@imx8mp-lpddr4-evk:~# ram_console_dump -a 0xc0100000 -r 1
RAM Console@0xc0100000:
*** Booting Zephyr OS build v3.7.0-37-g757815829721 ***
Cortex-A53: RTOS0: Hello World! Real-time Edge on imx8mp_evk
RTOS1: RAM console@0xc1100000
RTOS2: RAM console@0xc2100000
RTOS3: RAM console@0xc3100000
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac tic tac
```

Using commands `cat /proc/cpuinfo` can check that A53 Core0 and Core1 are used by SMP Linux as Core2 and Core3 are running RTOS.

3.3.2.3.2 Running use cases on i.MX 8M Mini LPDDR4 EVK

3.3.2.3.2.1 Overview

The following use cases can run on i.MX 8M Mini LPDDR4 EVK:

Table 27. Flexible Real-time System on i.MX 8M Mini

ID	M4	A53 Core0	A53 Core1	A53 Core2	A53 Core3
0	RTOS	SMP Linux			
1	RTOS	SMP Linux			RTOS
2	RTOS	SMP Linux		RTOS	RTOS
3	RTOS	Linux	RTOS	RTOS	RTOS
4	RTOS	RTOS	RTOS	RTOS	RTOS

By default, Cortex-M Core's RTOS uses UART4 as debug Console. Refer to [Section 2.8.1](#) for how to boot Cortex-M Core's RTOS image.

Cortex-A Core's RTOS could run with UART4 Console, RAM Console or UART2 Console, the following Cortex-A Core's RTOS images are provided:

```

/examples/heterogeneous-multicore/hello-world-ca/ddr_release/ #FreeRTOS A53
hello_world Images
├── hello_world_ca53_RTOS0_RAM_CONSOLE-0x94bff000.bin
├── hello_world_ca53_RTOS0_RAM_CONSOLE-0x94bff000.elf
├── hello_world_ca53_RTOS0_UART4.bin
├── hello_world_ca53_RTOS0_UART4.elf
├── hello_world_ca53_RTOS1_RAM_CONSOLE-0x95bff000.bin
├── hello_world_ca53_RTOS1_RAM_CONSOLE-0x95bff000.elf
├── hello_world_ca53_RTOS2_RAM_CONSOLE-0x96bff000.bin
├── hello_world_ca53_RTOS2_RAM_CONSOLE-0x96bff000.elf
├── hello_world_ca53_RTOS3_RAM_CONSOLE-0x97bff000.bin
├── hello_world_ca53_RTOS3_RAM_CONSOLE-0x97bff000.elf
├── hello_world_ca53_RTOS3_UART2.bin
└── hello_world_ca53_RTOS3_UART2.elf
/examples/heterogeneous-multicore/hello-world-ca-zephyr/ #Zephyr A53
hello_world Images
├── hello_world_ca53_RTOS0_RAM_CONSOLE-0x93d00000.bin
├── hello_world_ca53_RTOS0_RAM_CONSOLE-0x93d00000.elf
├── hello_world_ca53_RTOS0_UART4.bin
├── hello_world_ca53_RTOS0_UART4.elf
├── hello_world_ca53_RTOS1_RAM_CONSOLE-0x94d00000.bin
├── hello_world_ca53_RTOS1_RAM_CONSOLE-0x94d00000.elf
├── hello_world_ca53_RTOS2_RAM_CONSOLE-0x95d00000.bin
├── hello_world_ca53_RTOS2_RAM_CONSOLE-0x95d00000.elf
├── hello_world_ca53_RTOS3_RAM_CONSOLE-0x96d00000.bin
├── hello_world_ca53_RTOS3_RAM_CONSOLE-0x96d00000.elf
├── hello_world_ca53_RTOS3_UART2.bin
└── hello_world_ca53_RTOS3_UART2.elf
/examples/heterogeneous-multicore/hello-world-cm/release/ #FreeRTOS M4
hello_world Images
├── hello_world_cm4_UART4.bin
└── hello_world_cm4_UART4.elf

```

RAM Console address is appended at the end of image name, for example `hello_world_ca53_RTOS3_RAM_CONSOLE-0x97bff000.bin`, it uses RAM Console which buffer address is `0x97bff000`.

Refer to [Section 2.6.3](#) for how to boot RTOS on Cortex-A Core, and refer to [Section 3.3.1.3](#) for how to dump RAM Console log.

The following chapters will introduce how to use U-Boot commands or Linux remoteproc to boot FreeRTOS or Zephyr use cases.

3.3.2.3.2.2 Running use cases by using U-Boot Commands

- Run the FreeRTOS applications:

Follow the following steps to run FreeRTOS on M4 Core, two FreeRTOS instances on A53 Core2 and Core3, and a SMP Linux kernel on A53 Core0 and Core1.

1. Boot Cortex-M Core's RTOS

```
u-boot=> ext4load mmc 1:2 0x48000000 /examples/heterogeneous-multicore/
hello-world-cm/release/hello_world_cm4_UART4.bin;
u-boot=> cp.b 0x48000000 0x7e0000 20000;
u-boot=> bootaux 0x7e0000
```

Then, the following log is displayed from UART4:

```
Cortex-M4: RTOS0: Hello world! Real-time Edge on MIMX8MM-EVK
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac
```

2. Boot the first Cortex-A Core's RTOS on A53 Core2

```
u-boot=> ext4load mmc 1:2 0x93C00000 /examples/heterogeneous-multicore/
hello-world-ca/ddr_release/hello_world_ca53_RTOS0_RAM_CONSOLE-0x94bff000.bin
u-boot=> dcache flush; icache flush; cpu 2 release 0x93C00000
```

Then check the RAM Console's log as follows:

```
u-boot=> dcache flush; md 0x94bff000
94bff000: 5f4d4152 534e4f43 00454c4f 00000000 RAM_CONSOLE.....
94bff010: 94bff040 00000000 00000fc0 00000000 @.....
94bff020: 000000c2 00000000 00000000 00000000 .....
94bff030: 00000000 00000000 00000000 00000000 .....
94bff040: 6f430a0d 78657472 3335412d 5452203a ..Cortex-A53: RT
94bff050: 3a30534f 6c654820 77206f6c 646c726f OS0: Hello world
94bff060: 65522021 742d6c61 20656d69 65676445 ! Real-time Edge
94bff070: 206e6f20 584d494d 2d4d4d38 0d4b5645 on MIMX8MM-EVK.
94bff080: 4f54520a 203a3153 204d4152 736e6f63 .RTOS1: RAM cons
94bff090: 40656c6f 35397830 30666662 0a0d3030 ole@0x95bff000..
94bff0a0: 534f5452 52203a32 63204d41 6f736e6f RTOS2: RAM conso
94bff0b0: 3040656c 62363978 30306666 520a0d30 le@0x96bff000..R
94bff0c0: 33534f54 4152203a 6f63204d 6c6f736e TOS3: RAM consol
94bff0d0: 78304065 66623739 30303066 69740a0d e@0x97bff000..ti
94bff0e0: 61742063 69742063 61742063 69742063 c tac tic tac ti
94bff0f0: 61742063 69742063 61742063 69742063 c tac tic tac ti
```

3. Boot the second Cortex-A Core's RTOS on A53 Core3

```
u-boot=> ext4load mmc 1:2 0x94C00000 /examples/heterogeneous-multicore/
hello-world-ca/ddr_release/hello_world_ca53_RTOS1_RAM_CONSOLE-0x95bff000.bin
u-boot=> dcache flush; icache flush; cpu 3 release 0x94C00000
```

Then check the RAM Console's log as follows:

```
u-boot=> dcache flush; md 0x95bff000
95bff000: 5f4d4152 534e4f43 00454c4f 00000000 RAM_CONSOLE.....
95bff010: 95bff040 00000000 00000fc0 00000000 @.....
95bff020: 000000af 00000000 00000000 00000000 .....
95bff030: 00000000 00000000 00000000 00000000 .....
95bff040: 6f430a0d 78657472 3335412d 5452203a ..Cortex-A53: RT
95bff050: 3a31534f 6c654820 77206f6c 646c726f OS1: Hello world
95bff060: 65522021 742d6c61 20656d69 65676445 ! Real-time Edge
95bff070: 206e6f20 584d494d 2d4d4d38 0d4b5645 on MIMX8MM-EVK.
95bff080: 6369740a 63617420 63697420 63617420 .tic tac tic tac
```

```

95bfff090: 63697420 63617420 63697420 63617420  tic tac tic tac
95bfff0a0: 63697420 63617420 63697420 63617420  tic tac tic tac
95bfff0b0: 63697420 63617420 63697420 63617420  tic tac tic tac
95bfff0c0: 63697420 63617420 63697420 63617420  tic tac tic tac
95bfff0d0: 740a0d20 74206369 74206361 74206369  ..tic tac tic t
95bfff0e0: 74206361 74206369 74206361 00206369  ac tic tac tic .
95bfff0f0: 00000000 00000000 00000000 00000000  .....

```

4. Boot SMP Linux from A53 Core0 and Core1:

If M-Core is booting up, use `imx8mm-evk-multicore-rpmsg.dtb`

```
u-boot=> setenv fdtfile imx8mm-evk-multicore-rpmsg.dtb
```

Otherwise, can also use `imx8mm-evk-multicore-rtos.dtb`:

```
u-boot=> setenv fdtfile imx8mm-evk-multicore-rtos.dtb
```

Then, boot up the kernel using the command below:

```
u-boot=> setenv mmcargs $mmcargs clk_ignore_unused
u-boot=> boot
```

After Linux kernel boots up, check Cortex-A Core's RTOS log from the RAM console:

```

root@imx8mm-lpddr4-evk:~# ram_console_dump -a 0x94bfff000 -r 1
RAM Console@0x94bfff000:

Cortex-A53: RTOS0: Hello world! Real-time Edge on MIMX8MM-EVK
RTOS1: RAM console@0x95bfff000
RTOS2: RAM console@0x96bfff000
RTOS3: RAM console@0x97bfff000
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
> .....

root@imx8mm-lpddr4-evk:~# ram_console_dump -a 0x95bfff000 -r 1
RAM Console@0x95bfff000:

Cortex-A53: RTOS1: Hello world! Real-time Edge on MIMX8MM-EVK

```

[illegible]

- Run the Zephyr applications:

Follow the following steps to run two Zephyr instances on A53 Core2 and Core3, and run a SMP Linux kernel on A53 Core0 and Core1.

1. Boot the first Cortex-A Core's RTOS on A53 Core2 with UART4.

```
u-boot=> ext4load mmc 1:2 0x93C00000 /examples/heterogeneous-multicore/
hello-world-ca-zephyr/hello_world_ca53_RTOS0_UART4.bin
u-boot=> dcache flush; icache flush; cpu 2 release 0x93C00000
```

Then, the UART4 displays the logs as follows:

```
*** Booting Zephyr OS build zephyr-v3.5.0-27-g24a3599ab7a0 ***
Cortex-A53: RTOS0: Hello World! Real-time Edge on mimx8mm_evk_a53
RTOS1: RAM console@0x94d00000
RTOS2: RAM console@0x95d00000
RTOS3: RAM console@0x96d00000
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tic
```

2. Boot the second Cortex-A core RTOS on A53 Core3.

```
u-boot=> ext4load mmc 1:2 0x94C00000 /examples/heterogeneous-multicore/
hello-world-ca-zephyr/hello_world_ca53_RTOS1_RAM_CONSOLE-0x94d00000.bin
u-boot=> dcache flush; icache flush; cpu 3 release 0x94C00000
```

Then check the RAM Console log as follows:

```
u-boot=> dcache flush; md 94d00000
94d00000: 202a2a2a 746f6f42 20676e69 6870655a *** Booting Zeph
94d00010: 4f207279 75622053 20646c69 6870657a yr OS build zeph
94d00020: 762d7279 2e352e33 37322d30 3432672d yr-v3.5.0-27-g24
94d00030: 39353361 37626139 2a203061 430a2a2a a3599ab7a0 ***.C
94d00040: 6574726f 35412d78 52203a33 31534f54 ortex-A53: RTOS1
94d00050: 6548203a 206f6c6c 6c726f57 52202164 : Hello World! R
94d00060: 2d6c6165 656d6974 67644520 6e6f2065 eal-time Edge on
94d00070: 6d696d20 6d6d3878 6b76655f 3335615f mimx8mm evk a53
```

```

94d00080: 6369740a 63617420 63697420 63617420 .tic tac tic tac
94d00090: 63697420 63617420 63697420 63617420 tic tac tic tac
94d000a0: 63697420 63617420 63697420 63617420 tic tac tic tac
94d000b0: 63697420 63617420 63697420 63617420 tic tac tic tac
94d000c0: 63697420 63617420 63697420 63617420 tic tac tic tac
94d000d0: 69740a20 61742063 69742063 61742063 .tic tac tic ta
94d000e0: 69742063 61742063 69742063 61742063 c tic tac tic ta
94d000f0: 69742063 61742063 69742063 61742063 c tic tac tic ta

```

3. Boot SMP Linux from A53 Core0 and Core1:

```
u-boot=> setenv fdtfile imx8mm-evk-multicore-rtos.dtb
```

Then, boot up the kernel using the command below:

```
u-boot=> setenv mmcargs $mmcargs clk_ignore_unused
u-boot=> boot
```

3.3.2.3.2.3 Running use cases by using remoteproc

Apart from the U-Boot commands, users can also use remoteproc under Linux to run RTOS. Take the following use case as an example. It uses the flexible `hello_world` application images to set up the following use case on i.MX 8M Mini.

Users can use remoteproc to start and stop RTOS dynamically without rebooting the board. Similarly, demo switching from use case #1 to use case #2 can be done dynamically by stopping FreeRTOS on Cortex-A Core2 and starting a Zephyr on it.

Table 28. Dynamic Flexible Real-time System on i.MX 8M Mini

ID	M4	A53-Core0	A53-Core1	A53-Core2	A53-Core3
1	FreeRTOS	SMP Linux		FreeRTOS	FreeRTOS
2	FreeRTOS	SMP Linux		Zephyr	FreeRTOS

1. Boot SMP Linux from A53 Core0, Core1, Core2 and Core3:

```
u-boot=> run prepare_mcore
u-boot=> setenv fdtfile imx8mm-evk-multicore-rtos.dtb
u-boot=> setenv mmcargs $mmcargs clk_ignore_unused
u-boot=> run bsp_bootcmd
```

After Linux boot up, using commands `cat /proc/cpuinfo` can check that four A53 Core0, Core1, Core2 and Core3 are all used by SMP Linux.

2. Boot FreeRTOS on M4 Core

Use remoteproc to boot FreeRTOS image `hello_world_cm4_UART4.elf` on M4 Core, it uses UART4 as debug console.

```
root@imx8mm-lpddr4-evk:~# echo /examples/heterogeneous-multicore/hello-world-
cm/release/hello_world_cm4_UART4.elf > /sys/devices/platform/imx8mm-cm4/
remoteproc/remoteproc4/firmware
root@imx8mm-lpddr4-evk:~# echo start > /sys/devices/platform/imx8mm-cm4/
remoteproc/remoteproc4/state
```

Then the following log is displayed from UART4:

```
Cortex-M4: RTOS0: Hello world! Real-time Edge on MIMX8MM-EVK
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac
```

3. Boot FreeRTOS on A53 Core2

Using the following commands to boot Zephyr image `hello_world_ca53_RTOS0_RAM_CONSOLE-0x94bfff000.elf` on A53 Core2, it use RAM Console which buffer address is at `0x94bfff000`:

```
root@imx8mm-lpddr4-evk:~# echo /examples/heterogeneous-multicore/hello-world-ca/ddr_release/hello_world_ca53_RTOS0_RAM_CONSOLE-0x94bfff000.elf > /sys/devices/platform/remoteproc-ca53-2/remoteproc/remoteproc1/firmware
root@imx8mm-lpddr4-evk:~# echo start > /sys/devices/platform/remoteproc-ca53-2/remoteproc/remoteproc1/state
```

Then use the tool to dump RAM Console and log is as follows:

```
root@imx8mm-lpddr4-evk:/sys# ram_console_dump -a 0x94bfff000 -r 1
RAM Console@0x94bfff000:

Cortex-A53: RTOS0: Hello world! Real-time Edge on MIMX8MM-EVK
RTOS1: RAM console@0x95bfff000
RTOS2: RAM console@0x96bfff000
RTOS3: RAM console@0x97bfff000
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac
```

Using commands `cat /proc/cpuinfo` can check that A53 Core0, Core1 and Core3 are used by SMP Linux as Core2 is running FreeRTOS.

4. Boot FreeRTOS on A53 Core3

Using the following commands to boot FreeRTOS image `hello_world_ca53_RTOS1_RAM_CONSOLE-0x95bfff000.elf` on A53 Core3, it use RAM Console which buffer address is at `0x95bfff000`:

```
root@imx8mm-lpddr4-evk:~# echo /examples/heterogeneous-multicore/hello-world-ca/ddr_release/hello_world_ca53_RTOS1_RAM_CONSOLE-0x95bfff000.elf > /sys/devices/platform/remoteproc-ca53-3/remoteproc/remoteproc2/firmware
root@imx8mm-lpddr4-evk:~# echo start > /sys/devices/platform/remoteproc-ca53-3/remoteproc/remoteproc2/state
```

Then use the tool to dump RAM Console and log is as follows:

```
root@imx8mm-lpddr4-evk:~# ram_console_dump -a 0x95bfff000 -r 1
RAM Console@0xc1ffff000:

Cortex-A53: RTOS1: Hello world! Real-time Edge on MIMX8MM-EVK
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac
```

Using commands `cat /proc/cpuinfo` can check that A53 Core0 and Core1 are used by SMP Linux as Core2 and Core3 are running FreeRTOS.

5. Stop FreeRTOS and then start Zephyr on A53 Core2

Stop FreeRTOS on A53 Core2:

```
root@imx8mm-lpddr4-evk:~# echo stop > /sys/devices/platform/remoteproc-ca53-2/remoteproc/remoteproc1/state
```

Using commands `cat /proc/cpuinfo` can check that A53 Core0, Core1 and Core2 are used by SMP Linux as RTOS on Core2 has been stopped and Core2 is hot plugged back to Linux kernel.

Start Zephyr on A53 Core2:

```
root@imx8mm-lpddr4-evk:~# echo /examples/heterogeneous-multicore/hello-world-ca-zephyr/hello_world_ca53_RTOS0_RAM_CONSOLE-0x93d00000.elf > /sys/devices/platform/remoteproc-ca53-2/remoteproc/remoteproc1/firmware
```

```
root@imx8mm-lpddr4-evk:~# echo start > /sys/devices/platform/remoteproc-
ca53-2/remoteproc/remoteproc1/state
```

Then use the tool to dump RAM Console and log is as follows:

```
root@imx8mm-lpddr4-evk:~# ram_console_dump -a 0x93d00000 -r 1
RAM Console@0xc0100000:
*** Booting Zephyr OS build v3.7.0-37-g757815829721 ***
Cortex-A53: RTOS0: Hello World! Real-time Edge on imx8mm_evk
RTOS1: RAM console@0xc1100000
RTOS2: RAM console@0xc2100000
RTOS3: RAM console@0xc3100000
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac tic tac
```

Using commands `cat /proc/cpuinfo` can check that A53 Core0 and Core1 are used by SMP Linux as Core2 and Core3 are running RTOS.

3.3.2.3.3 Running use cases on i.MX 91 9x9 QSB and i.MX 91 11x11 EVK

3.3.2.3.3.1 Overview

The following use cases can run on i.MX 91 9x9 QSB and i.MX 91 11x11 EVK platforms:

Table 29. Flexible Real-time System on i.MX 91 boards

ID	A55 Core0
0	Linux
1	RTOS

The RTOS of the Cortex-A core can run with UART1 console. The following RTOS images are provided:

```
/examples/heterogeneous-multicore/hello-world-ca/ddr_release/ #FreeRTOS A55
hello_world Images
└─ hello_world_ca55_RTOS0_UART1.bin
└─ hello_world_ca55_RTOS0_UART1.elf
/examples/heterogeneous-multicore/hello-world-ca-zephyr/ #Zephyr A55
hello_world Images
└─ hello_world_ca55_RTOS0_UART1.bin
└─ hello_world_ca55_RTOS0_UART1.elf
```

Refer to [Section 2.6.3](#) for how to boot RTOS on Cortex-A core.

3.3.2.3.4 Running use cases on i.MX 93 11x11 EVK and i.MX 93 14x14 EVK

3.3.2.3.4.1 Overview

The following use cases can run on i.MX 93 11x11 EVK and i.MX 93 14x14 EVK:

Table 30. Flexible Real-time System on i.MX 93

ID	M33	A55 Core0	A55 Core1
0	RTOS	SMP Linux	
1	RTOS	Linux	RTOS

Table 30. Flexible Real-time System on i.MX 93...continued

ID	M33	A55 Core0	A55 Core1
2	RTOS	RTOS	RTOS

By default, RTOS of the Cortex-M core uses UART2 as debug console. Refer to [Section 2.8.1](#) for how to boot the RTOS image of Cortex-M core.

RTOS of the Cortex-A core can run with UART2 console, RAM console, or UART1 console. The following RTOS images are provided:

```

/examples/heterogeneous-multicore/hello-world-ca/ddr_release/ #FreeRTOS A55
hello_world Images
├── hello_world_ca55_RTOS0_RAM_CONSOLE-0xd0fff000.bin
├── hello_world_ca55_RTOS0_RAM_CONSOLE-0xd0fff000.elf
├── hello_world_ca55_RTOS0_UART2.bin
├── hello_world_ca55_RTOS0_UART2.elf
├── hello_world_ca55_RTOS1_RAM_CONSOLE-0xd1fff000.bin
├── hello_world_ca55_RTOS1_RAM_CONSOLE-0xd1fff000.elf
├── hello_world_ca55_RTOS1_UART1.bin
└── hello_world_ca55_RTOS1_UART1.elf
/examples/heterogeneous-multicore/hello-world-ca-zephyr/ #Zephyr A55
hello_world Images
├── hello_world_ca55_RTOS0_RAM_CONSOLE-0xd0100000.bin
├── hello_world_ca55_RTOS0_RAM_CONSOLE-0xd0100000.elf
├── hello_world_ca55_RTOS0_UART2.bin
├── hello_world_ca55_RTOS0_UART2.elf
├── hello_world_ca55_RTOS1_RAM_CONSOLE-0xd1100000.bin
├── hello_world_ca55_RTOS1_RAM_CONSOLE-0xd1100000.elf
├── hello_world_ca55_RTOS1_UART1.bin
└── hello_world_ca55_RTOS1_UART1.elf
/examples/heterogeneous-multicore/hello-world-cm/release/ #FreeRTOS M33
hello_world Images
├── hello_world_cm33_UART2.bin
└── hello_world_cm33_UART2.elf

```

RAM console address is appended at the end of image name. For example `hello_world_ca55_RTOS1_RAM_CONSOLE-0xd1fff000.bin`, it uses RAM Console whose buffer address is `0xd1fff000`.

Refer to [Section 2.6.3](#) for how to boot RTOS on Cortex-A core, and refer to [Section 3.3.1.3](#) for how to dump RAM console log.

The following chapters introduce how to use U-Boot commands or Linux remoteproc to boot FreeRTOS or Zephyr use cases.

3.3.2.3.4.2 Running use cases by using U-Boot Commands

Taking Use Case #1 as an example, it runs FreeRTOS on M33 Core, run Linux on A55 Core0 and run FreeRTOS or Zephyr on A55 Core1, below are the steps to boot multiple operating systems on the platform:

- Run the FreeRTOS applications:
 1. Boot Cortex-M Core's RTOS

```

u-boot=> ext4load mmc 1:2 0xd0000000 /examples/heterogeneous-multicore/
hello-world-cm/release/hello_world_cm33_UART2.bin;
u-boot=> cp.b 0xd0000000 0x201e0000 20000;
u-boot=> bootaux 0x1ffe0000

```

Then, the following log is displayed from UART2:

```
Cortex-M33: RTOS0: Hello world! Real-time Edge on MIMX93-EVK
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac
```

2. Boot the Cortex-A Core's RTOS on Core1

```
u-boot=> ext4load mmc 1:2 0xd0000000 /examples/heterogeneous-multicore/
hello-world-ca/ddr_release/hello_world_ca55_RTOS0_RAM_CONSOLE-0xd0fff000.bin
u-boot=> dcache flush; icache flush; cpu 1 release 0xd0000000
```

Then, check the RAM Console's log as follows:

```
u-boot=> dcache flush; md 0xd0fff000
d0fff000: 5f4d4152 534e4f43 00454c4f 00000000 RAM_CONSOLE.....
d0fff010: d0fff040 00000000 00000fc0 00000000 @.....
d0fff020: 000000c9 00000000 00000000 00000000 .....
d0fff030: 00000000 00000000 00000000 00000000 .....
d0fff040: 6f430a0d 78657472 3535412d 5452203a ..Cortex-A55: RT
d0fff050: 3a30534f 6c654820 77206f6c 646c726f OS0: Hello world
d0fff060: 65522021 742d6c61 20656d69 65676445 ! Real-time Edge
d0fff070: 206e6f20 584d494d 452d3339 0a0d4b56 on MIMX93-EVK..
d0fff080: 534f5452 52203a31 63204d41 6f736e6f RTOS1: RAM conso
d0fff090: 3040656c 66316478 30306666 740a0d30 le@0xd1fff000..t
d0fff0a0: 74206369 74206361 74206369 74206361 ic tac tic tac t
d0fff0b0: 74206369 74206361 74206369 74206361 ic tac tic tac t
d0fff0c0: 74206369 74206361 74206369 74206361 ic tac tic tac t
d0fff0d0: 74206369 74206361 74206369 74206361 ic tac tic tac t
d0fff0e0: 74206369 74206361 74206369 0d206361 ic tac tic tac .
d0fff0f0: 6369740a 63617420 63697420 63617420 .tic tac tic tac
```

3. Boot Linux from Cortex-A Core0:

```
# Running on i.MX 93 EVK
u-boot=> setenv fdtfile imx93-11x11-evk-multicore-rtos.dtb
# Running on i.MX 93 14x14 EVK
u-boot=> setenv fdtfile imx93-14x14-evk-multicore-rtos.dtb
```

After Linux kernel boots up, check Cortex-A Core's RTOS log from the RAM console:

```
root@imx93evk:~# ram_console_dump -a 0xd0fff000 -r 1
RAM Console@0xd0fff000:

Cortex-A55: RTOS0: Hello world! Real-time Edge on MIMX93-EVK
RTOS1: RAM console@0xd1fff000
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
```



```
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac
> .....
```

- Run the Zephyr applications:

1. Boot the Cortex-A Core's RTOS on Core1 with UART2

```
u-boot=> ext4load mmc 1:2 0xd0000000 /examples/heterogeneous-multicore/
hello-world-ca-zephyr/hello_world_ca55_RTOS0_UART2.bin
u-boot=> dcache flush; icache flush; cpu 1 release 0xd0000000
```

Then the UART2 shows the logs as follows:

```
*** Booting Zephyr OS build zephyr-v3.5.0-27-g24a3599ab7a0 ***
Cortex-A55: RTOS0: Hello World! Real-time Edge on m1mx93_evk_a53
RTOS1: RAM console@0xd1100000
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
```

2. Boot Linux from Cortex-A Core0:

```
# Running on i.MX 93 EVK
u-boot=> setenv fdtfile imx93-11x11-evk-multicore-rtos.dtb
# Running on i.MX 93 14x14 EVK
u-boot=> setenv fdtfile imx93-14x14-evk-multicore-rtos.dtb
```

Then, boot up the kernel using the command below:

```
u-boot=> setenv mmcargs $mmcargs clk_ignore_unused
u-boot=> boot
```

3.3.2.3.4.3 Running use cases by using remoteproc (i.MX 93)

In addition to using U-Boot commands, users can also use remoteproc under Linux to run RTOS. Take the following use case as an example, using the flexible `hello_world` application images to set up the following use case on i.MX 93.

Using the remoteproc allows to start and to stop the RTOS dynamically without rebooting the board. Users can switch the demo from use case #1 to use case #2 dynamically by stopping FreeRTOS on Cortex-A Core2 and start a Zephyr on it.

Table 31. Dynamic Flexible Real-time System on i.MX 93

ID	M4	A55-Core0	A55-Core1
1	FreeRTOS	Linux	FreeRTOS
2	FreeRTOS	Linux	Zephyr

1. Boot SMP Linux from A55 Core0 and Core1

```
u-boot=> run prepare_mcore
u-boot=> setenv fdtfile imx93-11x11-evk-multicore-rtos.dtb
u-boot=> setenv mmcargs $mmcargs clk_ignore_unused
u-boot=> run bsp_bootcmd
```

After Linux boots up, using commands `cat /proc/cpuinfo` can check that the A55 Core0 and Core1 are all used by SMP Linux.

2. Boot FreeRTOS on the M33 Core

Use remoteproc to boot FreeRTOS image `hello_world_cm33_UART2.elf` on M33 Core, it uses UART2 as debug console.

```
root@imx93evk:~# echo /examples/heterogeneous-multicore/hello-world-cm/
release/hello_world_cm33_UART2.elf > /sys/devices/platform/remoteproc-cm33/
remoteproc/remoteproc1/firmware
root@imx93evk:~# echo start > /sys/devices/platform/remoteproc-cm33/
remoteproc/remoteproc1/state
```

Then, the following log is displayed from UART2:

```
Cortex-M33: RTOS0: Hello world! Real-time Edge on MIMX93-EVK
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
```

3. Boot FreeRTOS on A55 Core1

Using the following commands to boot FreeRTOS image `hello_world_ca55_RTOS0_RAM_CONSOLE-0xd0fff000.elf` on A55 Core1; it uses RAM Console for which the buffer address is `0xd0fff000`:

```
root@imx93evk:~# echo /examples/heterogeneous-multicore/hello-world-ca/
ddr_release/hello_world_ca55_RTOS0_RAM_CONSOLE-0xd0fff000.elf > /sys/devices/
platform/remoteproc-ca55-1/remoteproc/remoteproc0/firmware
root@imx93evk:~# echo start > /sys/devices/platform/remoteproc-ca55-1/
remoteproc/remoteproc0/state
```

Then use the tool to dump RAM Console and log is as follows:

```
root@imx93evk:~# ram_console_dump -a 0xd0fff000 -r 1
RAM Console@0xd0fff000:

Cortex-A55: RTOS0: Hello world! Real-time Edge on MIMX93-EVK
RTOS1: RAM console@0xd1fff000
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac tic tac tic tac tic tac tic
```

Using commands `cat /proc/cpuinfo` can check that A55 Core0 is used by Linux as Core1 is running FreeRTOS.

4. Stop FreeRTOS and then start Zephyr on A55 Core1

Stop FreeRTOS on A55 Core1:

```
root@imx93evk:~# echo stop > /sys/devices/platform/remoteproc-ca55-1/
remoteproc/remoteproc0/state
```

Using commands `cat /proc/cpuinfo` can check that A55 Core0 and Core1 are used by SMP Linux as RTOS on Core1 has been stopped and Core1 is hot plugged back to Linux kernel.

Start Zephyr on A55 Core1:

```
root@imx93evk:~# echo /examples/heterogeneous-multicore/hello-world-ca-
zephyr/hello_world_ca55_RTOS0_RAM_CONSOLE-0xd0100000.elf > /sys/devices/
platform/remoteproc-ca55-1/remoteproc/remoteproc0/firmware
root@imx93evk:~# echo start > /sys/devices/platform/remoteproc-ca55-1/
remoteproc/remoteproc0/state
```

Then, use the tool to dump RAM console and the log is as follows:

```
root@imx93evk:~# ram_console_dump -a 0xd0100000 -r 1
RAM Console@0xd0100000:
*** Booting Zephyr OS build v3.7.0-37-g757815829721 ***
Cortex-A55: RTOS0: Hello World! Real-time Edge on imx93_evk
RTOS1: RAM console@0xd1100000
```

```
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac tic tac tic tac tic tac tic tac tic tac tic
```

Using commands `cat /proc/cpuinfo` can check that A55 Core0 is used by SMP Linux as Core1 is running Zephyr.

3.3.2.3.5 Running use cases on i.MX 95 EVK (15x15 and 19x19)

3.3.2.3.5.1 Overview

The following use cases can run on i.MX 95 15x15 EVK and i.MX 9519x19 EVK boards:

Table 32. Flexible Real-time System on i.MX 95

ID	A55 Core0~3	A55 Core4	A55 Core5
0	SMP Linux		RTOS
1	SMP Linux	RTOS	RTOS
2	RTOS	RTOS	RTOS

Note: For case ID:2, there isn't U-Boot/Linux to dump RAM console outputs, so you must leverage a JTAG debugger to dump the RAM console buffer memory.

Due to the limitation of the number of Message Units, there are up to 4 SCMI agents for the A core Logical Machine in the System Manager configuration, one is assigned to the ATF, and only 3 are assigned to the OS of the A core. So there are 3 combinations shown in the above table.

The RTOS can run with UART3 console, RAM console, or UART1 console. The following RTOS images are provided:

```
/examples/heterogeneous-multicore/hello-world-ca/ddr_release/ #FreeRTOS A55
hello_world Images
— hello_world_ca55_RTOS0_RAM_CONSOLE-0xd0fff000.bin
— hello_world_ca55_RTOS0_RAM_CONSOLE-0xd0fff000.elf
— hello_world_ca55_RTOS0_UART3.bin
— hello_world_ca55_RTOS0_UART3.elf
— hello_world_ca55_RTOS1_RAM_CONSOLE-0xd1fff000.bin
— hello_world_ca55_RTOS1_RAM_CONSOLE-0xd1fff000.elf
— hello_world_ca55_RTOS2_RAM_CONSOLE-0xd2fff000.bin
— hello_world_ca55_RTOS2_RAM_CONSOLE-0xd2fff000.elf
— hello_world_ca55_RTOS2_UART1.bin
— hello_world_ca55_RTOS2_UART1.elf
/examples/heterogeneous-multicore/hello-world-ca-zephyr/ #Zephyr A55
hello_world Images
— hello_world_ca55_RTOS0_RAM_CONSOLE-0xd0100000.bin
— hello_world_ca55_RTOS0_RAM_CONSOLE-0xd0100000.elf
— hello_world_ca55_RTOS0_UART3.bin
— hello_world_ca55_RTOS0_UART3.elf
— hello_world_ca55_RTOS1_RAM_CONSOLE-0xd1100000.bin
— hello_world_ca55_RTOS1_RAM_CONSOLE-0xd1100000.elf
— hello_world_ca55_RTOS2_RAM_CONSOLE-0xd2100000.bin
— hello_world_ca55_RTOS2_RAM_CONSOLE-0xd2100000.elf
— hello_world_ca55_RTOS2_UART1.bin
— hello_world_ca55_RTOS2_UART1.elf
```

The RAM console address is appended at the end of image name, for example `hello_world_ca55_RTOS1_RAM_CONSOLE-0xd1fff000.bin`, it uses the RAM console whose buffer address is `0xd1fff000`.

- For details on how to boot RTOS on Cortex-A core refer to [Section 2.6.3](#).
- For details on how to dump RAM Console log, refer to [Section 3.3.1.3](#).

The following chapters introduce how to use U-Boot commands or Linux remoteproc to boot FreeRTOS or Zephyr use cases.

3.3.2.3.5.2 Running use cases by using U-Boot Commands

Taking Use Case #1 as an example, it runs SMP Linux on A55 Core0 ~ Core3 and runs FreeRTOS or Zephyr on A55 Core4 and Core5. Follow the below steps to boot multiple operating systems on the platform:

- Run the FreeRTOS applications:

1. Boot the first Cortex-A Core's RTOS on Core5

```
u-boot=> ext4load mmc 1:2 0xd0000000 /examples/heterogeneous-multicore/
hello-world-ca/ddr_release/hello_world_ca55_RTOS0_UART3.bin;
u-boot=> dcache flush; icache flush; cpu 5 release 0xd0000000
```

Then, the following log is displayed from UART3:

```
Cortex-A55: RTOS0: Hello world! Real-time Edge on MIMX95-EVK
RTOS1: RAM console@0xd1fff000
RTOS2: RAM console@0xd2fff000
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac
```

2. Boot the second Cortex-A Core's RTOS on Core4

```
u-boot=> ext4load mmc 1:2 0xd1000000 /examples/heterogeneous-multicore/
hello-world-ca/ddr_release/hello_world_ca55_RTOS1_RAM_CONSOLE-0xd1fff000.bin
u-boot=> dcache flush; icache flush; cpu 4 release 0xd1000000
```

Then, check the RAM Console's log as follows:

```
u-boot=> dcache flush; md 0xd1fff000
d1fff000: 5f4d4152 534e4f43 00454c4f 00000000 RAM_CONSOLE.....
d1fff010: d1fff040 00000000 00000fc0 00000000 @.....
d1fff020: 00000088 00000000 00000000 00000000 .....
d1fff030: 00000000 00000000 00000000 00000000 .....
d1fff040: 6f430a0d 78657472 3535412d 5452203a ..Cortex-A55: RT
d1fff050: 3a31534f 6c654820 77206f6c 646c726f OS1: Hello world
d1fff060: 65522021 742d6c61 20656d69 65676445 ! Real-time Edge
d1fff070: 206e6f20 584d494d 452d3539 0a0d4b56 on MIMX95-EVK..
d1fff080: 20636974 20636174 20636974 20636174 tic tac tic tac
d1fff090: 20636974 20636174 20636974 20636174 tic tac tic tac
d1fff0a0: 20636974 20636174 20636974 20636174 tic tac tic tac
d1fff0b0: 20636974 20636174 20636974 20636174 tic tac tic tac
d1fff0c0: 20636974 20636174 00000000 00000000 tic tac .....
d1fff0d0: 00000000 00000000 00000000 00000000 .....
d1fff0e0: 00000000 00000000 00000000 00000000 .....
d1fff0f0: 00000000 00000000 00000000 00000000 .....
```

3. Boot SMP Linux:

```
# For i.MX 95 15x15 EVK
u-boot=> setenv fdtfile imx95-15x15-evk-multicore-rtos.dtb
# For i.MX 95 19x19 EVK
u-boot=> setenv fdtfile imx95-19x19-evk-multicore-rtos.dtb
u-boot=> setenv mmcargs $mmcargs clk_ignore_unused
u-boot=> boot
```

After Linux kernel boots up, check Cortex-A Core's RTOS log from the RAM console:

```
root@imx95-19x19-lpddr5-evk:~# ram_console_dump -a 0xd1fff000 -r 1
RAM Console@0xd1fff000:

Cortex-A55: RTOS1: Hello world! Real-time Edge on MIMX95-EVK
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac tic tac tic >
```

- Run the Zephyr applications:

1. Boot the first Cortex-A Core's RTOS on Core5

```
u-boot=> ext4load mmc 1:2 0xd0000000 /examples/heterogeneous-multicore/
hello-world-ca-zephyr/hello_world_ca55_RTOS0_UART3.bin
u-boot=> dcache flush; icache flush; cpu 5 release 0xd0000000
```

Then the UART3 shows the logs as follows:

```
*** Booting Zephyr OS build v3.7.0-37-g757815829721 ***
Cortex-A55: RTOS0: Hello World! Real-time Edge on imx95_evk
RTOS1: RAM console@0xd1100000
RTOS2: RAM console@0xd2100000
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
```

2. Boot the second Cortex-A Core's RTOS on Core4

```
u-boot=> ext4load mmc 1:2 0xd1000000 /examples/heterogeneous-multicore/
hello-world-ca-zephyr/hello_world_ca55_RTOS1_RAM_CONSOLE-0xd1100000.bin
u-boot=> dcache flush; icache flush; cpu 4 release 0xd1000000
```

Then, check the RAM Console's log as follows:

```
u-boot=> dcache flush; md 0xd1100000
d1100000: 5f4d4152 534e4f43 00454c4f 00000000 RAM_CONSOLE.....
d1100010: d1100040 00000000 00000fbf 00000000 @.....
d1100020: 000000ac 00000000 00000000 00000000 .....
d1100030: 00000000 00000000 00000000 00000000 .....
d1100040: 202a2a2a 746f6f42 20676e69 6870655a *** Booting Zeph
d1100050: 4f207279 75622053 20646c69 372e3376 yr OS build v3.7
d1100060: 332d302e 37672d37 31383735 39323835 .0-37-g757815829
d1100070: 20313237 0a2a2a2a 74726f43 412d7865 721 ***.Cortex-A
d1100080: 203a3535 534f5452 48203a31 6f6c6c65 55: RTOS1: Hello
d1100090: 726f5720 2021646c 6c616552 6d69742d World! Real-tim
d11000a0: 64452065 6f206567 6d69206e 5f353978 e Edge on imx95_
d11000b0: 0a6b7665 20636974 20636174 20636974 evk.tic tac tic
d11000c0: 20636174 20636974 20636174 20636974 tac tic tac tic
d11000d0: 20636174 20636974 20636174 20636974 tac tic tac tic
d11000e0: 20636174 20636974 20636174 00000000 tac tic tac ....
d11000f0: 00000000 00000000 00000000 00000000 .....
```

3. Boot SMP Linux:

```
# For i.MX 95 15x15 EVK
u-boot=> setenv fdtfile imx95-15x15-evk-multicore-rtos.dtb
# For i.MX 95 19x19 EVK
u-boot=> setenv fdtfile imx95-19x19-evk-multicore-rtos.dtb
```

```
u-boot=> setenv mmcargs $mmcargs clk_ignore_unused
u-boot=> boot
```

After Linux kernel boots up, check Cortex-A Core's RTOS log from the RAM console:

```
root@imx95-19x19-lpddr5-evk:~# ram_console_dump -a 0xd1100000 -r 1
RAM Console@0xd1100000:
*** Booting Zephyr OS build v3.7.0-37-g757815829721 ***
Cortex-A55: RTOS1: Hello World! Real-time Edge on imx95_evk
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tic
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tic
tic tac tic tac tic tac >
```

3.3.2.3.5.3 Running use cases by using remoteproc

In addition to U-Boot commands, users can leverage Linux's remoteproc framework to dynamically manage RTOS lifecycle on the i.MX 95. For example, the flexible hello_world application can be deployed to demonstrate RTOS switching without requiring a board reboot.

A typical workflow involve:

- 1. Stopping the **FreeRTOS** on Cortex-A55 Core4.
- 2. Launching a **Zephyr RTOS** on the same core.

This enables transitions between use cases (e.g., from Use Case #1 to #2 below) without rebooting by reallocating hardware resources on demand.

Table 33. Dynamically deploy Flexible Real-time System on i.MX 95

ID	A55 Core0~3	A55 Core4	A55 Core5
1	SMP Linux	FreeRTOS	FreeRTOS
2	SMP Linux	Zephyr RTOS	FreeRTOS

1. Booting SMP Linux from all 6 Cortex-A55 Cores

```
# For i.MX 95 15x15 EVK
u-boot=> setenv fdtfile imx95-15x15-evk-multicore-rtos.dtb
# For i.MX 95 19x19 EVK
u-boot=> setenv fdtfile imx95-19x19-evk-multicore-rtos.dtb
u-boot=> setenv mmcargs $mmcargs clk_ignore_unused
u-boot=> run bsp_bootcmd
```

After Linux booting up, using commands `cat /proc/cpuinfo` can check that all the 6 Cortex-A55 Cores are all used by SMP Linux.

2. Booting FreeRTOS on the Cortex-A55 Core4

Use remoteproc to boot FreeRTOS hello_world_ca55_RTOS0_UART3.elf on A55 Core4, it uses the UART3 as debug console.

```
root@imx95-19x19-lpddr5-evk:~# echo /examples/heterogeneous-multicore/
hello-world-ca/ddr_release/hello_world_ca55_RTOS0_UART3.elf > /sys/devices/
platform/remoteproc-ca55-4/remoteproc/remoteproc3/firmware
root@imx95-19x19-lpddr5-evk:~# echo start > /sys/devices/platform/remoteproc-
ca55-4/remoteproc/remoteproc3/state
```

Then, the following log is displayed from UART3:

```
Cortex-A55: RTOS0: Hello world! Real-time Edge on MIMX95-EVK
RTOS1: RAM console@0xd1fff000
```

```
RTOS2: RAM console@0xd2fff000
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
```

Using command `cat /proc/cpuinfo` can check that the A55 Core4 is not online anymore under Linux.

3. Booting FreeRTOS on Cortex-A55 Core5

Use the following commands to boot FreeRTOS `hello_world_ca55_RTOS1_RAM_CONSOLE-0xd1fff000.elf` on A55 Core5; it uses the RAM Console whose buffer address is `0xd1fff000`:

```
root@imx95-19x19-lpddr5-evk:~# echo /examples/heterogeneous-multicore/hello-
world-ca/ddr_release/hello_world_ca55_RTOS1_RAM_CONSOLE-0xd1fff000.elf > /
sys/devices/platform/remoteproc-ca55-5/remoteproc/remoteproc4/firmware
root@imx95-19x19-lpddr5-evk:~# echo start > /sys/devices/platform/remoteproc-
ca55-5/remoteproc/remoteproc4/state
```

Then use the `ram_console_dump` tool to dump RAM Console:

```
root@imx95-19x19-lpddr5-evk:/# ram_console_dump -a 0xd1fff000 -r 1
RAM console@0xd1fff000:

Cortex-A55: RTOS1: Hello world! Real-time Edge on MIMX95-EVK
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tac tic tac tic
tac
tic tac tic tac tic tac tic >
```

Using command `cat /proc/cpuinfo` can check that the A55 Core5 is also not online anymore under Linux.

4. Stopping the FreeRTOS on A55 Core4 and then start a Zephyr `hello_world_ca55_RTOS0_UART3.elf` on the same core

Stopping the FreeRTOS on A55 Core4:

```
root@imx95-19x19-lpddr5-evk:~# echo stop > /sys/devices/platform/remoteproc-
ca55-4/remoteproc/remoteproc3/state
```

Using commands `cat /proc/cpuinfo` can check that A55 Core4 has been released back to Linux and it's online again.

Starting Zephyr on A55 Core4:

```
root@imx95-19x19-lpddr5-evk:~# echo /examples/heterogeneous-multicore/hello-
world-ca-zephyr/hello_world_ca55_RTOS0_UART3.elf > /sys/devices/platform/
remoteproc-ca55-4/remoteproc/remoteproc3/firmware
root@imx95-19x19-lpddr5-evk:~# echo start > /sys/devices/platform/remoteproc-
ca55-4/remoteproc/remoteproc3/state
```

Then, the Zephyr `hello_world` log is displayed from UART3:

```
*** Booting Zephyr OS build v3.7.0-37-g757815829721 ***
Cortex-A55: RTOS0: Hello World! Real-time Edge on imx95_evk
RTOS1: RAM console@0xd1100000
RTOS2: RAM console@0xd2100000
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tac tic tac tic
tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
```

Using command `cat /proc/cpuinfo` can check that the A55 Core4 is not online anymore under Linux again.

3.3.3 IwIP Networking Stack

3.3.3.1 Overview

LwIP is a small independent implementation of the TCP/IP protocol suite, it is freely available under a BSD license.

3.3.3.2 Running LwIP Application

3.3.3.2.1 Overview

Heterogeneous Multicore Framework provides a **LwIP Ping application** in the repository [heterogeneous-multicore](#), which supports i.MX 8M Mini LPDDR4 EVK, i.MX 8M Plus LPDDR4 EVK, and i.MX 93 EVK, and i.MX 93 14x14 EVK.

The Heterogeneous-multicore LwIP Ping application runs on Cortex-A Core. It initializes the LWIP networking stack and configures ENET port with the default IP address "192.168.0.100" and default gateway address "192.168.0.254". Then ping the gateway.

Hardware Setup

- Connect the Ethernet port of a test board to another board using an Ethernet cable and configure the other board's Ethernet interface with the IP address "192.168.0.254".
- Run the LwIP Ping application.
Use U-Boot commands or `remoteproc` under Linux to start the application.

3.3.3.2.2 Running LwIP Networking Stack use cases by using U-Boot Commands

Follow the following steps to start the use cases by using U-Boot commands:

- **Running the application on i.MX 8M Mini LPDDR4 EVK board**

Boot up i.MX 8M Mini LPDDR4 EVK board, and boot LwIP application from U-Boot command line:

```
u-boot=> ext4load mmc 1:2 0x93C00000 /examples/heterogeneous-multicore/lwip-  
ping-ca/ddr_release/lwip_ping_ca53.bin  
u-boot=> dcache flush; icache flush; cpu 3 release 0x93C00000
```

- **Running the application on i.MX 8M Plus LPDDR4 EVK board**

Boot up i.MX 8M Plus LPDDR4 EVK board, and boot LwIP application from U-Boot command line:

```
u-boot=> ext4load mmc 1:2 0xC0000000 /examples/heterogeneous-multicore/lwip-  
ping-ca/ddr_release/lwip_ping_ca53.bin  
u-boot=> dcache flush; icache flush; cpu 2 release 0xC0000000
```

- **Running the application on i.MX 93 EVK and i.MX 93 14x14 EVK boards:**

Boot up i.MX 93 EVK board or i.MX 93 14x14 EVK, and boot LwIP application from U-Boot command line:

```
u-boot=> ext4load mmc 1:2 0xD0000000 /examples/heterogeneous-multicore/lwip-  
ping-ca/ddr_release/lwip_ping_ca55.bin  
u-boot=> dcache flush; icache flush; cpu 1 release 0xD0000000
```

Then, the following log is displayed on the FreeRTOS Console:

```
Initializing PHY...  
  
*****  
PING example  
*****
```



```

IPv4 Address      : 192.168.0.100
IPv4 Subnet mask  : 255.255.255.0
IPv4 Gateway      : 192.168.0.254
*****
ping: send
192.168.0.254

ping: recv
192.168.0.254
0 ms

ping: send
192.168.0.254

```

3.3.3.2.3 Running lwIP Networking Stack use cases using remoteproc

Follow the following steps to start the use cases by using remoteproc under Linux:

1. Boot up Linux

- For i.MX 8M Mini LPDDR4 EVK

```

u-boot=> setenv fdtfile imx8mm-evk-virtio-net-ca53.dtb
u-boot=> setenv mmcargs $mmcargs clk_ignore_unused
u-boot=> run bsp_bootcmd

```

- For i.MX 8M Plus LPDDR4 EVK

```

u-boot=> setenv fdtfile imx8mp-evk-multicore-lwip.dtb
u-boot=> setenv mmcargs $mmcargs clk_ignore_unused
u-boot=> run bsp_bootcmd

```

- For i.MX 93 EVK

```

u-boot=> setenv fdtfile imx93-11x11-evk-virtio-net-ca55.dtb
u-boot=> setenv mmcargs $mmcargs clk_ignore_unused
u-boot=> run bsp_bootcmd

```

- For i.MX 93 14x14 EVK

```

u-boot=> setenv fdtfile imx93-14x14-evk-virtio-net-ca55.dtb
u-boot=> setenv mmcargs $mmcargs clk_ignore_unused
u-boot=> run bsp_bootcmd

```

2. Run the application

- For i.MX 8M Mini LPDDR4 EVK

```

root@imx8mm-lpddr4-evk:~# echo /examples/heterogeneous-multicore/lwip-ping-
ca/ddr_release/lwip_ping_ca53.elf > /sys/devices/platform/remoteproc-
ca53-3/remoteproc/remoteproc2/firmware
root@imx8mm-lpddr4-evk:~# echo start > /sys/devices/platform/remoteproc-
ca53-3/remoteproc/remoteproc2/state

```

- For i.MX 8M Plus LPDDR4 EVK

```

root@imx8mp-lpddr4-evk:~# echo /examples/heterogeneous-multicore/lwip-ping-
ca/ddr_release/lwip_ping_ca53.elf > /sys/devices/platform/remoteproc-
ca53-3/remoteproc/remoteproc2/firmware
root@imx8mp-lpddr4-evk:~# echo start > /sys/devices/platform/remoteproc-
ca53-3/remoteproc/remoteproc2/state

```

- For i.MX93 EVK and i.MX 93 14x14 EVK

```
root@imx93evk:~# echo /examples/heterogeneous-multicore/lwip-ping-ca/
ddr_release/lwip_ping_ca55.elf > /sys/devices/platform/remoteproc-ca55-1/
remoteproc/remoteproc0/firmware
root@imx93evk:~# echo start > /sys/devices/platform/remoteproc-ca55-1/
remoteproc/remoteproc0/state
```

Then, the following log is displayed on the FreeRTOS Console:

```
Initializing PHY...

*****
PING example
*****
IPv4 Address      : 192.168.0.100
IPv4 Subnet mask  : 255.255.255.0
IPv4 Gateway      : 192.168.0.254
*****
ping: send
192.168.0.254

ping: recv
192.168.0.254
0 ms

ping: send
192.168.0.254
```

3.4 RPMSG data communication

3.4.1 Overview

RPMsg (Remote Processor Messaging) protocol defines a standardized binary interface and is used for inter-core communication between Heterogeneous AMP on i.MX MPU platforms.

Currently Real-time Edge supports the following Heterogeneous AMP:

- Linux on Cortex-A core(s)
- RTOS on Cortex-M core
- RTOS on Cortex-A core(s)

Between these OS running different processes, Real-time Edge supports inter-core communication between Cortex-M core and Cortex-A core. It also supports RPMSG between heterogeneous AMP on different Cortex-A cores.

3.4.2 RPMSG performance evaluation

This RPMSG performance application provides a method to evaluate the RPMSG channel's benchmarks between Linux as RPMSG master and FreeRTOS as RPMSG remote.

Note: The released image does not display the CPU load status of the RTOS side during the evaluation. To display the CPU load status, build the debug type image manually referring to the [Section 3.2.2](#).

3.4.2.1 Building RPMSG performance evaluation application

RPMSG performance evaluation application "rpmsg_perf" runs on FreeRTOS, it is an application in the repo: [heterogeneous-multicore](#).

Refer to "[Section 3.2 Building Heterogeneous Multicore RTOS Application](#)" for how to build RPMSG performance evaluation application.

3.4.2.2 Running RPMsg performance application

The FreeRTOS RPMsg remote can run on Cortex-M core or Cortex-A core to test the RPMsg performance between Cortex-A core and Cortex-M core or between 2 Cortex-A cores. Follow the steps below to run the application on i.MX 8M Plus LPDDR4 EVK:

1. Boot FreeRTOS and Linux

Use U-Boot commands or remoteproc to boot FreeRTOS on Cortex-A Core or Cortex-M Core

- Steps for using U-Boot commands to boot FreeRTOS, then start Linux

a. For FreeRTOS RPMsg remote on Cortex-M Core

```
u-boot=> load mmc 1:2 0x48000000 /examples/heterogeneous-multicore/rpmsg-
perf-cm/release/rpmsg_perf_cm7.bin
u-boot=> cp.b 48000000 7e0000 20000
u-boot=> bootaux 7e0000
```

Then boot up Linux:

```
u-boot=> setenv fdtfile imx8mp-evk-rpmsg.dtb
u-boot=> setenv mmcargs $mmcargs clk_ignore_unused
u-boot=> run bsp_bootcmd
```

b. Or for FreeRTOS RPMsg remote on Cortex-A core

```
u-boot=> load mmc 1:2 c0000000 /examples/heterogeneous-multicore/rpmsg-
perf-ca/ddr_release/rpmsg_perf_ca53.bin
u-boot=> dcache flush; icache flush
u-boot=> cpu 2 release c0000000
```

Then boot up Linux:

```
u-boot=> setenv fdtfile imx8mp-evk-multicore-rpmsg.dtb
u-boot=> setenv mmcargs $mmcargs clk_ignore_unused
u-boot=> run bsp_bootcmd
```

After Linux boot up, FreeRTOS shows RPMSG links up as following log:

```
Cortex-M7/A53: RPMSG performance with linux:
INFO: rpmsg_init           : RPMSG remote init ...
INFO: rpmsg_remote_init    : waiting for link establish ...
INFO: rpmsg_remote_init    : RPMSG link up
```

- Steps for using remoteproc under Linux to boot FreeRTOS

a. For FreeRTOS RPMsg remote on Cortex-M Core

Boot up Linux firstly:

```
u-boot=> run prepare_mcore
u-boot=> setenv fdtfile imx8mp-evk-rpmsg.dtb
u-boot=> setenv mmcargs $mmcargs clk_ignore_unused
u-boot=> run bsp_bootcmd
```

Then run FreeRTOS on Cortex-M Core:

```
root@imx8mp-lpddr4-evk:/# echo /examples/heterogeneous-multicore/rpmsg-
perf-cm/release/rpmsg_perf_cm7.elf > /sys/devices/platform/imx8mp-cm7/
remoteproc/remoteproc0/firmware
root@imx8mp-lpddr4-evk:/# echo start > /sys/devices/platform/imx8mp-cm7/
remoteproc/remoteproc0/state
```

- b. Or for FreeRTOS RPMsg remote on Cortex-A core

Boot up Linux firstly:

```
u-boot=> setenv fdtfile imx8mp-evk-multicore-rpmsg.dtb
u-boot=> setenv mmcargs $mmcargs clk_ignore_unused
u-boot=> run bsp_bootcmd
```

Then run FreeRTOS on Cortex-A Core:

```
root@imx8mp-lpddr4-evk:/# echo /examples/heterogeneous-multicore/
rpmsg-perf-ca/ddr_release/rpmsg_perf_ca53.elf > /sys/devices/platform/
remoteproc-ca53-2/remoteproc/remoteproc1/firmware
root@imx8mp-lpddr4-evk:/# echo start > /sys/devices/platform/remoteproc-
ca53-2/remoteproc/remoteproc1/state
```

FreeRTOS shows RPMMSG links up as following log:

```
Cortex-M7/A53: RPMsg performance with linux:
INFO: rpmsg_init           : RPMMSG remote init ...
INFO: rpmsg_remote_init    : waiting for link establish ...
INFO: rpmsg_remote_init    : RPMMSG link up
```

2. Install the Linux `rpmsg_perf` driver module using the commands below:

The driver creates a char device `rpmsg-perf30` as shown below, which is used by the user space tool `rpmsg_perf` to test the RPMsg benchmarks:

```
root@imx8mp-lpddr4-evk:~# modprobe rpmsg_perf
root@imx8mp-lpddr4-evk:~# ls /dev/rpmsg-perf30
/dev/rpmsg-perf30
```

3. Use the `rpmsg_perf` tool to test the RPMsg performance in Linux prompt, refer to the usage below:

```
root@imx8mp-lpddr4-evk:~# rpmsg_perf
usage: rpmsg_perf <dev> <as_sender> <no_copy> <packet_size> <test_time>
dev: specify rpmsg device, see /dev/rpmsg-perf<x>
as_sender: true for as_sender, false for as_receiver
no_copy: specify if use no_copy version API in remote side
packet_size: specify the packet size, the MAX value is 496
test_time: specify the test period in unit second
such as: rpmsg_perf /dev/rpmsg-perf<x> true true 64 60
```

The example runs `rpmsg_perf /dev/rpmsg-perf30 true true 64 60` in the usage:

```
root@imx8mp-lpddr4-evk:~# rpmsg_perf /dev/rpmsg-perf30 true true 64 60
[ 1643.799911] rpmsg_perf: packet size: 64, sent packets: 4075370, time: 60
s, rate: 67 kpps
```

It means that Linux sends 64 Bytes packets to the FreeRTOS side during the given period 60s. The FreeRTOS RPMsg remote receives these packets using `no_copy` version APIs of `rpmsg-lite`, and the performance is about 67 kpps.

3.4.3 RPMMSG between Cortex-A Core and Cortex-M Core

[Figure 29](#) shows RPMMSG communication between RTOS running on Cortex-M core and Linux running Cortex-A core.

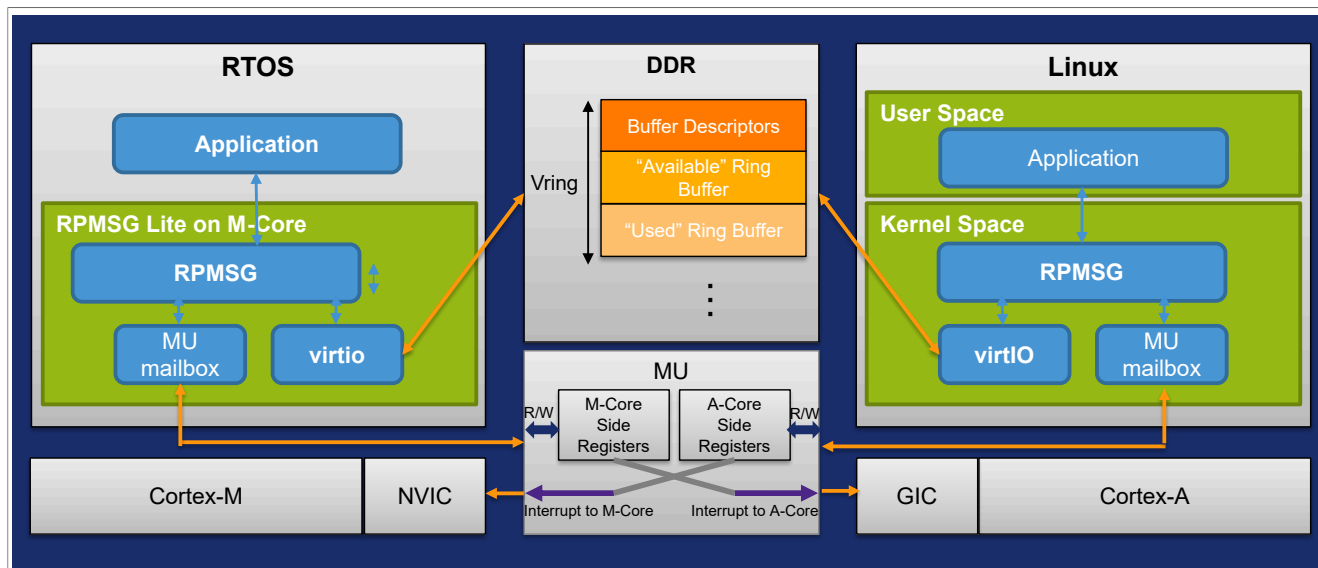


Figure 29. RPMSG between Cortex-A core and Cortex-M core

On i.MX MPU platforms, RPMSG builds virtual queue by leveraging Vring of VirtIO in shared memory of DDR. MU (Message Unit) is a hardware component in MPU platform that provides inter-core interrupt between Cortex-M core and Cortex-A core, so RPMSG uses MU as a mailbox notification mechanism.

In Linux, RPMSG communication is based on VirtIO driver and MU mailbox drivers. The RPMsg-Lite is an open-source component developed by NXP Semiconductors. It is a lightweight implementation of the RPMSG protocol. RPMsg-Lite is used on RTOS. It includes VirtIO driver, mailbox driver, and RPMSG driver. RPMsg-Lite is also enabled on RTOS running on Cortex-A cores.

Details about RPMsg-Lite can be found in the [RPMsg-Lite User's Guide](#).

3.4.3.1 RPMSG with enhanced 8MB Vring buffer

3.4.3.1.1 RPMSG merits

The RPMSG bus implements 2 virtqueues for transmitting and receiving respectively, and currently each virtqueue can support up to 256 RPMSG buffers with hardcode size 512B.

This feature increases the total number of RPMSG buffer to 8192 (4096 per direction) and extends the buffer size to 1024B.

3.4.3.1.2 Building and running the RPMSG demo (Cortex-A and Cortex-M core)

To build and run the demo for RPMSG between Cortex-A and Cortex-M cores, follow the steps listed below:

1. Enable RPMSG 8M buffer support in Real-time Edge software using the below commands:

```
$ cd yocto-real-time-edge/sources/meta-real-time-edge
# Open file "conf/distro/include/real-time-edge-base.inc" add "rpmsg_8m_buf"
to "DISTRO_FEATURES" like this:
DISTRO_FEATURES:append:mx8mm-nxp-bsp = " rpmsg_8m_buf"
```

2. Build the image using the commands below:

```
$ cd yocto-real-time-edge
$ DISTRO=nxp-real-time-edge MACHINE=imx8mm-lpddr4-evk source real-time-edge-
setup-env.sh -b build-imx8mm-real-time-edge
```

```
$ bitbake nxp-image-real-time-edge
```

3. Program the full SD card image. For this, use SD card with capacity of at least 4 GB.

```
$ bzip2 -d -c nxp-image-real-time-edge-imx8mm-lpddr4-evk.wic.bz2 | pv | sudo
dd of=/dev/sdx bs=1M && sync
# Note: find the right SD Card device name in your host machine and replace
the "sdx".
```

- #### 4. Start up M-core firmware under U-Boot:

```
=> ext4load mmc 1:2 0x48000000 /examples/heterogeneous-multicore/rpmsg-str-
echo-8m-cm/release/rpmsg_str_echo_cm4_8m_buf.bin
=> cp.b 0x48000000 0x7e0000 0x20000
=> bootaux 0x7e0000
```

- ### 5. Boot up Linux with RPMSG DTB:

```
=> setenv fdtfile imx8mm-evk-rpmsg-8m-buf.dtb
=> run bsp bootcmd
```

6. After Linux boots up, load `imx_rpmsg_tty.ko`

```
root@imx8mm-lpddr4-evk:~# modprobe imx rpmsg tty
```

Linux `imx_rpmmsg_tty` driver sends a “hello world!” message when probed, and it is displayed on the FreeRTOS console.

7. Test string transmitting through device “ttyRPMMSG30” from Linux prompt, the FreeRTOS console displays the received string. For example execute the following command:

```
root@imx8mm-lpddr4-evk:~# stty -F /dev/ttyRPMMSG30 -echo
root@imx8mm-lpddr4-evk:~# echo "any-string" > /dev/ttyRPMMSG30
```

8. In this demo, the single RPMSG buffer size is 1024B and the RPMSG header overhead is 16B, so the transmitting string will be split into up to 1008B fragments. Use the following commands to generate a file larger than 1KB to verify:

```
root@imx8mm-lpddr4-evk:~# for i in {1..300}; do echo -n `seq -s "" 0 1 9` >>
num.txt; done
root@imx8mm-lpddr4-evk:~# echo `cat num.txt` > /dev/ttyRPMMSG30
```

The log displays the message shown below:

[illegible]

[illegible]

3.4.4 RPMSG between Cortex-A Core and Cortex-A Core

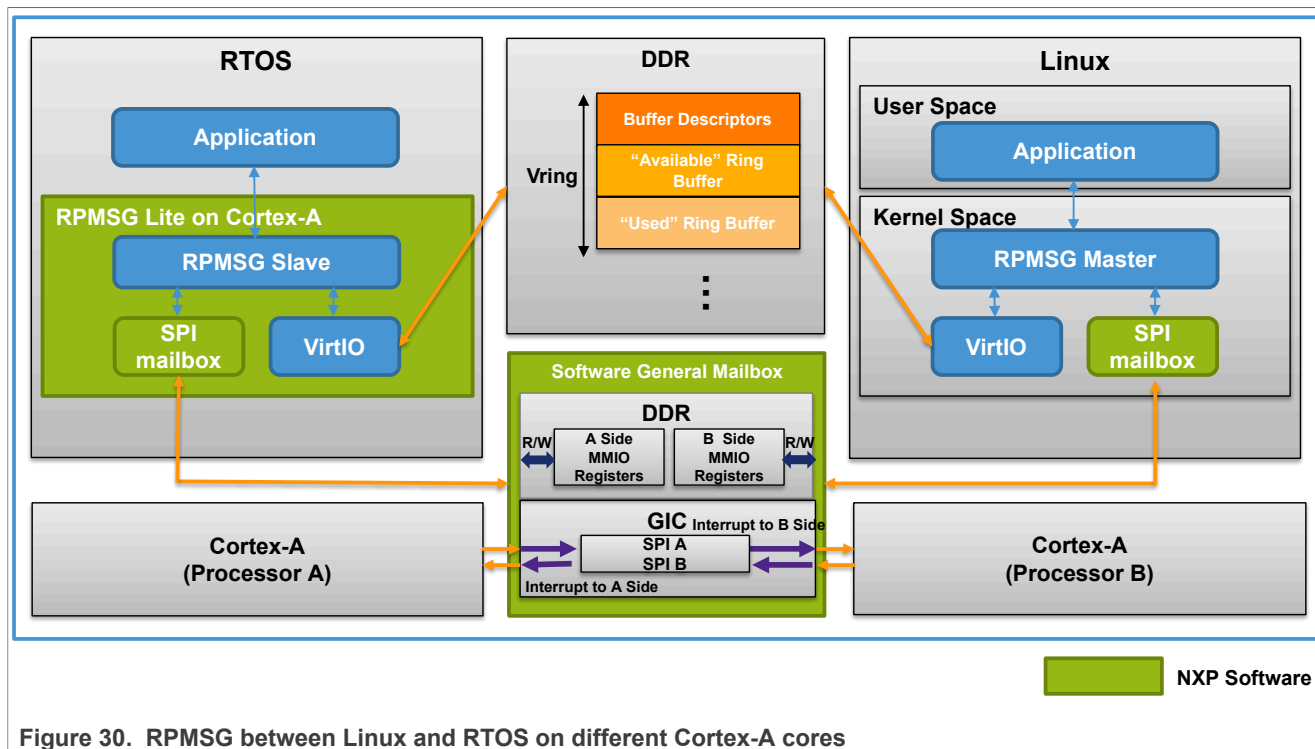
Heterogeneous Multicore Framework provides RPMSG communication between Cortex-A Core and Cortex-A Core:

- RPSMSG between Linux on Cortex-A Core and RTOS on Cortex-A Core
- RPSMSG between RTOS on Cortex-A Core and RTOS on Cortex-A Core

There is no MU hardware mailbox that can be used between different Cortex-A cores. Therefore, a Generic Software mailbox is created for message notification between Cortex-A cores. The Generic Software mailbox uses shared memory to simulate MMIO registers that are used by the mailbox driver. Two unused SPI interrupts in GIC are used as notification interrupts between Cortex-A cores. RPSmsg-Lite is also enabled on RTOS of Cortex-A cores.

3.4.4.1 RPMSG between Cortex-A Linux and Cortex-A RTOS

[Figure 30](#) illustrates the software setup for RPMSG between Linux on Cortex-A Core and RTOS on Cortex-A Core.



3.4.4.1.1 Building the rpmsg_str_echo demo

RPMSG "rpmsg_str_echo" application runs on FreeRTOS, it is an application named in the repo: [heterogeneous-multicore](#).

Refer to "[Building Heterogeneous Multicore RTOS Application](#)" for how to build RPMSG performance evaluation application.

3.4.4.1.2 Running the rpmsg_str_echo demo

- Setup UART Console
For i.MX 8M Mini LPDDR4 EVK and i.MX 8M Plus LPDDR4 EVK, open 2 terminal emulators to connect UART2 and UART4, respectively with the following setup:
 - 115200
 - No parity
 - 8 data bits
 - 1 stop bit
- Start up Linux and FreeRTOS
Can use U-Boot commands under U-Boot or remoteproc under Linux to start FreeRTOS, choose one of the methods below:
 - **Using U-Boot commands**
 - For i.MX 8M Mini LPDDR4 EVK
Start up FreeRTOS on the selected A-core under U-Boot:

```
u-boot=> ext4load mmc 1:2 93c00000 /examples/heterogeneous-multicore/rpmsg-str-echo-ca/ddr_release/rpmsg_str_echo_ca53_RTOS0_UART4.bin
u-boot=> dcache flush; icache flush;
u-boot=> cpu 2 release 93c00000
```


Boot up Linux with RPMSG DTB:

```
u-boot=> setenv fdtfile imx8mm-evk-multicore-rpmsg.dtb
u-boot=> run bsp_bootcmd
```

– **For i.MX 8M Plus LPDDR4 EVK**

```
u-boot=> ext4load mmc 1:2 0xC0000000 /examples/heterogeneous-multicore/
rpmsg-str-echo-ca/ddr_release/rpmsg_str_echo_ca53_RTOS0_UART4.bin
u-boot=> dcache flush; icache flush; cpu 2 release 0xC0000000
```

Boot up Linux with RPMSG DTB:

```
u-boot=> setenv fdtfile imx8mp-evk-multicore-rpmsg.dtb
u-boot=> run bsp_bootcmd
```

– **For i.MX 93 11x11 and 14x14 EVK**

```
u-boot=> ext4load mmc 1:2 0xD0000000 /examples/heterogeneous-multicore/
rpmsg-str-echo-ca/ddr_release/rpmsg_str_echo_ca55_RTOS0_UART2.bin
u-boot=> dcache flush; icache flush; cpu 1 release 0xD0000000
```

Boot up Linux with RPMSG DTB:

```
# Running on i.MX 93 11x11 EVK
u-boot=> setenv fdtfile imx93-11x11-evk-multicore-rpmsg.dtb
# Running on i.MX 93 14x14 EVK
u-boot=> setenv fdtfile imx93-14x14-evk-multicore-rpmsg.dtb
u-boot=> run bsp_bootcmd
```

– **Or using remoteproc**– **For i.MX 8M Mini LPDDR4 EVK****Boot up Linux with RPMSG DTB:**

```
u-boot=> setenv fdtfile imx8mm-evk-multicore-rpmsg.dtb
u-boot=> run bsp_bootcmd
```

Start up FreeRTOS on the selected A-core by using remoteproc:

```
root@imx8mm-lpddr4-evk:/# echo /examples/heterogeneous-multicore/rpmsg-
str-echo-ca/ddr_release/rpmsg_str_echo_ca53_RTOS0_UART4.elf > /sys/devices/
platform/remoteproc-ca53-2/remoteproc/remoteproc1/firmware
root@imx8mm-lpddr4-evk:/# echo start > /sys/devices/platform/remoteproc-
ca53-2/remoteproc/remoteproc1/state
```

– **For i.MX 8M Plus LPDDR4 EVK****Boot up Linux with RPMSG DTB:**

```
u-boot=> setenv fdtfile imx8mp-evk-multicore-rpmsg.dtb
u-boot=> run bsp_bootcmd
```

Start up FreeRTOS on the selected A-core by using remoteproc:

```
root@imx8mm-lpddr4-evk:/# echo /examples/heterogeneous-multicore/rpmsg-
str-echo-ca/ddr_release/rpmsg_str_echo_ca53_RTOS0_UART4.elf > /sys/devices/
platform/remoteproc-ca53-2/remoteproc/remoteproc1/firmware
root@imx8mm-lpddr4-evk:/# echo start > /sys/devices/platform/remoteproc-
ca53-2/remoteproc/remoteproc1/state
```

– **For i.MX 93 11x11 and 14x14 EVK****Boot up Linux with RPMSG DTB:**

```
# Running on i.MX 93 11x11 EVK
u-boot=> setenv fdtfile imx93-11x11-evk-multicore-rpmsg.dtb
# Running on i.MX 93 14x14 EVK
```

```
u-boot=> setenv fdtfile imx93-14x14-evk-multicore-rpmsg.dtb
u-boot=> run bsp_bootcmd
```

Start up FreeRTOS on the selected A-core by using remoteproc:

```
root@imx93evk:/# echo /examples/heterogeneous-multicore/rpmsg-str-echo-ca/
ddr_release/rpmsg_str_echo_ca55_RTOS0_UART2.elf > /sys/devices/platform/
remoteproc-ca55-1/remoteproc/remoteproc0/firmware
root@imx93evk:/# echo start > /sys/devices/platform/remoteproc-ca55-1/
remoteproc/remoteproc0/state
```

Then, the FreeRTOS application log is displayed as follows:

```
Cortex-A53: RTOS0: Multiple Endpoints RPMsg String Echo FreeRTOS Demo...
RTOS1: RAM console@0xc1fff000
INFO: rpmsg_init           : RPMMSG remote init ...
INFO: rpmsg_remote_init    : waiting for link establish ...
INFO: rpmsg_remote_init    : RPMMSG link up
INFO: rpmsg_create_ept     : Sending Nameservice string: rpmsg-virtual-tty-
channel-1
INFO: rpmsg_create_ept     : Sending Nameservice string: rpmsg-virtual-tty-
channel-1
INFO: rpmsg_create_ept     : Sending Nameservice string: rpmsg-virtual-tty-
channel-1
```

- Run the demo

1. After Linux boots up, load `imx_rpmsg_tty.ko`, it creates 3 `ttyRPMMSG` under `/dev`

```
root@imx8mm-lpddr4-evk:~# modprobe imx_rpmsg_tty
root@imx8mm-lpddr4-evk:~# ls /dev/ttyRPMMSG*
/dev/ttyRPMMSG3 /dev/ttyRPMMSG4 /dev/ttyRPMMSG5
```

2. Use minicom to open a console connecting one of the device `ttyRPMMSG` on the Linux prompt as shown below:

```
root@imx8mm-lpddr4-evk:~# minicom -D /dev/ttyRPMMSG3
```

Observe that the input string must then be echoed back on the console.

3.4.4.2 RPMMSG between Cortex-A RTOS and Cortex-A RTOS

[Figure 31](#) illustrates the software setup for RPMMSG between different RTOS on different Cortex-A Core.




Figure 31. RPMSG between RTOS and RTOS on different Cortex-A cores

3.4.4.2.1 Running the RPMSG ping-pong application

The RPMSG ping-pong application demonstrates the RPMSG communication between two FreeRTOS systems running respectively on two Cortex-A core islands.

Can use U-Boot commands under U-Boot or remoteproc under Linux to start FreeRTOS, choose one the method:

- Using U-Boot Commands to start RTOS

Follow the steps below to run the demo on i.MX 8M Plus LPDDR4 EVK:

1. First, boot the remote peer FreeRTOS on core2:

```
u-boot=> load mmc 1:2 0xC0000000 /examples/heterogeneous-multicore/rpmsg-
pingpong-remote-ca/ddr_release/rpmsg_pingpong_remote_ca53_UART4.bin
u-boot=> dcache flush; icache flush;
u-boot=> cpu 2 release 0xC0000000
```

The Remote peer boots up and waits for RPMSG link up, the logs on UART4 console are displayed as the following:

```
RPMsg Ping-Pong FreeRTOS Demo: remote: running at 0xc0000000
master RAM_CONSOLE at 0xc1fff000
INFO: rpmsg_init          : RPMSG remote init ...
INFO: rpmsg_remote_init   : waiting for link establish ...
```

2. As a second step, boot the Master peer FreeRTOS on core3, which uses the RAM console as output.

```
u-boot=> load mmc 1:2 0xC1000000 /examples/heterogeneous-
multicore/rpmsg-pingpong-master-ca/ddr_release/
rpmsg_pingpong_master_ca53_RAM_CONSOLE-0xc1fff000.bin
u-boot=> dcache flush; icache flush;
```

```
u-boot=> cpu 3 release 0xC1000000
```

The Master boots up and kicks the Remote peer to start the pingpong tests.

Then the remote peer FreeRTOS console displays the test result as the following:

```
RPMsng Ping-Pong FreeRTOS Demo: remote: running at 0xc0000000
master RAM_CONSOLE at 0xc1fff000
INFO: rpmsg_init           : RPMSG remote init ...
INFO: rpmsg_remote_init    : waiting for link establish ...
INFO: rpmsg_remote_init    : RPMSG link up
Waiting for ping ...
Received ping (0) ... sending pong (1)
Received ping (2) ... sending pong (3)
Received ping (4) ... sending pong (5)
Received ping (6) ... sending pong (7)
Received ping (8) ... sending pong (9)
Received ping (10) ... sending pong (11)
Received ping (12) ... sending pong (13)
Received ping (14) ... sending pong (15)
Received ping (16) ... sending pong (17)
Received ping (18) ... sending pong (19)
Received ping (20) ... sending pong (21)
Received ping (22) ... sending pong (23)
Received ping (24) ... sending pong (25)
Received ping (26) ... sending pong (27)
Received ping (28) ... sending pong (29)
Received ping (30) ... sending pong (31)
Received ping (32) ... sending pong (33)
Received ping (34) ... sending pong (35)
Received ping (36) ... sending pong (37)
Received ping (38) ... sending pong (39)
Received ping (40) ... sending pong (41)
Received ping (42) ... sending pong (43)
Received ping (44) ... sending pong (45)
Received ping (46) ... sending pong (47)
Received ping (48) ... sending pong (49)
Received ping (50) ... sending pong (51)
Received ping (52) ... sending pong (53)
Received ping (54) ... sending pong (55)
Received ping (56) ... sending pong (57)
Received ping (58) ... sending pong (59)
Received ping (60) ... sending pong (61)
Received ping (62) ... sending pong (63)
Received ping (64) ... sending pong (65)
Received ping (66) ... sending pong (67)
Received ping (68) ... sending pong (69)
Received ping (70) ... sending pong (71)
Received ping (72) ... sending pong (73)
Received ping (74) ... sending pong (75)
Received ping (76) ... sending pong (77)
Received ping (78) ... sending pong (79)
Received ping (80) ... sending pong (81)
Received ping (82) ... sending pong (83)
Received ping (84) ... sending pong (85)
Received ping (86) ... sending pong (87)
Received ping (88) ... sending pong (89)
Received ping (90) ... sending pong (91)
Received ping (92) ... sending pong (93)
Received ping (94) ... sending pong (95)
Received ping (96) ... sending pong (97)
Received ping (98) ... sending pong (99)
```

```
Received ping (100) ... sending pong (101)
```

```
RPMsg demo ends
```

The Master peer also displays the result on the RAM console:

```
u-boot=> dcache flush; md c1fff000 260
c1fff000: 734d5052 69502067 502d676e 20676e6f  RPMsg Ping-Pong
c1fff010: 65657246 534f5452 6d654420 6d203a6f  FreeRTOS Demo: m
c1fff020: 65747361 72203a72 696e6e75 6120676e  aster: running a
c1fff030: 78302074 30303163 30303030 4e490a0d  t 0xc1000000..IN
c1fff040: 203a4f46 736d7072 6e695f67 20207469  FO: rpmsg_init
c1fff050: 20202020 20202020 203a2020 534d5052  : RPMS
c1fff060: 616d2047 72657473 696e6920 2e2e2074  G master init ..
c1fff070: 490d0a2e 3a4f464e 6d707220 6d5f6773  ...INFO: rpmsg_m
c1fff080: 65747361 6e695f72 20207469 3a202020  aster_init :
c1fff090: 4d505220 6d204753 65747361 534e2072  RPMSG master NS
c1fff0a0: 72657320 65636976 61657220 0d0a7964  service ready..
c1fff0b0: 74696157 20676e69 20726f66 6120534e  Waiting for NS a
c1fff0c0: 756f6e6e 2065636e 0d2e2e2e 464e490a  nnounce .....INF
c1fff0d0: 72203a4f 67736d70 6d616e5f 72657365  O: rpmsg_nameser
c1fff0e0: 65636976 7361745f 63203a6b 74616572  vice_task: creat
c1fff0f0: 45206465 28205450 3a637273 202c3120  ed EPT (src: 1,
c1fff100: 3a747364 29303320 65530d0a 6e69646e  dst: 30)..Sendin
c1fff110: 69702067 2820676e 2e202930 72202e2e  g ping (0) ... r
c1fff120: 69656365 20646576 676e6f70 29312820  eceived pong (1)
c1fff130: 65530a0d 6e69646e 69702067 2820676e  ..Sending ping (
c1fff140: 2e202932 72202e2e 69656365 20646576  2) ... received
c1fff150: 676e6f70 29332820 65530a0d 6e69646e  pong (3)..Sendin
c1fff160: 69702067 2820676e 2e202934 72202e2e  g ping (4) ... r
c1fff170: 69656365 20646576 676e6f70 29352820  eceived pong (5)
c1fff180: 65530a0d 6e69646e 69702067 2820676e  ..Sending ping (
c1fff190: 2e202936 72202e2e 69656365 20646576  6) ... received
c1fff1a0: 676e6f70 29372820 65530a0d 6e69646e  pong (7)..Sendin
c1fff1b0: 69702067 2820676e 2e202938 72202e2e  g ping (8) ... r
c1fff1c0: 69656365 20646576 676e6f70 29392820  eceived pong (9)
c1fff1d0: 65530a0d 6e69646e 69702067 2820676e  ..Sending ping (
c1fff1e0: 20293031 202e2e2e 65636572 64657669  10) ... received
c1fff1f0: 6e6f7020 31282067 0a0d2931 646e6553  pong (11)..Send
c1fff200: 20676e69 676e6970 32312820 2e2e2029  ing ping (12) ..
c1fff210: 6572202e 76696563 70206465 20676e6f  . received pong
c1fff220: 29333128 65530a0d 6e69646e 69702067  (13)..Sending pi
c1fff230: 2820676e 20293431 202e2e2e 65636572  ng (14) ... rece
c1fff240: 64657669 6e6f7020 31282067 0a0d2935  ived pong (15)..
c1fff250: 646e6553 20676e69 676e6970 36312820  Sending ping (16
c1fff260: 2e2e2029 6572202e 76696563 70206465  ) ... received p
c1fff270: 20676e6f 29373128 65530a0d 6e69646e  ong (17)..Sendin
c1fff280: 69702067 2820676e 20293831 202e2e2e  g ping (18) ...
c1fff290: 65636572 64657669 6e6f7020 31282067  received pong (1
c1fff2a0: 0a0d2939 646e6553 20676e69 676e6970  9)..Sending ping
c1fff2b0: 30322820 2e2e2029 6572202e 76696563  (20) ... receiv
c1fff2c0: 70206465 20676e6f 29313228 65530a0d  ed pong (21)..Se
c1fff2d0: 6e69646e 69702067 2820676e 20293232  nding ping (22)
c1fff2e0: 202e2e2e 65636572 64657669 6e6f7020  ... received pon
c1fff2f0: 32282067 0a0d2933 646e6553 20676e69  g (23)..Sending
c1fff300: 676e6970 34322820 2e2e2029 6572202e  ping (24) ... re
c1fff310: 76696563 70206465 20676e6f 29353228  ceived pong (25)
c1fff320: 65530a0d 6e69646e 69702067 2820676e  ..Sending ping (
c1fff330: 20293632 202e2e2e 65636572 64657669  26) ... received
c1fff340: 6e6f7020 32282067 0a0d2937 646e6553  pong (27)..Send
c1fff350: 20676e69 676e6970 38322820 2e2e2029  ing ping (28) ..
```

```
clfff360: 6572202e 76696563 70206465 20676e6f . received pong
clfff370: 29393228 65530a0d 6e69646e 69702067 (29)..Sending pi
clfff380: 2820676e 20293033 202e2e2e 65636572 ng (30) ... rece
clfff390: 64657669 6e6f7020 33282067 0a0d2931 ived pong (31)..
clfff3a0: 646e6553 20676e69 676e6970 32332820 Sending ping (32
clfff3b0: 2e2e2029 6572202e 76696563 70206465 ) ... received p
clfff3c0: 20676e6f 29333328 65530a0d 6e69646e ong (33)..Sendin
clfff3d0: 69702067 2820676e 20293433 202e2e2e g ping (34) ...
clfff3e0: 65636572 64657669 6e6f7020 33282067 received pong (3
clfff3f0: 0a0d2935 646e6553 20676e69 676e6970 5)..Sending ping
clfff400: 36332820 2e2e2029 6572202e 76696563 (36) ... receiv
clfff410: 70206465 20676e6f 29373328 65530a0d ed pong (37)..Se
clfff420: 6e69646e 69702067 2820676e 20293833 nding ping (38)
clfff430: 202e2e2e 65636572 64657669 6e6f7020 ... received pon
clfff440: 33282067 0a0d2939 646e6553 20676e69 g (39)..Sending
clfff450: 676e6970 30342820 2e2e2029 6572202e ping (40) ... re
clfff460: 76696563 70206465 20676e6f 29313428 ceived pong (41)
clfff470: 65530a0d 6e69646e 69702067 2820676e ..Sending ping (
clfff480: 20293234 202e2e2e 65636572 64657669 42) ... received
clfff490: 6e6f7020 34282067 0a0d2933 646e6553 pong (43)..Send
clfff4a0: 20676e69 676e6970 34342820 2e2e2029 ing ping (44) ..
clfff4b0: 6572202e 76696563 70206465 20676e6f . received pong
clfff4c0: 29353428 65530a0d 6e69646e 69702067 (45)..Sending pi
clfff4d0: 2820676e 20293634 202e2e2e 65636572 ng (46) ... rece
clfff4e0: 64657669 6e6f7020 34282067 0a0d2937 ived pong (47)..
clfff4f0: 646e6553 20676e69 676e6970 38342820 Sending ping (48
clfff500: 2e2e2029 6572202e 76696563 70206465 ) ... received p
clfff510: 20676e6f 29393428 65530a0d 6e69646e ong (49)..Sendin
clfff520: 69702067 2820676e 20293035 202e2e2e g ping (50) ...
clfff530: 65636572 64657669 6e6f7020 35282067 received pong (5
clfff540: 0a0d2931 646e6553 20676e69 676e6970 1)..Sending ping
clfff550: 32352820 2e2e2029 6572202e 76696563 (52) ... receiv
clfff560: 70206465 20676e6f 29333528 65530a0d ed pong (53)..Se
clfff570: 6e69646e 69702067 2820676e 20293435 nding ping (54)
clfff580: 202e2e2e 65636572 64657669 6e6f7020 ... received pon
clfff590: 35282067 0a0d2935 646e6553 20676e69 g (55)..Sending
clfff5a0: 676e6970 36352820 2e2e2029 6572202e ping (56) ... re
clfff5b0: 76696563 70206465 20676e6f 29373528 ceived pong (57)
clfff5c0: 65530a0d 6e69646e 69702067 2820676e ..Sending ping (
clfff5d0: 20293835 202e2e2e 65636572 64657669 58) ... received
clfff5e0: 6e6f7020 35282067 0a0d2939 646e6553 pong (59)..Send
clfff5f0: 20676e69 676e6970 30362820 2e2e2029 ing ping (60) ..
clfff600: 6572202e 76696563 70206465 20676e6f . received pong
clfff610: 29313628 65530a0d 6e69646e 69702067 (61)..Sending pi
clfff620: 2820676e 20293236 202e2e2e 65636572 ng (62) ... rece
clfff630: 64657669 6e6f7020 36282067 0a0d2933 ived pong (63)..
clfff640: 646e6553 20676e69 676e6970 34362820 Sending ping (64
clfff650: 2e2e2029 6572202e 76696563 70206465 ) ... received p
clfff660: 20676e6f 29353628 65530a0d 6e69646e ong (65)..Sendin
clfff670: 69702067 2820676e 20293636 202e2e2e g ping (66) ...
clfff680: 65636572 64657669 6e6f7020 36282067 received pong (6
clfff690: 0a0d2937 646e6553 20676e69 676e6970 7)..Sending ping
clfff6a0: 38362820 2e2e2029 6572202e 76696563 (68) ... receiv
clfff6b0: 70206465 20676e6f 29393628 65530a0d ed pong (69)..Se
clfff6c0: 6e69646e 69702067 2820676e 20293037 nding ping (70)
clfff6d0: 202e2e2e 65636572 64657669 6e6f7020 ... received pon
clfff6e0: 37282067 0a0d2931 646e6553 20676e69 g (71)..Sending
clfff6f0: 676e6970 32372820 2e2e2029 6572202e ping (72) ... re
clfff700: 76696563 70206465 20676e6f 29333728 ceived pong (73)
clfff710: 65530a0d 6e69646e 69702067 2820676e ..Sending ping (
```

```

clfff720: 20293437 202e2e2e 65636572 64657669 74) ... received
clfff730: 6e6f7020 37282067 0a0d2935 646e6553 pong (75)..Send
clfff740: 20676e69 676e6970 36372820 2e2e2029 ing ping (76) ..
clfff750: 6572202e 76696563 70206465 20676e6f . received pong
clfff760: 29373728 65530a0d 6e69646e 69702067 (77)..Sending pi
clfff770: 2820676e 20293837 202e2e2e 65636572 ng (78) ... rece
clfff780: 64657669 6e6f7020 37282067 0a0d2939 ived pong (79)..
clfff790: 646e6553 20676e69 676e6970 30382820 Sending ping (80
clfff7a0: 2e2e2029 6572202e 76696563 70206465 ) ... received p
clfff7b0: 20676e6f 29313828 65530a0d 6e69646e ong (81)..Sendin
clfff7c0: 69702067 2820676e 20293238 202e2e2e g ping (82) ...
clfff7d0: 65636572 64657669 6e6f7020 38282067 received pong (8
clfff7e0: 0a0d2933 646e6553 20676e69 676e6970 3)..Sending ping
clfff7f0: 34382820 2e2e2029 6572202e 76696563 (84) ... receiv
clfff800: 70206465 20676e6f 29353828 65530a0d ed pong (85)..Se
clfff810: 6e69646e 69702067 2820676e 20293638 nding ping (86)
clfff820: 202e2e2e 65636572 64657669 6e6f7020 ... received pon
clfff830: 38282067 0a0d2937 646e6553 20676e69 g (87)..Sending
clfff840: 676e6970 38382820 2e2e2029 6572202e ping (88) ... re
clfff850: 76696563 70206465 20676e6f 29393828 ceived pong (89)
clfff860: 65530a0d 6e69646e 69702067 2820676e ..Sending ping (
clfff870: 20293039 202e2e2e 65636572 64657669 90) ... received
clfff880: 6e6f7020 39282067 0a0d2931 646e6553 pong (91)..Send
clfff890: 20676e69 676e6970 32392820 2e2e2029 ing ping (92) ..
clfff8a0: 6572202e 76696563 70206465 20676e6f . received pong
clfff8b0: 29333928 65530a0d 6e69646e 69702067 (93)..Sending pi
clfff8c0: 2820676e 20293439 202e2e2e 65636572 ng (94) ... rece
clfff8d0: 64657669 6e6f7020 39282067 0a0d2935 ived pong (95)..
clfff8e0: 646e6553 20676e69 676e6970 36392820 Sending ping (96
clfff8f0: 2e2e2029 6572202e 76696563 70206465 ) ... received p
clfff900: 20676e6f 29373928 65530a0d 6e69646e ong (97)..Sendin
clfff910: 69702067 2820676e 20293839 202e2e2e g ping (98) ...
clfff920: 65636572 64657669 6e6f7020 39282067 received pong (9
clfff930: 0a0d2939 646e6553 20676e69 676e6970 9)..Sending ping
clfff940: 30312820 2e202930 72202e2e 69656365 (100) ... recei
clfff950: 20646576 676e6f70 30312820 0a0d2931 ved pong (101)..
clfff960: 50520a0d 2067734d 6f6d6564 646e6520 ..RPMsg demo end
clfff970: 000a0d73 00000000 00000000 00000000 s.....

```

- Or using remoteproc to start RTOS

Start up Linux:

```

u-boot=> setenv fdtfile imx8mp-evk-multicore-rtos.dtb
u-boot=> setenv mmcargs $mmcargs clk_ignore_unused
u-boot=> run bsp_bootcmd

```

After Linux boot up, start remote peer FreeRTOS on core2, it use UART4 as debug console:

```

root@imx8mp-lpddr4-evk:~# echo /examples/heterogeneous-multicore/rpmsg-
pingpong-remote-ca/ddr_release/rpmsg_pingpong_remote_ca53_UART4.elf > /sys/
devices/platform/remoteproc-ca53-2/remoteproc/remoteproc1/firmware
root@imx8mp-lpddr4-evk:~# echo start > /sys/devices/platform/remoteproc-ca53-2/
remoteproc/remoteproc1/state

```

Then start the Master peer FreeRTOS on core3, which uses the RAM console as debug console:

```

root@imx8mp-lpddr4-evk:~# echo /examples/heterogeneous-
multicore/rpmsg-pingpong-master-ca/ddr_release/
rpmsg_pingpong_master_ca53_RAM_CONSOLE-0xc1fff000.elf > /sys/devices/platform/
remoteproc-ca53-3/remoteproc/remoteproc2/firmware

```



```
root@imx8mp-lpddr4-evk:~# echo start > /sys/devices/platform/remoteproc-ca53-3/remoteproc/remoteproc2/state
```

Can use the tool to check RAM Console log:

```
root@imx8mp-lpddr4-evk:~# ram_console_dump -a 0xc1fff000 -r 1
RAM Console@0xc1fff000:
RPMMsg Ping-Pong FreeRTOS Demo: master: running at 0xc1000000
INFO: rpmsg_init : RPMSG master init ...
INFO: rpmsg_master_init : RPMSG master NS service ready
Waiting for NS announce ...
INFO: rpmsg_nameservice_task: created EPT (src: 1, dst: 30)
Sending ping (0) ... received pong (1)
Sending ping (2) ... received pong (3)
Sending ping (4) ... received pong (5)
Sending ping (6) ... received pong (7)
Sending ping (8) ... received pong (9)
Sending ping (10) ... received pong (11)
Sending ping (12) ... received pong (13)
Sending ping (14) ... received pong (15)
Sending ping (16) ... received pong (17)
Sending ping (18) ... received pong (19)
Sending ping (20) ... received pong (21)
Sending ping (22) ... received pong (23)
Sending ping (24) ... received pong (25)
Sending ping (26) ... received pong (27)
Sending ping (28) ... received pong (29)
Sending ping (30) ... received pong (31)
Sending ping (32) ... received pong (33)
Sending ping (34) ... received pong (35)
Sending ping (36) ... received pong (37)
Sending ping (38) ... received pong (39)
Sending ping (40) ... received pong (41)
Sending ping (42) ... received pong (43)
Sending ping (44) ... received pong (45)
Sending ping (46) ... received pong (47)
Sending ping (48) ... received pong (49)
Sending ping (50) ... received pong (51)
Sending ping (52) ... received pong (53)
Sending ping (54) ... received pong (55)
Sending ping (56) ... received pong (57)
Sending ping (58) ... received pong (59)
Sending ping (60) ... received pong (61)
Sending ping (62) ... received pong (63)
Sending ping (64) ... received pong (65)
Sending ping (66) ... received pong (67)
Sending ping (68) ... received pong (69)
Sending ping (70) ... received pong (71)
Sending ping (72) ... received pong (73)
Sending ping (74) ... received pong (75)
Sending ping (76) ... received pong (77)
Sending ping (78) ... received pong (79)
Sending ping (80) ... received pong (81)
Sending ping (82) ... received pong (83)
Sending ping (84) ... received pong (85)
Sending ping (86) ... received pong (87)
Sending ping (88) ... received pong (89)
Sending ping (90) ... received pong (91)
Sending ping (92) ... received pong (93)
Sending ping (94) ... received pong (95)
Sending ping (96) ... received pong (97)
```



```
Sending ping (98) ... received pong (99)
Sending ping (100) ... received pong (101)
```

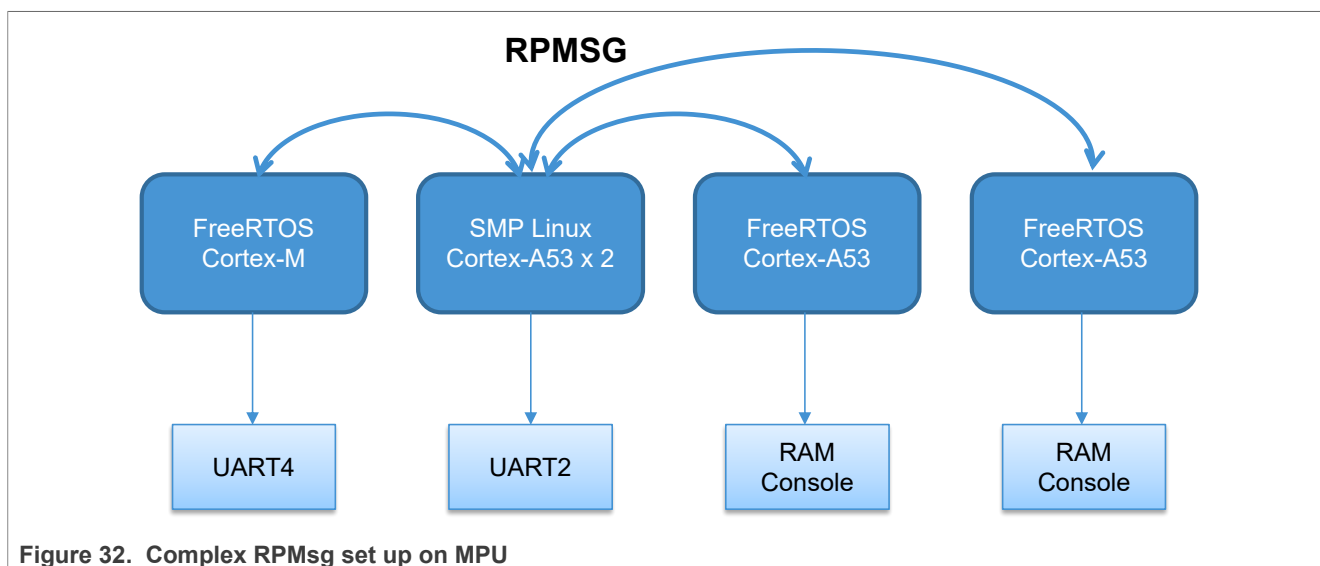
```
RPMsg demo ends
```

3.4.5 Complex RPMMSG on MPU

3.4.5.1 Overview

To demonstrate a typical RPMMSG use case on the MPU platform, a complex RPMMSG str-echo application is available in the repository: [heterogeneous-multicore](#).

[Figure 32](#) shows the application setup:



Rpmmsg-str-echo application setup enables RPMMSG communication between FreeRTOS and Linux. FreeRTOS runs RPMMSG master endpoint and Linux runs RPMMSG remote endpoint. By default, it creates three endpoints on both the master side and remote side. Each endpoint is one-to-one connected with the other side. Therefore, there are three RPMMSG channels between the master and remote side. Application on FreeRTOS receives data from the remote side and then sends it back to the same RPMMSG channel. On the Linux side, if data is sent to the master side, the same data is received or echoed back from the same channel.

3.4.5.2 Building the Complex str_echo application

RPMMSG complex str_echo application runs on FreeRTOS, it is an application named "rpmmsg_str_echo" in the repo: [heterogeneous-multicore](#).

Refer to "[Building Heterogeneous Multicore RTOS Application](#)" for how to build the application.

3.4.5.3 Running the Complex str_echo application

This section describes the steps for running the Complex str-echo application on i.MX 8M Plus LPDDR4 EVK and i.MX 8M Mini LPDDR4 EVK boards.

3.4.5.3.1 Running the application on i.MX 8M Plus LPDDR4 EVK

The following RTOS images are provided to run the application:

```
/examples/heterogeneous-multicore/rpmsg-str-echo-ca/ddr_release/
├── rpmsg_str_echo_ca53_RTOS0_RAM_CONSOLE-0xc0fff000.bin
├── rpmsg_str_echo_ca53_RTOS0_RAM_CONSOLE-0xc0fff000.elf
├── rpmsg_str_echo_ca53_RTOS0_UART4.bin
├── rpmsg_str_echo_ca53_RTOS0_UART4.elf
├── rpmsg_str_echo_ca53_RTOS1_RAM_CONSOLE-0xc1fff000.bin
├── rpmsg_str_echo_ca53_RTOS1_RAM_CONSOLE-0xc1fff000.elf
└── /examples/heterogeneous-multicore/rpmsg-str-echo-cm/release/
    ├── rpmsg_str_echo_cm7.bin
    └── rpmsg_str_echo_cm7.elf
```

There are two RTOS images provided for RTOS0 on Cortex-A53, one uses RAM console, the other uses UART4 console. Need to run A53 RAM Console image if running RTOS on Cortex M7 core simultaneously because Cortex M7 Core images use UART4 console by default.

Can use U-Boot commands under U-Boot or remoteproc under Linux to start FreeRTOS, choose one the method:

- Using U-Boot Commands to start RTOS, then start up Linux

1. Boot the First Cortex-A Core RTOS on A53 Core2:

```
u-boot=> ext4load mmc 1:2 0xC0000000 /examples/
heterogeneous-multicore/rpmsg-str-echo-ca/ddr_release/
rpmsg_str_echo_ca53_RTOS0_RAM_CONSOLE-0xc0fff000.bin
u-boot=> dcache flush; icache flush; cpu 2 release 0xC0000000
```

2. Boot the Second Cortex-A Core RTOS on A53 Core3:

```
u-boot=> ext4load mmc 1:2 0xC1000000 /examples/
heterogeneous-multicore/rpmsg-str-echo-ca/ddr_release/
rpmsg_str_echo_ca53_RTOS1_RAM_CONSOLE-0xc1fff000.bin
u-boot=> dcache flush; icache flush; cpu 3 release 0xC1000000
```

3. Boot M7 Core RTOS:

```
u-boot=> ext4load mmc 1:2 0x48000000 /examples/heterogeneous-multicore/
rpmsg-str-echo-cm/release/rpmsg_str_echo_cm7.bin
u-boot=> cp.b 0x48000000 0x7e0000 20000;
u-boot=> bootaux 0x7e0000
```

The below log is displayed for the UART4 console:

```
Cortex-M7: RTOS0: Multiple Endpoints RPMsg String Echo FreeRTOS Demo...
INFO: rpmsg_init          : RPMMSG init ...
INFO: rpmsg_init          : waiting for link establish ...
```

4. Boot up Linux using the commands:

```
u-boot=> setenv fdtfile imx8mp-evk-multicore-rpmsg.dtb
u-boot=> setenv mmcargs $mmcargs clk_ignore_unused
u-boot=> run bsp_bootcmd
```

5. When Linux is up, install tty driver module

```
root@imx8mp-lpddr4-evk:~# modprobe imx_rpmsg_tty
[ 21.770356] imx_rpmsg_tty virtio0.rpmsg-virtual-tty-channel-1.-1.3: new
channel: 0x400 -> 0x3!
[ 21.770576] Install rpmsg tty driver!
```

```
[ 21.773539] imx_rpmsg_tty virtio0.rpmsg-virtual-tty-channel-1.-1.4: new
channel: 0x401 -> 0x4!
[ 21.773804] Install rpmsg tty driver!
[ 21.774034] imx_rpmsg_tty virtio0.rpmsg-virtual-tty-channel-1.-1.5: new
channel: 0x402 -> 0x5!
[ 21.774156] Install rpmsg tty driver!
[ 21.774275] imx_rpmsg_tty virtio1.rpmsg-virtual-tty-channel-1.-1.6: new
channel: 0x400 -> 0x6!
[ 21.774360] Install rpmsg tty driver!
[ 21.774443] imx_rpmsg_tty virtio1.rpmsg-virtual-tty-channel-1.-1.7: new
channel: 0x401 -> 0x7!
[ 21.774530] Install rpmsg tty driver!
[ 21.774586] imx_rpmsg_tty virtio1.rpmsg-virtual-tty-channel-1.-1.8: new
channel: 0x402 -> 0x8!
[ 21.774663] Install rpmsg tty driver!
[ 21.774726] imx_rpmsg_tty virtio2.rpmsg-virtual-tty-channel-1.-1.0: new
channel: 0x400 -> 0x0!
[ 21.774810] Install rpmsg tty driver!
[ 21.774880] imx_rpmsg_tty virtio2.rpmsg-virtual-tty-channel-1.-1.1: new
channel: 0x401 -> 0x1!
[ 21.774960] Install rpmsg tty driver!
[ 21.775022] imx_rpmsg_tty virtio2.rpmsg-virtual-tty-channel-1.-1.2: new
channel: 0x402 -> 0x2!
[ 21.775111] Install rpmsg tty driver!
```

6. Then check the Cortex-A RTOS0's RAM Console Log:

```
root@imx8mp-lpddr4-evk:~# ram_console_dump -a 0xC0FFF000 -r 1
RAM Console@0xc0fff000:

Cortex-A53: RTOS0: Multiple Endpoints RPMsg String Echo FreeRTOS Demo...
RTOS1: RAM console@0xc1ffff00
INFO: rpmsg_init : RPMMSG remote init ...
INFO: rpmsg_remote_init : waiting for link establish ...
INFO: rpmsg_remote_init : RPMMSG link up
ept3: Get Message From Master Side : "hello world!" [len : 12]
ept4: Get Message From Master Side : "hello world!" [len : 12]
ept5: Get Message From Master Side : "hello world!" [len : 12]
> .....
```

7. Check Cortex-A RTOS1's RAM Console Log:

```
root@imx8mp-lpddr4-evk:~# ram_console_dump -a 0xC1FFF000 -r 1
RAM Console@0xc1fff000:

Cortex-A53: RTOS1: Multiple Endpoints RPMsg String Echo FreeRTOS Demo...
INFO: rpmsg_init : RPMMSG remote init ...
INFO: rpmsg_remote_init : waiting for link establish ...
INFO: rpmsg_remote_init : RPMMSG link up
ept6: Get Message From Master Side : "hello world!" [len : 12]
ept7: Get Message From Master Side : "hello world!" [len : 12]
ept8: Get Message From Master Side : "hello world!" [len : 12]
> .....
```

The following log for Cortex-M Core's UART4 Log:

```
Cortex-M7: RTOS0: Multiple Endpoints RPMsg String Echo FreeRTOS Demo...
INFO: rpmsg_init : RPMMSG remote init ...
INFO: rpmsg_remote_init : waiting for link establish ...
INFO: rpmsg_remote_init : RPMMSG link up
ept0: Get Message From Master Side : "hello world!" [len : 12]
ept1: Get Message From Master Side : "hello world!" [len : 12]
```

```
ept2: Get Message From Master Side : "hello world!" [len : 12]
```

It creates the following RPMsg devices:

```
root@imx8mp-lpddr4-evk:~# ls /dev/ttyRPMSG*
/dev/ttyRPMSG0 /dev/ttyRPMSG1 /dev/ttyRPMSG2 /dev/ttyRPMSG3 /dev/
ttyRPMSG4 /dev/ttyRPMSG5 /dev/ttyRPMSG6 /dev/ttyRPMSG7 /dev/ttyRPMSG8
```

- /dev/ttyRPMSG0 ~ 2 are three endpoints connected to Cortex-M Core RTOS;
- /dev/ttyRPMSG3 ~ 5 are three endpoints connected to Cortex-A Core RTOS0;
- /dev/ttyRPMSG6 ~ 8 are three endpoints connected to Cortex-A Core RTOS1.

- Or start up Linux, then using remoteproc to start RTOS

1. Start up Linux

```
u-boot=> setenv fdtfile imx8mp-evk-multicore-rpmsg.dtb
u-boot=> setenv mmcargs $mmcargs clk_ignore_unused
u-boot=> run bsp_bootcmd
```

2. Start the first A53 RTOS

After Linux boot up, start the first A53 Core RTOS on A53 Core2:

```
root@imx8mp-lpddr4-evk:~# echo /examples/heterogeneous-multicore/rpmsg-str-
echo-ca/ddr_release/rpmsg_str_echo_ca53_RTOS0_RAM_CONSOLE-0xc0fff000.elf >
/sys/devices/platform/remoteproc-ca53-2/remoteproc/remoteproc1/firmware
root@imx8mp-lpddr4-evk:~# echo start > /sys/devices/platform/remoteproc-
ca53-2/remoteproc/remoteproc1/state
```

The RAM Console log for the first A53 RTOS can be checked using the below command:

```
root@imx8mp-lpddr4-evk:~# ram_console_dump -a 0xc0fff000 -r 1
```

3. Start the second A53 RTOS

Then start the second A53 Core RTOS on A53 Core3:

```
root@imx8mp-lpddr4-evk:~# echo /examples/heterogeneous-multicore/rpmsg-str-
echo-ca/ddr_release/rpmsg_str_echo_ca53_RTOS1_RAM_CONSOLE-0xc1fff000.elf >
/sys/devices/platform/remoteproc-ca53-3/remoteproc/remoteproc2/firmware
root@imx8mp-lpddr4-evk:~# echo start > /sys/devices/platform/remoteproc-
ca53-3/remoteproc/remoteproc2/state
```

The RAM Console log for the second A53 RTOS can be checked with command:

```
root@imx8mp-lpddr4-evk:~# ram_console_dump -a 0xc1fff000 -r 1
```

4. Start M7 Core RTOS

```
root@imx8mp-lpddr4-evk:~# echo /examples/heterogeneous-multicore/rpmsg-str-
echo-cm/release/rpmsg_str_echo_cm7.elf > /sys//devices/platform/imx8mp-cm7/
remoteproc/remoteproc4/firmware
root@imx8mp-lpddr4-evk:~# echo start > /sys//devices/platform/imx8mp-cm7/
remoteproc/remoteproc4/state
```

5. Install tty driver module:

```
root@imx8mp-lpddr4-evk:~# modprobe imx_rpmsg_tty
```

It creates the following RPMsg devices:

```
root@imx8mp-lpddr4-evk:~# ls /dev/ttyRPMSG*
/dev/ttyRPMSG0 /dev/ttyRPMSG1 /dev/ttyRPMSG2 /dev/ttyRPMSG3 /dev/
ttyRPMSG4 /dev/ttyRPMSG5 /dev/ttyRPMSG6 /dev/ttyRPMSG7 /dev/ttyRPMSG8
```

- /dev/ttyRPMSG0 ~ 2 are three endpoints connected to Cortex-M Core RTOS;
- /dev/ttyRPMSG3 ~ 5 are three endpoints connected to Cortex-A Core RTOS0;

– /dev/ttyRPMMSG6 ~ 8 are three endpoints connected to Cortex-A Core RTOS1.

- Test RPMMSG Communication

Use "echo" or "minicom" to verify the RPMMsg communication between the two RTOS.

For example, use "echo" and send a sample string to the Cortex-M Core's endpoint:

```
root@imx8mp-lpddr4-evk:~# for i in {0..8}; do stty -F /dev/ttyRPMMSG$i -echo;
done
root@imx8mp-lpddr4-evk:~# echo "adfad" > /dev/ttyRPMMSG2
```

Then in the Cortex-M Core's Console, the below string is received:

```
ept2: Get Message From Master Side : "adfad" [len : 5]
ept2: Get New Line From Master Side
```

Or use a minicom to open one RPMMsg endpoint. It echoes back the character inputted in the minicom console:

```
root@imx8mp-lpddr4-evk:~# minicom -D /dev/ttyRPMMSG6
```

Then, input some characters typed in using the minicom console. This character is sent to the RTOS endpoint from Linux. Then the application running on RTOS sends the characters back to Linux. Finally all characters are echoed back in the minicom console. For example, if you input the characters "hello world!", the below log would be displayed:

```
Welcome to minicom 2.8

OPTIONS: I18n
Compiled on Jan  1 2021, 17:45:55.
Port /dev/ttyRPMMSG6, 08:36:41

Press CTRL-A Z for help on special keys

hello world!
```

Known Issues:

- Do not turn the data cache off in U-Boot while booting Cortex-M Core RTOS.

3.4.5.3.2 Running the application on i.MX 8M Mini LPDDR4 EVK

The following RTOS images are provided to run the application:

```
/examples/heterogeneous-multicore/rpmsg-str-echo-ca/ddr_release/
├── rpmsg_str_echo_ca53_RTOS0_RAM_CONSOLE-0x94bff000.bin
├── rpmsg_str_echo_ca53_RTOS0_RAM_CONSOLE-0x94bff000.elf
├── rpmsg_str_echo_ca53_RTOS0_UART4.bin
├── rpmsg_str_echo_ca53_RTOS0_UART4.elf
├── rpmsg_str_echo_ca53_RTOS1_RAM_CONSOLE-0x95bff000.bin
├── rpmsg_str_echo_ca53_RTOS1_RAM_CONSOLE-0x95bff000.elf
/examples/heterogeneous-multicore/rpmsg-str-echo-cm/release/
├── rpmsg_str_echo_cm4.bin
└── rpmsg_str_echo_cm4.elf
```

There are two RTOS images provided for RTOS0 on Cortex-A53, one uses RAM Console, the other uses UART4 Console, it needs to run RAM Console image if run RTOS on Cortex M4 Core simultaneously because Cortex M4 Core images uses UART4 Console by default.

Can use U-Boot commands under U-Boot or remoteproc under Linux to start FreeRTOS, choose one the method:

- Using U-Boot Commands to start RTOS, then start up Linux

1. Boot the First Cortex-A Core RTOS

```
u-boot=> ext4load mmc 1:2 0x93C00000 /examples/
heterogeneous-multicore/rpmsg-str-echo-ca/ddr_release/
rpmsg_str_echo_ca53_RTOS0_RAM_CONSOLE-0x94bff000.bin
u-boot=> dcache flush; icache flush; cpu 2 release 0x93C00000
```

2. Boot the Second Cortex-A Core RTOS

```
u-boot=> ext4load mmc 1:2 0x94C00000 /examples/
heterogeneous-multicore/rpmsg-str-echo-ca/ddr_release/
rpmsg_str_echo_ca53_RTOS1_RAM_CONSOLE-0x95bff000.bin
u-boot=> dcache flush; icache flush; cpu 3 release 0x94C00000
```

3. Boot Cortex-M Core RTOS

```
u-boot=> ext4load mmc 1:2 0x48000000 /examples/heterogeneous-multicore/
rpmsg-str-echo-cm/release/rpmsg_str_echo_cm4.bin
u-boot=> cp.b 0x48000000 0x7e0000 20000;
u-boot=> bootaux 0x7e0000
```

The log for UART4 Console will be the following log:

```
Cortex-M4: RTOS0: Multiple Endpoints RPMsg String Echo FreeRTOS Demo...
INFO: rpmsg_init          : RPMSG init ...
INFO: rpmsg_init          : waiting for link establish ...
```

4. Boot Linux up

```
u-boot=> setenv fdtfile imx8mm-evk-multicore-rpmsg.dtb
u-boot=> setenv mmcargs $mmcargs clk_ignore_unused
u-boot=> boot
```

5. When Linux is up, install tty driver module

```
root@imx8mm-lpddr4-evk:~# modprobe imx_rpmsg_tty
[ 21.770356] imx_rpmsg_tty virtio0.rpmsg-virtual-tty-channel-1.-1.3: new
channel: 0x400 -> 0x3!
[ 21.770576] Install rpmsg tty driver!
[ 21.773539] imx_rpmsg_tty virtio0.rpmsg-virtual-tty-channel-1.-1.4: new
channel: 0x401 -> 0x4!
[ 21.773804] Install rpmsg tty driver!
[ 21.774034] imx_rpmsg_tty virtio0.rpmsg-virtual-tty-channel-1.-1.5: new
channel: 0x402 -> 0x5!
[ 21.774156] Install rpmsg tty driver!
[ 21.774275] imx_rpmsg_tty virtio1.rpmsg-virtual-tty-channel-1.-1.6: new
channel: 0x400 -> 0x6!
[ 21.774360] Install rpmsg tty driver!
[ 21.774443] imx_rpmsg_tty virtio1.rpmsg-virtual-tty-channel-1.-1.7: new
channel: 0x401 -> 0x7!
[ 21.774530] Install rpmsg tty driver!
[ 21.774586] imx_rpmsg_tty virtio1.rpmsg-virtual-tty-channel-1.-1.8: new
channel: 0x402 -> 0x8!
[ 21.774663] Install rpmsg tty driver!
[ 21.774726] imx_rpmsg_tty virtio2.rpmsg-virtual-tty-channel-1.-1.0: new
channel: 0x400 -> 0x0!
[ 21.774810] Install rpmsg tty driver!
[ 21.774880] imx_rpmsg_tty virtio2.rpmsg-virtual-tty-channel-1.-1.1: new
channel: 0x401 -> 0x1!
[ 21.774960] Install rpmsg tty driver!
[ 21.775022] imx_rpmsg_tty virtio2.rpmsg-virtual-tty-channel-1.-1.2: new
channel: 0x402 -> 0x2!
```

```
[ 21.775111] Install rpmsg tty driver!
```

6. Then check Cortex-A RTOS0's RAM console log:

```
root@imx8mm-lpddr4-evk:~# ram_console_dump 0x94BFF000
RAM Console@0x94bff000:

Cortex-A53: RTOS0: Multiple Endpoints RPMsg String Echo FreeRTOS Demo...
RTOS1: RAM console@0x95Bff000
INFO: rpmsg_init          : RPMSG remote init ...
INFO: rpmsg_remote_init   : waiting for link establish ...
INFO: rpmsg_remote_init   : RPMSG link up
ept3: Get Message From Master Side : "hello world!" [len : 12]
ept4: Get Message From Master Side : "hello world!" [len : 12]
ept5: Get Message From Master Side : "hello world!" [len : 12]
> .....
```

7. And check the Cortex-A RTOS1's RAM console Log:

```
root@imx8mm-lpddr4-evk:~# ram_console_dump -a 0x95BFF000 -r 1
RAM Console@0x95bff000:

Cortex-A53: RTOS1: Multiple Endpoints RPMsg String Echo FreeRTOS Demo...
INFO: rpmsg_init          : RPMSG remote init ...
INFO: rpmsg_remote_init   : waiting for link establish ...
INFO: rpmsg_remote_init   : RPMSG link up
ept6: Get Message From Master Side : "hello world!" [len : 12]
ept7: Get Message From Master Side : "hello world!" [len : 12]
ept8: Get Message From Master Side : "hello world!" [len : 12]
> .....
```

The following is the log displayed for Cortex-M Core's UART4 console:

```
Cortex-M4: RTOS0: Multiple Endpoints RPMsg String Echo FreeRTOS Demo...
INFO: rpmsg_init          : RPMSG remote init ...
INFO: rpmsg_remote_init   : waiting for link establish ...
INFO: rpmsg_remote_init   : RPMSG link up
ept0: Get Message From Master Side : "hello world!" [len : 12]
ept1: Get Message From Master Side : "hello world!" [len : 12]
ept2: Get Message From Master Side : "hello world!" [len : 12]
```

The following RPMsg devices are created:

```
root@imx8mm-lpddr4-evk:~# ls /dev/ttyRPMSG*
/dev/ttyRPMSG0 /dev/ttyRPMSG1 /dev/ttyRPMSG2 /dev/ttyRPMSG3 /dev/
ttyRPMSG4 /dev/ttyRPMSG5 /dev/ttyRPMSG6 /dev/ttyRPMSG7 /dev/ttyRPMSG8
```

/dev/ttyRPMSG0 ~ 2 are three endpoints connected to Cortex-M Core RTOS, /dev/ttyRPMSG3 ~ 5 are three endpoints connected to Cortex-A Core RTOS0, /dev/ttyRPMSG6 ~ 8 are three endpoints connected to Cortex-A Core RTOS1.

• Or start up Linux, then using remoteproc to start RTOS

1. Start up Linux

```
u-boot=> setenv fdtfile imx8mm-evk-multicore-rpmsg.dtb
u-boot=> setenv mmcargs $mmcargs clk_ignore_unused
u-boot=> run bsp_bootcmd
```

2. Start the first A53 RTOS

After Linux boot up, start the first A53 Core RTOS on A53 Core2:

```
root@imx8mm-lpddr4-evk:~# echo /examples/heterogeneous-multicore/rpmsg-str-
echo-ca/ddr_release/rpmsg_str_echo_ca53_RTOS0_RAM_CONSOLE-0x94bff000.elf >
/sys/devices/platform/remoteproc-ca53-2/remoteproc/remoteproc1/firmware
```



```
root@imx8mm-lpddr4-evk:~# echo start > /sys/devices/platform/remoteproc-
ca53-2/remoteproc/remoteproc1/state
```

The RAM Console log for the first A53 RTOS can be checked with command:

```
root@imx8mm-lpddr4-evk:~# ram_console_dump -a 0x94bff000 -r 1
```

3. Start the second A53 RTOS

Then start the second A53 Core RTOS on A53 Core3:

```
root@imx8mm-lpddr4-evk:~# echo /examples/heterogeneous-multicore/rpmsg-str-
echo-ca/ddr_release/rpmsg_str_echo_ca53_RTOS1_RAM_CONSOLE-0x95bff000.elf >
/sys/devices/platform/remoteproc-ca53-3/remoteproc/remoteproc2/firmware
root@imx8mm-lpddr4-evk:~# echo start > /sys/devices/platform/remoteproc-
ca53-3/remoteproc/remoteproc2/state
```

The RAM Console log for the second A53 RTOS can be checked with command:

```
root@imx8mm-lpddr4-evk:~# ram_console_dump -a 0x95bff000 -r 1
```

4. Start M4 Core RTOS

```
root@imx8mm-lpddr4-evk:~# echo /examples/heterogeneous-multicore/rpmsg-str-
echo-cm/release/rpmsg_str_echo_cm4.elf > /sys//devices/platform/imx8mm-cm4/
remoteproc/remoteproc4/firmware
root@imx8mm-lpddr4-evk:~# echo start > /sys//devices/platform/imx8mm-cm4/
remoteproc/remoteproc4/state
```

5. Install tty driver module:

```
root@imx8mm-lpddr4-evk:~# modprobe imx_rpmsg_tty
```

It creates the following RPMSG devices:

```
root@imx8mm-lpddr4-evk:~# ls /dev/ttyRPMSG*
/dev/ttyRPMSG0 /dev/ttyRPMSG1 /dev/ttyRPMSG2 /dev/ttyRPMSG3 /dev/
ttyRPMSG4 /dev/ttyRPMSG5 /dev/ttyRPMSG6 /dev/ttyRPMSG7 /dev/ttyRPMSG8
```

/dev/ttyRPMSG0 ~ 2 are three endpoints connected to Cortex-M Core RTOS, /dev/ttyRPMSG3 ~ 5 are three endpoints connected to Cortex-A Core RTOS0, /dev/ttyRPMSG6 ~ 8 are three endpoints connected to Cortex-A Core RTOS1.

• Test RPMSG Communication

Use "echo" or "minicom" to verify the RPMSG communication between the two RTOS.

For example, use "echo" and send a sample string to the Cortex-M Core's endpoint:

```
root@imx8mm-lpddr4-evk:~# for i in {0..8}; do stty -F /dev/ttyRPMSG$i -echo;
done
root@imx8mm-lpddr4-evk:~# echo "adfad" > /dev/ttyRPMSG2
```

Then in the Cortex-M Core's Console, the below string is received:

```
ept2: Get Message From Master Side : "adfad" [len : 5]
ept2: Get New Line From Master Side
```

Or use a minicom to open one RPMSG endpoint. It echoes back the character inputted in the minicom console:

```
root@imx8mm-lpddr4-evk:~# minicom -D /dev/ttyRPMSG6
```

Then input some characters typed in using the minicom console. This character is sent to the RTOS endpoint from Linux. Then the application running on RTOS sends the characters back to Linux. Finally all characters are echoed back in the minicom console. For example, if you input the characters "hello world!", the below log would be displayed:

```
Welcome to minicom 2.8
```



```
OPTIONS: I18n
Compiled on Jan  1 2021, 17:45:55.
Port /dev/ttyRPMMSG6, 08:36:41

Press CTRL-A Z for help on special keys

hello world!
```

Known Issues:

- Do not turn the data cache off in U-Boot while booting Cortex-M Core RTOS.

3.5 RPMMSG based resource sharing

3.5.1 Overview

On NXP MPU platforms, in general, RTOS runs on Cortex-M Core(s) and Linux runs on Cortex-A Core(s). In some use cases, considering power management and real-time performance, Cortex-M Core owns and controls physical resources or peripherals, but shares these physical resources or peripherals with Cortex-A Core(s). The `rpmsg_lite_uart_sharing_rtos` is a FreeRTOS example to share physical UART owned by Cortex-M Core with Cortex-A Core.

3.5.2 Software architecture and design

This chapter describes different software architectures based on different technologies.

3.5.3 Resource sharing based on SRTM

This example uses the Simplified Real-Time Messaging (SRTM) protocol to communicate between Cortex-A and Cortex-M Cores. SRTM is used for communication among SoCs/processors in the same SoC. The figure below shows the software architecture for resource sharing based on SRTM.




Figure 33. Resource sharing software architecture

SRTM runs on Cortex-M Core which owns the hardware resources. To share these, it provides an application protocol based on RPMSG.

Virtual Device Drivers run on the resource user. The drivers provide standard device service on Cortex-A, which needs to use the hardware resources shared by SRTM.

3.5.3.1 UART sharing design details

The UART sharing example is designed with the following features:

- RTOS on Cortex-M Core owns and fully controls the physical UART ports.
- SRTM service runs on RTOS and provides physical device sharing service to Linux.
- The virtual UART driver on Linux provides a standard UART device service to applications.
- Multiple virtual UART ports are provided in Linux.
- Each virtual UART port in Linux can map to a dedicated physical UART on FreeRTOS.
- Multiple virtual UART ports can be mapped to the same physical UART.

Supported Platforms: i.MX 8M Mini LPDDR4 EVK, i.MX 93 EVK

It includes the following software components:

- Physical UART driver on FreeRTOS
- SRTM UART sharing service on FreeRTOS
- `rpmsg_lite_uart_sharing_rtos` application on FreeRTOS
- Virtual UART driver in Linux

[Figure 34](#) illustrates the software architecture of a UART sharing design.

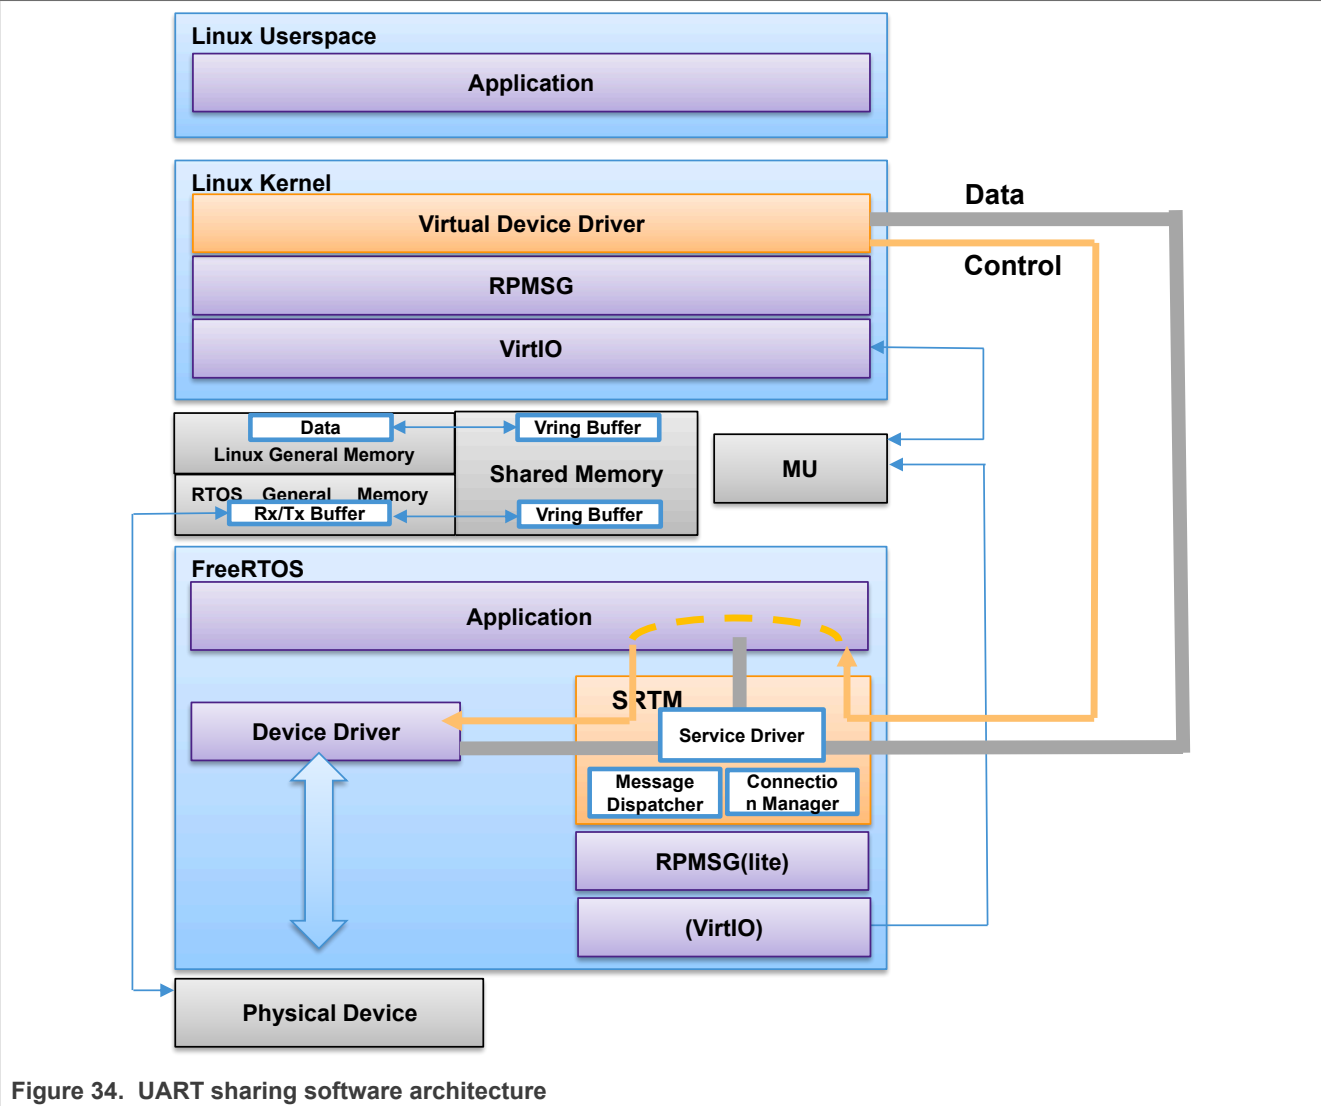


Figure 34. UART sharing software architecture

In order to support multiple virtual UART on a single physical UART, a multiple virtual UART protocol is used. The example described in this document follows the packet format described in the [Table 34](#).

Table 34. Packet format for UART sharing

Fields	Start Flags (4 bytes)				Address (1 bytes)	Payload Size (1 bytes)	Payload (n bytes)
HEX	24	55	54	2C	x	n	xxxxxx...
ASCII	\$	U	T	,			

- The packet header includes fields that indicate start flags, address, and payload size. It is 6 bytes by default.
- The “Start flags” field is used to figure out the start of data packets. Users can configure start flags with specified characters and size. The default start flags are 4 bytes: “\$UT,”.
 - The “Address” field is reused by receive from transmit directions. For the receive direction (blue colored path in [Figure 34](#)), it is the destination address or ID of target virtual device. For the transmit direction (orange colored path in [Figure 34](#)), it is the source address or ID from which virtual device is transmitted.

- The “Payload Size” is the size of the payload data. It is one byte, so the maximum payload size is 255 bytes. "Payload" is the actual data exchanged within protocol and it follows the packet header.

3.5.4 Source code files and configuration

1. Source code files:

The source files for different software components are listed in the [Table 35](#):

Table 35. Software source code list

Name	Software component	Source Files/Directory
FreeRTOS application: rpmsg_lite_uart_sharing_rtos	mcux-sdk-examples	evkmimx8mm/multicore_examples/ rpmsg_lite_uart_sharing_rtos/
SRTM Service	mcu-sdk	components/srtm/services/ srtm_uart_service.c srtm_uart_service.h srtm_uart_adapter.c srtm_uart_adapter.h
Virtual UART driver	real-time-edge-linux	drivers/tty/rpmsg_tty.c

2. Linux Virtual UART driver

By default, the Real-time Edge kernel builds the virtual UART driver as module (rpmsg_tty.ko) by enabling the configurable item: CONFIG_RPMSG_TTY=m.

3. Virtual UART and physical UART mapping

The UART Sharing Service supports **three modes** of mapping between **virtual UART** and **physical UART**:

a. Virtual UART to physical UART 1:1 mapping

- Virtual UARTs on the A-core have 1:1 mapping to physical UART on the M-core.
- Each physical UART connects to a different device.
- Each virtual UART uses a dedicated RPMsg endpoint.

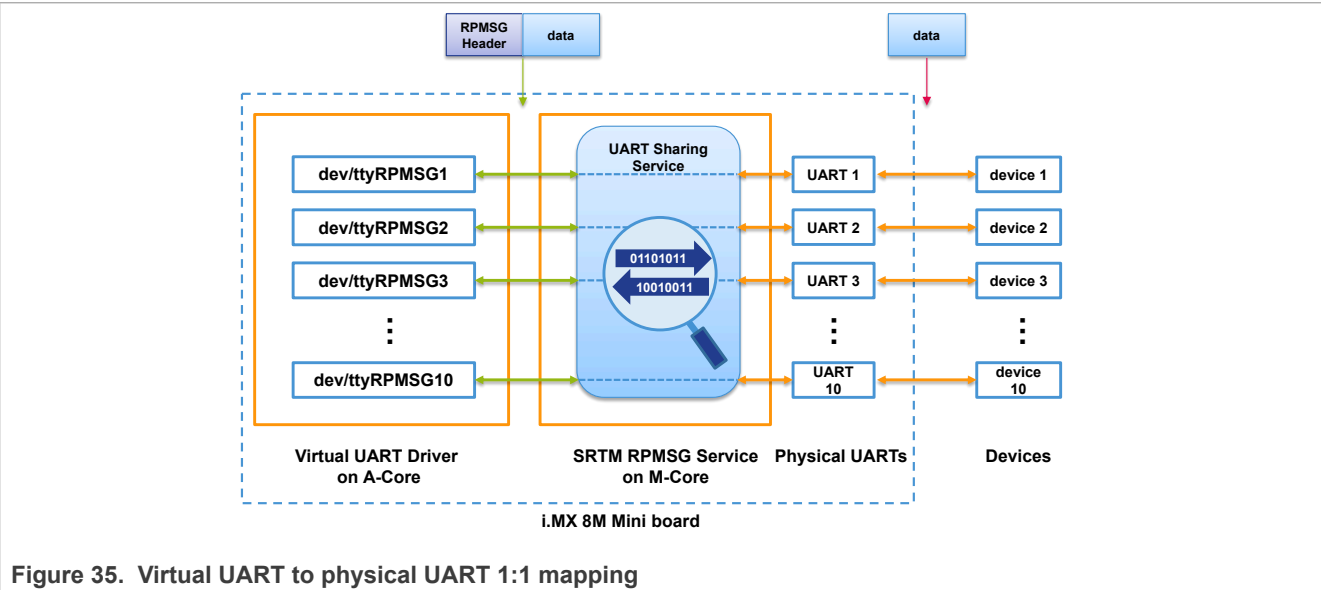


Figure 35. Virtual UART to physical UART 1:1 mapping

b. Virtual UART to physical UART n:1 mapping

- Multiple Virtual UARTs on A-core maps to a single physical UARTs on M-core.
- Physical UART connects to a device or another board.
- Each virtual UART uses a dedicated RPMsg Endpoint.

- A multiple UART Header is used to establish multiple virtual UART channels on a single physical UART connect. For details about multiple UART Headers, refer to the section [Section 3.5.3.1](#).

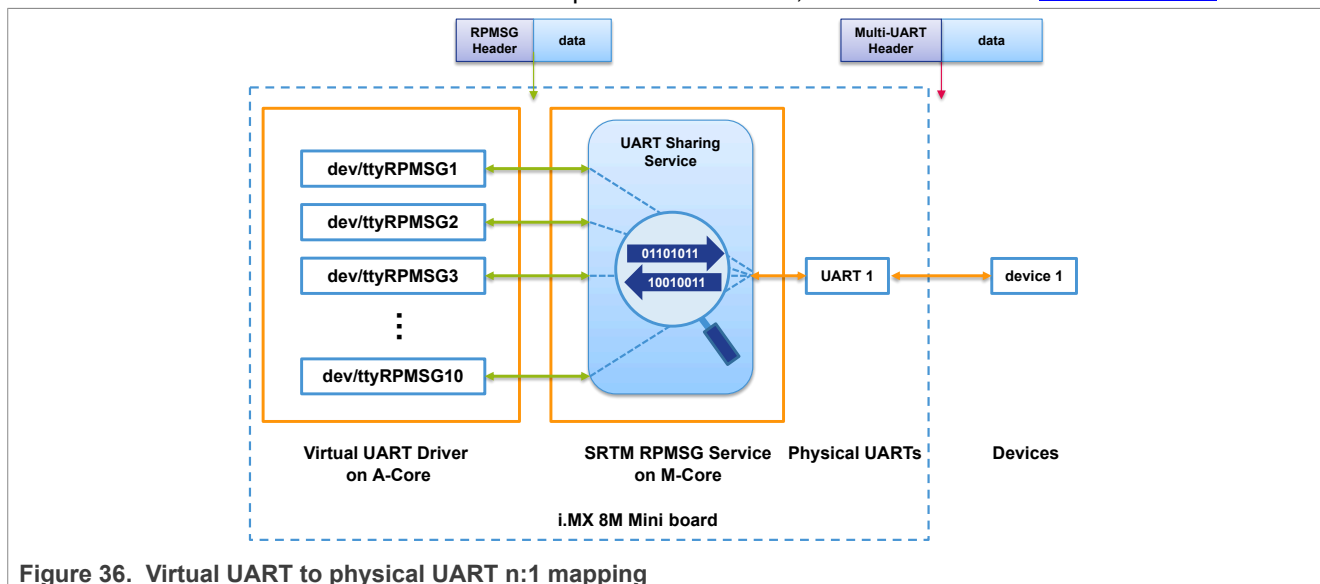


Figure 36. Virtual UART to physical UART n:1 mapping

- c. Virtual UART to physical UART flexible mapping: This mapping mode can support virtual UART to physical UART 1:1 mapping and n:1 mapping simultaneously. The following figure shows flexible mapping between two i.MX 8M Mini boards.

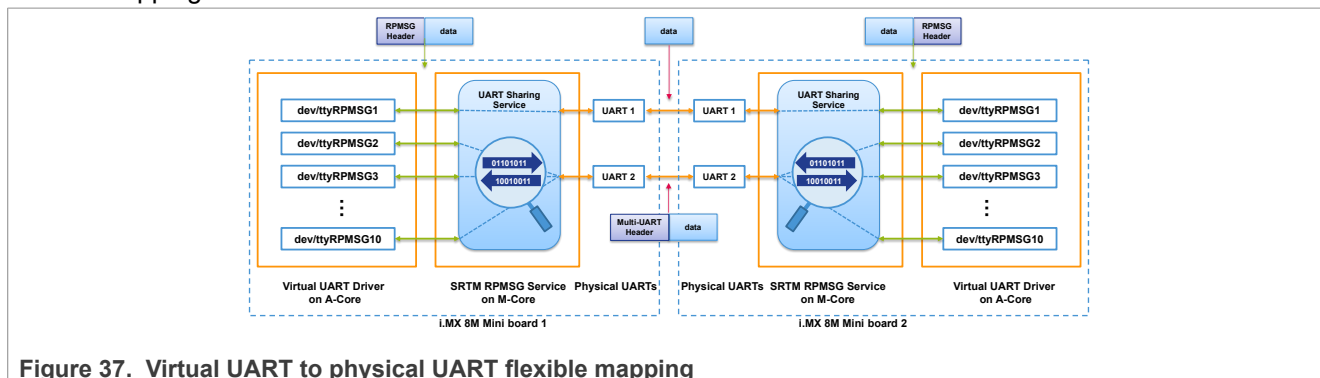


Figure 37. Virtual UART to physical UART flexible mapping

The mapping between virtual UART and physical UART is configured in Linux device tree, as shown in a dts node example below:

```
uart_rpbus_3: uart-rpbus-3 {
    compatible = "fsl,uart-rpbus";
    bus_id = <3>; /* use uart3 */
    flags=<IMX_SRTM_UART_SUPPORT_MULTI_UART_MSG_FLAG>;
    status = "okay";
};
```

This dts node is configured for virtual UART3.

Note:

- The “bus_id” specifies the physical UART instance ID that this virtual UART maps to. If the property of “bus_id” is not configured, the message sent from Linux to this virtual UART is display on M-core’s debug console directly.
- Physical UART ID is configured in the FreeRTOS application “rpbmsg_lite_uart_sharing_rtos”.
- On i.MX 8M Mini LPDDR4 EVK, physical UART3 can be used, so all virtual UART ports are mapped to physical UART3 by default.

- On i.MX 93 EVK, physical LPUART5 can be used, so all virtual UART ports are mapped to physical LPUART5 by default.
- If *flags* is set with the value `IMX_SRTM_UART_SUPPORT_MULTI_UART_MSG_FLAG`, the multiple virtual UART is mapped to a single physical UART instance specified by *bus_id* (that implies that multiple virtual UART protocol packet headers are used).
- If *flags* is not set, this virtual UART is mapped 1:1 with physical UART instance specified by *bus_id*.
By default, there are 11 virtual UARTs in the dtb file `imx8mm-evk-rpmsg.dtb` for i.MX 8M Mini LPDDR4 EVK and `imx93-11x11-evk-rpmsg.dtb` for i.MX 93 EVK.
- The virtual UART 0 to 9 are n:1 mapped to physical UART.
- The virtual UART 10 has no *bus_id* and displays messages sent from Linux to M-core's debug console.

3.5.5 Building and running the demo on i.MX 8M Mini LPDDR4 EVK

3.5.5.1 Hardware setup for i.MX 8M Mini EVK

Use flying wire to connect UART3 between two i.MX 8M Mini EVK boards. UART3's pin is provided in J1003 connector; use the following pin connection between the two boards.

Table 36. PIN connection between two i.MX 8M Mini boards

i.MX 8M Mini Board1		Connection	i.MX 8M Mini Board2	
Pin	Function		Pin	Function
6	GND	<->	6	GND
8	UART3_TXD	<->	10	UART3_RXD
10	UART3_RXD	<->	8	UART3_TXD

3.5.5.2 Building the demo images

The demo images "rpmsg_uart_sharing.bin" are by default compiled with the i.MX 8M Mini LPDDR4 EVK target image compiling, and are installed into the "/examples" directory of the target rootfs.

Or the image can be built separately by using the following Yocto command:

```
DISTRO=nxp-real-time-edge MACHINE=imx8mm-lpddr4-evk bitbake rpmsg-uart-sharing-cm
```

The image can be found on directory "<image-build-dir>/tmp/deploy/images/imx8mmevk/examples/" on building host.

3.5.5.3 Running the i.MX 8M Mini EVK demo

1. Connect two i.MX 8M Mini EVK boards by following the steps in section of "[Hardware Setup](#)".
2. Connect two i.MX 8M Mini EVK boards to your PC via USB cable between the USB-UART connector and the PC USB connector.
3. Open the terminal application on the PC, such as PuTTY or TeraTerm, and connect to the debug serial port number, two debug consoles for each board, one for the Linux debug console and another for the FreeRTOS debug console.
4. Deploy Real-time Edge release root files in SD card and modify the on-board switch to boot from MicroSD card.
5. Power on the board and enter into U-Boot command line, then execute the following command:

```
u-boot => setenv fdtfile imx8mm-evk-rpmsg.dtb
```

To make the above change permanent, execute the following command once:

```
u-boot => saveenv
```

6. Then, use the following commands to download and run FreeRTOS image:

```
u-boot => ext4load mmc 1:2 0x48000000 /examples/heterogeneous-multicore/
rpmsg-uart-sharing-cm/release/rpmsg_uart_sharing.bin;
u-boot => cp.b 0x48000000 0x7e0000 20000; bootaux 0x7e0000
```

Then, FreeRTOS debug console displays the following log:

```
##### RPMSG UART SHARING DEMO #####
Build Time: Mar 2 2022--09:38:19
*****
Wait the Linux kernel boot up to create the link between M core and A
core.
```

7. Then, boot the Linux kernel by executing the following commands:

```
u-boot => setenv jh_clk clk_ignore_unused
u-boot => boot
```

After the Linux kernel boots up, the below line on the FreeRTOS debug console indicates that RPMSG connection between Cortex-A core and Cortex-M core has been established:

```
Task A is working now.
```

Execute the above steps (1 to 7) on each i.MX 8M Mini EVK board.

8. After Linux boots up, enter Linux command line, use the following commands to test the demo:
 - a. Check device files are available:

```
root@imx8mm-lpddr4-evk:~# ls /dev/ttyRPMSG*
```

If the default dtb file `imx8mm-evk-rpmsg.dtb` is used, 11 device files from `/dev/ttyRPMSG0` to `/dev/ttyRPMSG10` must be displayed. The `/dev/ttyRPMSG0` to `/dev/ttyRPMSG9` have n:1 mapping to physical UART3. The `/dev/ttyRPMSG10` is without `bus_id` and displays the message sent from Linux to M-core's debug console.

- b. Check each virtual UART from `/dev/ttyRPMSG0` to `/dev/ttyRPMSG9` is connected to peer virtual UART between two boards, for example, use "minicom" to open and configure the same virtual UART on both boards:

```
root@imx8mm-lpddr4-evk:~# minicom -s -D /dev/ttyRPMSG6
```

Then, configure the UART as shown in the following figures:

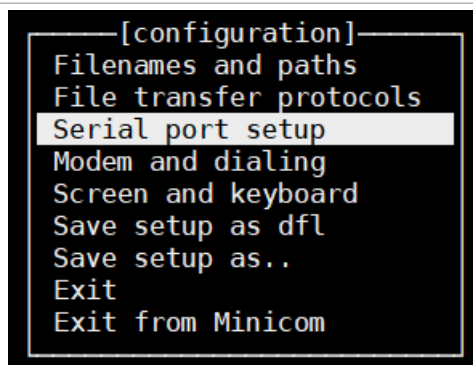


Figure 38. Configuring RPMSG Virtual UART Step1




Figure 39. Configuring RPMSG Virtual UART Step2

Save the settings and back to minicom main window, then input any characters in one board's minicom window. These characters would be displayed on the other board's minicom window.

3.5.6 Building and running the demo on i.MX 93 EVK and i.MX 93 14x14 EVK

3.5.6.1 Hardware setup for i.MX 93 EVK and i.MX 93 14x14 EVK

Use flying wire to connect LPUART5 between two i.MX 93 EVK or i.MX 93 14x14 EVK boards.

- On i.MX 93 EVK, LPUART5's pin is provided in J1001 connector.
- On i.MX 93 14x14 EVK, LPUART5's pin is provided in J16 connector.

Use the following pin connection between the two boards.

Table 37. Pin connections between two i.MX 93 EVK or i.MX 93 14x14 EVK boards

Board 1		Connection	Board 2	
Pin	Function		Pin	Function
30	GND	<->	30	GND
28	LPUART5_RX	<->	27	LPUART5_TX
27	LPUART5_TX	<->	28	LPUART5_RX

3.5.6.2 Building the demo images

The demo image "rpmsg_uart_sharing.bin" is by default compiled with the i.MX 93 EVK and and i.MX 93 14x14 EVK target image compiling, and are installed into the "/examples" directory of the target rootfs.

Or the image can be built separately by using the following Yocto command:

```
# For i.MX 93 EVK
DISTRO=nxp-real-time-edge MACHINE=imx93evk source real-time-edge-setup-env.sh -b
<build_dir>
# For i.MX93 14x14 EVK
```

```
DISTRO=nxp-real-time-edge MACHINE=imx93-14x14-lpddr4x-evk source real-time-edge-
setup-env.sh -b <build_dir>
```

```
bitbake rpmsg-uart-sharing-cm
```

The image can be found on directory "<image-build-dir>/tmp/deploy/images/<MACHINE>/examples/" on building host.

3.5.6.3 Running the demo on i.MX 93 EVK and i.MX 93 14x14 EVK

1. Connect two i.MX 93 EVK or i.MX 93 14x14 EVK boards by following the steps listed in [Section 3.5.6.1](#).
2. Connect two i.MX 93 EVK or i.MX 93 14x14 EVK boards to your PC via USB cable between the USB-UART connector and the PC USB connector.
3. Open the terminal application on the PC, such as PuTTY or TeraTerm, and connect to the debug serial port number, four debug consoles for each board. Use the third one for the Linux debug console and the fourth one for the FreeRTOS debug console.
4. Deploy Real-time Edge release root files in SD card and modify the on-board switch to boot from MicroSD card.
5. Power on the board and enter into U-Boot command line. Then, execute the following command:

```
# Running on i.MX 93 EVK
u-boot => setenv fdtfile imx93-11x11-evk-uart-sharing-cm33.dtb
# Running on i.MX 93 14x14 EVK
u-boot => setenv fdtfile imx93-14x14-evk-uart-sharing-cm33.dtb
```

To make changes permanent, execute the following commands once (after `setenv` above):

```
u-boot => saveenv
```

6. Then, use the following command to download and run FreeRTOS image:

```
u-boot => ext4load mmc 1:2 0x80000000 /examples/heterogeneous-multicore/
rpmsg-uart-sharing-cm/release/rpmsg_uart_sharing.bin
u-boot => cp.b 0x80000000 0x201e0000 0x10000
u-boot => bootaux 0x1ffe0000 0
```

Then, FreeRTOS debug console would display the following log:

```
##### RPMSG UART SHARING DEMO #####
Build Time: Apr 5 2011 23:00:00
Start SRTM communication
*****
Wait for the Linux kernel boot up to create the link between M core and A
core.
*****
```

7. And then, boot Linux kernel by executing the following command:

```
u-boot => setenv jh_clk clk_ignore_unused
u-boot => boot
```

8. After the Linux kernel boots up, in the FreeRTOS console, an extra line of log as shown below indicates that RPMSG connection between Cortex-A core and Cortex-M core has been established:

```
Task A is working now.
```

Execute the above steps (1 to 7) on each of the two boards.

9. After Linux boots up, enter Linux command line, use the following commands to test the demo:

a. Check device files are available:

```
root @imx93evk:~# ls /dev/ttyRPMMSG*
```

There should be 11 device files from “/dev/ttyRPMMSG0” to “/dev/ttyRPMMSG10”. The “/dev/ttyRPMMSG0” to “/dev/ttyRPMMSG9” have n:1 mapping to physical LPUART5, “/dev/ttyRPMMSG10” is without “bus_id” and displays the message sent from Linux to M-core’s debug console.

b. Check each virtual UART from “/dev/ttyRPMMSG0” to “/dev/ttyRPMMSG9” is connected to peer virtual UART between two boards, for example, execute the following on the first board:

• Use “minicom” to open and configure the same virtual UART on both boards:

```
root@imx93evk:~# minicom -s -D /dev/ttyRPMMSG6
```

Then configure the UART as shown in the following figures:




Figure 40. Configure RPMMSG Virtual UART Step1




Figure 41. Configure RPMMSG Virtual UART Step2

Save the settings and go back to minicom main window, then input any characters in one board's minicom window. Then, the characters would be displayed on the other board's minicom window.

3.5.7 Building and running the demo on i.MX 93 QSB

3.5.7.1 Hardware setup for i.MX 93 QSB

Use flying wire to connect LPUART5 between two i.MX 93 QSB boards. LPUART5's pin is provided in J1401 connector. Use the following pin connection between the two boards.

Table 38. PIN connection between two i.MX 93 QSB boards

i.MX 93 QSB Board1		Connection	i.MX 93 QSB Board2	
Pin	Function		Pin	Function
30	GND	<->	30	GND
28	LPUART5_RX	<->	27	LPUART5_TX
27	LPUART5_TX	<->	28	LPUART5_RX

3.5.7.2 Building the demo images

The demo image "rpmsg_lite_uart_sharing_rtos.bin" is by default compiled with the i.MX 93 QSB target image compiling, and are installed into the "/examples" directory of the target rootfs.

Or the image can be built separately by using the following Yocto command:

```
DISTRO=nxp-real-time-edge MACHINE=imx93-9x9-lpddr4-qsb source real-time-edge-setup-env.sh -b <build_dir>

bitbake rpmsg-lite-uart-sharing-rtos-mcimx93qsb
```

The image can be found on directory "<image-build-dir>/tmp/deploy/images/imx93-9x9-lpddr4-qsb/examples/" on building host.

3.5.7.3 Running the demo on i.MX 93 QSB

- 1. Connect two i.MX 93 QSB boards by following the steps listed in [Section 3.5.7.1](#).
- 2. Connect two i.MX 93 QSB boards to your PC via USB cable between the USB-UART connector and the PC USB connector.
- 3. Open the terminal application on the PC, such as PuTTY or TeraTerm, and connect to the debug serial port number, four debug consoles for each board. Use the third one for the Linux debug console and the fourth one for the FreeRTOS debug console.
- 4. Deploy Real-time Edge release root files in SD card and modify the on-board switch to boot from MicroSD card.
- 5. Power on the board and enter into U-Boot command line. Then, execute the following command:

```
u-boot => setenv fdtfile imx93-9x9-qsb-uart-sharing-cm33.dtb
```

To make changes permanent, execute the following commands once (after setenv above):

```
u-boot => saveenv
```

- 6. Then, use the following command to download and run FreeRTOS image:

```
u-boot => ext4load mmc 1:2 0x80000000 /examples/rpmsg-lite-uart-sharing-rtos-mcimx93qsb/release/rpmsg_lite_uart_sharing_rtos.bin
u-boot => cp.b 0x80000000 0x201e0000 0x10000
u-boot => bootaux 0x1ffe0000 0
```

Then, FreeRTOS debug console would display the following log:

```
##### RPMSG UART SHARING DEMO #####
Build Time: Oct 30 2023 11:20:34
Start SRTM communication
*****
Wait for the Linux kernel boot up to create the link between M core and A
core.
*****
```

7. And then, boot Linux kernel by executing the following command:

```
u-boot => setenv jh_clk clk_ignore_unused
u-boot => boot
```

8. After the Linux kernel boots up, in the FreeRTOS console, an extra line of log as shown below indicates that RPMSG connection between Cortex-A core and Cortex-M core has been established:

```
Task A is working now.
```

Execute the above steps (1 to 7) on each i.MX 93 QSB board.

9. After Linux boots up, enter Linux command line using the following commands to test the demo:
 - a. Check the device files are available:

```
root@imx93-9x9-lpddr4-qsb:~# ls /dev/ttyRPMSG*
```

There should be 11 device files from “/dev/ttyRPMSG0” to “/dev/ttyRPMSG10” if the default dtb file `imx93-9x9-qsb-uart-sharing-cm33.dtb` is used. The “/dev/ttyRPMSG0” to “/dev/ttyRPMSG9” have n:1 mapping to physical LPUART5, “/dev/ttyRPMSG10” is without “bus_id” and displays the message sent from Linux to M-core’s debug console.

- b. Check each virtual UART from “/dev/ttyRPMSG0” to “/dev/ttyRPMSG9” is connected to peer virtual UART between two boards, for example, ues "minicom" to open and configure the same virtual UART on both boards:

```
root@imx93-9x9-lpddr4-qsb:~# minicom -s -D /dev/ttyRPMSG6
```

Then configure the UART as shown in the following figures:




Figure 42. Configure RPMSG Virtual UART Step1

```
A - Serial Device      : /dev/ttyRPMSG6
B - Lockfile Location  : /var/lock
C - Callin Program     :
D - Callout Program    :
E - Bps/Par/Bits       : 115200 8N1
F - Hardware Flow Control : No
G - Software Flow Control : No
H - RS485 Enable       : No
I - RS485 Rts On Send  : No
J - RS485 Rts After Send : No
K - RS485 Rx During Tx  : No
L - RS485 Terminate Bus : No
M - RS485 Delay Rts Before: 0
N - RS485 Delay Rts After : 0

Change which setting? █
```

Figure 43. Configure RPMSG Virtual UART Step2

Save the settings and go back to minicom main window, then input any characters in one board's minicom window. Then, the characters would be displayed on the other board's minicom window.

3.6 Heterogeneous Multicore VirtIO and networking sharing

3.6.1 Heterogeneous Multicore VirtIO

Heterogeneous Multicore VirtIO applies para-virtualization VirtIO technology to build resource sharing between Heterogeneous asymmetric multiprocessing (AMP). The main difference from para-virtualization VirtIO is that Heterogeneous Multicore VirtIO does not use and depend on any hypervisor. Therefore, it can be used for resource sharing between Cortex-A and Cortex-M cores, or between multiple Cortex-A cores.

The VirtIO is a standard for para-virtualization to provide high-performance IO device virtualization for VM. The [Figure 44](#) shows the architecture of the VirtIO solution.




Figure 44. Para-virtualization VirtIO on hypervisor

The front-end VirtIO driver runs in the guest kernel space, and the backend VirtIO device runs in the Hypervisor. The Virtqueue and Vring provide data transfer capability via shared memory. The hypervisor emulates the device logic by VMExit and injecting vCPU IRQ. Therefore, para-virtualization VirtIO is used for resource sharing between Virtual Machine guest OS and host OS, and it depends on the Hypervisor to run the VirtIO backend.

The Heterogeneous Multicore VirtIO in Real Time Edge uses VirtIO technology. However, it runs VirtIO backend on any CPU core including Cortex-A core and Cortex-M core. The VirtIO front-end runs on any other CPU core. This technology uses VirtIO to establish communication between the front-end and backend, which run on different CPU cores. Therefore, VirtIO can be used to share hardware resources between different CPU cores.

In the current implementation, the backend runs on RTOS and owns the hardware resources, such as peripherals. The front-end runs on Linux. As there is no hypervisor providing VMExit and vCPU IRQ injecting mechanism, the hardware or software mailbox between the front-end and back-end is needed. [Figure 45](#) shows the architecture of the Heterogeneous Multicore VirtIO.

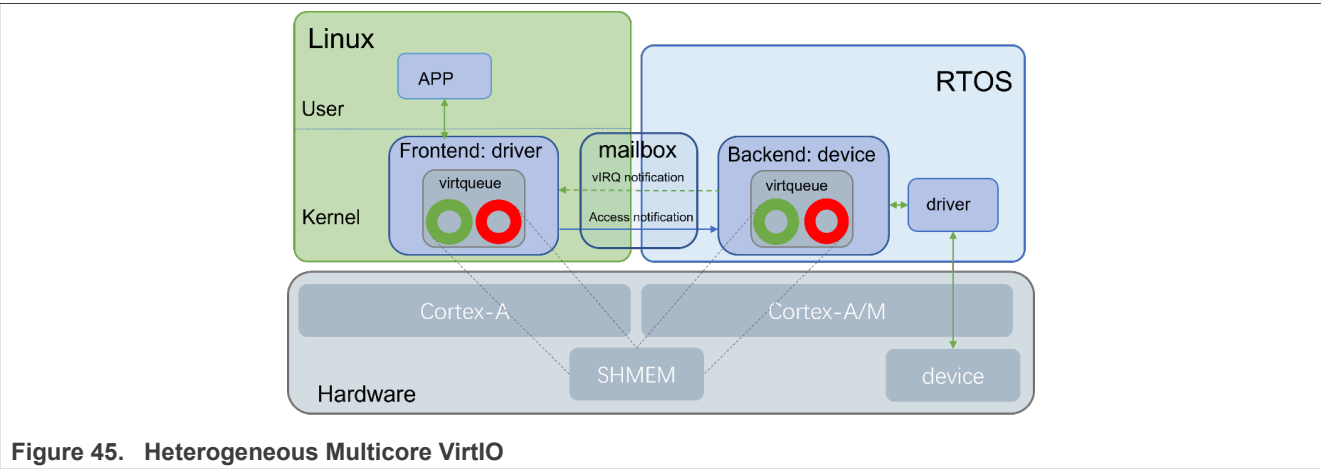


Figure 45. Heterogeneous Multicore VirtIO

Heterogeneous Multicore VirtIO uses shared memory to build Vring structure and data buffers. The shared memory must be coherent between front end and Backend.

3.6.2 Heterogeneous Multicore VirtIO performance evaluation

3.6.2.1 Overview

A VirtIO transmission device is introduced for evaluating the performance between the frontend and backend through virtqueues. There are 2 virtqueues for transmitting and receiving directions separately, and the device configuration registers are used to configure and control the test cases.

The [Table 39](#) lists the combination of supported cases.

Table 39. Heterogeneous Multicore VirtIO performance evaluation

Direction	Pkt size	Frontend Buffer copy	Backend Buffer copy
TX (Linux -> FreeRTOS)	Max 2KB	Y/N	Y/N
RX (FreeRTOS -> Linux)	Max 2KB	Y/N	Y/N

3.6.2.2 Building VirtIO performance evaluation backend application

VirtIO performance evaluation backend application runs on FreeRTOS, it is an application named "virtio_perf" in the repo: [heterogeneous-multicore](#).

Refer to "[Building Heterogeneous Multicore RTOS Application](#)" for steps to build the application.

3.6.2.3 Running VirtIO performance testing

The VirtIO performance testing supports running on i.MX 8M Mini LPDDR4 EVK and i.MX 8M Plus LPDDR4 EVK.

Note: The released image does not display the CPU load status of the RTOS side during the evaluation, to display the CPU load status, build the debug type image manually referring to the steps listed in [Section 3.2.2 Build with Standalone Mode](#).

Perform the following steps for VirtIO performance testing:

1. Set up the UART console for the frontend and backend

Connect the DEBUG UART slot on the board to your PC through the USB Cable. On the PC, this step creates two USB serial ports (port0 and port1) for i.MX 8M Mini EVK board, and four USB serial ports (port0 ~ port3) for i.MX 8M Plus EVK board. Open two UART consoles for UART port0 and port1 on i.MX 8M Mini EVK board or port2 and port3 on i.MX 8M Plus EVK board with the following setup:

- 115200
- No parity
- 8 data bits
- 1 stop bit

The first UART console is used for Linux that runs the VirtIO frontend, the another is used for RTOS that runs the VirtIO backend.

2. Boot up VirtIO backend FreeRTOS and frontend Linux

Heterogeneous Multicore VirtIO Backend can run on Cortex-A core or Cortex-M core to evaluate different use cases.

- For i.MX 8M Mini LPDDR4 EVK

Users can use U-Boot commands or Linux remoteproc to start the backend RTOS. Choose one method to start RTOS and Linux:

– Using U-Boot Commands to start RTOS, then start Linux

a. Run the backend on Cortex-M Core

On the U-Boot command line, execute the following commands to boot Cortex-M core with backend firmware:

```
u-boot => ext4load mmc 1:2 0x48000000 /examples/heterogeneous-
multicore/virtio-perf-cm/release/virtio_perf_cm4.bin
u-boot => cp.b 0x48000000 0x7e0000 0x20000
u-boot => bootaux 0x7e0000
```

Then, boot Linux kernel:

```
u-boot => setenv fdtfile imx8mm-evk-virtio-perf-cm4.dtb
u-boot => setenv mmcargs $mmcargs clk_ignore_unused
u-boot => run bsp_bootcmd
```

b. Or run the backend on Cortex-A Core

Execute the following command in U-Boot command line:

```
u-boot => ext4load mmc 1:2 0x93c00000 /examples/heterogeneous-
multicore/virtio-perf-ca/ddr_release/virtio_perf_ca53.bin
u-boot => dcache flush; icache flush; cpu 3 release 0x93c00000
```

Then, boot Linux kernel:

```
u-boot => setenv fdtfile imx8mm-evk-virtio-perf-ca53.dtb
u-boot => setenv mmcargs $mmcargs clk_ignore_unused
u-boot => run bsp_bootcmd
```

– Or start Linux, then use remoteproc to start RTOS

a. Run the backend on Cortex-M Core

Boot up Linux:

```
u-boot => run prepare_mcore; mw.l 0xb8400000 0 1
u-boot => setenv fdtfile imx8mm-evk-virtio-perf-cm4.dtb
u-boot => setenv mmcargs $mmcargs clk_ignore_unused
u-boot => run bsp_bootcmd
```

After Linux boot up, start backend RTOS running on Cortex-M Core by using remoteproc under Linux:

```
root@imx8mm-lpddr4-evk:~# modprobe -r virtio_trans
root@imx8mm-lpddr4-evk:~# modprobe -r virtio_mmio
root@imx8mm-lpddr4-evk:~# echo /examples/heterogeneous-multicore/
virtio-perf-cm/release/virtio_perf_cm4.elf > /sys/devices/platform/
imx8mm-cm4/remoteproc/remoteproc0/firmware
root@imx8mm-lpddr4-evk:~# echo start > /sys/devices/platform/imx8mm-
cm4/remoteproc/remoteproc0/state
root@imx8mm-lpddr4-evk:~# modprobe virtio_mmio
root@imx8mm-lpddr4-evk:~# modprobe virtio_trans
```

b. Or run the backend on Cortex-A Core

Boot up Linux:

```
u-boot => setenv fdtfile imx8mm-evk-virtio-perf-ca53.dtb ; mw.l
0xb8400000 0 1
u-boot => setenv mmcargs $mmcargs clk_ignore_unused
u-boot => run bsp_bootcmd
```

After Linux boot up, start backend RTOS running on Cortex-A Core by using remoteproc under Linux:

```
root@imx8mm-lpddr4-evk:~# modprobe -r virtio_trans
root@imx8mm-lpddr4-evk:~# modprobe -r virtio_mmio
root@imx8mm-lpddr4-evk:~# echo /examples/heterogeneous-multicore/
virtio-perf-ca/ddr_release/virtio_perf_ca53.elf > /sys/devices/
platform/remoteproc-ca53-3/remoteproc/remoteproc2/firmware
root@imx8mm-lpddr4-evk:~# echo start > /sys/devices/platform/
remoteproc-ca53-3/remoteproc/remoteproc2/state
root@imx8mm-lpddr4-evk:~# modprobe virtio_mmio
root@imx8mm-lpddr4-evk:~# modprobe virtio_trans
```

- **For i.MX 8M Plus LPDDR4 EVK**

Can use U-Boot commands or Linux remoteproc to start backend RTOS, choose one method to start RTOS and Linux:

- **Using U-Boot Commands to start RTOS, then start Linux**

- a. Run the backend on Cortex-M Core

On the U-Boot command line, execute the following commands to boot Cortex-M core with backend firmware:

```
u-boot => ext4load mmc 1:2 0x48000000 /examples/heterogeneous-
multicore/virtio-perf-cm/release/virtio_perf_cm7.bin
u-boot => cp.b 0x48000000 0x7e0000 0x20000
u-boot => bootaux 0x7e0000
```

Then, boot Linux kernel:

```
u-boot => setenv fdtfile imx8mp-evk-virtio-perf-cm7.dtb
u-boot => setenv mmcargs $mmcargs clk_ignore_unused
u-boot => run bsp_bootcmd
```

- b. Run the backend on Cortex-A Core

Execute the following command in U-Boot command line:

```
u-boot => ext4load mmc 1:2 0xc0000000 /examples/heterogeneous-
multicore/virtio-perf-ca/ddr_release/virtio_perf_ca53.bin
u-boot => dcache flush; icache flush; cpu 3 release 0xc0000000
```

Then, boot Linux kernel:

```
u-boot => setenv fdtfile imx8mp-evk-virtio-perf-ca53.dtb
u-boot => setenv mmcargs $mmcargs clk_ignore_unused
u-boot => run bsp_bootcmd
```

- **Or start Linux, then use remoteproc to start RTOS**

- a. Run the backend on Cortex-M Core

Boot up Linux:

```
u-boot => run prepare_mcore; mw.l 0xb8400000 0 1
u-boot => setenv fdtfile imx8mp-evk-virtio-perf-cm7.dtb
u-boot => setenv mmcargs $mmcargs clk_ignore_unused
u-boot => run bsp_bootcmd
```

After Linux boot up, start backend RTOS running on Cortex-M Core by using remoteproc under Linux:

```
root@imx8mp-lpddr4-evk:~# modprobe -r virtio_trans
root@imx8mp-lpddr4-evk:~# modprobe -r virtio_mmio
root@imx8mp-lpddr4-evk:~# echo /examples/heterogeneous-multicore/
virtio-perf-cm/release/virtio_perf_cm7.elf > /sys/devices/platform/
imx8mp-cm7/remoteproc/remoteproc0/firmware
```

```
root@imx8mp-lpddr4-evk:~# echo start > /sys/devices/platform/imx8mp-
cm7/remoteproc/remoteproc0/state
root@imx8mp-lpddr4-evk:~# modprobe virtio_mmio
root@imx8mp-lpddr4-evk:~# modprobe virtio_trans
```

b. Run the backend on Cortex-A Core

Boot up Linux:

```
u-boot => setenv fdtfile imx8mp-evk-virtio-perf-ca53.dtb; mw.l
0xfc700000 0 1
u-boot => setenv mmcargs $mmcargs clk_ignore_unused
u-boot => run bsp_bootcmd
```

After Linux boots up, start backend RTOS running on Cortex-A Core by using remoteproc under Linux:

```
root@imx8mp-lpddr4-evk:~# modprobe -r virtio_trans
root@imx8mp-lpddr4-evk:~# modprobe -r virtio_mmio
root@imx8mp-lpddr4-evk:~# echo /examples/heterogeneous-multicore/
virtio-perf-ca/ddr_release/virtio_perf_ca53.elf > /sys/devices/
platform/remoteproc-ca53-3/remoteproc/remoteproc2/firmware
root@imx8mp-lpddr4-evk:~# echo start > /sys/devices/platform/
remoteproc-ca53-3/remoteproc/remoteproc2/state
root@imx8mp-lpddr4-evk:~# modprobe virtio_mmio
root@imx8mp-lpddr4-evk:~# modprobe virtio_trans
```

3. Use the “vt_test.sh” tool in Linux to start the performance testing use case.

The following is the help information of the tool.

```
root@imx8mm-lpddr4-evk:~# vt_test.sh -h
USAGE: vt_test.sh [-h] [-s pkt_size] [-r regression] [-t type] [-b backend
copy] [-f frontend copy]
-s: Packet size: max 2048 Bytes, default: 64 Bytes
-r: Regression times: default: 1000
-t: Test type: 0: TX (frontend to backend); 1: RX (backend to frontend)
-b: Backend copy buffer option: 0: not copy; 1: copy
-f: Frontend copy buffer option: 0: not copy; 1: copy
-h: This USAGE info
```

- a. “-s” specifies the packet size to be used for testing, such as “-s 64”, it uses 64-byte packets for testing.
- b. “-t” specifies the testing direction:
 - “-t 0” means the test sends packets from frontend (Linux) to backend (RTOS on A-Core or M-Core),
 - “-t 1” means the test sends packets from backend (RTOS on A-Core or M-Core) to frontend (Linux).
- c. “-r” specifies the regression times:
 - “-r 100000” indicates that the test case sends 100000 packets with the direction from backend to frontend or frontend to backend, which is specified by the “-t” parameter.
- d. “-b” specifies whether there is a memory copy in the backend:
 - For the memory copy case, use “-b 1”, for enabling memory copy from Vring buffer to user application buffer when receiving packet from frontend, or copy from user application buffer to Vring buffer when transmitting packets to frontend.
 - For no memory copy case “-b 0”, specifies there will be no memory copy in backend for each packet receiving or transmitting.
- e. “-f” specifies whether there is a memory copy in the frontend.

For example, see the command below:

```
vt_test.sh -s 64 -r 1000000 -t 0 -b 0 -f 0
```

The above test case transmits 1000000 packets from frontend (Linux) to backend (RTOS on A-Core or M-Core), each packet size is 64 bytes, there is no memory copy both on frontend and backend. The test log is as follows:

```
root@imx8mm-lpddr4-evk:~# vt_test.sh -s 64 -r 1000000 -t 0 -b 0 -f 0
[20.561527] *****
[20.561539] Front-end: interrupt mode
[20.561543] Back-end: interrupt mode
[20.561544] Front-end: do NOT copy buffer
[20.561546] Back-end: do NOT copy buffer
[20.561547] Test case: TX
[20.561547] pkt_size: <64>
[20.561547] regress times: <1000000>
[21.868494] tx_test: pkt_size (64 B), pkt_cnt (1000000), period (1298108 µs)
```

The log shows that 1000000 packets are transmitted from frontend to backend in 1298108 µs. Therefore, the performance is 770 kpps or 394 Mbit/s.

3.6.3 Heterogeneous Multicore VirtIO network sharing

3.6.3.1 Overview

Figure 46 shows the heterogeneous Multicore VirtIO network sharing architecture.

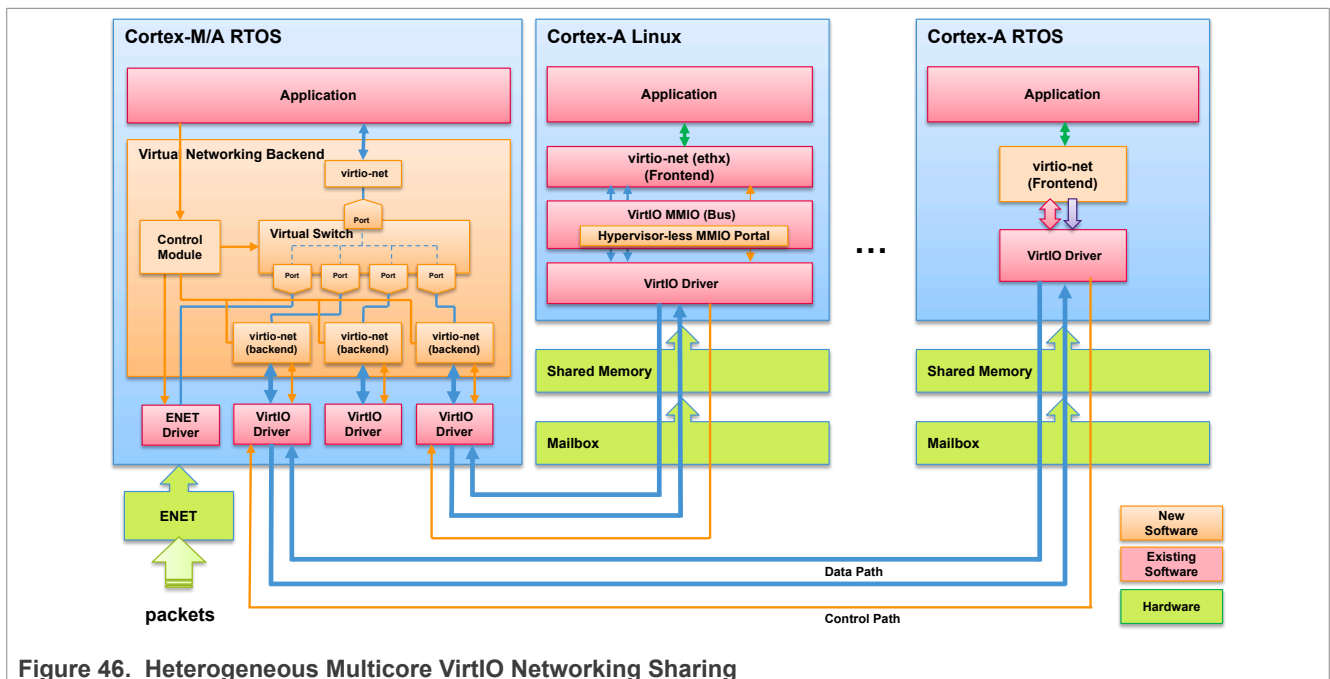


Figure 46. Heterogeneous Multicore VirtIO Networking Sharing

The virtual networking frontend runs on Cortex-A core, the frontend in Linux reuses the existing “drivers/net/virtio_net.c” driver by selecting kernel configuration item “CONFIG_VIRTIO_NET”. RTOS frontend driver is not enabled in this release.

The Virtual Networking Backend can run in RTOS on Cortex-A core or Cortex-M core. The virtio-net backend drivers use Heterogeneous Multicore VirtIO to communicate with virtio-net frontend. It includes two data paths to handle data packets receiving/transmitting of frontend and one control path to handle control requirements from the frontend.

A Virtual Switch in the backend is used to switch packets from different ports. The switch ports include one “remote port” and many “local ports”. In general, a “remote port” is a physical Ethernet port such as physical ENET port, which is used to receive/transmit packets from/to the physical Ethernet port. Here, “local port” refers to the virtual software port, such as the local port for virtio-net backend used to receive/transmit from/to virtio-net backend. In fact, the packet is from/to virtio-net frontend through data path of Heterogeneous Multicore VirtIO and another type of “local port” is used to connect to “virt-net” on RTOS to provide virtual Ethernet interface for RTOS locally.

In the current implementation, an Ethernet L2 switch acts as a virtual switch. Each “local port” has a different MAC address, so the packets received from a “remote port” can be switched to the destination “local port” according to destination MAC address in the networking packets. The packets whose destination MAC address does not match any “local port” are discarded, except broadcast packets. For the packets received from the “local port”, the switch tries to check whether destination MAC address matches the address of other “local ports”. If a matching entry is found, the packet is switched to the matched “local port”, so that the virtual switch can implement switching packets between “local ports” locally. In case a match is not found, the packets are sent to the external by “remote port”. That is to say Virtual Switch supports “local switch” and “remote switch”.

The Virtual Switch can connect to multiple “local port”, that is to say single physical Ethernet port, controlled by CPU Core running backend, can be shared with multiple OS running on different CPU Core by though multiple front-end, and local virtual Ethernet driver in back-end also provide Ethernet service for back-end CPU core locally.

3.6.3.2 Building VirtIO network sharing backend application

VirtIO network sharing backend application runs on FreeRTOS, it is an application named “virtio_net_backend” in the repo: [heterogeneous-multicore](#).

Refer to “[Building Heterogeneous Multicore RTOS Application](#)” for how to build the application.

3.6.3.3 Running VirtIO network sharing

VirtIO Network Sharing supports running on i.MX 8M Mini, i.MX 8M Plus, and i.MX 93 platforms.

1. Set up the UART console for the frontend and backend

Connect the DEBUG UART slot on the board to your PC through the USB Cable. On the PC, this step creates two USB serial ports (port0 and port1) for i.MX 8M Mini EVK board, and four USB serial ports (port0 ~ port3) for i.MX 8M Plus EVK board and i.MX 93 EVK board. Open two UART consoles for UART port0 and port1 on i.MX 8M Mini EVK board or port2 and port3 on i.MX 8M Plus EVK board and i.MX 93 EVK board UART with the following setup:

- 115200
- No parity
- 8 data bits
- 1 stop bit

The first UART console is used for Linux that runs the VirtIO frontend. The UART console is used for RTOS, which runs the VirtIO backend.

2. Hardware setup

Connect the ENET port on the selected EVK board to a networking switch or another board by using an Ethernet cable.

For i.MX 8M Mini EVK board, there is a single Ethernet port on the board. So, use this port for testing.

For i.MX 8M Plus EVK board and i.MX 93 EVK board, there are two Ethernet ports are on the board. The first one, which is close to DEBUG USB port is ENET port and it is used for VirtIO Networking Sharing. So connect this port to the networking link. Another Ethernet port is an ENET QoS port and it is not used for this demo.

3. Booting backend and frontend

Heterogeneous Multicore VirtIO backend can run on Cortex-A core or Cortex-M core.

- **For i.MX 8M Mini LPDDR4 EVK**

Can use U-Boot commands or Linux remoteproc to start backend RTOS, choose one method to start RTOS and Linux:

- **Use U-Boot Commands to start RTOS, then start Linux**

- a. Run the backend on a Cortex-M core

On U-Boot command prompt, execute the following commands to boot Cortex-M core with firmware:

```
u-boot => ext4load mmc 1:2 0x48000000 /examples/heterogeneous-
multicore/virtio-net-backend-cm/release/virtio_net_backend_cm4.bin
u-boot => cp.b 0x48000000 0x7e0000 0x20000
u-boot => bootaux 0x7e0000
```

Then boot Linux kernel:

```
u-boot => setenv fdtfile imx8mm-evk-virtio-net-cm4.dtb
u-boot => setenv mmcargs $mmcargs mem=2048MB clk_ignore_unused
u-boot => run bsp_bootcmd
```

- b. Run the backend on a Cortex-A core

Execute the following command in the U-Boot command line:

```
u-boot => ext4load mmc 1:2 0x93c00000 /examples/heterogeneous-
multicore/virtio-net-backend-ca/ddr_release/virtio_net_backend_ca53.bin
u-boot => dcache flush; icache flush;
u-boot => cpu 3 release 0x93c00000
```

Then, boot the Linux kernel:

```
u-boot => setenv fdtfile imx8mm-evk-virtio-net-ca53.dtb
u-boot => setenv mmcargs $mmcargs maxcpus=3 clk_ignore_unused
u-boot => run bsp_bootcmd
```

- **Or start Linux, then use remoteproc to start RTOS**

- a. Run the backend on Cortex-M core

Boot up Linux:

```
u-boot => run prepare_mcore; mw.l 0xb8400000 0 1
u-boot => setenv fdtfile imx8mm-evk-virtio-net-cm4.dtb
u-boot => setenv mmcargs $mmcargs mem=2048MB clk_ignore_unused
u-boot => run bsp_bootcmd
```

After Linux boot up, start backend RTOS running on Cortex-M Core by using remoteproc under Linux:

```
root@imx8mm-lpddr4-evk:/# modprobe -r virtio_net
root@imx8mm-lpddr4-evk:/# modprobe -r virtio_mmio
root@imx8mm-lpddr4-evk:/# echo /examples/heterogeneous-multicore/
virtio-net-backend-cm/release/virtio_net_backend_cm4.elf > /sys/
devices/platform/imx8mm-cm4/remoteproc/remoteproc0/firmware
root@imx8mm-lpddr4-evk:/# echo start > /sys/devices/platform/imx8mm-
cm4/remoteproc/remoteproc0/state
root@imx8mm-lpddr4-evk:/# modprobe virtio_mmio
root@imx8mm-lpddr4-evk:/# modprobe virtio_net
```

- b. Run the backend on Cortex-A core

Boot up Linux:

```
u-boot => setenv fdtfile imx8mm-evk-virtio-net-ca53.dtb; mw.l
0xb8400000 0 1
u-boot => setenv mmcargs $mmcargs maxcpus=3 clk_ignore_unused
```

```
u-boot => run bsp_bootcmd
```

After Linux boot up, start backend RTOS running on Cortex-A Core by using remoteproc under Linux:

```
root@imx8mm-lpddr4-evk:/# modprobe -r virtio_net
root@imx8mm-lpddr4-evk:/# modprobe -r virtio_mmio
root@imx8mm-lpddr4-evk:/# echo /examples/heterogeneous-multicore/
virtio-net-backend-ca/ddr_release/virtio_net_backend_ca53.elf > /sys/
devices/platform/remoteproc-ca53-3/remoteproc/remoteproc2/firmware
root@imx8mm-lpddr4-evk:/# echo start >/sys/devices/platform/remoteproc-
ca53-3/remoteproc/remoteproc2/state
root@imx8mm-lpddr4-evk:/# modprobe virtio_mmio
root@imx8mm-lpddr4-evk:/# modprobe virtio_net
```

• For i.MX 8M Plus LPDDR4 EVK

Can use U-Boot commands or Linux remoteproc to start backend RTOS, choose one method to start RTOS and Linux:

– Using U-Boot Commands to start RTOS, then start Linux

- a. Run the backend on Cortex-M core

U-Boot command prompt, execute the following commands to boot Cortex-M core with firmware:

```
u-boot => ext4load mmc 1:2 0x48000000 /examples/heterogeneous-
multicore/virtio-net-backend-cm/release/virtio_net_backend_cm7.bin
u-boot => cp.b 0x48000000 0x7e0000 0x20000
u-boot => bootaux 0x7e0000
```

Then, boot Linux kernel:

```
u-boot => setenv fdtfile imx8mp-evk-virtio-net-cm7.dtb
u-boot => setenv mmcargs $mmcargs mem=2048MB clk_ignore_unused
u-boot => run bsp_bootcmd
```

- b. Run the backend on Cortex-A core

Execute the following command in U-Boot command line:

```
u-boot => ext4load mmc 1:2 0xC0000000 /examples/heterogeneous-
multicore/virtio-net-backend-ca/ddr_release/virtio_net_backend_ca53.bin
u-boot => dcache flush; icache flush;
u-boot => cpu 3 release 0xC0000000
```

Then boot Linux kernel:

```
u-boot => setenv fdtfile imx8mp-evk-virtio-net-ca53.dtb
u-boot => setenv mmcargs $mmcargs maxcpus=3 clk_ignore_unused
u-boot => run bsp_bootcmd
```

– Or start Linux, then use remoteproc to start RTOS

- a. Run the backend on Cortex-M core

Boot up Linux:

```
u-boot => run prepare_mcore; mw.l 0xb8400000 0 1
u-boot => setenv fdtfile imx8mp-evk-virtio-net-cm7.dtb
u-boot => setenv mmcargs $mmcargs mem=2048MB clk_ignore_unused
u-boot => run bsp_bootcmd
```

After Linux boot up, start backend RTOS running on Cortex-M Core by using remoteproc under Linux:

```
root@imx8mp-lpddr4-evk:/# modprobe -r virtio_net
root@imx8mp-lpddr4-evk:/# modprobe -r virtio_mmio
```

```

root@imx8mp-lpddr4-evk:/# echo /examples/heterogeneous-multicore/
virtio-net-backend-cm/release/virtio_net_backend_cm7.elf > /sys/
devices/platform/imx8mp-cm7/remoteproc/remoteproc0/firmware
root@imx8mp-lpddr4-evk:/# echo start > /sys/devices/platform/imx8mp-
cm7/remoteproc/remoteproc0/state
root@imx8mp-lpddr4-evk:/# modprobe virtio_mmio
root@imx8mp-lpddr4-evk:/# modprobe virtio_net

```

b. Run the backend on Cortex-A core

Boot up Linux:

```

u-boot => setenv fdtfile imx8mp-evk-virtio-net-ca53.dtb; mw.l
0xfc700000 0 1
u-boot => setenv mmcargs $mmcargs maxcpus=3 clk_ignore_unused
u-boot => run bsp_bootcmd

```

After Linux boot up, start backend RTOS running on Cortex-A Core by using remoteproc under Linux:

```

root@imx8mp-lpddr4-evk:/# modprobe -r virtio_net
root@imx8mp-lpddr4-evk:/# modprobe -r virtio_mmio
root@imx8mp-lpddr4-evk:/# echo /examples/heterogeneous-multicore/
virtio-net-backend-ca/ddr_release/virtio_net_backend_ca53.elf > /sys/
devices/platform/remoteproc-ca53-3/remoteproc/remoteproc2/firmware
root@imx8mp-lpddr4-evk:/# echo start > /sys/devices/platform/remoteproc-
ca53-3/remoteproc/remoteproc2/state
root@imx8mp-lpddr4-evk:/# modprobe virtio_mmio
root@imx8mp-lpddr4-evk:/# modprobe virtio_net

```

• for i.MX 93 EVK

Can use U-Boot commands or Linux remoteproc to start backend RTOS, choose one method to start RTOS and Linux:

– Using U-Boot Commands to start RTOS, then start Linux

a. Run the backend on Cortex-M core

On U-Boot command prompt, execute the following commands to boot Cortex-M core with firmware:

```

u-boot => ext4load mmc 1:2 0xd0000000 /examples/heterogeneous-
multicore/virtio-net-backend-cm/release/virtio_net_backend_cm33.bin
u-boot => cp.b 0xd0000000 0x201e0000 20000
u-boot => bootaux 0x1ffe0000

```

Then boot Linux kernel:

```

# For i.MX 93 EVK
u-boot => setenv fdtfile imx93-11x11-evk-virtio-net-cm33.dtb
# For i.MX 93 14x14 EVK
u-boot => setenv fdtfile imx93-14x14-evk-virtio-net-cm33.dtb

u-boot => setenv mmcargs $mmcargs clk_ignore_unused
u-boot => run bsp_bootcmd

```

b. Run the backend on Cortex-A core

Execute the following command in the U-Boot command line:

```

u-boot => ext4load mmc 1:2 0xd0000000 /examples/heterogeneous-
multicore/virtio-net-backend-ca/ddr_release/virtio_net_backend_ca55.bin
u-boot => dcache flush && icache flush
u-boot => cpu 1 release 0xd0000000

```

Then boot Linux kernel:

```

# For i.MX 93 EVK

```



```

u-boot => setenv fdtfile imx93-11x11-evk-virtio-net-ca55.dtb
# For i.MX 93 14x14 EVK
u-boot => setenv fdtfile imx93-14x14-evk-virtio-net-ca55.dtb

u-boot => setenv mmcargs $mmcargs maxcpus=1 clk_ignore_unused
u-boot => run bsp_bootcmd

```

– Or start Linux, then use remoteproc to start RTOS

a. Run the backend on Cortex-M core

Boot up Linux:

```

# For i.MX 93 EVK
u-boot => setenv fdtfile imx93-11x11-evk-virtio-net-cm33.dtb
# For i.MX 93 14x14 EVK
u-boot => setenv fdtfile imx93-14x14-evk-virtio-net-cm33.dtb

u-boot => run prepare_mcore; mw.l 0xa8400000 0 1
u-boot => setenv mmcargs $mmcargs clk_ignore_unused
u-boot => run bsp_bootcmd

```

After Linux boot up, start backend RTOS running on Cortex-M Core by using remoteproc under Linux:

```

root@imx93evk:/sys# modprobe -r virtio_net
root@imx93evk:/sys# modprobe -r virtio_mmio
root@imx93evk:/sys# echo /examples/heterogeneous-multicore/virtio-
net-backend-cm/release/virtio_net_backend_cm33.elf > /sys/devices/
platform/remoteproc-cm33/remoteproc/remoteproc0/firmware
root@imx93evk:/sys# echo start > /sys/devices/platform/remoteproc-cm33/
remoteproc/remoteproc0/state
root@imx93evk:/sys# modprobe virtio_mmio
root@imx93evk:/sys# modprobe virtio_net

```

b. Run the backend on Cortex-A core

Boot up Linux:

```

# For i.MX 93 EVK
u-boot => setenv fdtfile imx93-11x11-evk-virtio-net-ca55.dtb
# For i.MX 93 14x14 EVK
u-boot => setenv fdtfile imx93-14x14-evk-virtio-net-ca55.dtb

u-boot => setenv mmcargs $mmcargs maxcpus=1 clk_ignore_unused; mw.l
0xfc700000 0 1
u-boot => run bsp_bootcmd

```

After Linux boot up, start backend RTOS running on Cortex-A Core by using remoteproc under Linux:

```

root@imx93evk:/sys# modprobe -r virtio_net
root@imx93evk:/sys# modprobe -r virtio_mmio
root@imx93evk:/sys# echo /examples/heterogeneous-multicore/virtio-net-
backend-ca/ddr_release/virtio_net_backend_ca55.elf > /sys/devices/
platform/remoteproc-ca55-1/remoteproc/remoteproc0/firmware
root@imx93evk:/sys# echo start > /sys/devices/platform/remoteproc-
ca55-1/remoteproc/remoteproc0/state
root@imx93evk:/sys# modprobe virtio_mmio
root@imx93evk:/sys# modprobe virtio_net

```

4. Evaluate Networking Sharing

After backend starts, the second UART console displays the following backend log:

```
Starting Virtio networking backend...
```

```
virtio network device initialization succeed!
Switch enabled with enet remote port succeed!
ENET: PHY link is up with speed 1000M full-duplex
```

After the kernel boots up, use “ifconfig” and “ping” commands to verify the virtual networking interface. In the following log, “eth0” is virtio_net interface, but it may be different on different platform. So, use “ethtool” to find out the virtio_net interface that is using “virtio_net” driver, and the default MAC address of virtio_net interface is “00:04:9f:00:01:02”.

```
root@imx8mm-lpddr4-evk:~# ifconfig
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
    inet 192.168.1.107 netmask 255.255.255.0 broadcast 192.168.1.255
    inet6 fd08:d7d5:e652::733 prefixlen 128 scopeid 0x0<global>
    inet6 fd08:d7d5:e652:0:201:2ff:fe03:405 prefixlen 64 scopeid
0x0<global>
    inet6 fe80::201:2ff:fe03:405 prefixlen 64 scopeid 0x20<link>
    ether 00:04:9f:00:01:02 txqueuelen 1000 (Ethernet)
    RX packets 54 bytes 5544 (5.4 KiB)
    RX errors 0 dropped 0 overruns 0 frame 0
    TX packets 143 bytes 20887 (20.3 KiB)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

root@imx8mm-lpddr4-evk:~# ethtool -i eth0
driver: virtio_net
version: 1.0.0
firmware-version:
expansion-rom-version:
bus-info: b8400000.virtio_net
supports-statistics: yes
supports-test: no
supports-EEPROM-access: no
supports-register-dump: no
supports-priv-flags: no

root@imx8mm-lpddr4-evk:~# ping 192.168.1.1
PING 192.168.1.1 (192.168.1.1) 56(84) bytes of data.
64 bytes from 192.168.1.1: icmp_seq=1 ttl=64 time=0.888 ms
64 bytes from 192.168.1.1: icmp_seq=2 ttl=64 time=0.541 ms
64 bytes from 192.168.1.1: icmp_seq=3 ttl=64 time=2.13 ms
64 bytes from 192.168.1.1: icmp_seq=4 ttl=64 time=2.29 ms
64 bytes from 192.168.1.1: icmp_seq=5 ttl=64 time=1.73 ms
```

Use the following command to change the MAC address of virtio_net:

```
root@imx8mm-lpddr4-evk:~# ifconfig eth0 hw ether 00:04:9f:00:01:03
```

Attention: While running M-core backend on i.MX 93 EVK and i.MX 93 14x14 EVK platforms, changing the MAC address causes the M-core to be stuck.

3.7 Unified Life Cycle Management

3.7.1 Overview

Heterogeneous Multicore Framework provides Unified Life Cycle Management for both the Cortex-A and the Cortex-M cores.

Real-time Edge supports bootstrapping the **native Zephyr** and **native FreeRTOS** on the **Cortex-A core** and **Cortex-M core** with **U-Boot** command and with the **RemoteProc** under Linux as listed in the below table.

Table 40. Bootstrapping options for Cortex-A and CortexM cores

Core Type	U-Boot	RemoteProc on Linux
NativeZephyr on M core	Y	Y
NativeFreeRTOS on M core	Y	Y
NativeZephyr on A core	Y	Y
NativeFreeRTOS on A core	Y	Y

3.7.2 Using U-Boot Commands to Manage Life Cycle of RTOS on Cortex-A Core

Multiple U-Boot commands can be used to start or stop the Cortex-A Core RTOS.

- "cpu" command

The following is the help information of the U-Boot `cpu` command:

```
u-boot=> cpu
cpu - Multiprocessor CPU boot manipulation and release

Usage:
cpu <num> reset           - Reset cpu <num>
cpu status                - Status of all cpus
cpu <num> status          - Status of cpu <num>
cpu <num> disable         - Disable cpu <num>
cpu <num> release <addr> [args] - Release cpu <num> at <addr> with [args]
```

Except `cpu <num> reset` command, all other `cpu` commands have been implemented on i.MX 8M Mini, i.MX 8M Plus, i.MX 93, and i.MX 95 platforms.

Note: `cpu <num> disable` is not implemented on i.MX 95.

In general, U-Boot runs on the master Core (Core0) of Cortex-A Cores. Therefore, users can

- Use the command `cpu <num> release <addr> [args]` to start RTOS on any other slave Cortex-A Cores (Cores except Core0) from a specified memory address. For this, you must load the RTOS binary images into the corresponding memory space.
 - Use the `cpu <num> disable` command to power off any slave Core which runs RTOS.
 - Use the `cpu status` or the `cpu <num> status` commands to check the status (running or power off) for all or a specified Cortex-A Core.
- `go` command

Following is the help information of the U-Boot `go` command:

```
u-boot=> go
go - start application at address 'addr'

Usage:
go addr [arg ...]
    - start application at address 'addr'
      passing 'arg' as arguments
```

The `go` command can be used to start RTOS on the master Core (Core0) but it does not return to the U-Boot command line.

In summary, use the U-Boot `go` command to boot the RTOS from Core0. Use the U-Boot command `cpu` to boot or to power off the RTOS running on other Cortex-A Cores except the Core0.

The below example shows how to run the `hello_world` example on the following boards:

- i.MX 8M Mini EVK
- i.MX 8M Plus EVK
- i.MX 91 EVK and i.MX 91 QSB

- i.MX 93 11x11 EVK and i.MX 93 14x14 EVK
- i.MX 95 15x15 EVK and i.MX 95 19x19 EVK

1. Setup UART console for Native RTOS

Connect DEBUG UART slot on the board to your PC through the USB Cable. This step creates some USB serial ports on the PC:

- i.MX 8M Mini EVK (port 0~1): port1 is for U-Boot/Linux or RTOS; port0 for RTOS.
- i.MX 8M Plus EVK (port 0~3): port2 is for U-Boot/Linux or RTOS; port3 for RTOS.
- i.MX 91 EVK and QSB (port 0~3): port2 is for U-Boot/Linux or RTOS.
- i.MX 93 11x11 and 14x14 EVK (port 0~3): port2 is for U-Boot/Linux or RTOS; port3 for RTOS.
- i.MX95 15x15 and 19x19 EVK (port 0~3): port2 is for U-Boot/Linux or RTOS; port0 for RTOS.

Open the UART consoles needed with the following setup:

- 115200
- No parity
- 8 data bits
- 1 stop bit

2. Booting Native FreeRTOS Image

After powering up the board and entering U-Boot command line, execute the following U-Boot commands to run the `hello_world` example.

- For i.MX 8M Mini EVK

Using the commands below to start RTOS from Core3 (the first Core is Core0):

```
=> ext4load mmc 1:2 0x93C00000 /examples/heterogeneous-multicore/hello-world-
ca/ddr_release/hello_world_ca53_RTOS0_UART4.bin
=> dcache flush; icache flush
=> cpu 3 release 0x93C00000
```

After the preceding steps are followed, the UART4 console displays the following RTOS log:

```
Cortex-A53: RTOS0: Hello world! Real-time Edge on MIMX8MM-EVK
RTOS1: RAM console@0x95bff000
RTOS2: RAM console@0x96bff000
RTOS3: RAM console@0x97bff000
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
```

Use the following command to power off the RTOS running on Core3 (the first Core is Core0):

```
u-boot=> cpu 3 disable
```

- For i.MX 8M Plus EVK

Using the commands below to start RTOS from Core3 (the first Core is Core0):

```
=> ext4load mmc 1:2 0xC0000000 /examples/heterogeneous-multicore/hello-world-
ca/ddr_release/hello_world_ca53_RTOS0_UART4.bin
=> dcache flush; icache flush;
=> cpu 3 release 0xC0000000
```

After the preceding steps are followed, the UART4 console displays the following RTOS log:

```
Cortex-A53: RTOS0: Hello world! Real-time Edge on MIMX8MP-EVK
RTOS1: RAM console@0xc1fff000
RTOS2: RAM console@0xc2fff000
RTOS3: RAM console@0xc3fff000
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac
```

Use the following command to power off the RTOS running on Core3 (the first Core is Core0):

```
u-boot=> cpu 3 disable
```

- For i.MX 91 EVK and QSB board:
Using the commands below to start RTOS:

```
=> ext4load mmc 1:2 0x80000000 /examples/heterogeneous-multicore/hello-world-
ca/ddr_release/hello_world_ca55_RTOS0_UART1.bin
=> dcache flush; icalache flush;
=> go 0x80000000
```

After the preceding steps are followed, the UART1 console displays the following RTOS log:

```
Cortex-A55: RTOS0: Hello world! Real-time Edge on MIMX91-EVK
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
```

- For i.MX 93 11x11 and 14x14 EVK board:
Using the commands below to start RTOS from Core1 (the first Core is Core0):

```
=> ext4load mmc 1:2 0xD0000000 /examples/heterogeneous-multicore/hello-world-
ca/ddr_release/hello_world_ca55_RTOS0_UART2.bin
=> dcache flush; icalache flush;
=> cpu 1 release 0xD0000000
```

After the preceding steps are followed, the UART2 console displays the following RTOS log:

```
Cortex-A55: RTOS0: Hello world! Real-time Edge on MIMX93-EVK
RTOS1: RAM console@0xd1fff000
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
```

Use the following command to poweroff the RTOS running on Core1 (the first Core is Core0):

```
u-boot=> cpu 1 disable
```

- For i.MX95 15x15 and 19x19 EVK board:
Using the commands below to start RTOS from Core5 (the first Core is Core0):

```
=> ext4load mmc 1:2 0xD0000000 /examples/heterogeneous-multicore/hello-world-
ca/ddr_release/hello_world_ca55_RTOS0_UART3.bin
=> dcache flush; icalache flush;
=> cpu 5 release 0xD0000000
```

After the preceding steps are followed, the UART3 console displays the following RTOS log:

```
Cortex-A55: RTOS0: Hello world! Real-time Edge on MIMX95-EVK
RTOS1: RAM console@0xd1fff000
RTOS2: RAM console@0xd2fff000
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
```

- go command example
The Native RTOS image can also be booted from the first Cortex-A Core, which is called master Core. Use the same command but use the "go" command to replace "cpu" command. For example, use the command below to boot hello_world example on the first Cortex-A Core on i.MX 8M Mini EVK board:

```
=> ext4load mmc 1:2 0x93C00000 /examples/heterogeneous-multicore/hello-world-
ca/ddr_release/hello_world_ca53_RTOS0_UART4.bin
=> dcache flush; icalache flush;
=> go 0x93C00000
```

After the preceding steps are followed, the UART4 console displays the following RTOS log:

```
Cortex-A53: RTOS0: Hello world! Real-time Edge on MIMX8MM-EVK
RTOS1: RAM console@0x95bfff000
RTOS2: RAM console@0x96bfff000
RTOS3: RAM console@0x97bfff000
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac
tic
tic tac tic tac tic tac tic tac tic tac tic tac tic tac
```

3. Booting Native Zephyr Image

After powering up the board and entering U-Boot command line, execute the following U-Boot commands to run the `hello_world` example.

- For i.MX 8M Mini EVK

Using below commands to boot `hello_world` example on the Cortex-A Core3 (the first Core is Core0):

```
=> ext4load mmc 1:2 0x93C00000 /examples/heterogeneous-multicore/hello-world-
ca-zephyr/hello_world_ca53_RTOS0_UART4.bin
=> dcache flush; icache flush
=> cpu 3 release 0x93C00000
```

After the preceding steps are followed, the UART4 console displays the following RTOS log:

```
*** Booting Zephyr OS build zephyr-v3.5.0-27-g24a3599ab7a0 ***
Cortex-A53: RTOS0: Hello World! Real-time Edge on mimx8mm_evk_a53
RTOS1: RAM console@0x94d00000
RTOS2: RAM console@0x95d00000
RTOS3: RAM console@0x96d00000
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac
tic
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac
```

Use the following command to power off the RTOS running on Core3 (the first Core is Core0):

```
u-boot=> cpu 3 disable
```

- For i.MX 8M Plus EVK

Using below commands to boot `hello_world` example on the Cortex-A Core3 (the first Core is Core0):

```
=> ext4load mmc 1:2 0xC0000000 /examples/heterogeneous-multicore/hello-world-
ca-zephyr/hello_world_ca53_RTOS0_UART4.bin
=> dcache flush; icache flush;
=> cpu 3 release 0xC0000000
```

After the preceding steps are followed, the UART4 console displays the following RTOS log:

```
*** Booting Zephyr OS build zephyr-v3.5.0-27-g24a3599ab7a0 ***
Cortex-A53: RTOS0: Hello World! Real-time Edge on mimx8mp_evk_a53
RTOS1: RAM console@0xc1100000
RTOS2: RAM console@0xc2100000
RTOS3: RAM console@0xc3100000
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac
tic
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac
```

Use the following command to poweroff the RTOS running on Core3 (the first Core is Core0):

```
u-boot=> cpu 3 disable
```

- For i.MX 91 EVK and QSB board:

Using the commands below to start RTOS:

```
=> ext4load mmc 1:2 0x80000000 /examples/heterogeneous-multicore/hello-world-
ca-zephyr/hello_world_ca55_RTOS0_UART1.bin
=> dcache flush; icache flush;
=> go 0x80000000
```

After the preceding steps are followed, the UART1 console displays the following RTOS log:

```
*** Booting Zephyr OS build zephyr-v3.5.0-27-g24a3599ab7a0 ***
Cortex-A55: RTOS0: Hello World! Real-time Edge on mimx91_evk_a55
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac
tic
tic tac tic tac tic tac tic tac tic tac tic tac tic tac
```

- For i.MX 93 11x11 and 14x14 EVK board:

Using the commands below to start RTOS from Core1 (the first Core is Core0):

```
=> ext4load mmc 1:2 0xD0000000 /examples/heterogeneous-multicore/hello-world-
ca-zephyr/hello_world_ca55_RTOS0_UART2.bin
=> dcache flush; icache flush;
=> cpu 1 release 0xD0000000
```

After the preceding steps are followed, the UART2 console displays the following RTOS log:

```
*** Booting Zephyr OS build zephyr-v3.5.0-27-g24a3599ab7a0 ***
Cortex-A55: RTOS0: Hello World! Real-time Edge on mimx93_evk_a55
RTOS1: RAM console@0xd1100000
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac
tic
tic tac tic tac tic tac tic tac tic tac tic tac tic tac
```

Use the following command to poweroff the RTOS running on Core1 (the first Core is Core0):

```
u-boot=> cpu 1 disable
```

- For i.MX95 15x15 and 19x19 EVK board:

Using the commands below to start RTOS from Core5 (the first Core is Core0):

```
=> ext4load mmc 1:2 0xD0000000 /examples/heterogeneous-multicore/hello-world-
ca-zephyr/hello_world_ca55_RTOS0_UART3.bin
=> dcache flush; icache flush;
=> cpu 5 release 0xD0000000
```

After the preceding steps are followed, the UART3 console displays the following RTOS log:

```
*** Booting Zephyr OS build v3.7.0-37-g757815829721 ***
Cortex-A55: RTOS0: Hello World! Real-time Edge on imx95_evk
RTOS1: RAM console@0xd1100000
RTOS2: RAM console@0xd2100000
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac
tic
```

- go command example

The Native RTOS image can also be booted from the first Cortex-A Core, which is called master Core. Use the same command but use the "go" command to replace "cpu" command. For example, use the command below to boot `hello_world` example on the first Cortex-A Core on i.MX 8M Mini EVK board:

```
=> ext4load mmc 1:2 0x93C00000 /examples/heterogeneous-multicore/hello-world-
ca-zephyr/hello_world_ca53_RTOS0_UART4.bin
=> dcache flush; icache flush
=> go 0x93C00000
```

After the preceding steps are followed, the UART4 console displays the following RTOS log:

```
*** Booting Zephyr OS build zephyr-v3.5.0-27-g24a3599ab7a0 ***
Cortex-A53: RTOS0: Hello World! Real-time Edge on mimx8mm_evk_a53
RTOS1: RAM console@0x94d00000
RTOS2: RAM console@0x95d00000
RTOS3: RAM console@0x96d00000
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac
tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac
tac
```

3.7.3 Using RemoteProc under Linux to Manage Life Cycle of RTOS on Cortex-A Core

Apart from the U-Boot commands in U-Boot, another Cortex-A Core RTOS (native FreeRTOS and native Zephyr) management tool is remoteproc under Linux. Using remoteproc can dynamically start or stop native FreeRTOS or native Zephyr under Linux. This Cortex-A Core RTOS life cycle management method has limitations. The limitation is that Linux must run on at least a single Cortex-A Core to manage RTOS running on all the other Cortex-A Cores.

Remoteproc (Remote Processor Framework) under Linux is used to control remote processors on which different instances of operating systems are running. For example, it can be used to power on, load firmware, or to power off the remote processors. For SMP Linux, which it runs on multiple Cortex-A Cores, each Cortex-A Core can be used as a remote processor, and can use remoteproc to bring up or bring down another RTOS instance on it. To run another RTOS on it, remoteproc first removes this Cortex-A Core from SMP Linux by using CPU hotplug, and brings it to be down. Then, remoteproc brings it up again with a new RTOS instance to run it. Remoteproc can also bring this CPU Core running RTOS to be down and then plug it back to SMP Linux by using CPU hotplug.

• Linux Device Tree Configuration

To use remoteproc to manage life cycle on Cortex-A Core, you must add device nodes in Linux dts. For example, the following dts nodes are defined in `imx8m-rproc-ca53.dtsi` which defines four remoteproc instances. In each instance `fsl,cpus-bits` defines CPU Core bitmask to specify which CPU Core(s) can be managed by this instance. The `memory-region` is used to specify the reserved memory space used by RTOS to be run on this instance. So, the first instance in the following example runs on Core1, the second one runs on Core2, the third one runs on Core3. The fourth one runs on Core2 and Core3 and it can be used to run SMP RTOS. However, since the CPU Core used by the fourth one is duplicated with the second one and the third one, so we cannot run RTOS simultaneously on the second one or third one with the fourth one.

```
ca53_1: remoteproc-ca53-1 {
    compatible = "fsl,imx-rproc-psci";
    /* bitmask:0b0010, assign A53 Core 1 */
    fsl,cpus-bits = <0x2>;
    memory-region = <&rtos_ca53_reserved>;
};

ca53_2: remoteproc-ca53-2 {
    compatible = "fsl,imx-rproc-psci";
    /* bitmask:0b0100, assign A53 Core 2 */
    fsl,cpus-bits = <0x4>;
    memory-region = <&rtos_ca53_reserved>;
};

ca53_3: remoteproc-ca53-3 {
    compatible = "fsl,imx-rproc-psci";
    /* bitmask:0b1000, assign A53 Core 3 */
    fsl,cpus-bits = <0x8>;
};
```



```

        memory-region = <&rtos_ca53_reserved>;
    };

    ca53_2_3: remoteproc-ca53-2-3 {
        compatible = "fsl,imx-rproc-psci";
        /* bitmask:0b1100, assign A53 Core 2 and Core 3 */
        fsl,cpus-bits = <0xc>;
        memory-region = <&rtos_ca53_reserved>;
    };

```

• Hardware Resource Allocation between RTOS and Linux

To run a flexible AMP system with multiple RTOSes and Linux running together on the single SoC simultaneously, the hardware resources must be allocated to different OSes carefully. This allocation is necessary to avoid conflicts and the hardware resources include memory and peripherals. Refer to `imx8mp-evk-multicore-rtos.dts` in Real-time Edge Linux as an example. The memory used by RTOS must be added to the `reserved-memory` in Linux dts. In addition, the Linux dts node 'UART4' must be disabled as it is used as the RTOS debug console.

• Using remoteproc to start or stop RTOS on Cortex-A Core

When running multiple OSes on different Cortex-A Cores on one platform, you must avoid GIC to be re-configured to crash the OS that has already been started. For Zephyr, you need to enable "`CONFIG_GIC_SAFE_CONFIG=y`". For the Linux kernel, you must enable "`CONFIG_GIC_GENTLE_CONFIG=y`", and for FreeRTOS in Real-time Edge, a similar feature has already enabled by default.

Take Heterogeneous Multicore FreeRTOS application `hello_world` as example, follow the following steps to start or stop the RTOS by using remoteproc.

1. Setup UART console for Native RTOS

Refer to the steps described in **Setup UART console** in the Section [Using U-Boot Commands to Manage Life Cycle of RTOS on Cortex-A Core](#).

2. Starting the Linux kernel

To run RTOS, a device tree that enables remoteproc and has compatible resource allocation between RTOS and Linux must be used.

To do this, when U-Boot is executing, stop at the U-Boot prompt with a terminal emulator connected to the serial port and execute the following commands (based on the board and the application):

• For i.MX 8M Mini EVK

```

=> setenv fdtfile imx8mm-evk-multicore-rtos.dtb
=> setenv mmcargs $mmcargs clk_ignore_unused
=> run bsp_bootcmd

```

• For i.MX 8M Plus EVK

```

=> setenv fdtfile imx8mp-evk-multicore-rtos.dtb
=> setenv mmcargs $mmcargs clk_ignore_unused
=> run bsp_bootcmd

```

• For i.MX 93 11x11 and 14x14 EVK

```

# For i.MX 93 11x11 EVK
=> setenv fdtfile imx93-11x11-evk-multicore-rtos.dtb
# For i.MX 93 14x14 EVK
=> setenv fdtfile imx93-14x14-evk-multicore-rtos.dtb
=> setenv mmcargs $mmcargs clk_ignore_unused
=> run bsp_bootcmd

```

- For i.MX 95 15x15 19x19 EVK

```
# For i.MX 95 15x15 EVK
=> setenv fdtfile imx95-15x15-evk-multicore-rtos.dtb
# For i.MX 95 19x19 EVK
=> setenv fdtfile imx95-19x19-evk-multicore-rtos.dtb
=> setenv mmcargs $mmcargs clk_ignore_unused
=> run bsp_bootcmd
```

3. Using remoteproc

When Linux boots up, log in Linux.

The `hello_world` application is located at the following directory of the root file system:

```
/examples/heterogeneous-multicore/hello-world-ca/ddr_release #FreeRTOS Images
/examples/heterogeneous-multicore/hello-world-ca-zephyr #Zephyr Images
```

The remoteproc sys portal is available in the following directories:

- For i.MX 8M Mini EVK and i.MX 8M Plus EVK:

```
/sys/devices/platform/remoteproc-ca53-1/remoteproc/remoteproc0
/sys/devices/platform/remoteproc-ca53-2/remoteproc/remoteproc1
/sys/devices/platform/remoteproc-ca53-3/remoteproc/remoteproc2
/sys/devices/platform/remoteproc-ca53-2-3/remoteproc/remoteproc3
```

These four remoteproc portals correspond to the remoteproc instances defined in the Linux device tree (`imx8m-rproc-ca53.dtsi`).

- For i.MX 93 11x11 and 14x14 EVK:

```
/sys/devices/platform/remoteproc-ca55-1/remoteproc/remoteproc0
```

This remoteproc portal corresponds to the remoteproc instance defined in the Linux device tree (`imx93-rproc-ca55.dtsi`).

- For i.MX 95 (15x15 and 19x19) EVK:

```
/sys/devices/platform/remoteproc-ca55-1/remoteproc/remoteproc0
/sys/devices/platform/remoteproc-ca55-2/remoteproc/remoteproc1
/sys/devices/platform/remoteproc-ca55-3/remoteproc/remoteproc2
/sys/devices/platform/remoteproc-ca55-4/remoteproc/remoteproc3
/sys/devices/platform/remoteproc-ca55-5/remoteproc/remoteproc4
/sys/devices/platform/remoteproc-ca55-4-5/remoteproc/remoteproc5
```

This remoteproc portal is the corresponding remoteproc instance defined in the Linux device tree (`imx95-rproc-ca55.dtsi`).

a. Start RTOS

To start RTOS on a specified CPU Core, first you must specify the elf RTOS image by echoing the image path and name. If the RTOS image is located in the `/lib/firmware/` directory, only specify the image name without the image path to the corresponding remoteproc instance's `firmware` portal. Then, echo `start` to the corresponding remoteproc instance's `state` portal.

For example, use the following commands to start FreeRTOS `hello_world` application on Cortex-A Core3 (the first Core is Core0) on i.MX 8M Mini EVK and i.MX 8M Plus EVK boards:

```
root@imx8mp-lpddr4-evk:~# echo /examples/heterogeneous-multicore/hello-world-ca/ddr_release/
hello_world_ca53_RTOS0_UART4.elf > /sys/devices/platform/remoteproc-ca53-3/remoteproc/
remoteproc2/firmware
root@imx8mp-lpddr4-evk:~# echo start > /sys/devices/platform/remoteproc-ca53-3/remoteproc/
remoteproc2/state
```

Then, the following RTOS log displays on the RTOS UART console on i.MX 8M Plus EVK. (A similar log is obtained for the i.MX 8M Mini EVK board):

```
Cortex-A53: RTOS0: Hello world! Real-time Edge on MIMX8MP-EVK
RTOS1: RAM console@0xc1fff000
```

```
RTOS2: RAM console@0xc2fff000
RTOS3: RAM console@0xc3fff000
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac
tic tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac
tic tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac
```

Or Use the following commands to start FreeRTOS `hello_world` application on Cortex-A Core1 (first Core is Core0) on i.MX 93 11x11 and 14x14 EVK:

```
root@imx93evk:~# echo /examples/heterogeneous-multicore/hello-world-ca/
ddr_release/hello_world_ca55_RTOS0_UART2.elf > /sys/devices/platform/
remoteproc-ca55-1/remoteproc/remoteproc0/firmware
root@imx93evk:~# echo start > /sys/devices/platform/remoteproc-ca55-1/
remoteproc/remoteproc0/state
```

Then, the following RTOS log is displayed on the RTOS UART console on i.MX 93 11x11 and 14x14 EVK:

```
Cortex-A55: RTOS0: Hello world! Real-time Edge on MIMX93-EVK
RTOS1: RAM console@0xd1fff000
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac
tic tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac
```

b. Stop RTOS

To stop RTOS on a specified CPU core, echo `stop` to the corresponding `state` portal of the remoteproc instance.

For example, use the following command to stop RTOS running on Cortex-A Core3 (first Core is Core0) on i.MX 8M Mini EVK and i.MX 8M Plus EVK:

```
root@imx8mp-lpddr4-evk:~# echo stop > /sys/devices/platform/remoteproc-
ca53-3/remoteproc/remoteproc2/state
```

Or use the following command to stop RTOS running on Cortex-A Core1 (the first Core is Core0) on i.MX 93 11x11 and 14x14 EVK:

```
root@imx93evk:~# echo stop > /sys/devices/platform/remoteproc-ca55-1/
remoteproc/remoteproc0/state
```

• Cortex-A Core dynamically allocated between Linux and RTOS

Before running RTOS by using remoteproc, SMP Linux kernel uses all Cortex-A Cores by default. After running RTOS on the specified Cortex-A Core by using remoteproc, these Cortex-A core(s) are hot removed from SMP Linux kernel and RTOS runs on them. After using the `remoteproc stop` command to stop the RTOS running on Cortex-A Core, these CPU Cores once used by RTOS are hot added back to SMP Linux kernel.

Can use the following command to check which CPU Core are used by SMP Linux kernel currently:

```
root@imx93evk:~# cat /proc/cpuinfo
```

Note: The remoteproc portal instance may change according to the different device tree configuration used. Therefore, use the `name` under each remoteproc instance portal to check which instance needs to be used.

3.7.4 Using U-Boot Commands to Manage Life Cycle of RTOS on Cortex-M Core

U-boot command "`bootaux`" is used to boot Cortex-M Core RTOS Image from U-Boot, for example, after the board is booted into the U-Boot console.

- Use the following command to boot Arm Cortex-M core on i.MX 8M Mini LPDDR4 EVK board:

```
=> ext4load mmc 1:2 0x48000000 /examples/heterogeneous-multicore/hello-world-cm/
release/hello_world_cm4_UART4.bin; cp.b 0x48000000 0x7e0000 20000;
=> bootaux 0x7e0000
```

- Use the following command on i.MX 8M Plus LPDDR4 EVK board:

```
=> ext4load mmc 1:2 0x48000000 /examples/heterogeneous-multicore/hello-world-cm/
release/hello_world_cm7_UART4.bin; cp.b 0x48000000 0x7e0000 20000;
=> bootaux 0x7e0000
```

- Use the following command on i.MX 93 EVK board:

```
=> ext4load mmc 1:2 0xd0000000 /examples/heterogeneous-multicore/hello-world-cm/
release/hello_world_cm33_UART2.bin; cp.b 0xd0000000 0x201e0000 20000;
=> bootaux 0x1ffe0000
```

- Then, the below RTOS log is displayed on the UART console (taking i.MX 8M Mini EVK as an example)

```
Cortex-M4: RTOS0: Hello world! Real-time Edge on MIMX8MM-EVK
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac
```

3.7.5 Using RemoteProc under Linux to Manage Life Cycle of RTOS on Cortex-M Core

Use Linux remoteproc to start and stop RTOS on Cortex-M Core, taking the Heterogeneous Multicore `hello_world` application as an example:

- Booting Arm Cortex-M core on i.MX 8M Mini EVK, i.MX 8M Plus EVK and i.MX 93 EVK boards:

1. Starting Linux kernel:

```
=> run prepare_mcore
=> setenv fdtfile imx8mm-evk-multicore-rtos.dtb # for i.MX 8M Mini EVK
=> setenv fdtfile imx8mp-evk-multicore-rtos.dtb # for i.MX 8M Plus EVK
=> setenv fdtfile imx93-11x11-evk-multicore-rtos.dtb # for i.MX 93 EVK
=> boot
```

2. After Linux booting up, boot up Cortex-M core RTOS application:

```
# i.MX 8M Mini EVK
root@imx8mm-lpddr4-evk:~# echo -n /examples/heterogeneous-multicore/hello-
world-cm/release/hello_world_cm4_UART4.elf > /sys/devices/platform/imx8mm-
cm4/remoteproc/remoteproc4/firmware
root@imx8mm-lpddr4-evk:~# echo start > /sys/devices/platform/imx8mm-cm4/
remoteproc/remoteproc4/state

# i.MX 8M Plus EVK
root@imx8mp-lpddr4-evk:~# echo -n /examples/heterogeneous-multicore/hello-
world-cm/release/hello_world_cm7_UART4.elf > /sys/devices/platform/imx8mp-
cm7/remoteproc/remoteproc4/firmware
root@imx8mp-lpddr4-evk:~# echo start > /sys/devices/platform/imx8mp-cm7/
remoteproc/remoteproc4/state

# i.MX 93 EVK
root@imx93evk:~# echo -n /examples/heterogeneous-multicore/hello-world-cm/
release/hello_world_cm33_UART2.elf > /sys/devices/platform/remoteproc-cm33/
remoteproc/remoteproc1/firmware
```

```
root@imx93evk:~# echo start > /sys/devices/platform/remoteproc-cm33/
remoteproc/remoteproc1/state
```

Then, the below RTOS log is displayed on the UART console (taking i.MX 8M Mini EVK as an example):

```
Cortex-M4: RTOS0: Hello world! Real-time Edge on MIMX8MM-EVK
tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic tac tic
tac
tic tac tic tac tic tac tic tac tic tac tic tac tic tac
```

3. Then the Cortex-M core RTOS application can be shut down:

```
# i.MX 8M Mini EVK
root@imx8mm-lpddr4-evk:~# echo stop > /sys/devices/platform/imx8mm-cm4/
remoteproc/remoteproc4/state

# i.MX 8M Plus EVK
root@imx8mp-lpddr4-evk:~# echo stop > /sys/devices/platform/imx8mp-cm7/
remoteproc/remoteproc4/state

# i.MX 93 EVK
root@imx93evk:~# echo stop > /sys/devices/platform/remoteproc-cm33/
remoteproc/remoteproc1/state
```

Note: The remoteproc portal instance ID /sys/devices/platform/remoteproc-cm7/remoteproc/remoteproc4 might be different, depending on the device tree configuration. Therefore, use the name under each remoteproc instance portal to check which instance must be used.

3.8 Industrial applications

3.8.1 EtherCAT MainDevice Stack SOEM

3.8.1.1 Overview

The Heterogeneous Multicore framework provides a flexible mechanism to run Preemp-RT Linux or RTOS on different cores with Industrial applications. SOEM can run on FreeRTOS to scale down to the Cortex-M core and scale up to the Cortex-A core. Taking i.MX 8M Plus platform as an example, it is equipped with four Cortex-A53 cores and one Cortex-M7 core. [Table 41](#) describes the use cases that can be executed on this MPU platform.

Table 41. SOEM use cases

SOEM examples	M7	A53	A53	A53	A53
digital_io	baremetal/FreeRTOS	FreeRTOS			
servo_motor	baremetal/FreeRTOS	FreeRTOS			
servo_motor_rt1180	baremetal/FreeRTOS	FreeRTOS			

SOEM is supported on the below platforms:

- i.MX 8M Plus LPDDR4 EVK platform
- i.MX 8M Mini LPDDR4 EVK platform
- i.MX 93 EVK platform

Three examples are provided in this release:

- SOEM `digital_io` example for IO control using BECKHOFF EK1100 as EtherCAT SubDevice

- SOEM `servo_motor` example using Inovance SV680 servo as EtherCAT SubDevice
 - SOEM `servo_motor_rt1180` example using i.MX RT1180-EVK as EtherCAT SubDevice
- For detailed information of SOEM, refer to [Section 6.1.4](#).

3.8.1.2 Running SOEM Application on Cortex-A core

For detailed information, refer to [Section 6.1.4](#) for steps to run the SOEM stack on the Cortex-A core for the supported platforms.

3.8.1.3 Running SOEM Application on Cortex-M core

For detailed information, please refer to [Section 6.1.4](#) to run the SOEM stack on the Cortex-M core for the supported platforms.

4 Heterogeneous Multi-SoC Framework

This section describes the features of heterogeneous Multi-SoC framework and how to implement it using NXP hardware platforms.

4.1 Introduction

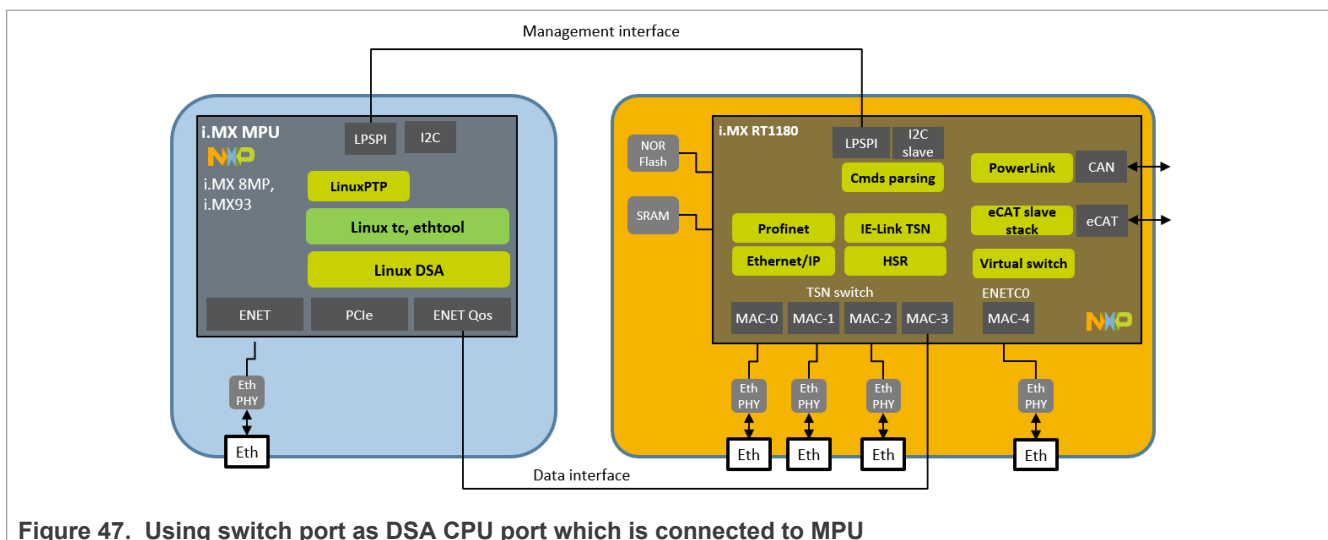
Heterogeneous Multi-SoC Framework enables the usage of a combination of MPU and i.MX RT1180 as an Industrial Switch. This extends the MPU hardware capability with the i.MX RT1180 hardware capability, thereby providing switch functionality, TSN functionality, and the capability of supporting different Industrial Protocols. The i.MX RT1180 can be used to run real-time tasks such as industrial protocols in the real-time domain. On the other hand, the MPU can process compute-heavy tasks in the non-real-time domain.

The external Ethernet ports of i.MX RT1180 can be exposed to the MPU side as standard Ethernet interfaces as data path. Different interfaces such as SPI or I2C can be used as the management interface between MPU and i.MX RT1180.


4.2 Software architecture

The Linux Distributed Switch Architecture (DSA) framework is used to expose the i.MX RT1180 NETC switch ports to MPU side. In this architecture, one of the NETC switch ports or ENETC port is used as the data interface. Different interfaces (for example LPSPi, I2C or message unit) can be used as management interfaces.

For more information regarding Linux DSA, refer to <https://docs.kernel.org/networking/dsa/dsa.html>.



Note: The industrial protocols listed in the box of i.MX RT1180 are not supported yet.



The software architecture implementation includes:

- NETC DSA switch driver on Linux
- Device driver of DSA control interface on Linux DSA
- Service driver of DSA control interface on i.MX RT1180
- NETC DSA switch configuration on i.MX RT1180
 - Using NETC switch port as DSA CPU port
 - Using ENETC port as DSA CPU port

The external Ethernet ports of i.MX RT1180 are exposed to MPU with Linux DSA framework. The ports can be viewed as standard Ethernet ports which could support the below operations:

- Binding Linux IP address to a specific port
- Broadcast on a specific port and others

To identify the source traffic of incoming traffic, the outer VLAN tag is used as the DSA tag to indicate which port the traffic is coming from. The VID field of the VLAN tag is encoded to include the source/destination port. Below is the description of the 12-bit VID field:

Table 42. VID field description

11	10	9	8	7	6	5	4	3	2	1	0
RSV		VBID	SWITCH_ID			VBID		PORT			

Note: The above custom VID definition is coming from the Linux `tag_8021q` source code (`net/dsa/tag_8021q.c`) which is used by the Linux NETC DSA driver. As a result, VLAN ID 3072 to 3076 and 3088 are reserved by the Linux NETC DSA driver and will not be allowed to be used by the user.

4.2.1 Using one of the i.MX RT1180 switch ports as DSA CPU port

Consider the case when one of the i.MX RT1180 switch ports is used as DSA CPU port that is connected to MPU. In such a situation, the packets received on the external ports of i.MX RT1180 are filtered using Ingress Port Filter Table. The incoming packets are filtered according to the ingress port ID, and forwarded to the DSA CPU port by either stream forwarding (by-passing the 802.1Q bridge forwarding) or 802.1Q bridge forwarding using Ingress Stream Table. The outer VLAN tag is added by using the Egress Treatment Table on egress DSA CPU port.

The packets received on the DSA CPU port are filtered according to the DSA tag, which is viewed as a VID from i.MX RT1180 perspective. This is done by using Ingress Port Filter Table and the packets are directed to the port mapped to the VID by using Ingress Stream Table. The outer VLAN tag is removed on ingress DSA CPU port by also using Ingress Stream Table.

4.2.2 Using i.MX RT1180 ENETC port as DSA CPU port

For the case using the ENETC port as the DSA CPU port that is connected to MPU, the packets received on the external ports of i.MX RT1180 are filtered using the Ingress Port Filter Table. The incoming packets are filtered according to the ingress port ID, and forwarded to the internal switch port (also called switch management port). The forwarding is done by either stream forwarding (by-passing the 802.1Q bridge forwarding) or 802.1Q bridge forwarding using Ingress Stream Table. The software virtual switch running on the CPU of i.MX RT1180 receives the packets and the software virtual switch adds the DSA tag and forwards the packets to the ENETC port.

The packets received on the ENETC port are forwarded by the software virtual switch to the internal switch port, where the packets are filtered according to the DSA tag, which is viewed as VID from i.MX RT1180 perspective. This is done by using Ingress Port Filter Table and the packets are directed to the port mapped to the VID by using Ingress Stream Table. The outer VLAN tag is removed by using Egress Treatment Table on egress switch port.

4.3 Running the Heterogeneous Multi-SoC Framework

This section describes the process to implement and run the heterogeneous Multi-SoC framework on supported MPU platforms and i.MX RT1180 EVK.

4.3.1 Building and running on MPU + i.MX RT1180 EVK

The currently supported MPU platforms include i.MX 8M Plus EVK and i.MX 93 EVK.

Note: For i.MX RT1180 EVK, the supported board revision is SCH-50577 REV C2 (700-50577 REVC), which has i.MX RT1180 Rev. B0 chip.

4.3.1.1 Hardware setup for i.MX 8M Plus EVK and i.MX RT1180 EVK

1. SPI connection between i.MX 8M Plus EVK and i.MX RT1180 EVK

On i.MX 8M Plus EVK, ECSPI2 pins are available on **J21** connector, but due to signal incompatibility with the i.MX RT1180 EVK LPSPI3 pins on **J44** connector, board rework is needed on i.MX 8M Plus EVK to replace a translating transceiver NBT0104 **U55** which has limited capacitive loading for 70 pF with NTS0104GU12, which has internal pull-up of 10K and bigger capacitive loading. Refer to the schematic shown in the [Figure 49](#) below:




Figure 49. Hardware rework for SPI signals on i.MX 8M Plus EVK

Connect ECSPi2 pins on J21 connector to i.MX RT1180 EVK LPSPi3 pins by following the connection in [Table 43](#).

Table 43. Pin connection between i.MX 8M Plus EVK and i.MX RT1180 EVK

i.MX 8M Plus EVK		Connection	i.MX RT1180 EVK	
Pin	Function		Pin	Function
19	ECSPi2_MOSI	<->	10	LPSPi3_SIN
21	ECSPi2_MISO	<->	8	LPSPi3_SOUT
23	ECSPi2_SCLK	<->	12	LPSPi3_CLK
24	ECSPi2_SS0	<->	6	LPSPi3_PCS0
6	GND	<->	14	GND

2. Ethernet connection between i.MX 8M Plus EVK and i.MX RT1180 EVK

For the case using one of the i.MX RT1180 external switch ports as DSA CPU port, use Ethernet cable to connect i.MX 8M Plus EVK ENET2 port (eth1 in Linux) on J8 RJ45 connector and i.MX RT1180 EVK NETC switch port 3 (ENET3) on J31 RJ45 connector.

For the case using ENETC port as DSA CPU port, use Ethernet cable to connect i.MX 8M Plus EVK ENET2 port (eth1 in Linux) on J8 RJ45 connector and i.MX RT1180 EVK ENETC port (ENET4) on J32 RJ45 connector.

4.3.1.2 Hardware setup for i.MX 93 EVK and i.MX RT1180 EVK

1. LPSPi connection between i.MX 93 EVK and i.MX RT1180 EVK

Using flying wire to connect i.MX 93 EVK LPSPi3 pins on J1001 connector and i.MX RT1180 EVK LPSPi3 pins on J44 connector by following the pin connection in below table:

Table 44. Pin connection between i.MX 93 EVK and i.MX RT1180 EVK

i.MX 93 EVK		Connection	i.MX RT1180 EVK	
Pin	Function		Pin	Function
19	LPSPi3_MOSI	<->	10	LPSPi3_SIN
21	LPSPi3_MISO	<->	8	LPSPi3_SOUT
23	LPSPi3_CLK	<->	12	LPSPi3_CLK
24	LPSPi3_PCS0	<->	6	LPSPi3_PCS0
25	GND	<->	14	GND

2. Ethernet connection between i.MX 93 EVK and i.MX RT1180 EVK

For the case using one of the i.MX RT1180 external switch ports as DSA CPU port, use Ethernet cable to connect i.MX 93 EVK ENET_QOS port (eth1 in Linux) on **J501** RJ45 connector and i.MX RT1180 EVK NETC switch port 3 (ENET3) on **J31** RJ45 connector.

For the case using ENETC port as DSA CPU port, use Ethernet cable to connect i.MX 93 EVK ENET_QOS port (eth1 in Linux) on **J501** RJ45 connector and i.MX RT1180 EVK ENETC port (ENET4) on **J32** RJ45 connector.

4.3.1.3 Building the DSA switch application for i.MX RT1180 EVK

The demo application images `dsa_switch.elf/dsa_switch.bin` and `dsa_enetc.elf/dsa_enetc.bin` for i.MX RT1180 EVK are installed into target rootfs with the MPU Yocto image build. These images are available in the `/examples/heterogeneous-multi-soc/dsa-switch-evkmimxrt1180-cm33/` directory.

Also, the demo application images can be deployed into the Yocto build directory by using the following Yocto commands:

```
# For i.MX 8M Plus EVK
$ DISTRO=nxp-real-time-edge MACHINE=imx8mp-lpddr4-evk source real-time-edge-
setup-env.sh -b <build_dir>
$ bitbake dsa-switch-evkmimxrt1180-cm33

# For i.MX 93 EVK
$ DISTRO=nxp-real-time-edge MACHINE=imx93evk source real-time-edge-setup-env.sh
-b <build_dir>
$ bitbake dsa-switch-evkmimxrt1180-cm33
```

After the Yocto `bitbake` command is used, the demo application images can be found in the below directory on the build host:

```
<build_dir>/tmp/deploy/images/<MACHINE>/examples/heterogeneous-multi-soc/dsa-
switch-evkmimxrt1180-cm33
```

The demo applications are available in both `release` and `ram_release` mode. The `release` image boots from the external QuadSPI NOR flash and then relocates to internal memory (Code TCM, System TCM, and OCRAM) and needs to be flashed in the external NOR flash. The `ram_release` image boots from internal memory and must be loaded directly into internal memory by host debug tools.

4.3.1.4 Burning the image to flash on i.MX RT1180 EVK

This section describes how to burn DSA image to i.MX RT1180 EVK board.

4.3.1.4.1 Hardware setup

To burn the DSA image to i.MX RT1180 EVK board, use the hardware setup described in this section.

1.4 Board pictures

Figure 2 shows the connectors on the top-side view of the MIMXRT1180-EVK board.




Figure 2. MIMXRT1180-EVK top view connectors

Figure 3 shows the connectors on the bottom side of the MIMXRT1180-EVK board.

Figure 50. MIMXRT1180-EVK top view connectors

Figure 50 shows the connectors on the top-side view of the MIMXRT1180-EVK board.

To burn the image to i.MX RT1180, connect the PC to the MCU-Link USB J53 using a micro-USB cable.




Figure 51. Connection between PC and i.MX RT1180

4.3.1.4.2 Preparing the DSA images

There are two ways to get the DSA (distributed switch architecture) images, which are described in the following sections:

- Using pre-built images
- Building the DSA image

4.3.1.4.2.1 Using pre-built image

Real-time Edge release package includes DSA images in root file system. The folder is */example/heterogeneous-multi-soc*. The content of this package is as follows:

```
/examples/heterogeneous-multi-soc/  
├── dsa-switch-evkmimxrt1180-cm33  
│   ├── ram_release  
│   │   ├── dsa_enetc.elf  
│   │   └── dsa_switch.elf  
│   └── release  
│       ├── dsa_enetc.bin  
│       ├── dsa_enetc.elf  
│       ├── dsa_switch.bin  
│       └── dsa_switch.elf
```

Images under the `ram_release` folder are for RAM boot. Images under `release` folder are for flash boot. The `dsa_enetc` image supports enetc as DSA CPU port. The `dsa_switch` image supports one of switch ports as DSA CPU port.

In this example we need to copy `release/dsa_switch.elf` to the PC.

4.3.1.4.2.2 Building the DSA image

To build the DSA images, refer to the section [Section 4.3.1.3](#).

4.3.1.4.2.3 Prepare SPT tool

Download and install MCUXpresso Secure Provisioning Tool (SPT) from the below link on the host PC: <https://www.nxp.com/webapp/Download?colCode=MCUXPRESSO-SECURE-PROVISIONING-V9.0.1-WIND&appType=license> Then, install MCUXpresso Secure Provisioning to the PC.

4.3.1.4.2.4 Burning the DSA image

1. Set “**SW5[1..4]**” of board to “**0000**” to enable SDP mode.
2. Open SPT. Then, click “**USB**” to change to “**UART**”, select the “**Port**” of the board, set “**Baud rate**” to “**115200**”.

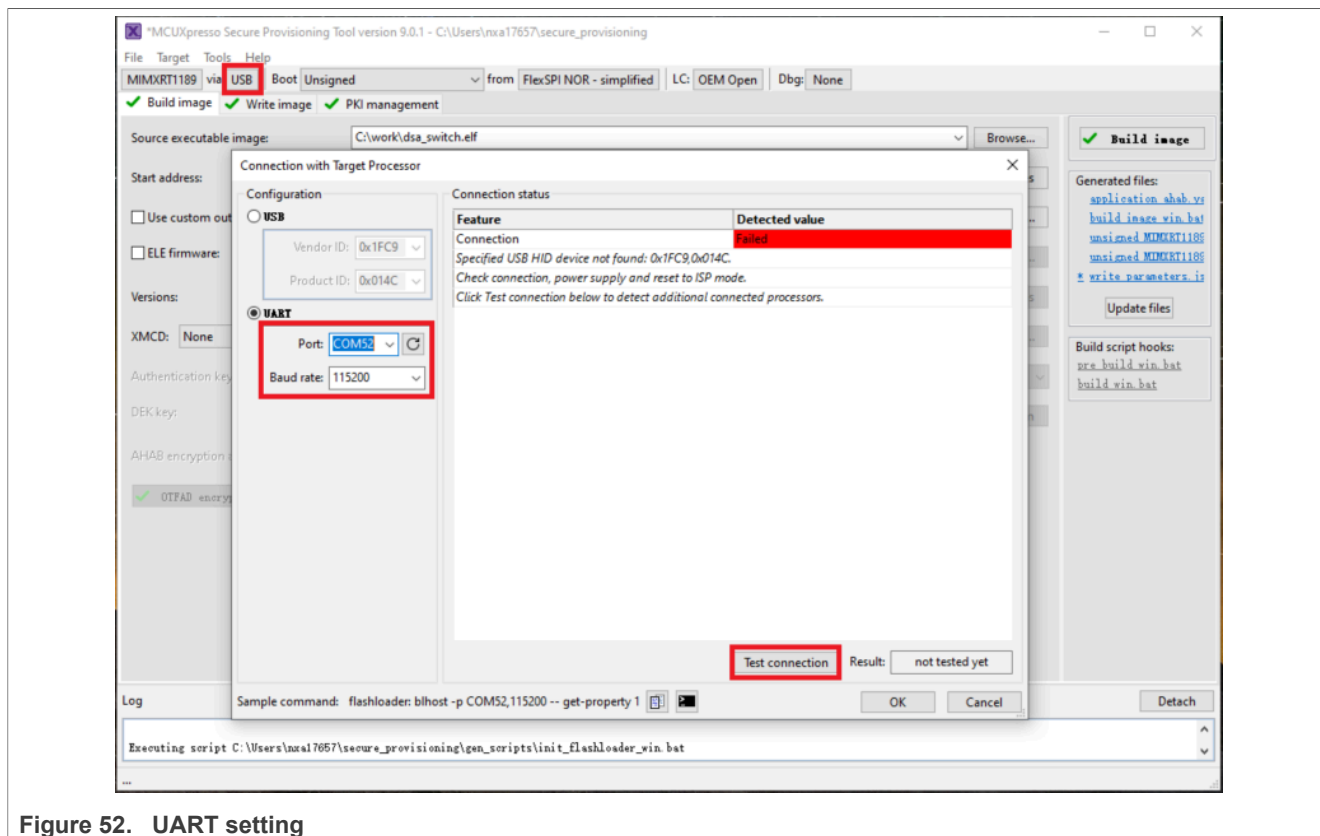


Figure 52. UART setting

3. Click the button “**Test connection**” to check the UART connection.

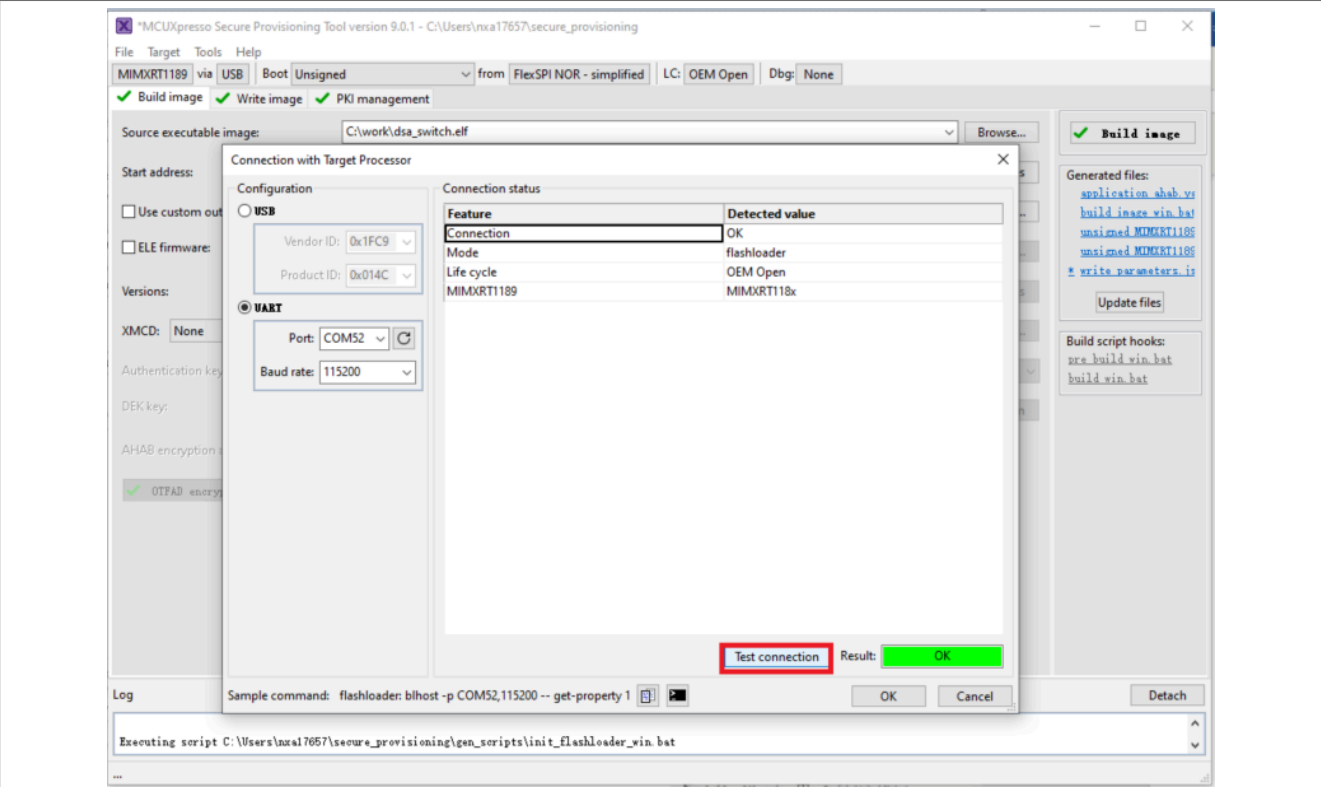


Figure 53. UART connection

4. Click “Browse” to select “release/dsa_switch.elf”

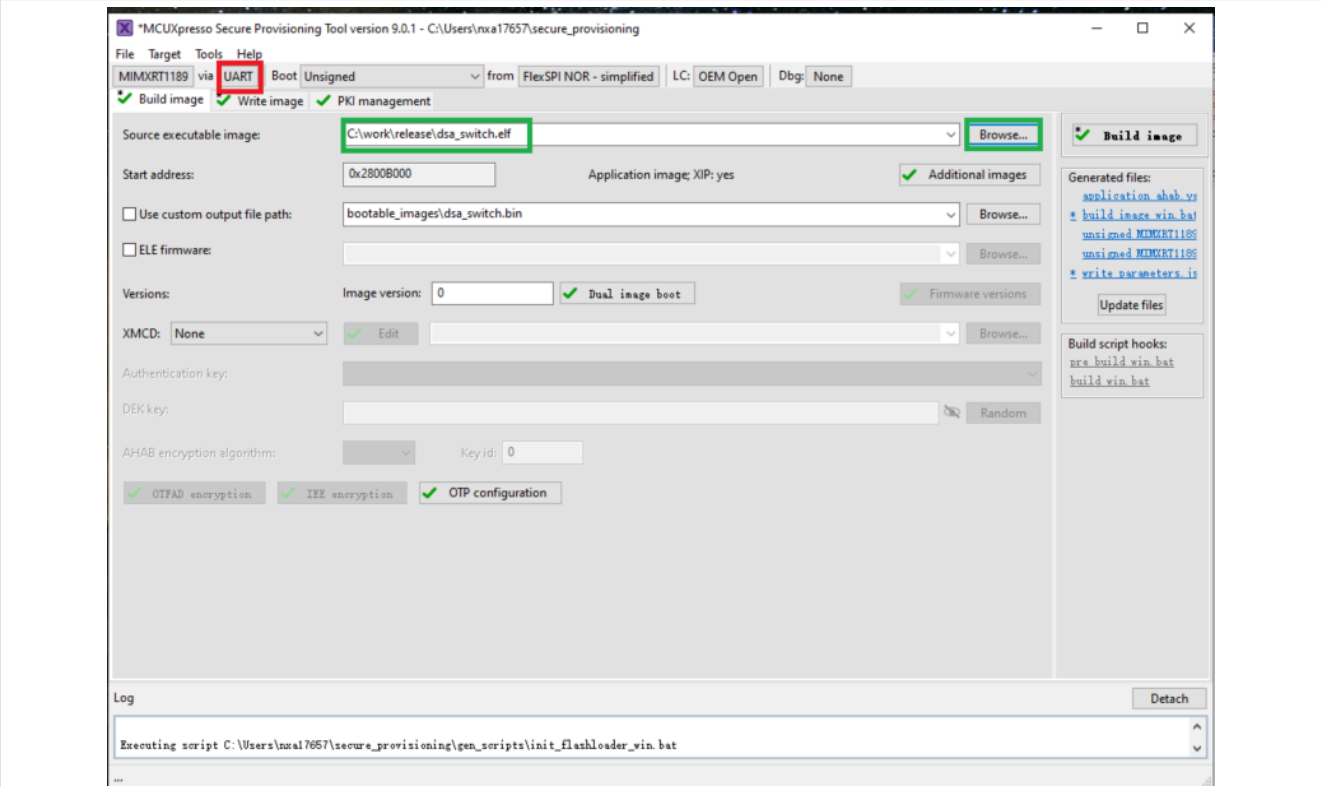


Figure 54. Select “release/dsa_switch.elf”

5. Click the “build image” button to build the images.

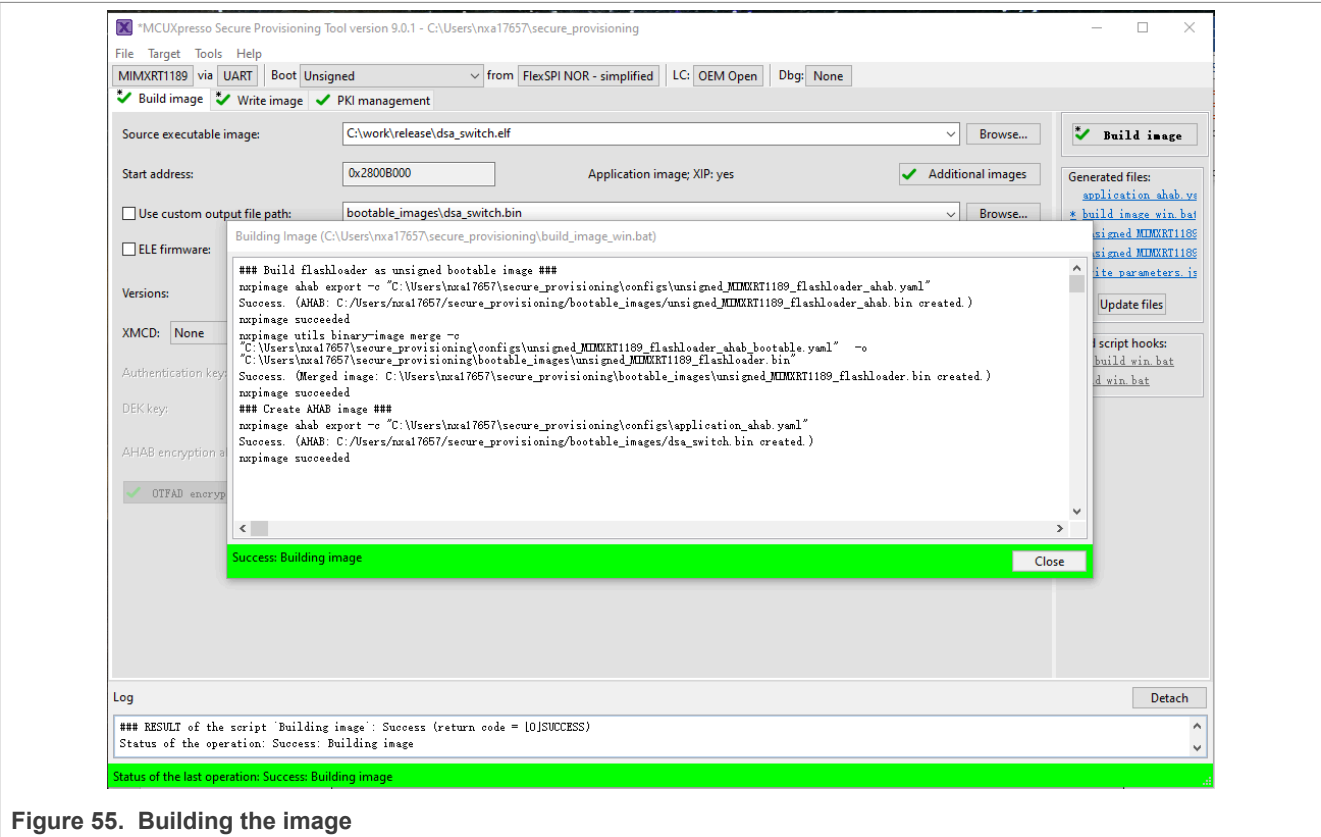


Figure 55. Building the image

6. Click the “write image” button to burn the image to the flash.

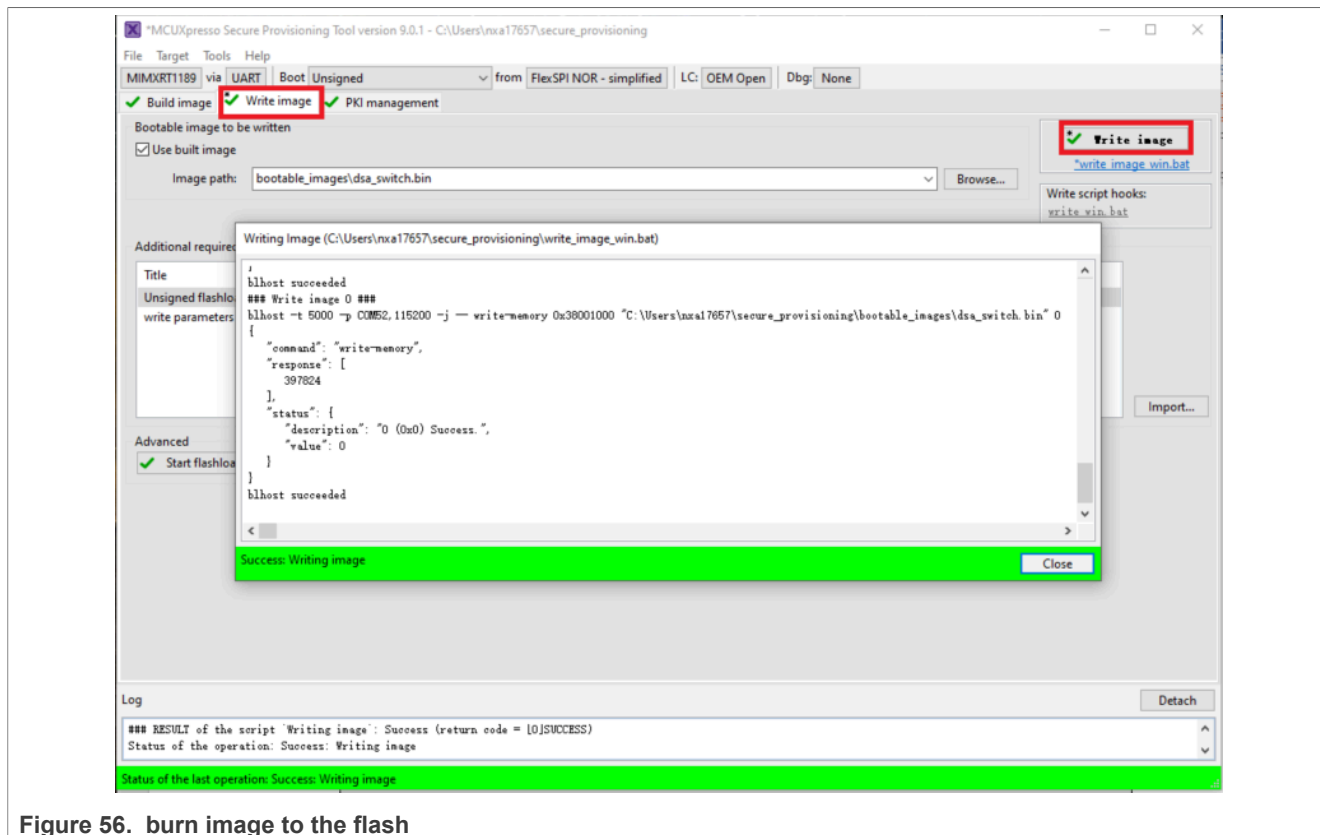


Figure 56. burn image to the flash

7. Change **SW5** to **"0100"**- flash boot mode.

4.3.1.4.2.5 Checking the DSA image

1. Start Putty or other serial terminal tools to access i.MX RT1180.
The configuration of the terminal is as follows:
 - Speed (bit/s): 115200
 - Data bits: 8
 - Stop bits: 1
 - Parity: None
 - Flow control: None
2. Here is a partial log at startup

```

COMS2 - PuTTY
started
INFO      0 app_dsa_ctrl_task      : DSA control task started

Copyright 2024 NXP

INFO      0 app_main_task          : main_task started

BRIDGE>>INIT      0.000000000 os      genavb_init      : NXP's GenAVB/TSN stack version 6_5-20e0ld7e
INFO      0.000000000 os      _netc_1588_init      : timerClk period (ns): 4 genClk period (ns): 100000
INFO      0.000000000 os      msgintr_msiX_msgaddr : msgintr addr return for the MSI-X Entry
INFO      0.000000000 os      _netc_1588_init      : refClkKHz: 240000000 defaultPpb: 0
INFO      0.000000000 os      _netc_1588_init      : Alarm2 is configured
INFO      0.000000000 os      _netc_1588_init      : Alarm1 is configured for Fiper
INIT      0.000000000 os      hw_clock_init      : rate: 1000000000, period: 0, mult(to ns): 1, shift(to ns): 0, mult(to cy
cles): 1, shift(to cycles): 0
INIT      0.000000000 os      hw_clock_register   : hw_clock(3) registered
INIT      0.000000000 os      hw_clock_init      : rate: 240000000, period: 0, mult(to ns): 2236962133, shift(to ns): 29, m
ult(to cycles): 1030792151, shift(to cycles): 32
INIT      0.000000000 os      hw_clock_register   : hw_clock(4) registered
INIT      0.000000000 os      hw_timer_register   : hw_timer(FFE9478) of hw_clock(3) registered, pps support
INIT      0.000000000 os      hw_clock_init      : rate: 800000000, period: 1000000000, mult(to ns): 208715200, shift(to ns):
24, mult(to cycles): 343587384, shift(to cycles): 32
INIT      0.000000000 os      hw_clock_register   : hw_clock(1) registered
INIT      0.000000000 os      hw_timer_register   : hw_timer(FFE4F70) of hw_clock(1) registered
INIT      0.000000000 os      hw_timer_register   : hw_timer(FFE4F98) of hw_clock(1) registered
INIT      0.000000000 os      hw_timer_register   : hw_timer(FFE4FC0) of hw_clock(1) registered
ERR      0.000000000 os      gpt_driver_init     : failed to register GPT media clock recovery
INIT      0.000000000 os      logical_port_init   : logical_port(0) enabled, endpoint(0)
INIT      0.000000000 os      logical_port_init   : logical_port(1) enabled, endpoint(1)
INIT      0.000000000 os      logical_port_init   : logical_port(2) enabled, bridge(0, 0)
INIT      0.000000000 os      logical_port_init   : logical_port(3) enabled, bridge(0, 1)
INIT      0.000000000 os      logical_port_init   : logical_port(4) enabled, bridge(0, 2)
INIT      0.000000000 os      logical_port_init   : logical_port(5) enabled, bridge(0, 3)
INIT      0.000000000 os      logical_port_init   : logical_port(6) enabled, bridge(0, 4)
INIT      0.000000000 os      logical_ports_configuration: 2 endpoints, 5 bridge ports
INIT      0.000000000 os      _port_init         : port(0): FFE7998
INFO      0.000000000 os      enetc_ep_init       : port(0) driver handle (FFE8718)

```

Figure 57. Figure 5partial log at startup

4.3.1.5 Bringing up MPU and i.MX RT1180 EVK

1. Make SPI and Ethernet connections between i.MX 8M Plus EVK or i.MX 93 EVK and i.MX RT 1180 EVK by following the steps in the Section [Section 4.3.1.1](#) or [Section 4.3.1.2](#).
2. Make a serial connection between **J23** connector on i.MX 8M Plus EVK or **J1401** connector on i.MX 93 EVK and USB connector on host PC using USB cable. Open the terminal application on host PC, such as PuTTY or TeraTerm, and connect to the 3rd debug console which is used by Linux running on MPU.
3. Make a serial connection between **J53** connector on i.MX RT1180 EVK and USB connector on host PC using USB cable. Open another terminal application on host PC, and connect to the debug console which is used by FreeRTOS running on i.MX RT1180 EVK. Note that the terminal settings 'Implicit CR in every LF' should be enabled for correct display of the logs.
4. Power on i.MX RT1180 EVK board, and flash the binary image `release/dsa_switch.bin` or `release/dsa_enetc.bin` on i.MX RT1180 EVK and power cycle the board. Alternatively, one can download the elf image `ram_release/dsa_switch.elf` or `ram_release/dsa_enetc.elf` by host debug tools (e.g. using JTAG debugger connected to **J37** connector). After the demo application is running, there should be continuous logs displayed in the console.
5. Power on the MPU board and enter U-Boot command line. Then execute the corresponding commands depending on which port on i.MX RT1180 is used as DSA CPU port.

For the case using one of the i.MX RT1180 external switch ports as DSA CPU port:

```

# For i.MX 8M Plus EVK
u-boot=> setenv fdtfile imx8mp-evk-revb4-dsa.dtb
u-boot=> boot

# For i.MX 93 EVK
u-boot=> setenv fdtfile imx93-11x11-evk-dsa.dtb
u-boot=> boot

```

For the case using ENETC port as DSA CPU port:

```

# For i.MX 8M Plus EVK
u-boot=> setenv fdtfile imx8mp-evk-revb4-dsa-enetc.dtb
u-boot=> boot

```

```
# For i.MX 93 EVK
u-boot=> setenv fdtfile imx93-11x11-evk-dsa-enetc.dtb
u-boot=> boot
```

After Linux is up on MPU, refer to the [Section 4.3.2](#) for the detailed usage on how to show and configure various NETC switch features on i.MX RT1180.

4.3.2 Runtime usage on MPU and i.MX RT1180 EVK

This section describes the major i.MX RT1180 NETC switch features that can be queried or configured on MPU + i.MX RT1180 EVK heterogeneous multi-SoC architecture using Linux DSA (Distributed Switch Architecture).

4.3.2.1 i.MX RT1180 NETC switch interface under Linux

On MPU, after Linux is up, on successful initialization of the Linux NETC DSA driver, the i.MX RT1180 NETC front panel switch port 0 (ENET0), port 1 (ENET1), port 2 (ENET2), port 3 (ENET3) should have a network device interface attached with the `swpX` name format.

Table 45. i.MX RT1180 EVK NETC switch port mapping in Linux on MPU

i.MX RT1180 EVK NETC switch port	Label on i.MX RT1180 EVK	Linux network device name on MPU
Port 0	ENET0 (J28)	swp0
Port 1	ENET1 (J29)	swp1
Port 2	ENET2 (J30)	swp2
Port 3	ENET3 (J31)	N/A or swp3 (see note)
Port 4 (internal)	N/A	N/A

Note: Linux DSA does not currently create user network device for CPU port. When using one of the i.MX RT1180 external switch ports as CPU port (i.e. the i.MX RT1180 NETC switch port 3 is used as CPU port), it is not visible to the user as a normal network device in Linux. When using ENETC port as CPU port, the i.MX RT1180 NETC switch port 3 is shown as `swp3` in Linux.

One can use `ifconfig -a` command to show these network device interfaces.

Example - NETC switch port interfaces `swpX` shown by `ifconfig -a`

```
# ifconfig -a
[...]
swp0: flags=4098<BROADCAST,MULTICAST> mtu 1500
    ether 00:04:9f:08:49:31 txqueuelen 1000  (Ethernet)
    RX packets 0  bytes 0 (0.0 B)
    RX errors 0  dropped 0  overruns 0  frame 0
    TX packets 0  bytes 0 (0.0 B)
    TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0

swp1: flags=4098<BROADCAST,MULTICAST> mtu 1500
    ether 00:04:9f:08:49:31 txqueuelen 1000  (Ethernet)
    RX packets 0  bytes 0 (0.0 B)
    RX errors 0  dropped 0  overruns 0  frame 0
    TX packets 0  bytes 0 (0.0 B)
    TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0

swp2: flags=4098<BROADCAST,MULTICAST> mtu 1500
    ether 00:04:9f:08:49:31 txqueuelen 1000  (Ethernet)
    RX packets 78  bytes 19158 (18.7 KiB)
    RX errors 0  dropped 78  overruns 0  frame 0
    TX packets 0  bytes 0 (0.0 B)
    TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0

/* swp3 is only available when using ENETC port as DSA CPU port */
swp3: flags=4098<BROADCAST,MULTICAST> mtu 1500
    ether 00:04:9f:08:49:31 txqueuelen 1000  (Ethernet)
    RX packets 0  bytes 0 (0.0 B)
```

```
RX errors 0   dropped 0   overruns 0   frame 0
TX packets 0   bytes 0 (0.0 B)
TX errors 0   dropped 0 overruns 0   carrier 0   collisions 0
```

Also, the `ip link show` command uses the `swpX@eth1` name format to also indicate the associated master Ethernet interface (`eth1`) on MPU for the DSA switch port.

Example - NETC switch port interfaces `swpX` and DSA master interface `eth1` shown by `ip link show`.

```
# ip link show
[...]
3: eth1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1504 qdisc mq state UP mode
  DEFAULT group default qlen 1000
    link/ether 00:04:9f:08:49:31 brd ff:ff:ff:ff:ff:ff
4: swp0@eth1: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN mode DEFAULT
  group default qlen 1000
    link/ether 00:04:9f:08:49:31 brd ff:ff:ff:ff:ff:ff
5: swp1@eth1: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN mode DEFAULT
  group default qlen 1000
    link/ether 00:04:9f:08:49:31 brd ff:ff:ff:ff:ff:ff
6: swp2@eth1: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN mode DEFAULT
  group default qlen 1000
    link/ether 00:04:9f:08:49:31 brd ff:ff:ff:ff:ff:ff
/* swp3@eth1 is only available when using ENETC port as DSA CPU port */
7: swp3@eth1: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN mode DEFAULT
  group default qlen 1000
    link/ether 00:04:9f:08:49:31 brd ff:ff:ff:ff:ff:ff
```

4.3.2.2 Maximum frame size configuration

The Linux NETC DSA driver correlates between Layer-2 maximum frame size and Layer-3 MTU. Typical settings of Layer-2 maximum frame size is 1518 bytes for standard untagged frames and the corresponding Layer-3 MTU is 1500 bytes.

Example 1: Setting Layer-3 MTU of NETC switch port 1 (`swp1`) to 1000 bytes

```
# ifconfig swp1 mtu 1000
or
# ip link set dev swp1 mtu 1000
```

Note:

Setting the Layer-3 MTU for an external NETC switch port on i.MX RT1180 EVK to a value larger than 1500 bytes also changes the Layer-3 MTU of both the CPU port on i.MX RT1180 EVK and the master Ethernet interface (`eth1`) on MPU accordingly.

As the maximum MTU of `eth1` is 1600, the maximum MTU value that can be configured on an external NETC switch port is 1596 bytes. However, the NETC Ethernet MAC supports configurable maximum frame size up to 2000 bytes. The 4 extra bytes in MTU for both the CPU port on i.MX RT1180 EVK and the master Ethernet interface (`eth1`) on MPU are used to store the 4-bytes 802.1Q DSA tag.

4.3.2.3 Single port mode for i.MX RT1180 NETC switch ports

After Linux is up on MPU, by default all the i.MX RT1180 NETC switch ports available in Linux works in single port mode. In this configuration mode the traffic received on other switch ports is forwarded to the CPU port

(i.MX RT1180 NETC switch port 3). Each of the other external switch port interface can be used independently to send and receive packets.

Example 1: Single port configuration of the Linux NETC DSA driver

```
/* configure IP address on external switch interfaces */
# ip addr add 192.168.1.1/24 dev swp0
# ip addr add 192.168.2.1/24 dev swp1
# ip addr add 192.168.3.1/24 dev swp2
/* swp3 is only available when using ENETC port as DSA CPU port */
# ip addr add 192.168.4.1/24 dev swp3
/* master interface to be brought up first - up by default */
# ip link set eth1 up
/* bring up the switch slave interfaces */
# ip link set swp0 up
# ip link set swp1 up
# ip link set swp2 up
/* swp3 is only available when using ENETC port as DSA CPU port */
# ip link set swp3 up
```

Assuming there is one remote host connected to each of the external switch ports, one can ping the remote host.

```
/* Assuming the IP is 192.168.1.2 for the remote host connected to swp0 */
# ping 192.168.1.2
PING 192.168.1.2 (192.168.1.2) 56(84) bytes of data.
64 bytes from 192.168.1.2: icmp_seq=1 ttl=64 time=1.75 ms
64 bytes from 192.168.1.2: icmp_seq=2 ttl=64 time=1.77 ms
64 bytes from 192.168.1.2: icmp_seq=3 ttl=64 time=1.76 ms
[...]

/* Assuming the IP is 192.168.2.2 for the remote host connected to swp1 */
# ping 192.168.2.2
PING 192.168.2.2 (192.168.2.2) 56(84) bytes of data.
64 bytes from 192.168.2.2: icmp_seq=1 ttl=64 time=3.23 ms
64 bytes from 192.168.2.2: icmp_seq=2 ttl=64 time=1.75 ms
64 bytes from 192.168.2.2: icmp_seq=3 ttl=64 time=1.75 ms
[...]

/* Assuming the IP is 192.168.3.2 for the remote host connected to swp2 */
# ping 192.168.3.2
PING 192.168.3.2 (192.168.3.2) 56(84) bytes of data.
64 bytes from 192.168.3.2: icmp_seq=1 ttl=64 time=2.80 ms
64 bytes from 192.168.3.2: icmp_seq=2 ttl=64 time=1.76 ms
64 bytes from 192.168.3.2: icmp_seq=3 ttl=64 time=1.77 ms
[...]

/* swp3 is only available when using ENETC port as DSA CPU port */
/* Assuming the IP is 192.168.4.2 for the remote host connected to swp3 */
# ping 192.168.4.2
PING 192.168.4.2 (192.168.4.2) 56(84) bytes of data.
64 bytes from 192.168.4.2: icmp_seq=1 ttl=64 time=2.80 ms
64 bytes from 192.168.4.2: icmp_seq=2 ttl=64 time=1.76 ms
64 bytes from 192.168.4.2: icmp_seq=3 ttl=64 time=1.77 ms
[...]
```

4.3.2.4 Bridge mode for i.MX RT1180 NETC switch ports

In this configuration mode, the external switch ports (DSA slave interfaces) are added to a bridge for L2 forwarding. The `eth1` interface is brought up as the DSA master interface.

Example 1: Bridge configuration of the Linux NETC DSA driver

```
/* bring up master interface before the slave ports - up by default */
# ip link set eth1 up

/* bring up the switch slave interfaces */
# ip link set swp0 up
# ip link set swp1 up
# ip link set swp2 up
/* swp3 is only available when using ENETC port as DSA CPU port */
# ip link set swp3 up

/* create bridge */
# ip link add name br0 type bridge

/* add the external switch ports to the bridge */
# ip link set dev swp0 master br0
# ip link set dev swp1 master br0
# ip link set dev swp2 master br0
/* swp3 is only available when using ENETC port as DSA CPU port */
# ip link set dev swp3 master br0

/* configure IP address and bring up the bridge */
# ip addr add 192.168.200.1/24 dev br0
# ip link set dev br0 up
```

Execute the ping command on remote host machine

```
/* Assuming the IP is 192.168.200.2 and net device is eth1 on the remote host
 * connected to one of the external switch ports (swp0, swp1 and swp2)
 * when using one of the i.MX RT1180 switch ports as DSA CPU port or
 * (swp0, swp1, swp2 and swp3) when using ENETC port as DSA CPU port.
 */
/* Set the Remote host's IP and ping the br0 */
# ip addr add 192.168.200.2/24 dev eth1
# ping 192.168.200.1
```

Br0 pings the remote host machine

```
# ping 192.168.200.2
PING 192.168.200.2 (192.168.200.2) 56(84) bytes of data:
64 bytes from 192.168.200.2: icmp_seq=1 ttl=64 time=1.76 ms
64 bytes from 192.168.200.2: icmp_seq=2 ttl=64 time=1.80 ms
64 bytes from 192.168.200.2: icmp_seq=3 ttl=64 time=1.77 ms
[...]
```

At the same time, the L2 forwarding works using the external switch ports (`swp0`, `swp1` and `swp2`) when using one of the i.MX RT1180 external switch ports as DSA CPU port or (`swp0`, `swp1`, `swp2` and `swp3`) when using ENETC port as DSA CPU port. Assuming there is one remote host connected to each of the external switch ports, each remote host can ping other hosts.

4.3.2.5 i.MX RT1180 NETC switch port statistics counters

The NETC DSA switch driver supports `ethtool -S swpX` statistics reporting for each external DSA slave switch port through the associated net devices. Note that the first 4 stats (`tx_packets`, `tx_bytes`, `rx_packets`, `rx_bytes`) are counted by Linux networking stack and the other stats are directly read from NETC switch hardware.

Example 1: Query the port statistics counters of NETC switch port 1 (`swp1`)

```
# ethtool -S swp1
NIC statistics:
  tx_packets: 232
  tx_bytes: 30985
  rx_packets: 1178
  rx_bytes: 76913
  in-bytes: 150749
  in-valid-bytes: 150749
  in-pause-frames: 0
  in-invalid-frames: 1846
  in-vlan-frames: 0
  in-uc-frames: 814
  in-mc-frames: 85
  in-bc-frames: 947
  in-frames: 1846
[...]
```

```
  out-bytes: 381591
  out-valid-bytes: 381591
  out-pause-frames: 0
  out-valid-frames: 4539
  out-vlan-frames: 0
  out-uc-frames: 810
  out-mc-frames: 3673
  out-bc-frames: 56
  out-frames: 4539
[...]
```

```
  q0-rejected-bytes: 0
  q0-rejected-frames: 0
  q0-dequeue-bytes: 0
  q0-dequeue-frames: 0
  q0-dropped-bytes: 0
  q0-dropped-frames: 0
  q0-frames: 0
[...]
```

```
  q7-rejected-bytes: 0
  q7-rejected-frames: 0
  q7-dequeue-bytes: 0
  q7-dequeue-frames: 0
  q7-dropped-bytes: 0
  q7-dropped-frames: 0
  q7-frames: 0
```


4.3.2.6 VLAN configuration

By default, the DSA switch application running on the i.MX RT1180 EVK hardware platform configures the NETC switch as VLAN filtering enabled. The `bridge` tool from the `iproute2` package can be used to manipulate the VLAN filter table.

To make the bridge VLAN aware in Linux, run the below command to toggle the `vlan_filtering` property on the bridge that already exists.

```
# ip link set dev br0 type bridge vlan_filtering 1
```

By default, only the default `pvid` (1) of the bridge is installed on all the switch ports. So, all VLAN-tagged traffic, except that tagged with VID 1, is dropped.

Note: Since the default `pvid` (port-based VLAN) is 1, all untagged traffic also gets internally processed by the switch as having VID 1. So

- Untagged traffic is treated the same as traffic tagged with VID 1, or any other value that the `pvid` has.
- Deleting VID 1 from the VLAN table effectively blocks untagged traffic too.

Note: VLAN IDs '3072' to '3076' and '3088' are reserved by the Linux NETC DSA driver and are not allowed to be used by the user.

Example 1: To add a new VLAN filter entry by which both the NETC switch port 1 (`swp1`) and port 2 (`swp2`) can accept and transmit tagged traffic with VLAN ID 100:

```
# bridge vlan add dev swp1 vid 100
# bridge vlan add dev swp2 vid 100
```

Example 2: To delete the VLAN filter entry added in previous example.

```
# bridge vlan delete dev swp1 vid 100
# bridge vlan delete dev swp2 vid 100
```

Example 3: To display the current VLAN filter table:

```
# bridge vlan show
```

It is also possible for the switch to tag untagged traffic with a different VLAN ID on ingress using the `pvid` option.

Example 4: To add new VLAN filter entries by which both the NETC switch port 1 (`swp1`) and port 2 (`swp2`) can accept and transmit tagged traffic with VLAN ID 100 and 200, also configure the PVID of NETC switch port 1 (`swp1`) to 100 and configure the PVID of NETC switch port 2 (`swp2`) to 200.

```
# bridge vlan add dev swp1 vid 100 pvid
# bridge vlan add dev swp2 vid 100
# bridge vlan add dev swp1 vid 200
# bridge vlan add dev swp2 vid 200 pvid
```

Using the above configuration, when an untagged packet enters `swp1`, it gets internally processed by the switch as having VID 100. If it is forwarded to `swp2` as a result of FDB lookup and exits `swp2`, it is tagged with VLAN ID 100. Similarly, when an untagged packet enters `swp2`, it gets internally processed by the switch as having VID 200. If it is forwarded to `swp1` as a result of FDB lookup and exits `swp1`, it is tagged with VLAN ID 200.

In the above example, if the VLAN tag added by PVID on the original untagged packet is not desired, it can be stripped on egress using the `untagged` option.

Example 5: To strip the VLAN tag added by PVID on untagged source traffic:

```
# bridge vlan add dev swp1 vid 100 pvid untagged
# bridge vlan add dev swp2 vid 100 untagged
# bridge vlan add dev swp1 vid 200 untagged
# bridge vlan add dev swp2 vid 200 pvid untagged
```

Using the above configuration, when an untagged packet enters `swp1`, it gets internally processed by the switch as having VID 100. If it is forwarded to `swp2` as a result of FDB lookup and exits `swp2`, it is still untagged. Similarly, when an untagged packet enters `swp2`, it gets internally processed by the switch as having VID 200. If it is forwarded to `swp1` as a result of FDB lookup and exits `swp1`, it is still untagged.

4.3.2.7 FDB configuration

By default, hardware MAC learning is enabled. FDB table entries <MAC, VID, PORT> are added or updated in the FDB table when packets with new unique <MAC + VID> are received on the NETC switch port.

Example 1: To display the current FDB table entries related to NETC switch port `swp1`

```
# bridge fdb show | grep swp1
```

Example 2: To add a static FDB entry associated with NETC switch port `swp1`

```
# bridge fdb add dev swp1 11:22:33:44:55:66 vlan 1 master static
```

Note that the VLAN filter entry for the VLAN ID specified for `vlan` option must have been added before adding the FDB entry for that VLAN ID.

Example 3: To delete a static FDB entry

```
# bridge fdb del dev swp1 11:22:33:44:55:66 vlan 1 master static
```

4.3.2.8 PTP Stack

PTP driver based on DSA has been added into Real-time Edge Linux. Linuxptp stack is supported.

For the detailed information about IEEE 1588/802.1AS, refer to [Section 5.3](#).

1. Hardware setup




Figure 58. PTP Boards Connection

PTP stack only supports those i.MX RT1180 switch ports that are used as DSA CPU ports. To set up DSA, refer to [Section 4.3.1](#)

For PTP peer, any board such as LS1028ARDB, i.MX 93, or i.MX 8M Plus can be used. Select an Ethernet port that supports PTP to connect with one of i.MX RT1180 switch ports. Here it is the SWP2 port of i.MX RT1180.

2. Running DSA switch elf on i.MX RT1180

Flash image and RAM images are provided by Real-time Edge. To burn the flash image, refer to the *i.MX RT1180 Getting started User Guide*.

3. Running ptp4l on boards

The below command can be used to start the ptp4l program on DSA switch (i.MX 8M Plus / i.MX 93) side.

- The "-i interface" option specifies the PTP port.
- The "-p phc-device" option specifies the PTP hardware clock device.

In the above hardware setup, the SWP2 port of i.MX RT1180 is used for clock synchronization, so "-i swp2" is used here.

```
# ptp4l -i swp2 -p /dev/ptpX -f /etc/ptp4l_cfg/gPTP.cfg -m
```

The ethtool utility can be used to identify the PHC (PTP hardware clock) device number as shown below. The PTP Hardware Clock value displayed by *ethtool* is the index of the PTP hardware clock. It corresponds to the naming of the */dev/ptp** devices. Here, the PTP Hardware Clock is 1, which implies that "-p /dev/ptp1" must be used to specify the PHC device.

```
~# ethtool -T swp2
Time stamping parameters for swp2:
Capabilities:
    hardware-transmit
    hardware-receive
    hardware-raw-clock
PTP Hardware Clock: 1
Hardware Transmit Timestamp Modes:
    off
    on
Hardware Receive Filter Modes:
    none
    ptpv2-l2-event
```

Scenario 1: MPU + i.MX RT1180 running as PTP master clock: MPU as PTP master

Starting ptp4l program on PTP peer to be running as PTP slave clock.

- Provide the network interface with hardware time stamping capability with the "-i" option and replace the "<network interface>" with the actual network interface name, such as eth1, in the "-i <network interface>" option.

- Use the `ethtool` utility as mentioned above on PTP peer side to identify the PHC device number and replace the "X" in "`-p /dev/ptpX`" with it.

```
# ptp4l -i <network interface> -p /dev/ptpX -f /etc/ptp4l_cfg/gPTP.cfg -m -s
ptp4l[25687.144]: ioctl SIOCEHTOOL failed: Operation not supported
ptp4l[25687.144]: selected /dev/ptp1 as PTP clock
ptp4l[25687.162]: port 1: get_ts_info not supported
ptp4l[25687.208]: port 1: INITIALIZING to LISTENING on INIT_COMPLETE
ptp4l[25687.209]: port 0: INITIALIZING to LISTENING on INIT_COMPLETE
ptp4l[25690.716]: port 1: new foreign master 2ecbe2.ffff.7c9e22-1
ptp4l[25690.717]: selected best master clock 2ecbe2.ffff.7c9e22
ptp4l[25690.717]: port 1: assuming the grand master role
ptp4l[25690.717]: port 1: LISTENING to GRAND_MASTER on RS_GRAND_MASTER
```

Starting `ptp4l` program on DSA switch side to be running as PTP master clock.

```
# ptp4l -i swp2 -p /dev/ptp1 -f /etc/ptp4l_cfg/gPTP.cfg -m
ptp4l[290529.407]: selected /dev/ptp1 as PTP clock
ptp4l[290529.444]: port 1: INITIALIZING to LISTENING on INIT_COMPLETE
ptp4l[290529.444]: port 0: INITIALIZING to LISTENING on INIT_COMPLETE
ptp4l[290532.815]: port 1: LISTENING to MASTER on
ANNOUNCE_RECEIPT_TIMEOUT_EXPIRES
ptp4l[290532.815]: selected local clock 2ecbe2.ffff.7c9e22 as best master
ptp4l[290532.815]: port 1: assuming the grand master role
ptp4l[290537.820]: port 1: new foreign master 00049f.ffff.06fe8b-1
ptp4l[290537.820]: selected best master clock 00049f.ffff.06fe8b
ptp4l[290537.821]: port 1: MASTER to SLAVE on RS_SLAVE
ptp4l[290539.448]: rms 133343990721727 max 266687981443494 freq -36907 +/- 3430
delay 365 +/- 0
ptp4l[290540.451]: rms 903 max 1395 freq -34021 +/- 1209 delay 365 +/- 0
ptp4l[290541.453]: rms 1496 max 1570 freq -31334 +/- 385 delay 364 +/- 0
ptp4l[290542.454]: rms 939 max 1242 freq -30825 +/- 33 delay 365 +/- 0
ptp4l[290543.456]: rms 296 max 487 freq -31125 +/- 106 delay 364 +/- 0
ptp4l[290544.458]: rms 58 max 92 freq -31448 +/- 73 delay 364 +/- 0
ptp4l[290545.458]: rms 88 max 98 freq -31596 +/- 24 delay 364 +/- 0
ptp4l[290546.459]: rms 51 max 70 freq -31618 +/- 6 delay 364 +/- 0
ptp4l[290547.459]: rms 19 max 36 freq -31606 +/- 10 delay 363 +/- 0
ptp4l[290548.460]: rms 8 max 18 freq -31589 +/- 12 delay 361 +/- 0
ptp4l[290549.462]: rms 9 max 16 freq -31577 +/- 10 delay 361 +/- 0
ptp4l[290550.464]: rms 6 max 8 freq -31573 +/- 5 delay 363 +/- 0
ptp4l[290551.465]: rms 5 max 10 freq -31576 +/- 7 delay 365 +/- 0
ptp4l[290552.467]: rms 7 max 13 freq -31583 +/- 8 delay 365 +/- 0
ptp4l[290553.468]: rms 6 max 11 freq -31583 +/- 8 delay 365 +/- 0
ptp4l[290554.468]: rms 6 max 11 freq -31572 +/- 5 delay 364 +/- 0
ptp4l[290555.469]: rms 8 max 14 freq -31570 +/- 10 delay 364 +/- 0
ptp4l[290556.469]: rms 7 max 13 freq -31574 +/- 9 delay 364 +/- 0
ptp4l[290557.470]: rms 6 max 10 freq -31573 +/- 8 delay 364 +/- 0
ptp4l[290558.471]: rms 6 max 12 freq -31570 +/- 8 delay 364 +/- 0
ptp4l[290559.472]: rms 5 max 9 freq -31569 +/- 7 delay 364 +/- 0
ptp4l[290560.472]: rms 4 max 7 freq -31576 +/- 3 delay 364 +/- 0
ptp4l[290561.473]: rms 5 max 11 freq -31577 +/- 6 delay 363 +/- 0
ptp4l[290562.473]: rms 6 max 12 freq -31569 +/- 7 delay 363 +/- 0
```

Scenario 2: MPU + i.MX RT1180 running as PTP slave clock

Start `ptp4l` program on PTP peer side to be running as PTP master clock.

- Providing the network interface with hardware time stamping capability with the "-i" option and replace the "<network interface>" with actual network interface name, such as `eth1`, in the "-i <network interface>" option.

- Use the ethtool utility as mentioned above on PTP peer side to identify the PHY device number and replace the "X" in "-p /dev/ptpX" with it.

```
# ptp4l -i <network interface> -p /dev/ptpX -f /etc/ptp4l_cfg/gPTP.cfg -m
ptp4l[290699.108]: selected /dev/ptp1 as PTP clock
ptp4l[290699.144]: port 1: INITIALIZING to LISTENING on INIT_COMPLETE
ptp4l[290699.144]: port 0: INITIALIZING to LISTENING on INIT_COMPLETE
ptp4l[290702.668]: port 1: LISTENING to MASTER on
ANNOUNCE_RECEIPT_TIMEOUT_EXPIRES
ptp4l[290702.668]: selected local clock 2ecbe2.ffff.7c9e22 as best master
ptp4l[290702.668]: port 1: assuming the grand master role
```

Starting ptp4l program on MPU side to be running as PTP slave clock. DSA switch

```
# ptp4l -i swp2 -p /dev/ptp1 -f /etc/ptp4l_cfg/gPTP.cfg -m -s
ptp4l[25866.504]: ioctl SIOCEHTOOL failed: Operation not supported
ptp4l[25866.504]: selected /dev/ptp1 as PTP clock
ptp4l[25866.536]: port 1: get_ts_info not supported
ptp4l[25866.580]: port 1: INITIALIZING to LISTENING on INIT_COMPLETE
ptp4l[25866.580]: port 0: INITIALIZING to LISTENING on INIT_COMPLETE
ptp4l[25870.400]: selected local clock 00049f.ffff.06fe8b as best master
ptp4l[25870.566]: port 1: new foreign master 2ecbe2.ffff.7c9e22-1
ptp4l[25870.566]: selected best master clock 2ecbe2.ffff.7c9e22
ptp4l[25870.566]: port 1: LISTENING to UNCALIBRATED on RS_SLAVE
ptp4l[25870.835]: port 1: UNCALIBRATED to SLAVE on MASTER_CLOCK_SELECTED
ptp4l[25871.587]: rms 1685 max 2235 freq -1906 +/- 1164 delay 362 +/- 0
ptp4l[25872.588]: rms 311 max 477 freq -893 +/- 421 delay 363 +/- 0
ptp4l[25873.589]: rms 525 max 552 freq +60 +/- 139 delay 362 +/- 0
ptp4l[25874.585]: rms 326 max 437 freq +233 +/- 11 delay 363 +/- 0
ptp4l[25875.591]: rms 107 max 179 freq +135 +/- 42 delay 362 +/- 0
ptp4l[25876.592]: rms 19 max 28 freq +20 +/- 23 delay 362 +/- 0
ptp4l[25877.592]: rms 31 max 35 freq -31 +/- 12 delay 362 +/- 0
ptp4l[25878.592]: rms 18 max 31 freq -39 +/- 7 delay 362 +/- 0
ptp4l[25879.592]: rms 6 max 10 freq -30 +/- 6 delay 363 +/- 0
ptp4l[25880.593]: rms 4 max 6 freq -25 +/- 6 delay 364 +/- 0
ptp4l[25881.591]: rms 3 max 6 freq -27 +/- 4 delay 365 +/- 0
ptp4l[25882.593]: rms 4 max 6 freq -23 +/- 4 delay 364 +/- 0
ptp4l[25883.594]: rms 4 max 5 freq -29 +/- 4 delay 365 +/- 0
ptp4l[25884.599]: rms 4 max 7 freq -22 +/- 4 delay 365 +/- 0
ptp4l[25885.592]: rms 4 max 6 freq -27 +/- 5 delay 365 +/- 0
ptp4l[25886.597]: rms 3 max 7 freq -25 +/- 5 delay 366 +/- 0
ptp4l[25887.597]: rms 5 max 7 freq -27 +/- 7 delay 366 +/- 0
ptp4l[25888.598]: rms 3 max 5 freq -26 +/- 4 delay 366 +/- 0
ptp4l[25889.602]: rms 5 max 9 freq -22 +/- 6 delay 365 +/- 0
ptp4l[25890.601]: rms 3 max 6 freq -26 +/- 3 delay 364 +/- 0
ptp4l[25891.605]: rms 5 max 9 freq -22 +/- 7 delay 364 +/- 0
ptp4l[25892.604]: rms 5 max 7 freq -31 +/- 3 delay 365 +/- 0
ptp4l[25893.606]: rms 7 max 13 freq -18 +/- 6 delay 365 +/- 0
ptp4l[25894.608]: rms 4 max 7 freq -26 +/- 5 delay 365 +/- 0
ptp4l[25895.604]: rms 5 max 9 freq -24 +/- 7 delay 365 +/- 0
ptp4l[25896.602]: rms 5 max 9 freq -27 +/- 7 delay 365 +/- 0
ptp4l[25897.609]: rms 5 max 7 freq -28 +/- 6 delay 365 +/- 0
ptp4l[25898.611]: rms 5 max 9 freq -25 +/- 7 delay 365 +/- 0
ptp4l[25899.602]: rms 3 max 6 freq -22 +/- 4 delay 365 +/- 0
```

4.3.2.9 802.1CB usage

4.3.2.9.1 Network configuration

Use the below steps to configure the 802.1CB usage.

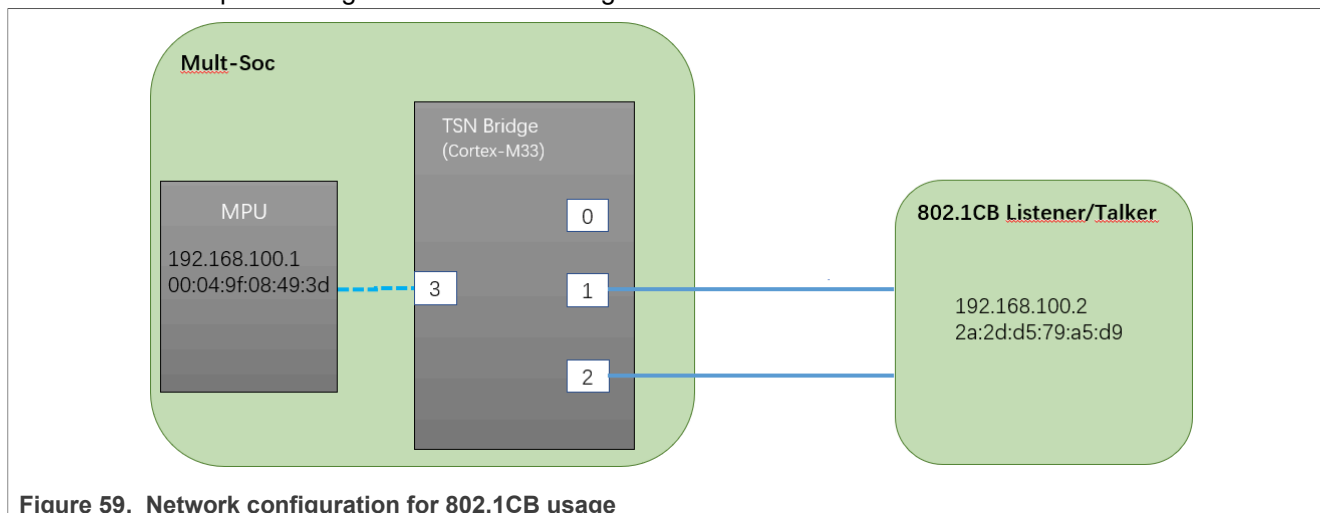


Figure 59. Network configuration for 802.1CB usage

1. Add switch ports into the bridge:

```
ip link add name br0 type bridge vlan_filtering 1
ip link set br0 up
ip link set swp0 master br0 && ip link set swp0 up
ip link set swp1 master br0 && ip link set swp1 up
ip link set swp2 master br0 && ip link set swp2 up
```

2. Connect swp1 to the network, add IP on the br0 and test the connection:

```
ifconfig br0 192.168.100.1
ping 192.168.100.2
```

3. Add destination MAC in FDB if the MAC is not learned:

```
bridge fdb add dev swp1 2a:2d:d5:79:a5:d9 vlan 1 master static
```

4. Drop streams to prevent loop back

```
tc qdisc add dev swp1 clsact
tc qdisc add dev swp2 clsact
tc filter add dev swp1 ingress protocol 802.1Q flower skip_sw \
src_mac 00:04:9f:08:49:3d vlan_id 0 action police mtu 1 \
conform-exceed drop/ok
tc filter add dev swp2 ingress protocol 802.1Q flower skip_sw \
src_mac 00:04:9f:08:49:3d vlan_id 0 action police mtu 1 \
conform-exceed drop/ok
tc filter add dev swp1 ingress protocol 802.1Q flower skip_sw \
src_mac 00:aa:bb:cc:dd:11 vlan_id 0 action police mtu 1 \
conform-exceed drop/ok
tc filter add dev swp2 ingress protocol 802.1Q flower skip_sw \
src_mac 00:aa:bb:cc:dd:11 vlan_id 0 action police mtu 1 \
conform-exceed drop/ok
```

Note: “00:04:9f:08:49:3d” is the local bridge MAC on Linux. “00:aa:bb:cc:dd:11” is the local pseudo MAC on i.MX RT1180.

5. Connect swp2 to the network to obtain redundancy paths.

4.3.2.9.2 802.1CB sequence generator

1. Add 802.1CB sequence generator:

```
tc qdisc add dev swp1 clsact
tc filter add dev swp1 egress protocol 802.1Q flower skip_sw \
dst_mac 2a:2d:d5:79:a5:d9 vlan_id 1 action frer rtag tag-action tag-push
```

2. Split the stream to swp2

```
tc filter add dev swp1 egress protocol 802.1Q flower skip_sw \
dst_mac 2a:2d:d5:79:a5:d9 vlan_id 1 action mirrored egress mirror dev swp2
```

3. Check the R-tag in listener side # tcpdump -i eno0

```
tcpdump: verbose output suppressed, use -v[v]... for full protocol decode
listening on eno0, link-type EN10MB (Ethernet), snapshot length 262144 bytes
00:31:27.536183 00:04:9f:08:49:3d (oui Freescale) > 2a:2d:d5:79:a5:d9 (oui
Unknown), ethertype Unknown (0xf1c1), length 104:
 0x0000: 0000 0000 0800 4500 0054 4f2a 4000 4001 .....E..TO*@.@.
 0x0010: a229 c0a8 6401 c0a8 6403 0800 7a3c 0003 ..).d...d...z<..
 0x0020: 0001 bf55 a866 0000 0000 4a30 0d00 0000 ...U.f....J0....
 0x0030: 0000 1011 1213 1415 1617 1819 1a1b 1c1d .....
 0x0040: 1e1f 2021 2223 2425 2627 2829 2a2b 2c2d ...!"#$%&'()*+,-
 0x0050: 2e2f 3031 3233 3435 3637 ...../01234567
00:31:28.556234 00:04:9f:08:49:3d (oui Freescale) > 2a:2d:d5:79:a5:d9 (oui
Unknown), ethertype Unknown (0xf1c1), length 104:
 0x0000: 0000 0001 0800 4500 0054 4fe6 4000 4001 .....E..TO.@.@.
 0x0010: a16d c0a8 6401 c0a8 6403 0800 8cec 0003 ..m..d...d.....
 0x0020: 0002 c055 a866 0000 0000 367f 0d00 0000 ...U.f....6.....
 0x0030: 0000 1011 1213 1415 1617 1819 1a1b 1c1d .....
 0x0040: 1e1f 2021 2223 2425 2627 2829 2a2b 2c2d ...!"#$%&'()*+,-
 0x0050: 2e2f 3031 3233 3435 3637 ...../01234567
```

4. Drop streams to prevent loop back

```
tc qdisc add dev swp1 clsact
tc qdisc add dev swp2 clsact
tc filter add dev swp1 ingress protocol 802.1Q flower skip_sw \
src_mac 00:04:9f:08:49:3d vlan_id 0 action police mtu 1 \
conform-exceed drop/ok
tc filter add dev swp2 ingress protocol 802.1Q flower skip_sw \
src_mac 00:04:9f:08:49:3d vlan_id 0 action police mtu 1 \
conform-exceed drop/ok
tc filter add dev swp1 ingress protocol 802.1Q flower skip_sw \
src_mac 00:aa:bb:cc:dd:11 vlan_id 0 action police mtu 1 \
conform-exceed drop/ok
tc filter add dev swp2 ingress protocol 802.1Q flower skip_sw \
src_mac 00:aa:bb:cc:dd:11 vlan_id 0 action police mtu 1 \
conform-exceed drop/ok
```

5. Connect swp2 to the network to obtain redundancy paths.

4.3.2.9.3 802.1CB sequence recovery

1. Add 802.1CB sequence recovery:

```
tc filter add dev swp1 ingress protocol 802.1Q flower skip_sw \
dst_mac 00:04:9f:08:49:3d vlan_id 0 action frer rtag recover alg \
```

```
vector history-length 32 reset-time 10000 tag-action tag-pop
```

2. Ping on the remote terminal and check the packets received on MPU side:

```
# tcpdump -i br0
04:04:01.030771 IP aontst11.nl-nym01.nxp.com > imx93evk: ICMP echo request,
  id 51345, seq 2, length 64
04:04:01.030873 IP imx93evk > aontst11.nl-nym01.nxp.com: ICMP echo reply, id
  51345, seq 2, length 64
04:04:02.054815 IP aontst11.nl-nym01.nxp.com > imx93evk: ICMP echo request,
  id 51345, seq 3, length 64
04:04:02.054914 IP imx93evk > aontst11.nl-nym01.nxp.com: ICMP echo reply, id
  51345, seq 3, length 64
```

4.3.2.10 802.1 Qci usage

4.3.2.10.1 Network configuration

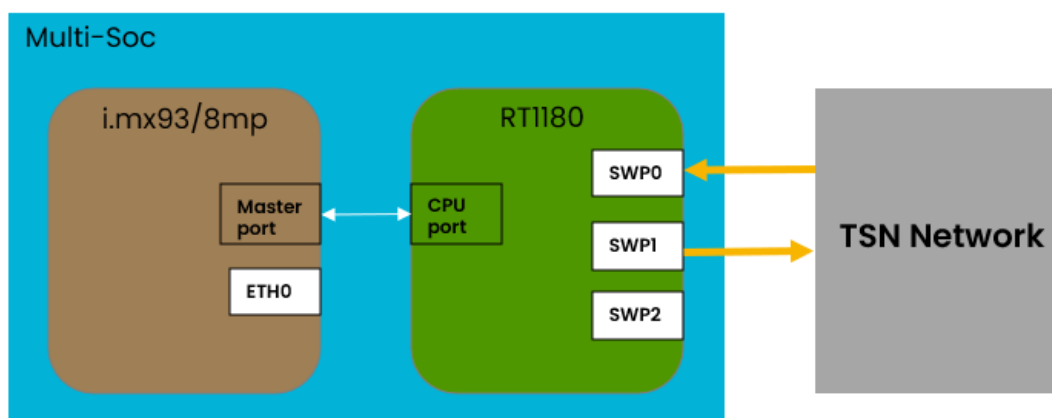


Figure 60. 802.1 Qci usage

Bridge mode

```
ip link add name br0 type bridge vlan_filtering 1
ip link set br0 up
ip link set swp0 master br0 && ip link set swp0 up
ip link set swp1 master br0 && ip link set swp1 up
ip link set swp2 master br0 && ip link set swp2 up
```

Standalone mode

```
ifconfig swp0 up
```

4.3.2.10.2 802.1 Qci configuration

1. Set Qci gate

```
tc qdisc add dev swp0 clsact
tc filter add dev swp0 ingress protocol 802.1Q flower skip_sw \
dst_mac 00:04:9f:08:49:3d vlan_id 0 action gate index 1 \
```



```
base-time 0 priority 3 sched-entry OPEN 6000 -1 -1 \
sched-entry CLOSE 6000 -1 -1
```

2. Set Qci flow meter

```
tc filter add dev swp0 ingress protocol 802.1Q flower skip_sw \
dst_mac 00:04:9f:08:49:3d vlan_id 0 action police index 1 \
rate 10Mbit burst 10000 conform-exceed drop/ok
```

3. Set Qci SFI priority

```
tc filter add dev swp0 ingress protocol 802.1Q flower skip_sw \
dst_mac 00:04:9f:08:49:3d vlan_id 0 vlan_prio 1 action gate \
index 1 base-time 0 priority 3 sched-entry CLOSE 6000 -1 -1
```

4. Set Qci SFI maxsdu

```
tc filter add dev swp0 ingress protocol 802.1Q flower skip_sw \
dst_mac 00:04:9f:08:49:3d vlan_id 0 vlan_prio 1 action police \
mtu 1000 conform-exceed drop/ok
```

5. Set both gate and flow meter

```
tc filter add dev swp0 ingress protocol 802.1Q flower skip_sw \
dst_mac 00:04:9f:08:49:3d vlan_id 0 action gate index 1 \
base-time 0 priority 3 sched-entry OPEN 6000 -1 -1 action police \
index 1 rate 10Mbit burst 10000 conform-exceed drop/ok
```

6. Show the Qci rules configured on swp0:

```
tc -s filter show dev swp0 ingress
```

7. Delete one of the Qci rule:

```
tc filter del dev swp0 ingress pref 49152
```

4.3.2.11 802.1 Qbv usage

4.3.2.11.1 Network configuration

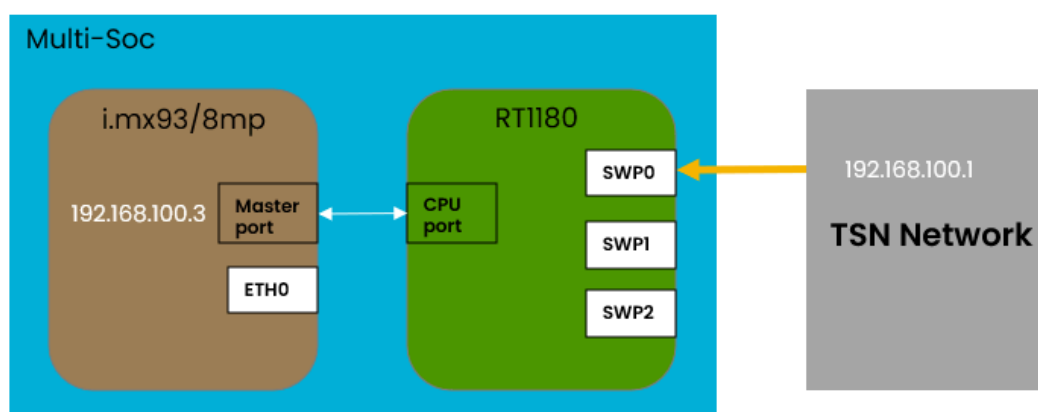


Figure 61. Network configuration for 802.1 Qbv usage

TSN 802.1 Qbv is an IEEE specification that allows to define transmission time slots for networking traffic queues. It adds enhancements to traffic scheduling with the introduction of time-aware shaper. For DSA switch,

both bridge mode and single port mode support TSN 802.1 Qbv setting. The following Qbv configuration example takes the network card swp0 as an example. Refer to the following command to select a work mode for swp0 - bridge mode or single port mode.

Bridge mode

```
ip link add name br0 type bridge vlan_filtering 1
ip link set br0 up
ip link set swp0 master br0 && ip link set swp0 up
ip link set swp1 master br0 && ip link set swp1 up
ip link set swp2 master br0 && ip link set swp2 up
ifconfig br0 192.168.100.3 up
```

Single port mode

```
ifconfig swp0 192.168.100.3 up
```

4.3.2.11.2 Configuring 802.1 Qbv via tc

1. Open the gate

```
tc qdisc replace dev swp0 parent root handle 100 taprio num_tc 8 \
map 0 1 2 3 4 5 6 7 queues 100 101 102 103 104 105 106 107 \
base-time 0 sched-entry S 01 300000 sched-entry S 01 300000 flags 0x2
```

2. Ping remote IP successfully

```
ping 192.168.100.1
```

3. Close the gate

```
tc qdisc replace dev swp0 parent root handle 100 taprio num_tc 8 \
map 0 1 2 3 4 5 6 7 queues 100 101 102 103 104 105 106 107 \
base-time 0 sched-entry S 01 300000 sched-entry S fe 300000 flags 0x2
```

4. ping remote IP failed

```
ping 192.168.100.1
```

5. Disable Qbv configuration

```
tc qdisc del dev swp0 parent root handle 100
```

4.3.2.11.3 Configuring 802.1 Qbv via netconf

1. For installing Netopeer2-cli on Ubuntu22.04, refer to "[Section 6.4.2.2](#)".
2. Create Qbv configuration file qbv.xml

```
<interfaces xmlns="urn:ietf:params:xml:ns:yang:ietf-interfaces"
  xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0"
  xmlns:dot1q="urn:ieee:std:802.1Q:yang:ieee802-dot1q-bridge"
  xmlns:sched="urn:ieee:std:802.1Q:yang:ieee802-dot1q-sched"
  xmlns:sched-bridge="urn:ieee:std:802.1Q:yang:ieee802-dot1q-sched-bridge">
  <interface>
    <name>swp0</name>
```

```

    <enabled>true</enabled>
    <type xmlns:ianaift="urn:ietf:params:xml:ns:yang:iana-if-
type">ianaift:ethernetCsmacd</type>
    <dot1q:bridge-port>
      <sched-bridge:gate-parameter-table>
        <sched-bridge:queue-max-sdu-table>
          <sched-bridge:traffic-class>0</sched-bridge:traffic-class>
          <sched-bridge:queue-max-sdu>1024</sched-bridge:queue-max-sdu>
        </sched-bridge:queue-max-sdu-table>
        <sched-bridge:gate-enabled>true</sched-bridge:gate-enabled>
        <sched-bridge:admin-gate-states>127</sched-bridge:admin-gate-
states>
        <sched-bridge:admin-control-list>
          <sched-bridge:gate-control-entry>
            <sched-bridge:index>0</sched-bridge:index>
            <sched-bridge:operation-name>sched:set-gate-states</
sched-bridge:operation-name>
            <sched-bridge:gate-states-value>5</sched-bridge:gate-
states-value>
            <sched-bridge:time-interval-value>400000</sched-
bridge:time-interval-value>
          </sched-bridge:gate-control-entry>
          <sched-bridge:gate-control-entry>
            <sched-bridge:index>1</sched-bridge:index>
            <sched-bridge:operation-name>sched:set-gate-states</
sched-bridge:operation-name>
            <sched-bridge:gate-states-value>7</sched-bridge:gate-
states-value>
            <sched-bridge:time-interval-value>300000</sched-
bridge:time-interval-value>
          </sched-bridge:gate-control-entry>
          <sched-bridge:gate-control-entry>
            <sched-bridge:index>2</sched-bridge:index>
            <sched-bridge:operation-name>sched:set-gate-states</
sched-bridge:operation-name>
            <sched-bridge:gate-states-value>7</sched-bridge:gate-
states-value>
            <sched-bridge:time-interval-value>300000</sched-
bridge:time-interval-value>
          </sched-bridge:gate-control-entry>
        </sched-bridge:admin-control-list>
        <sched-bridge:config-change>true</sched-bridge:config-change>
        <sched-bridge:admin-cycle-time>
          <sched-bridge:numerator>1</sched-bridge:numerator>
          <sched-bridge:denominator>10000</sched-bridge:denominator>
        </sched-bridge:admin-cycle-time>
        <sched-bridge:admin-base-time>
          <sched-bridge:seconds>0</sched-bridge:seconds>
          <sched-bridge:nanoseconds>5000</sched-bridge:nanoseconds>
        </sched-bridge:admin-base-time>
        <sched-bridge:supported-list-max>50</sched-bridge:supported-list-
max>
        <sched-bridge:supported-cycle-max>
          <sched-bridge:numerator>1</sched-bridge:numerator>
          <sched-bridge:denominator>1</sched-bridge:denominator>
        </sched-bridge:supported-cycle-max>
        <sched-bridge:supported-interval-max>10000000000</sched-
bridge:supported-interval-max>
      </sched-bridge:gate-parameter-table>
    </dot1q:bridge-port>

```

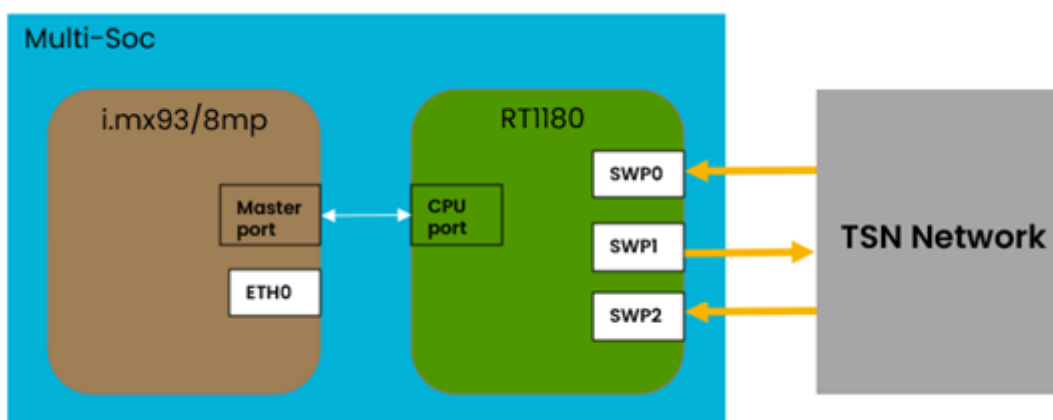
```
</interface>
</interfaces>
```

3. Configure Qbv by netopeer2-cli

```
$ netopeer2-cli
> connect --login root --host 192.168.1.1
> edit-config --target running --config=qbv.xml
```

4.3.2.12 802.1 Qbu usage

4.3.2.12.1 Network configuration



1. Bridge mode

```
ip link add name br0 type bridge vlan_filtering 1
ip link set br0 up
ip link set swp0 master br0 && ip link set swp0 up
ip link set swp1 master br0 && ip link set swp1 up
ip link set swp2 master br0 && ip link set swp2 up
ifconfig br0 192.168.100.3 up
```

4.3.2.12.2 Qbu configuration

1. Enable pmac

```
ethtool --set-mm swp0 tx-enabled on pmac-enabled on verify-enabled off \
tx-min-frag-size 60
ethtool --set-mm swp1 tx-enabled on pmac-enabled on verify-enabled off \
tx-min-frag-size 60
ethtool --set-mm swp2 tx-enabled on pmac-enabled on verify-enabled off \
tx-min-frag-size 60
```

2. Configure preemption queues

```
tc qdisc add dev swp1 root handle 1: mqprio num_tc 8 map 0 1 2 3 4 5 6 7
queues 1@0 1@1 1@2 1@3 1@4 1@5 1@6 1@7 fp E P E E E P E hw 1
```

3. Send packets from TSN network

Send priority 0 and priority 1 traffic on swp0 and swp2. The preemption queue is then preempted, and fragment packets are obtained on swp1.

5 Real-time networking

5.1 Time Sensitive Networking (TSN) on NXP platforms

Time Sensitive Networking (TSN) is an extension to traditional Ethernet networks, providing a set of standards compatible with IEEE 802.1 and 802.3. These extensions intend to address the limitations of standard Ethernet in sectors ranging from industrial and automotive applications to live audio and video systems. Applications running over traditional Ethernet must be designed to be very robust in order to withstand corner cases such as packet loss, delay, or even reordering. TSN aims to provide guarantees for deterministic latency and packet loss under congestion. Therefore, it allows critical and non-critical traffic to be converged in the same network.

This chapter describes the process and use cases for implementing TSN features on the i.MX 8M LPDDR4 Plus, i.MX 8DXL LPDDR4 EVK, i.MX 93 EVK, i.MX 93 9x9 QSB, i.MX 93 14x14 EVK, and LS1028ARDB boards.

5.1.1 TSN hardware capability

Table 46. TSN hardware capability on different platforms

Platform	802.1Qbv (Enhancements for Scheduled Traffic)	802.1Qbu and 802.3br (Frame Preemption)	802.1Qav (Credit Based Shaper)	802.1AS (Precision Time Protocol)	802.1CB (Frame Replication and Elimination for Reliability)	802.1Qci (Per Stream Filtering and Policing)
ENETC (LS1028 A)	Y	Y	Y	Y	N	Y
Felix switch (LS1028A)	Y	Y	Y	Y	Y	Y
Stmac (i.MX 8DXL, i.MX 8M Plus, i.MX 93)	Y	Y	Y	Y	N	N

5.1.2 TSN configuration

The table below describes the TSN configuration tools support on different platforms

Table 47. TSN configuration tool support on different hardware platforms

Platform	802.1Qbv (Enhancements for Scheduled Traffic)	802.1Qbu and 802.3br (Frame Preemption)	802.1Qav (Credit Based Shaper)	802.1AS (Precision Time Protocol)	802.1CB (Frame Replication and Elimination for Reliability)	802.1Qci (Per Stream Filtering and Policing)
ENETC (LS1028A)	tc-taprio tsntool	ethtool tsntool	tc-cbs tsntool	ptp4l	N/A	tc-flower tsntool
Felix switch (LS1028A)	tc-taprio tsntool	ethtool tsntool	tc-cbs tsntool	ptp4l, Gen AVB/TSN stack	tsntool	tc-flower tsntool
Stmac (i.MX 8DXL, i.MX 8M Plus, i.MX 93)	tc-taprio	ethtool	tc-cbs	ptp4l, Gen AVB/TSN stack	N/A	N/A

5.1.2.1 Using Linux traffic control (tc) to configure TSN

Enable the following configurations in kernel when using Linux traffic control (tc):

```
Symbol: NET_SCH_MQPRIO [=y] && NET_SCH_CBS [=y] && NET_SCH_TAPRIO [=y]
[*] Networking support --->
  Networking options --->
    [*] QoS and/or fair queueing --->
      <*> Credit Based Shaper (CBS)
      <*> Time Aware Priority (taprio) Scheduler
      <*> Multi-queue priority scheduler (MQPRIO)
    [*] Actions --->
      <*> Traffic Policing
      <*> Generic actions
      <*> Redirecting and Mirroring
      <*> SKB Editing
      <*> Vlan manipulation
      <*> Frame gate entry list control tc action
```

On LS1028A platform, ENETC QoS driver needs to be set to support tc configuration.

```
Symbol: FSL_ENETC_QOS [=y]
Device Drivers--->
  [*] Network device support --->
    [*] Ethernet driver support --->
      [*] Freescale devices
      [*] ENETC hardware Time-sensitive Network support
```

1. The below link provides details for using tc-taprio to set Qbv:

<https://man7.org/linux/man-pages/man8/tc-taprio.8.html>

2. The below link provides details for using tc-cbs to set Qav:

<https://man7.org/linux/man-pages/man8/tc-cbs.8.html>

3. The below link provides details for using tc-flower to set Qci and ACL:

<https://man7.org/linux/man-pages/man8/tc-flower.8.html>

5.1.2.2 Using Tstntool to configure TSN

Tstntool is a tool to set the TSN capability of the Ethernet ports of TSN Endpoint and TSN switch. It is used on LS1028A platform. You should enable TSN, ENETC_TSN, and MSCC_FELIX_SWITCH_TSN to support tstntool configuration on LS1028A.

```
Symbol: TSN [=y]
  [*] Networking support --->
    Networking options --->
      [*] 802.1 Time-Sensitive Networking support
Symbol: ENETC_TSN [=y] && FSL_ENETC_PTP_CLOCK [=y] && FSL_ENETC_HW_TIMESTAMPING [=y]
Device Drivers --->
  [*] Network device support --->
    [*] Ethernet driver support --->
      [*] Freescale devices
      <*> ENETC PF driver
      <*> ENETC VF driver
      -*- ENETC MDIO driver
      <*> ENETC PTP clock driver
```

```

[*]      ENETC hardware timestamping support
[*]      TSN Support for NXP ENETC driver
Symbol: MSCC_FELIX_SWITCH_TSN [=y]
Device Drivers --->
  [*] Network device support --->
    Distributed Switch Architecture drivers --->
      <*> Ocelot / Felix Ethernet switch support --->
      <*> TSN on FELIX switch driver

```

Enable PKTGEN in the kernel to use pktgen for testing using the commands below:

```

Symbol: NET_PKTGEN [=y]
[*] Networking support --->
  Networking options --->
    Network testing --->
      <*> Packet Generator (USE WITH CAUTION)

```

Refer to *Tsntool User Manual* for the details.

5.1.2.2.1 Tsntool User Manual

Tsntool is a tool to set the TSN capability of the Ethernet ports of TSN Endpoint and TSN switch. This document describes how to use tsntool for NXP's LS1028ARDB hardware platform.

Note: *Tsntool supports only the LS1028ARDB platform.*

5.1.2.2.1.1 Getting the source code

Github of the tsntool code is mentioned below.

<https://github.com/nxp-qorik/tsntool>

5.1.2.2.1.2 Tsn tool commands

The [Table 48](#) lists the TSN tool commands and their description.

Table 48. TSN tool commands and their description

Command	Description
help	Lists commands support
version	Shows software version
verbose	Debugs on/off for tsntool
quit	Quits prompt mode
qbvset	Sets time gate scheduling config for <ifname>
qbvget	Gets time scheduling entries for <ifname>
cbstreamidset	Sets stream identification table
cbstreamidget	Gets stream identification table and counters
qcisfiset	Sets stream filter instance
qcisfiget	Gets stream filter instance
qcisgiset	Sets stream gate instance
qcisgiget	Gets stream gate instance

Table 48. TSN tool commands and their description...continued

Command	Description
qcisficounterget	Gets stream filter counters
qcifmisset	Sets flow metering instance
qcifmiget	Gets flow metering instance
cbsset	Sets TCs credit-based shaper configuration
cbsget	Gets TCs credit-based shaper status
qbuset	Sets one 8-bits vector showing the preemptable traffic class
qbugetstatus	Not supported
tsdset	Not supported
tsdget	Not supported
ctset	Sets cut through queue status (specific for Is1028 switch)
cbgen	Sets sequence generate configure (specific for Is1028 switch)
cbrec	Sets sequence recover configure (specific for Is1028 switch)
dscpset	Sets queues map to DSCP of Qos tag (specific for Is1028 switch)
sendpkt	Not supported
regtool	Registers read/write of bar0 of PFs (specific for Is1028 enetc)
ptptool	ptptool get/set ptp timestamp. Useful commands: #get ptp0 clock time ptptool -g #get ptp1 clock time ptptool -g -d /dev/ptp1
dscpset	Set queues map to DSCP of QoS tag (specific for Is1028 switch)
qcicapget	Gets max capability of the qci instance
tsncapget	Gets tsn capability of the device

5.1.2.2.1.3 Tsntool commands and parameters

This section lists the tsntool commands along with the parameters and arguments, with which they can be used.

[Table 49](#) lists the qbvset parameter and its arguments.

Table 49. qbvset

Parameter <argument>	Description
--device <ifname>	An interface such as eno0/swp0.
--entryfile <filename>	A file script to input gatelist format. It has the following arguments: # 'NUMBER' 'GATE_VALUE' 'TIME_LONG' <ul style="list-style-type: none"> NUMBER: # 't' or 'T' head. Plus entry number. Duplicate entry number will result in an error. GATE_VALUE: # format: xxxxxxxb . # The MSB corresponds to traffic class 7. The LSB corresponds to traffic class 0. # A bit value of 0 indicates closed, whereas, a bit value of 1 indicates open. TIME_LONG: # nanoseconds. Do not input 0 time long. t0 11101111b 10000 t1 11011111b 10000 Note: Entryfile parameter must be set. If not set, there will be a vi text editor prompt, "require to input the gate list".
--basetime <value>	AdminBaseTime A 64-bit hex value means nanosecond until now.

Table 49. qbvset...continued

Parameter <argument>	Description
	OR a value input format as: Seconds.decimalSecond Example: 115.000125 means 115 seconds and 125 µs.
--cycletime <value>	AdminCycleTime
--cycleextend <value>	AdminCycleTimeExtension
--enable --disable	<ul style="list-style-type: none"> enable: enables the qbv for this port. disable: disables the qbv for this port. By default, the value is set to enable, if user does not provide any input.
--maxsdu <value>	queueMaxSDU
--initgate <value>	AdminGateStates
--configchange	ConfigChange. Default set to 1.
--configchangetime <value>	ConfigChangeTime

Table 50. qbvget

Parameter <argument>	Description
--device <ifname>	An interface such as eno0/swp0

Table 51. cbstreamidset

Parameter <argument>	Description
--enable --disable	<ul style="list-style-type: none"> enable: Enables the entry for this index. disable: Disables the entry for this index. By default, this field is set to enable if there is no enable or disable input.
--index <value>	Index entry number in this controller. Mandatory parameter. This value corresponds to tsnStreamIdHandle on switch configuration.
--device <string>	An interface such as eno0/swp0
--streamhandle <value>	tsnStreamIdHandle
--infacoutport <value>	tsnStreamIdInFacOutputPortList
--outfacoutport <value>	tsnStreamIdOutFacOutputPortList
--infacinport <value>	tsnStreamIdInFacInputPortList
--outfacinport <value>	tsnStreamIdOutFacInputPortList
--nullstreamid - -sourcemacvid - -destmacvid -- ipstreamid	tsnStreamIdIdentificationType: <ul style="list-style-type: none"> -nullstreamid: Null Stream identification -sourcemacvid: Source MAC and VLAN Stream identification -destmacvid: not supported -ipstreamid: not supported
--nulldmac <value>	tsnCpeNullDownDestMac

Table 51. `cbstreamidset` ...continued

Parameter <argument>	Description
<code>--nulltagged <value></code>	<code>tsnCpeNullDownTagged</code>
<code>--nullvid <value></code>	<code>tsnCpeNullDownVlan</code>
<code>--sourcemac <value></code>	<code>tsnCpeSmacVlanDownSrcMac</code>
<code>--sourcetagged <value></code>	<code>tsnCpeSmacVlanDownTagged</code>
<code>--sourcevid <value></code>	<code>tsnCpeSmacVlanDownVlan</code>

Table 52. `cbstreamidset`

Parameter <argument>	Description
<code>--device <ifname></code>	An interface such as <code>eno0/swp0</code>
<code>--index <value></code>	Index entry number in this controller. Mandatory to have.

Table 53. `qcisfiset`

Parameter <argument>	Description
<code>--device <ifname></code>	An interface such as <code>eno0/swp0</code>
<code>--enable --disable</code>	<ul style="list-style-type: none"> <code>enable</code>: enable the entry for this index <code>disable</code>: disable the entry for this index By default, this field is set to <code>enable</code> if there is no enable or disable input.
<code>--maxsdu <value></code>	Maximum SDU size.
<code>--flowmeterid <value></code>	Flow meter instance identifier index number.
<code>--index <value></code>	StreamFilterInstance. index entry number in this controller. This value corresponds to <code>tsnStreamIdHandle</code> of <code>cbstreamidset</code> command on switch configuration.
<code>--streamhandle <value></code>	StreamHandleSpec This value corresponds to <code>tsnStreamIdHandle</code> of <code>cbstreamidset</code> command.
<code>--priority <value></code>	PrioritySpec
<code>--gateid <value></code>	StreamGateInstanceID
<code>--oversizeenable</code>	StreamBlockedDueToOversizeFrameEnable
<code>--oversize</code>	StreamBlockedDueToOversizeFrame

Table 54. `qcisfiget`

parameter <argument>	Description
<code>--device <ifname></code>	An interface such as <code>eno0/swp0</code>
<code>--index <value></code>	Index entry number in this controller. Mandatory to have.

Table 55. qcisgiset

Parameter <argument>	Description
--device <ifname>	An interface such as eno0/swp0
--index <value>	Index entry number in this controller. Mandatory to have.
--enable --disable	<ul style="list-style-type: none"> enable: enable the entry for this index. PSFPGateEnabled disable: disable the entry for this index. <p>By default, this field is set to enable if there is no enable or disable input.</p>
--configchange	configchange
--enblkinvr	PSFPGateClosedDueToInvalidRxEnable
--blkinvr	PSFPGateClosedDueToInvalidRx
--initgate	PSFPAdminGateStates
--initip	AdminIPV
--cycletime	Default not set. Get by gatelistfile.
--cycletimeext	PSFPAdminCycleTimeExtension
--basetime	<p>PSFPAdminBaseTime</p> <p>A 64-bit hex value means nanosecond until now.</p> <p>OR a value input format as: Seconds.decimalSecond</p> <p>Example: 115.000125means 115 seconds and 125 µs.</p>
--gatelistfile	<p>PSFPAdminControlList. A file input the gate list: 'NUMBER' 'GATE_VALUE' 'IPV' 'TIME_LONG' 'OCTET_MAX'</p> <ul style="list-style-type: none"> NUMBER: # 't' or 'T' head. Plus entry number. Duplicate entry number will result in an error. GATE_VALUE: format: xb: The MSB corresponds to traffic class 7. The LSB corresponds to traffic class 0. A bit value of 0 indicates closed, A bit value of 1 indicates open. IPV: # 0~7 TIME_LONG: in nanoseconds. Do not input time long as 0. OCTET_MAX: The maximum number of octets that are permitted to pass the gate. If zero, there is no maximum. <p>Example:</p> <pre>t0 1b -1 50000 10</pre>

Table 56. qcisgiget

Parameter <argument>	Description
--device <ifname>	An interface such as eno0/swp0
--index <value>	Index entry number in this controller. Mandatory to have.

Table 57. qcifmiset

Parameter <argument>	Description
--device <ifname>	An interface such as eno0/swp0
--index <value>	Index entry number in this controller. Mandatory to have.
--disable	If not set disable, then to be set enable.
--cir <value>	cir. kbit/s.

Table 57. qcifmisset ...continued

Parameter <argument>	Description
--cbs <value>	cbs. octets.
--eir <value>	eir.kbit/s.
--ebs <value>	ebs.octets.
--cf	cf. couple flag.
--cm	cm. color mode.
--dropyellow	drop yellow.
--markred_enable	mark red enable.
--markred	mark red.

Table 58. qcifmigset parameter

Parameter <argument>	Description
--device <ifname>	An interface such as eno0/swp0
--index <value>	Index entry number in this controller. Mandatory to have.

Table 59. qbuset parameter

Parameter <argument>	Description
--device <ifname>	An interface such as eno0/swp0
--preemptable <value>	8-bit hex value. Example: 0xfe The MS bit corresponds to traffic class 7. The LS bit to traffic class 0. A bit value of 0 indicates express. A bit value of 1 indicates preemptable.

Table 60. cbsset command

Parameter <argument>	Description
--device <ifname>	An interface such as eno0/swp0
--tc <value>	Traffic class number.
--percentage <value>	Set percentage of tc limitation.
--all <tc-percent:tc-percent...>	Not supported.

Table 61. cbsget

Parameter <argument>	Description
--device <ifname>	An interface such as eno0/swp0
--tc <value>	Traffic class number.

Table 62. regtool

Parameter <argument>	Description
Usage: regtool { pf number } { offset } [data]	pf number: pf number for the pci resource to act on
	offset: offset into pci memory region to act upon
	data: data to be written

Table 63. ctset

Parameter <argument>	Description
--device <ifname>	An interface such as swp0
--queue_stat <value>	Specifies which priority queues have to be processed in cut-through mode of operation. Bit 0 corresponds to priority 0, Bit 1 corresponds to priority 1 so-on.

Table 64. cbgen

Parameter <argument>	Description
--device <ifname>	An interface such as swp0
--index <value>	Index entry number in this controller. Mandatory to have. This value corresponds to <code>tsnStreamIdHandle</code> of <code>cbstreamidset</code> command.
--iport_mask <value>	INPUT_PORT_MASK: If the packet is from input port belonging to this port mask, then it's a known stream and Sequence generation parameters can be applied
--split_mask <value>	SPLIT_MASK: Port mask used to add redundant paths (or ports). If split is enabled (STREAM_SPLIT) for a stream. This is OR'ed with the final port mask determined by the forwarding engine.
--seq_len <value>	SEQ_SPACE_LOG2: Minimum value is 1 and maximum value is 28. $tsnSeqGenSpace = 2^{**}SEQ_SPACE_LOG2$ For example, if this value is 12, then valid sequence numbers are from 0x0 to 0xFFF.
--seq_num <value>	GEN_REC_SEQ_NUM: The sequence number to be used for outgoing packet passed to <code>SEQ_GEN</code> function. Note: Only lower 16-bits are sent in RED_TAG.

Table 65. cbrec

Parameter <argument>	Description
--device <ifname>	An interface such as swp0
--index <value>	Index entry number in this controller. Mandatory to have. This value corresponds to <code>tsnStreamIdHandle</code> of <code>cbstreamidset</code> command.
--seq_len <value>	SEQ_SPACE_LOG2: Min value is 1 and maximum value is 28. $tsnSeqRecSeqSpace = 2^{**}SEQ_REC_SPACE_LOG2$ For example, if this value is 12, then valid sequence numbers are from 0x0 to 0xFFF.
--his_len <value>	SEQ_HISTORY_LEN: Refer to SEQ_HISTORY, Min 1 and Max 32.
--rtag_pop_en	REDTAG_POP: If True, then the redundancy tag is popped by rewriter.

Table 66. dscpset

Parameter <argument>	Description
--device <ifname>	An interface such as swp0
--disable	Disables DSCP to traffic class for frames
--index	DSCP value
--cos	Priority number of queue which is mapped to
--dpl	Drop level which is mapped to

Table 67. qcicapget

Parameter <argument>	Description
--device <ifname>	An interface such as swp0

Table 68. tsncapget

Parameter <argument>	Description
--device <ifname>	An interface such as swp0

5.1.2.2.1.4 Input tips

While providing the command input, user can use the following shortcut keys to make the input faster:

- When a user inputs a command, use the **TAB** key to help list the related commands.
For example:

```
tsntool> qbv
```

Then press **TAB** key, to get all related `qbv*` start commands.

If there is only one choice, it is filled as the whole command automatically.

- When you want to input parameters and do not remember the parameter name, you can just input "--". Then pressing **TAB** key displays all the parameters.
- If you input only half the parameter's name, pressing the **TAB** key lists all the related names.
- History:** Press the up arrow "↑". User gets the command history and can re-use the command.

5.1.2.2.1.5 Non-interactive mode

Tsntool also supports non-interactive mode.

For example:

In the interactive mode:

```
tsntool> qbuset --device eno0 --preemptable 0xfe
```

In non-interactive mode:

```
tsntool qbuset --device eno0 --preemptable 0xfe
```

5.1.2.3 Using NETCONF to configure TSN remotely

1. Overview

The NETCONF protocol defines a mechanism for device management and configuration retrieval and modification. It enables a client to adjust to the specific features of any network equipment by using a remote procedure call (RPC) paradigm and a system to expose device (server) capabilities.

YANG is a standards-based, extensible, hierarchical data modeling language. YANG is used to model the configuration and state data used by NETCONF operations, RPCs, and server event notifications.

2. Support for different platforms in Real-time Edge

Table 70. Real-time Edge platform support

TSN offload	Real-time Edge		
	LS1028A		i.MX 8DXL / i.MX 8M Plus / i.MX 93
	libtsn	tc	tc
802.1Qbv (Time Aware Shaper)	Y	Y	Y
802.1Qbu/802.3br (Frame Preemption)	Y	Y	Y
802.1Qav (Credit Based Shaper)	-	-	-
802.1CB (Frame Replication and Elimination for Reliability)	-	-	N/A
802.1Qci (Per-Stream Filtering and Policing)	Y	Y	N/A
IP config	Y	Y	Y
MAC config	Y	Y	Y
VLAN config	Y	Y	Y

3. NETCONF tools Installation and configuration

The command `netopeer2-cli` can be used as a NETCONF tool, which is present in the Netopeer2 software package. Refer to [Section 6.4.2](#) for more information. The example XML configuration files can be found in the `Instances` directory of the git repo (<https://github.com/nxp-real-time-edge-sw/real-time-edge-sysrepo.git>), including Qbv, Qbu, Qci, and stream identification via network.

- For Netopeer2 installation, refer to [Section 6.4.2.2](#).
- For the usage of `netopeer2-cli`, refer to [Section 6.4.3.3](#).
- For a configuration example, refer to [Section 6.4.3.7](#).

5.1.2.4 TSN configuration demo

In this demo, a network administrator can configure and monitor the TSN network on a web page. It can manage multiple TSN devices. Before starting the demo, set up the device network as shown in the [Figure 62](#).

This section describes the steps for remote configuration.

• Overview

The **Web UI** allows the remote control of the YANG model. The user can connect to the http server and input TSN parameters on the web UI.

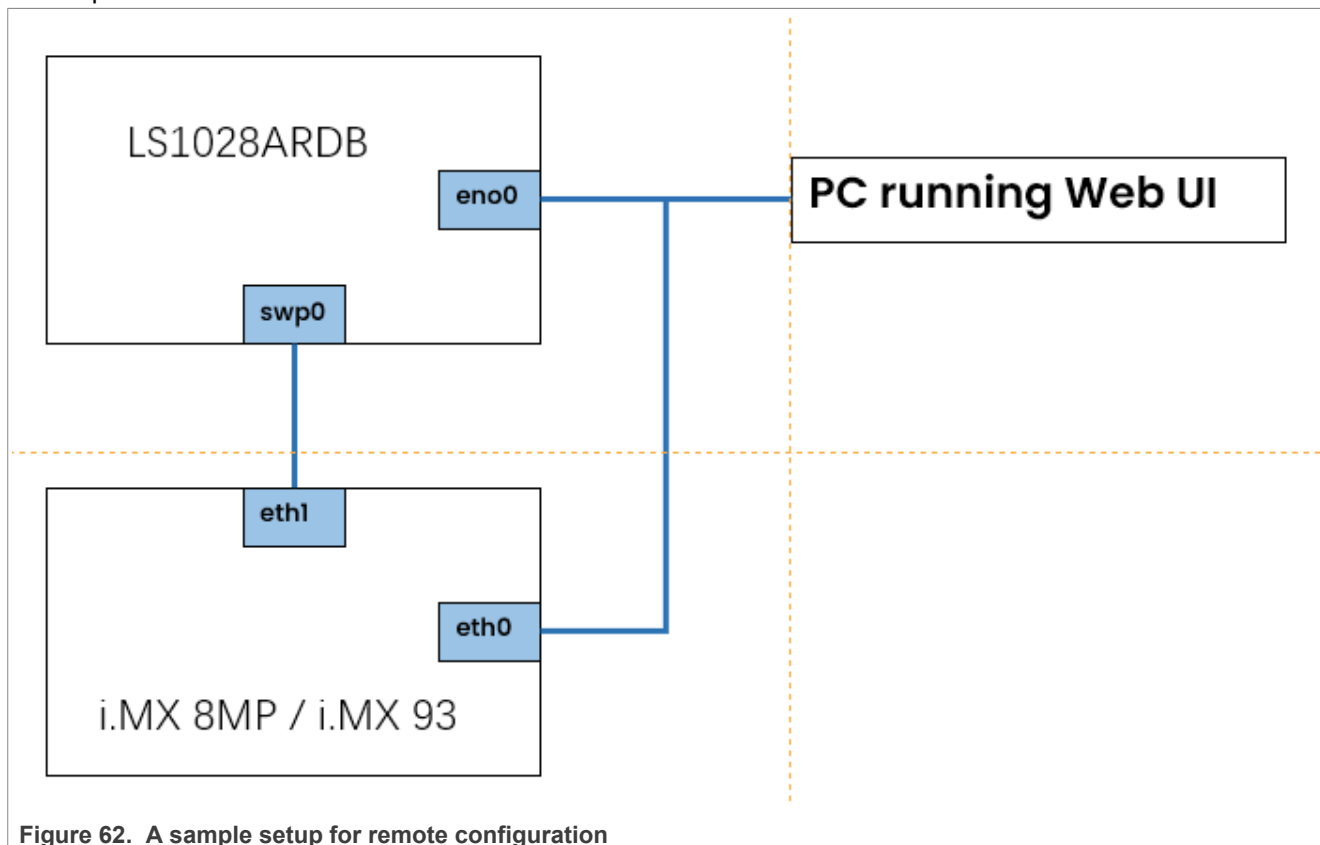


Figure 62. A sample setup for remote configuration

5.1.2.4.1 Software setup

The Web UI allows the remote control of the YANG model and also get devices information by websockets. This demo is already added to tsntool in the folder `tsntool/demos/cnc/`.

Set up a Web server on an Ubuntu PC which has connected to the TSN network. The following steps show how to set up the software and configure devices:

1. Install dependent packages on Ubuntu 22.04.

For Ubuntu:

```
$ sudo apt update
$ sudo apt install -y libtool libtool-bin libffi-dev \
libxslt1-dev libcurl4-openssl-dev xsltproc zlib1g-dev \
libssl-dev libaugeas-dev libreadline-dev \
pkg-config libxml2-dev cmake openssh-server \
libnss-mdns avahi-utils sphinx-doc sphinx-common

$ sudo apt install -y python3-sphinx python3-setuptools python3-libxml2
$ sudo apt install -y python3-pip python3-dev python3-flask python3-pexpect

$ sudo pip3 install flask-restful websockets argparse
```

2. Install pyang

```
$ git clone https://github.com/mbj4668/pyang.git
$ cd pyang
$ git checkout b92b17718de53758c4c8a05b6818ea66fc0cd4d8 -b fornetconfl
```



```
$ sudo python setup.py install
```

3. Install libssh:

```
$ git clone https://git.libssh.org/projects/libssh.git
$ cd libssh
$ git checkout fe18ef279881b65434e3e44fc4743e4b1c7cb891 -b fornetconf1
$ mkdir build; cd build/
$ cmake ..
$ make
$ sudo make install
```

Note: There is a version issue for libssh installation on Ubuntu below version 16.04. Apt-get install libssh may get version 0.6.4. But libnetconf needs a version of 0.7.3 or later. Remove the default one and reinstall by downloading the source code and installing it manually.

4. Get tsntool source code on the web server PC:

```
git clone https://github.com/nxp-qorIQ/tsntool
cd tsntool/demos/cnc/
```

5. Install libnetconf:

In the below command segments,

- PATH-to-tsntool is the path to the tsntool source code.

```
$ git clone https://github.com/CESNET/libnetconf.git
$ cd libnetconf
$ git checkout 8e934324e4b1e0ba6077b537e55636e1d7c85aed -b fornetconf1
$ cp PATH-to-tsntool/demos/cnc/0003-update-the-hostkeys-of-ssh-session-connection.patch .
$ git am 0003-update-the-hostkeys-of-ssh-session-connection.patch
$ autoreconf --force --install
$ ./configure
$ make
$ sudo make install
```

6. Install python library:

In the below command segments,

- PATH-to-libnetconf is the path to the libnetconf source code.
- PATH-to-tsntool is the path to the tsntool source code.

```
$ cd PATH-to-libnetconf/
```

The libnetconf needs to add two patches based on the below commit point to fix the demo python support. Ensure that the commit id is 313fdadd15427f7287801b92fe81ff84c08dd970.

```
$ git checkout 313fdadd15427f7287801b92fe81ff84c08dd970 -b cnc-server
$ cp PATH-to-tsntool/demos/cnc/*patch .
$ git am 0001-lnctool-to-make-install-transapi-yang-model-proper.patch
$ git am 0002-automatic-python3-authorizing-with-root-password-non.patch
$ cd PATH-to-libnetconf/python
$ python3 setup.py build; sudo python3 setup.py install
```

Note:

If rebuilding python lib, user need to remove the build folder by command `rm build -rf` before rebuilding. On the boards Real-time Edge supports, avahi-daemon and netopeer server are required. Remember to also add the netopeer2-server run at boards.

7. To start the web server on webserver PC, input the command below at shell into the folder: `PATH-to-tsntool/demos/cnc/`:

```
sudo python3 cnc.py
```

8. Start the topoagent server on the boards supported

- Make sure the netopeer2-server run at boards(Not necessary for topology discovery).
- Make sure the lldpd daemon is running at boards.
- Make sure the avahi-daemon is running at boards.
 - Start the topology server on LS1028ardb board (execute the following commands on the device console):

```
ifconfig eno2 up
ip link add name switch type bridge vlan_filtering 1
ip link set switch up
ip link set swp0 master switch && ip link set swp0 up
ip link set swp1 master switch && ip link set swp1 up
ip link set swp2 master switch && ip link set swp2 up
ip link set swp3 master switch && ip link set swp3 up

#Stop lldpd service:
pkill lldpd

#Start lldpd and limit interfaces to use. Use all ports except the control
port.
lldpd -I swp0,swp1,swp2,swp3

# Set the host name as below:

avahi-set-host-name real-time-edge-ls1028ardb
cd /home/root/samples/cncdemo/
python3 topoagent.py
```

- Start the topology server on the i.MX 8M Plus or the i.MX 93 board:

```
#Stop lldpd service.
pkill lldpd
#Start lldpd and limit interfaces to use.. Use all ports except the
control port.
lldpd -I eth1
#If the hostname is not real-time-edge-$boardname, change to real-time-
edge-$boardname.
avahi-set-host-name real-time-edge-imx93evk
cd /home/root/samples/cncdemo/
python3 topoagent.py
```

9. Use the web browser to track the topology and configuration of the devices. Input the IP of web server with the port 8180 at browser. For example:

```
http://10.193.20.147:8180
```

Note:

TSN configuration debug:

- It is recommended to track the boards using tsntool to check the real tsn configuration for comparison.
- For tsn configuration, it is also recommended to track if the netopeer2-server is running at board or not.

Limitations of Web UI are:

- The server setup on a Ubuntu PC could be more compatible.
- Supports Qbv, Qbu, and Qci in current version.

- For Qci setting, Stream-gate entry should be set ahead of setting the Stream-filter as sysrepo required. Or else, user will get failure for setting Stream-filter without a stream gate id link to.
- The boards and the web server PC are required to be in the same IP domain since the bridge may block the probe frames.

5.1.2.4.2 Remote configuration

This section describes the steps for remote configuration.

• Overview

The **Web UI** allows the remote control of the YANG model. The user can connect to the http server and input TSN parameters on the web UI. Click **"Yes, confirm"** button to send the parameters to the board as shown in [Figure 63](#).

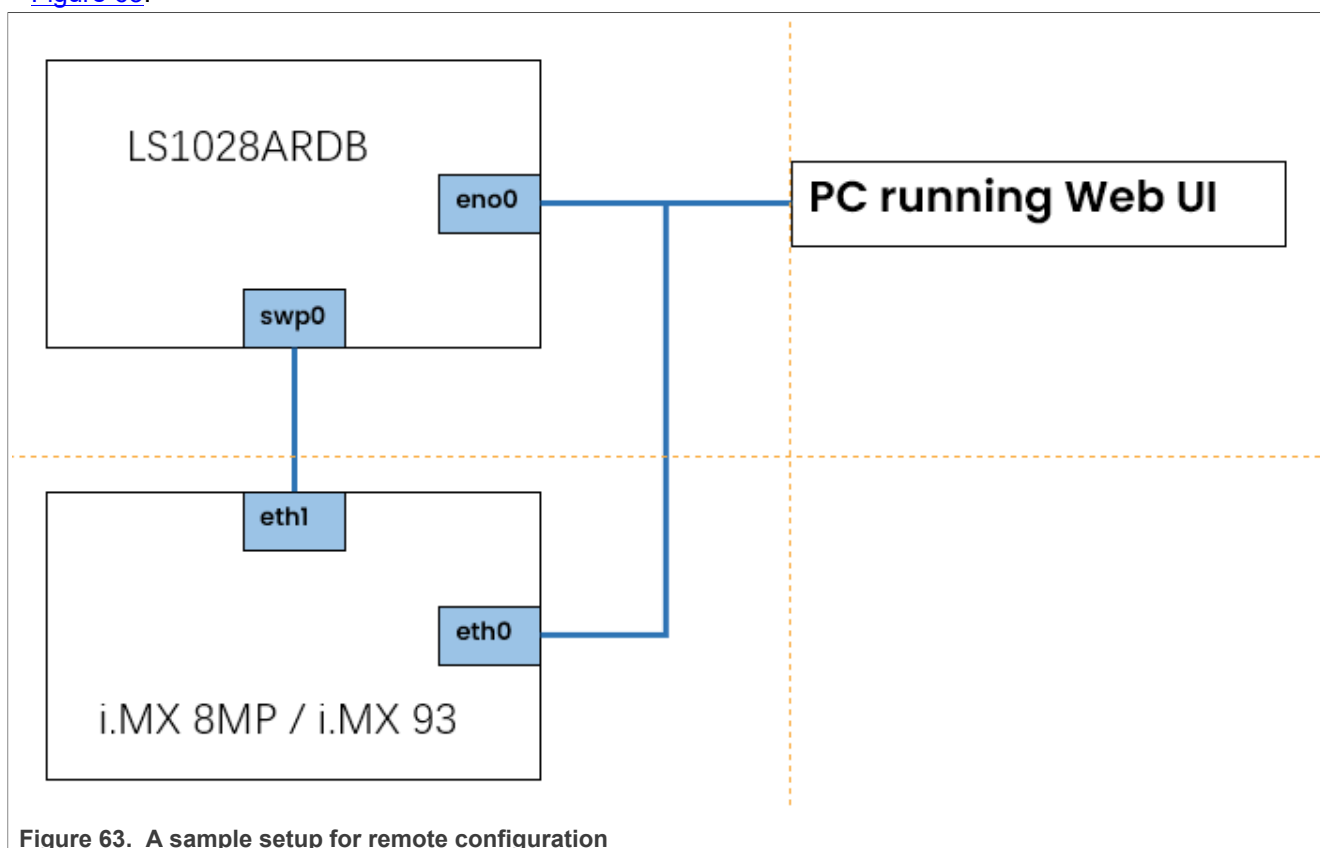


Figure 63. A sample setup for remote configuration

• User Interface

Click the device displayed on the home page, and an interface description table appears. Click the interface to jump to the configuration page.

- **Qbv Configuration:** Selecting 'qvb' option setting displays the options as shown in [Figure 64](#).



Figure 64. Qbv Configuration

- **Qbu Configuration:** Selecting 'qbu' option setting displays the options as shown in [Figure 65](#).



Figure 65. Qbu configuration

- **Qci Configuration:** Selecting 'qci' option setting displays the options as shown in [Figure 66](#).



Figure 66. Qci configuration

The **qci** interface allows user to select the configuration for "stream identify", "stream filter", "stream gate", and "flow metering".

Note:

1. Configure the "stream identify" first, then configure the "stream gate" and "flow metering", configure the "stream filter" at last.
2. "index" in "stream filter" configuration and "streamhandle" in "stream identify" should be the same value.
3. "flow meter index" in "stream filter" and "index" in "flow metering" should be the same value (63-246).

5.1.2.4.3 Dynamic remote configuration

The dynamic TSN configuration is used for the TSN configuration dynamically. Users do not need to log into each TSN node to specify the TSN parameters for TSN configuration. They only need to select the path, the base time, and then specify the cycle time. Then, the schedule mapping component calculates the TSN configuration parameters according to the user input and the path selected. The configuration parameters are delivered to each node by YANG models.

5.1.2.4.3.1 TSN working flow

This section provides an example of the TSN configuration working flow, which is described below:

After topology discovery and device registration, the network topology can be displayed over the web-browser. The user should select the nodes, specify the stream, input the timing requirement through the stream reservation component, and schedule configuration component. The results are passed down to the schedule mapping component to calculate the mapping from customer input to the TSN configuration. The configuration is instantiated using the YANG model and is delivered to different nodes for actual configuration.

The major components include:

- TSN network topology discovery
- Schedule mapping
- NETCONF/YANG configuration

- TSN Protocol Driver and TSN configuration
- Dashboard for stream management and customer input parse

Figure 67 shows the architecture diagram.

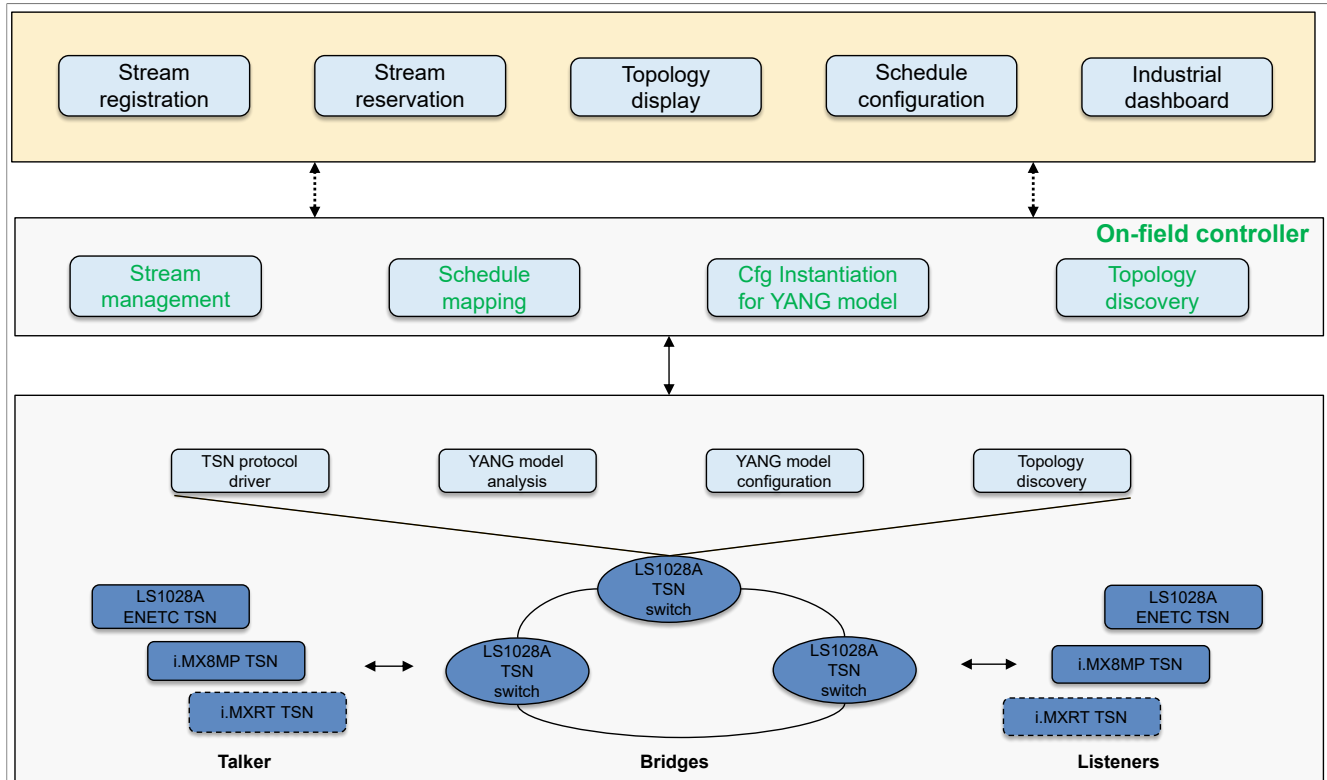


Figure 67. TSN architecture diagram

The Figure 67 illustrates the three layers of the TSN architecture. The first layer is the TSN network layer. The second layer is the service layer, which runs on the on-field controller/server. The third layer is an optional service that runs in the cloud or on-field server.

The TSN network layer includes TSN switches and endpoints that form the TSN network. For example, it includes the LS1028A TSN switch and TSN endpoints such as LS1028 ENETC TSN and i.MX 8M Plus TSN endpoint. The different components run on each of the nodes. For example, the topology discovery component collects the network topology, YANG model performs the TSN register configuration, and NETCONF server parses the YANG model for TSN configuration.

The second layer is the on-field controller layer. It is the server running on-field to host the services of the industrial board, topology discovery and schedule mapping.

The third layer runs on the cloud, which could host the services running on the on-field controller. This layer is an optional layer.

5.1.2.4.3.2 Topology discovery

The topology discovery component is used to discover network connections by running LLDP on each TSN network node. The connection information is delivered to topology discovery service running on the on-field server.

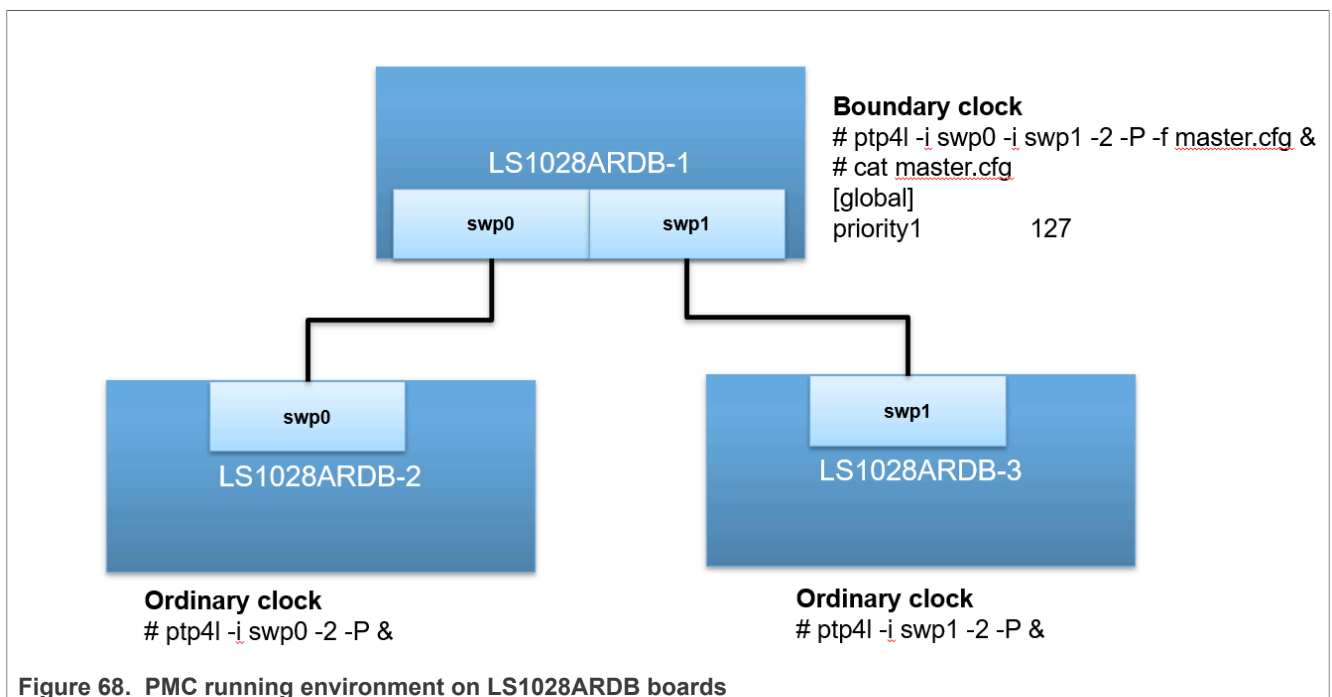
5.1.2.4.3.3 Path selection

Path selection implements an algorithm to select the path between the selected talker and listener. If there are multiple paths, the dashboard displays all paths and the user can select one of the paths for the stream. Set a different VLAN ID for the selected path, and the stream with this VID can flow in the path.

5.1.2.4.3.4 Path delay

Clock synchronization and path delay calculation are two prerequisites for schedule mapping. Clock synchronization uses gPTP to synchronize the clock of the system. The example described in this document uses linuxptp PMC tool to get the path delay.

[Figure 68](#) shows a sample configuration to show the PMC running environment on LS1028ARDB boards.



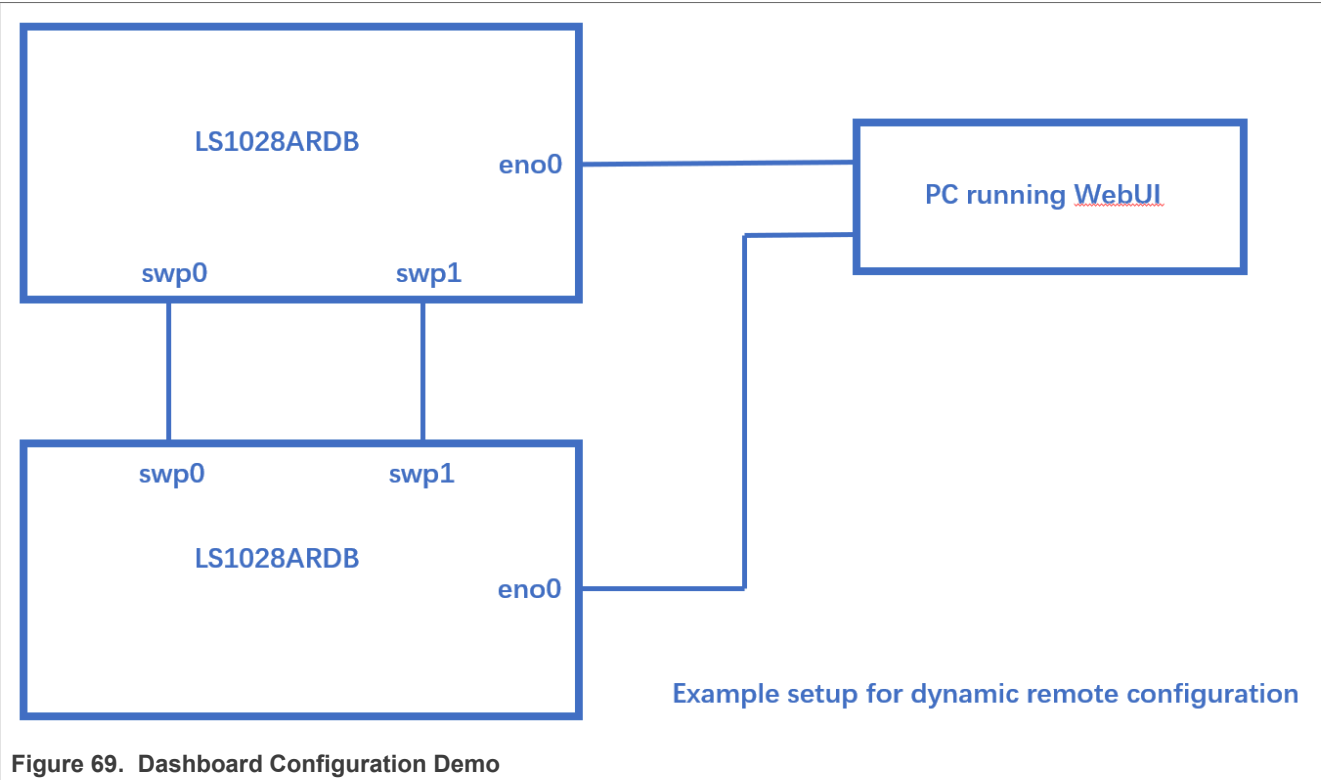
5.1.2.4.3.5 Schedule mapping

The schedule mapping component is a critical component to convert the customer requirement to TSN register configuration. This component performs the following:

- Gets the user input and converting the input into TSN parameters.
- Gets the path and path delay from the link object of the NetworkGraph file.
- Gets the old TSN configuration for each node and calculates a new configuration to meet the user's requirements.

5.1.2.4.3.6 Dashboard configuration demo

The figure below shows the dashboard for the configuration demo.



Registering a stream

Click **“Check Path”** button, input the start device in **“first device”** input box, and end device in **“Second device”** input box. Then click the **“submit”** button, path is described in the [Figure 70](#).



Click **“Register Stream”** button, then select the path in path select. Fill **VLAN ID**, **Stream ID**, **Priority**, and then click **“Add”** button. The [Figure 71](#) shows the output a stream table.

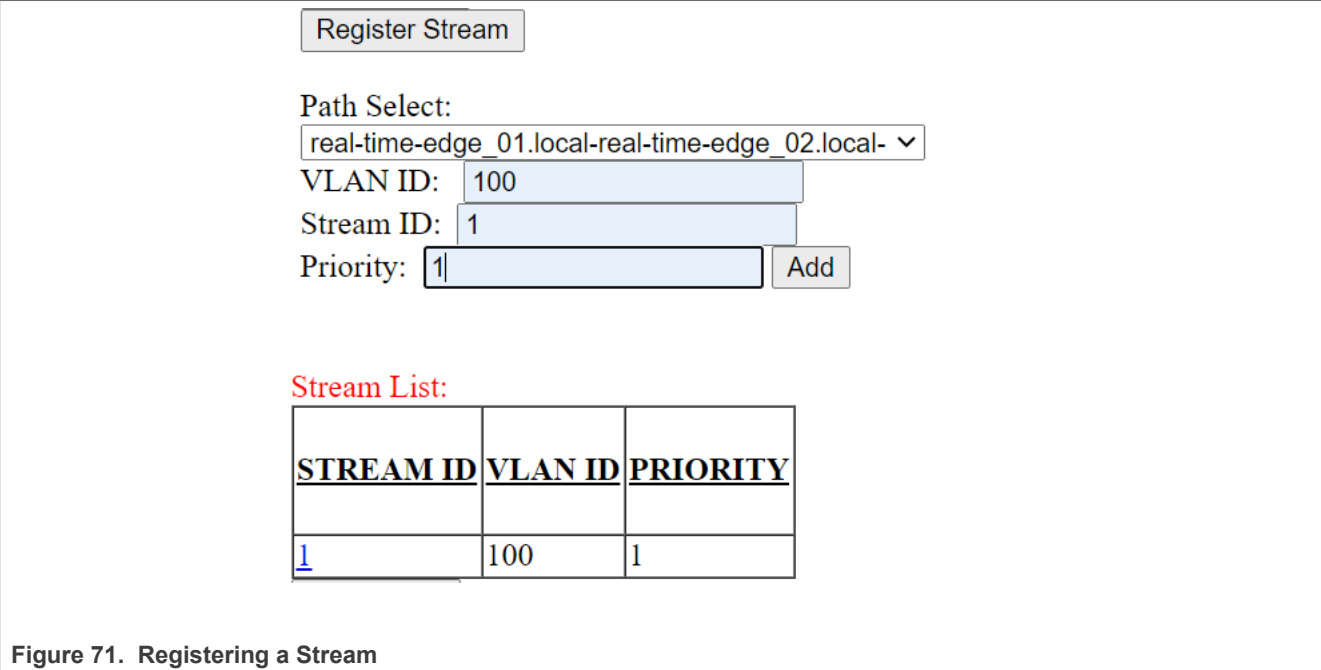


Figure 71. Registering a Stream

Configure stream identification

Click one stream ID in stream table, jump to stream configuration page. Select streamidentify and fill information in input boxes. The stream MAC information and VLAN ID identify a stream according to the 802.1CB definition. This information is used by the PSFP configuration. Therefore, the **streamidentify** page should be configured before configuring Qci and CFQ. Refer [Figure 72](#).




Figure 72. Configuring stream identification

Configure Qbv and Qci On Stream

Select Qbv, and then fill basetime, cycletime, and gate open time in the respective input textboxes. Select **enable Qci** button to configure both Qci gate control on input port and Qbv control list on the output port. The CNC server calculates the gate open time slot on each board and get a minimum time delay. Each path node tries to open gate with a minimum time delay. Refer to [Figure 73](#).




Figure 73. Configuring Qbv and Qci On Stream

Configuring CQF

The CQF configuration is based on the 802.1Qch definition to configure Qbv and Qci. The CQF configuration cannot be mixed with the previous Qbv configuration. In CQF configuration, the cycle time and gate open time for all streams must be the same, and cycle time must be an integer multiple of gate open time. Packets are delayed for a gate open time on each path node. Refer [Figure 74](#).

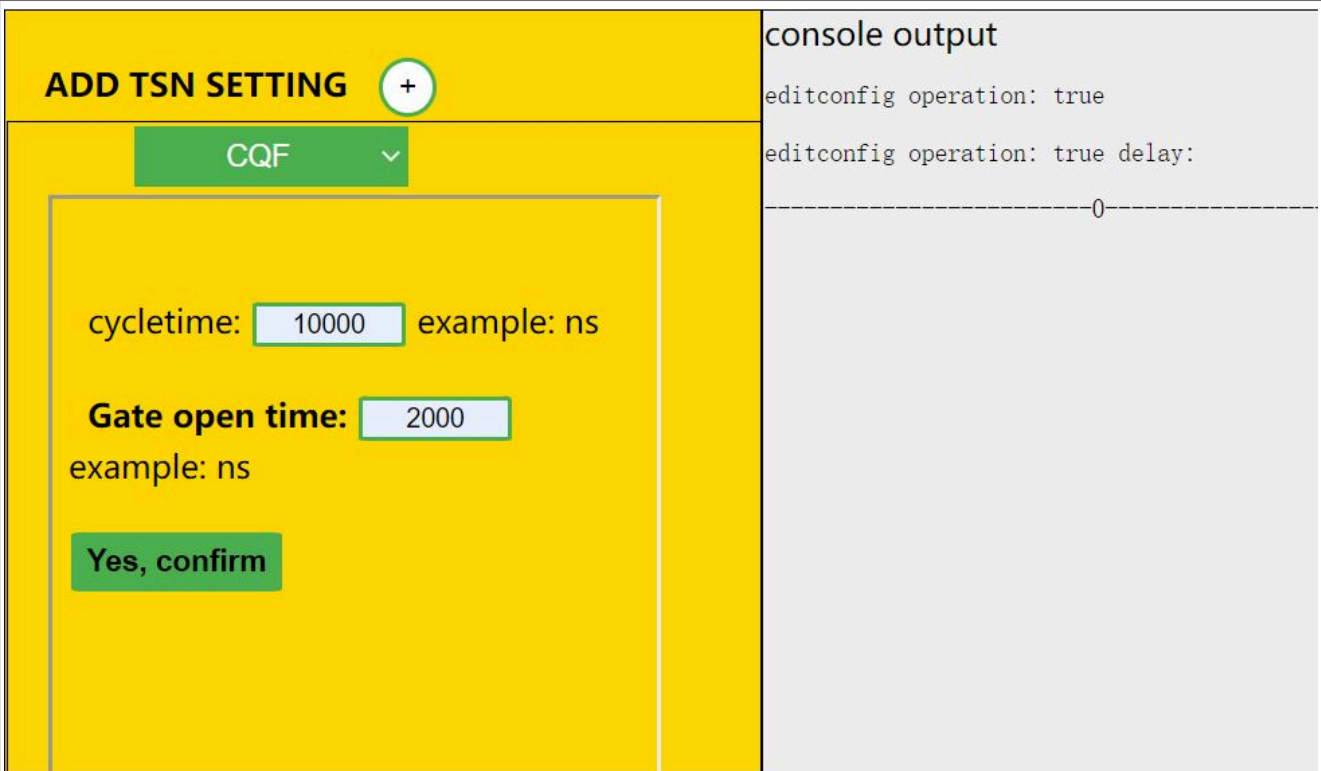


Figure 74. Configuring CQF

5.1.3 TSN on i.MX 8DXL / i.MX 8M Plus / i.MX 93

The following sections describe TSN configuration on i.MX 8DXL, i.MX 8M Plus, or i.MX 93 hardware platforms

5.1.3.1 Test environment

On i.MX 8M Plus EVK / i.MX 93 EVK / i.MX 93 14x14 EVK platform, the interface name of ENET_QOS port which supports TSN is eth1. On i.MX 8DXL EVK / i.MX 93 9x9 LPDDR4 QSB, the interface name of ENET_QOS port which supports TSN is eth0.

Note: For i.MX 93 14x14 EVK, in order to use ENET_QOS interface, TJA1103SDB ENET PHY daughter card is needed to be connected on J9 connector. To verify the TSN features, two i.MX 93 14x14 EVK boards are connected back-to-back on ENET_QOS interface via TJA1103ADB ENET PHY daughter card. Also on one of the TJA1103SDB, connect jumper to short pin 2-3 of J14 (CONFIG 6).

Connect ENET_QOS port to the TestCenter to test TSN features. The commands in this section use the i.MX 8M Plus EVK platform as example:

Use the following command to check the TSN Ethernet device name:

```
#ls /sys/devices/platform/soc@0/30800000.bus/30bf0000.ethernet/net/ eth1
```

The [Figure 75](#) shows the TSN test environment setup.

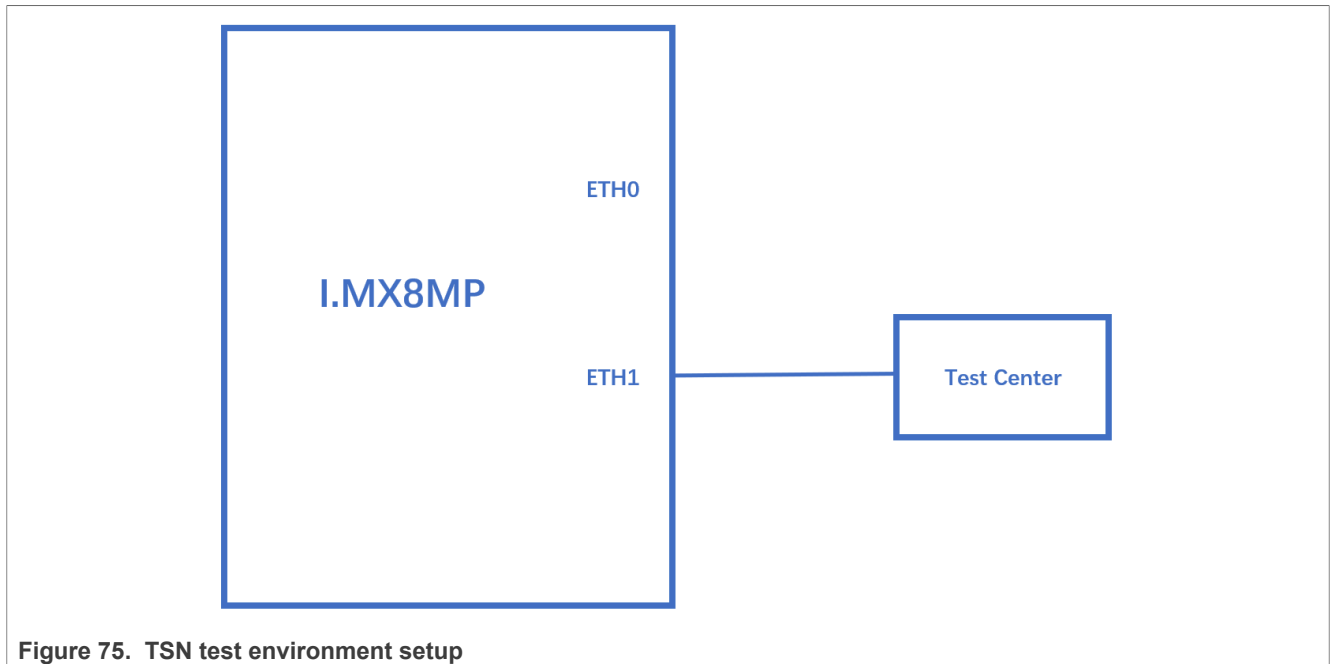


Figure 75. TSN test environment setup

Note: TestCenter is a device used to capture streams from eth1 of i.MX 8M Plus board. For this example, the Spirent TestCenter is used to capture preemptable frames in the Qbu test case.

5.1.3.2 Clock synchronization

To test 1588 synchronization on dwcMAC interfaces, use the following procedure:

1. Connect eth1 interfaces on two boards in a back-to-back manner. The Linux booting log is as follows:

```
...
pps pps0: new PPS source ptp0
...
```

2. Configure the IP address using the command below:

```
ifconfig eth1 192.168.3.1
```

3. Check PTP clock and time stamping capability:

```
# ethtool -T eth1
Time stamping parameters for eth1:
Capabilities:
    hardware-transmit      (SOF_TIMESTAMPING_TX_HARDWARE)
    software-transmit      (SOF_TIMESTAMPING_TX_SOFTWARE)
    hardware-receive       (SOF_TIMESTAMPING_RX_HARDWARE)
    software-receive       (SOF_TIMESTAMPING_RX_SOFTWARE)
    software-system-clock  (SOF_TIMESTAMPING_SOFTWARE)
    hardware-raw-clock     (SOF_TIMESTAMPING_RAW_HARDWARE)
PTP Hardware Clock: 1
Hardware Transmit Timestamp Modes:
    off                    (HWTSTAMP_TX_OFF)
    on                    (HWTSTAMP_TX_ON)
Hardware Receive Filter Modes:
    none                  (HWTSTAMP_FILTER_NONE)
    all                   (HWTSTAMP_FILTER_ALL)
    ptpv1-l4-event       (HWTSTAMP_FILTER_PTP_V1_L4_EVENT)
    ptpv1-l4-sync        (HWTSTAMP_FILTER_PTP_V1_L4_SYNC)
```

ptpv1-l4-delay-req	(HWTSTAMP_FILTER_PTP_V1_L4_DELAY_REQ)
ptpv2-l4-event	(HWTSTAMP_FILTER_PTP_V2_L4_EVENT)
ptpv2-l4-sync	(HWTSTAMP_FILTER_PTP_V2_L4_SYNC)
ptpv2-l4-delay-req	(HWTSTAMP_FILTER_PTP_V2_L4_DELAY_REQ)
ptpv2-event	(HWTSTAMP_FILTER_PTP_V2_EVENT)
ptpv2-sync	(HWTSTAMP_FILTER_PTP_V2_SYNC)
ptpv2-delay-req	(HWTSTAMP_FILTER_PTP_V2_DELAY_REQ)

4. Run ptp4l on two boards:

```
ptp4l -i eth1 -p /dev/ptp1 -m -2
```

5. After running, one board is automatically selected as the master, and the slave board displays synchronization messages.

6. For 802.1AS testing, use the configuration file gPTP.cfg in linuxptp source. Run the below command on the boards, instead:

```
ptp4l -i eth1 -p /dev/ptp1 -f /etc/ptp4l_cfg/gPTP.cfg -m
```

Or use GenAVB/TSN Stack with the following command: 'avb.sh start'. Note that the configuration file /etc/genavb/fgptp.cfg is automatically used.

Note: i.MX 8M Plus current dwmac driver (eth1) initializes few hardware functions while opening net device, including PTP initialization. Before that, the operations such as ethtool queries, and PTP operations might not work. So, the workaround is to do operations on the eth1 and PTP of dwmac only after "ifconfig eth1 up".

Note: If Qbu preemption is enabled on remote device and the PTP packets are sent as preemption frames, run clock synchronization using the ptp4l command along with the parameter --hwts_filter=full. For example:

```
ptp4l -i eth1 -p /dev/ptp1 -f /etc/ptp4l_cfg/gPTP.cfg -m --hwts_filter=full
```

5.1.3.3 Qbv

Note: The Qbu frame preemption capability is enabled automatically if the link partner also supports frame preemption. This is achieved by LLDP protocol by which the link partner announces its support for the preemption capability via an Additional Ethernet Capabilities TLV in an LLDPDU addressed to the Nearest Bridge group address (see IEEE Std 802.1Q).. So firstly user needs to check whether frame preemption is active by command `ethtool --show-frame-preemption eth1`. If so, user needs to disable frame preemption using command `ethtool --set-frame-preemption eth1 disabled` before configuring Qbv using `s (SetGateStates)` command in `sched-entry`.

1. Enable the ptp device, and get the current ptp time.

```
ptp4l -i eth1 -p /dev/ptp1 -m
#Get current time(seconds)
devmem2 0x30bf0b08
0x5E01F9B2
```

2. Get the basetime to be 2 minutes later.

```
#Basetime = (currenttime + 120) * 1000000000 = 1577187882000000000
```

3. Set time schedule, open queue 1 in 100 µs and open queue 2 in 100 µs.

```
tc qdisc replace dev eth1 parent root handle 100 taprio \
```

```
num_tc 5 map 0 1 2 3 4 queues 1@0 1@1 1@2 1@3 1@4 base-time
15771878820000000000 \
    sched-entry S 1 100000 \
    sched-entry S 2 100000 \
    sched-entry S 4 100000 flags 2
```

4. Send two streams into queue 1 and queue 2.

```
/home/root/samples/pktgen/pktgen_twoqueue.sh -i eth1 -q 1 -s 1000 -n 0 -m
90:e2:ba:ff:ff:ff
```

5. Capture the streams on TestCenter, 100 μ s queue 1 frames (length=1004) and 100 μ s queue 2 frames (length=1504) will be got. Or if the Ethernet port is connected to another board, the frames can be captured on that board by using Linux tcpdump command as shown below:

```
tcpdump -i eth0 -e -n -t -xx -c 10000 -w tsn.pcap
```

Then Wireshark can be used to analyze the pcap file on host PC.

Note:

- *More than one entry needs to be set on each tc taprio command.*
- *Use “devmem2 0x30bf0c58” to get Qbv status and check if qbv status is active. Refer to the MTL_EST_Status register.*

5.1.3.4 Qbu

1. Using ethtool to enable Qbu on eth1, set queue 2 to be preemptable.

```
ethtool --set-frame-preemption eth1 preemptible-queues-mask 0x04 min-frag-
size 60
```

Note: Once Qbu is enabled, queue 0 is always preemptable queue. To support preemption, the MAC should have at least 1 queue designated as an express queue.

Note: On a back-to-back setup using two i.MX 8M Plus EVK boards connected via eth1, ensure that Qbu is enabled on eth1 of both boards.

2. Send two streams into queue 1 and queue 2.

```
/home/root/samples/pktgen/pktgen_twoqueue.sh -i eth1 -q 1 -s 150 -n 0 -m
90:e2:ba:ff:ff:ff
```

3. Capture the mPacket on Spirent TestCenter. Users can observe that Q2 frames are preempted into fragments.

Note: Spirent TestCenter can capture the preamble of mPacket. Refer to Section 99.3, "MAC Merge Packet (mPacket)" of IEEE standard for Ethernet 802.3-2018 for the mPacket format.

- Below is an example mPacket that contains an express packet, which has SMD value of 0xD5.




Figure 76. Sample mPacket that contains an express packet

- b. Below is an example mPacket containing an initial fragment of a preemptable packet, which has SMD-S1 value of 0x4C.




Figure 77. Sample mPacket containing an initial fragment of a preemptable packet

- c. Below is an example mPacket containing a continuation fragment of a preemptable packet, which has SMD-C1 value of 0x52, as well as frag_count value of 0xE6.

[illegible]

Figure 78. Sample mPacket containing a continuation fragment of a preemptable packet

4. User can also check the below counter for the number of fragments transmitted.

```
ethtool -S eth1 | grep "mmc tx fpe fragment cntr"
```

5. Qbu combined with Qbv test

Once a queue is set to be a preemptable queue and the gate open/close is invalid in Qbv gate control list, the queue is considered as always "Open". Use **Hold/Release** to control all preemptable queues. When the GCL entry is set from **Hold** to **Release**, preemptable queues begin transmitting. When GCL entry is set from Release to Hold, preemptable queues are held.

```
tc qdisc replace dev eth1 parent root handle 100 taprio \
    num_tc 5 map 0 1 2 3 4 queues 1@0 1@1 1@2 1@3 1@4 base-time
    1577187882000000000 \
    sched-entry H 2 100000 \
    sched-entry R 4 100000 flags 2
```

5.1.3.4.1 Preemption verification

The preemption capability is enabled only if the link partner announces its support for the preemption capability via an Additional Ethernet Capabilities TLV in an LLDPDU addressed to the Nearest Bridge group address (see IEEE Std 802.1Q). The preemption capability is disabled if the MAC Merge sublayer receives indication of link failure.

Connect `eth1` of two boards back to back. Then, run below command to enable the hardware verification:

```
ethtool --set-frame-preemption eth1 fp on preemptible-queues-mask 0x04 min-frag-
size 60
```

After that, make the link down and then up to enable the verify mPacket and response mPacket to be exchanged between two boards. Run below command to show the preemption status.

```
ethtool --show-frame-preemption eth1
```

Capture the LLDP frames on eth1 port to check the Additional Ethernet Capabilities TLV.

5.1.3.5 Qav

To perform a Qav test, perform the steps listed below:

1. Set a queue map handle.

```
tc qdisc add dev eth1 root handle 1: mqprio num_tc 5 map 0 1 2 3 4 queues 1@0
1@1 1@2 1@3 1@4 hw 0
```

2. Set the bandwidth of queue 3 to be 20 Mbit/s.

```
tc qdisc replace dev eth1 parent 1:4 cbs locredit -1470 hicredit 30 sendslope
-980000 idleslope 20000 offload 1
```

3. Send a stream into queue 3:

```
/home/root/samples/pktgen/pktgen_sample01_simple.sh -i eth1 -q 3 -s 500 -n
3000
```

4. Get the result. The bandwidth is 19 Mbit/s. For a workaround to the deviation, refer to the [Note 1](#) after point 8 of this section.

```
WARN : Missing destination MAC address
WARN : Missing destination IP address
Running... ctrl^C to stop
Done
Result device: eth1
Params: count 3000 min_pkt_size: 500 max_pkt_size: 500
       frags: 0 delay: 0 clone_skb: 0 ifname: eth1
       flows: 0 flowlen: 0
       queue_map_min: 3 queue_map_max: 3
       dst_min: 198.18.0.42 dst_max:
       src_min: src_max:
       src_mac: a6:85:82:fc:89:bf dst_mac: 02:5d:ae:ba:e0:00
       udp_src_min: 9 udp_src_max: 109 udp_dst_min: 9 udp_dst_max: 9
       src_mac_count: 0 dst_mac_count: 0
       Flags: UDPSRC_RND NO_TIMESTAMP QUEUE_MAP_RND
Current:
       pkts-sofar: 3000 errors: 0
       started: 5631940023us stopped: 5632560030us idle: 79984us
       seq_num: 3001 cur_dst_mac_offset: 0 cur_src_mac_offset: 0
       cur_saddr: 0.0.0.0 cur_daddr: 198.18.0.42
       cur_udp_dst: 9 cur_udp_src: 41
       cur_queue_map: 3
       flows: 0
Result: OK: 620007(c540023+d79984) usec, 3000 (500byte,0frags)
       4838pps 19Mb/sec (19352000bps) errors: 0
```

5. Set the bandwidth of queue 4 to be 40 Mbit/s.

```
tc qdisc replace dev eth1 parent 1:5 cbs locredit -1440 hicredit 60 sendslope
-960000 idleslope 40000 offload 1
```

6. Send a stream into queue 4 and get the result.

```
/home/root/samples/pktgen/pktgen_sample01_simple.sh -i eth1 -q 4 -s 500 -n
3000
WARN : Missing destination MAC address
WARN : Missing destination IP address
Running... ctrl^C to stop
Done
Result device: eth1
Params: count 3000 min_pkt_size: 500 max_pkt_size: 500
       frags: 0 delay: 0 clone_skb: 0 ifname: eth1
       flows: 0 flowlen: 0
       queue_map_min: 4 queue_map_max: 4
       dst_min: 198.18.0.42 dst_max:
       src_min: src_max:
       src_mac: a6:85:82:fc:89:bf dst_mac: 02:5d:ae:ba:e0:00
       udp_src_min: 9 udp_src_max: 109 udp_dst_min: 9 udp_dst_max: 9
       src_mac_count: 0 dst_mac_count: 0
       Flags: UDPSRC_RND NO_TIMESTAMP QUEUE_MAP_RND
Current:
pkts-sofar: 3000 errors: 0
started: 6113136017us stopped: 6113443758us idle: 38457us
seq_num: 3001 cur_dst_mac_offset: 0 cur_src_mac_offset: 0
cur_saddr: 0.0.0.0 cur_daddr: 198.18.0.42
cur_udp_dst: 9 cur_udp_src: 17
cur_queue_map: 4
flows: 0
Result: OK: 307741(c269283+d38457) usec, 3000 (500byte,0frags)
       9748pps 38Mb/sec (38992000bps) errors: 0
```

7. Send two streams into queue 3 and queue 4 using the command below:

```
/home/root/samples/pktgen/pktgen_twoqueue.sh -i eth1 -q 3 -s 1500 -n 0
```

8. Capture the streams on TestCenter; sort the frames by one Q3 frame and two Q4 frames.

Note: According to errata TKT0649448, the following formula holds true:

Additional/Extra bandwidth = ((# of packets x 12 bytes) / (Total number of bytes that is transmitted in that window, including the preamble bytes of each packet)) x Fractional bandwidth.

actual bandwidth = goal_bandwidth * (1 - (12 / (8 + average_packet_length)))

5.1.4 TSN on LS1028A

The **tsntool** is an application configuration tool to configure the TSN capability on LS1028ARDB. The files **/usr/bin/tsntool** and **/usr/lib/libtsn.so** are located in the rootfs. Run **tsntool** to start the setting shell.

5.1.4.1 TSN configuration on ENETC

The **tsntool** is an application configuration tool to configure the TSN capability. Users can find the files **/usr/bin/tsntool** and **/usr/lib/libtsn.so** in the rootfs. Run **tsntool** to start the setting shell. The following sections describe the TSN configuration examples on the ENETC Ethernet driver interfaces.

Before testing the ENETC TSN test cases, you must enable **mqprio** by using the command below:

```
tc qdisc add dev eno0 root handle 1: mqprio num_tc 8 map 0 1 2 3 4 5 6 7 queues
100 101 102 103 104 105 106 107 hw 1
```

5.1.4.1.1 Clock synchronization

To test 1588 synchronization on ENETC interfaces, use the following procedure:

1. Connect ENETC interfaces on two boards in a back-to-back manner. (For example, eno0 to eno0.)
The linux booting log is as follows:

```
...
pps pps0: new PPS source ptp0
...
```

2. Check PTP clock and timestamping capability:

```
# ethtool -T eno0
Time stamping parameters for eno0:
Capabilities:
  hardware-transmit          (SOF_TIMESTAMPING_TX_HARDWARE)
  hardware-receive          (SOF_TIMESTAMPING_RX_HARDWARE)
  hardware-raw-clock        (SOF_TIMESTAMPING_RAW_HARDWARE)
PTP Hardware Clock: 0
Hardware Transmit Timestamp Modes:
  off      (HWTSTAMP_TX_OFF)
  on       (HWTSTAMP_TX_ON)
Hardware Receive Filter Modes:
  none     (HWTSTAMP_FILTER_NONE)
  all      (HWTSTAMP_FILTER_ALL)
```

3. Configure the IP address and run ptp4l on two boards:

```
# ifconfig eno0 <ip_addr>
# ptp4l -i eno0 -p /dev/ptp0 -m
```

4. After running, one board would be automatically selected as the master, and the slave board would print synchronization messages.
5. For 802.1AS testing, just use the configuration file gPTP.cfg in linuxptp source. Run the below command on the boards, instead:

```
# ptp4l -i eno0 -p /dev/ptp0 -f /etc/ptp4l_cfg/gPTP.cfg -m
```

5.1.4.1.2 Qbv

This test includes the Basic Gates Closing test, Basetime test, and the Qbv performance test. These are described in the following sections.

5.1.4.1.2.1 Basic gates closing

The commands below describe the steps for closing the basic gates:

```
cat > qbv0.txt << EOF
t0      000000000b      20000
EOF
```

```
#Explanation:
# 'NUMBER'      : t0
# 'GATE_VALUE'  : 000000000b
# 'TIME_LONG'   : 20000 ns
```

```
tsntool
tsntool> verbose
```

```
tsntool> qbvset --device eno0 --entryfile ./qbv0.txt
ethtool -S eno0
ping 192.168.0.2 -c 1    #Should not pass any frame since gates are all off.
```

5.1.4.1.2.2 Basetime test

Base on case 1 qbv1.txt gate list.

```
#create ls gate
cat > qbv1.txt << EOF
t0  11111111b      10000
t1  00000000b      99990000
EOF
#ENETC Qbv basetime can be set any past time or future time.
#For the past time, hardware calculate by:
# effective-base-time = base-time + N x cycle-time
#where N is the smallest integer number of cycles such that effective-base-time
#>= now.
#If you want a future time, you can get current time by:
tsntool> ptptool -g
#Below example shows basetime start at 260.666 s (start of 1 January 1970):
tsntool> qbvset --device eno0 --entryfile qbv1.txt --basetime 260.666
tsntool> qbvget --device eno0 #User can check configchange time
tsntool> regtool 0 0x11a10 #Check pending status, 0x1 means time gate is working
#Waiting to change state, ping remote computer
ping 192.168.0.2 -A -s 1000
#The reply time will be about 100 ms
```

Since 10000 ns is the maximum limit for package size 1250 B.

```
ping 192.168.0.2 -c 1 -s 1300 #frame should not pass
```

5.1.4.1.2.3 Qbv performance test

Use the setup described in the [Figure 79](#) for testing ENETC port0 (MAC0).

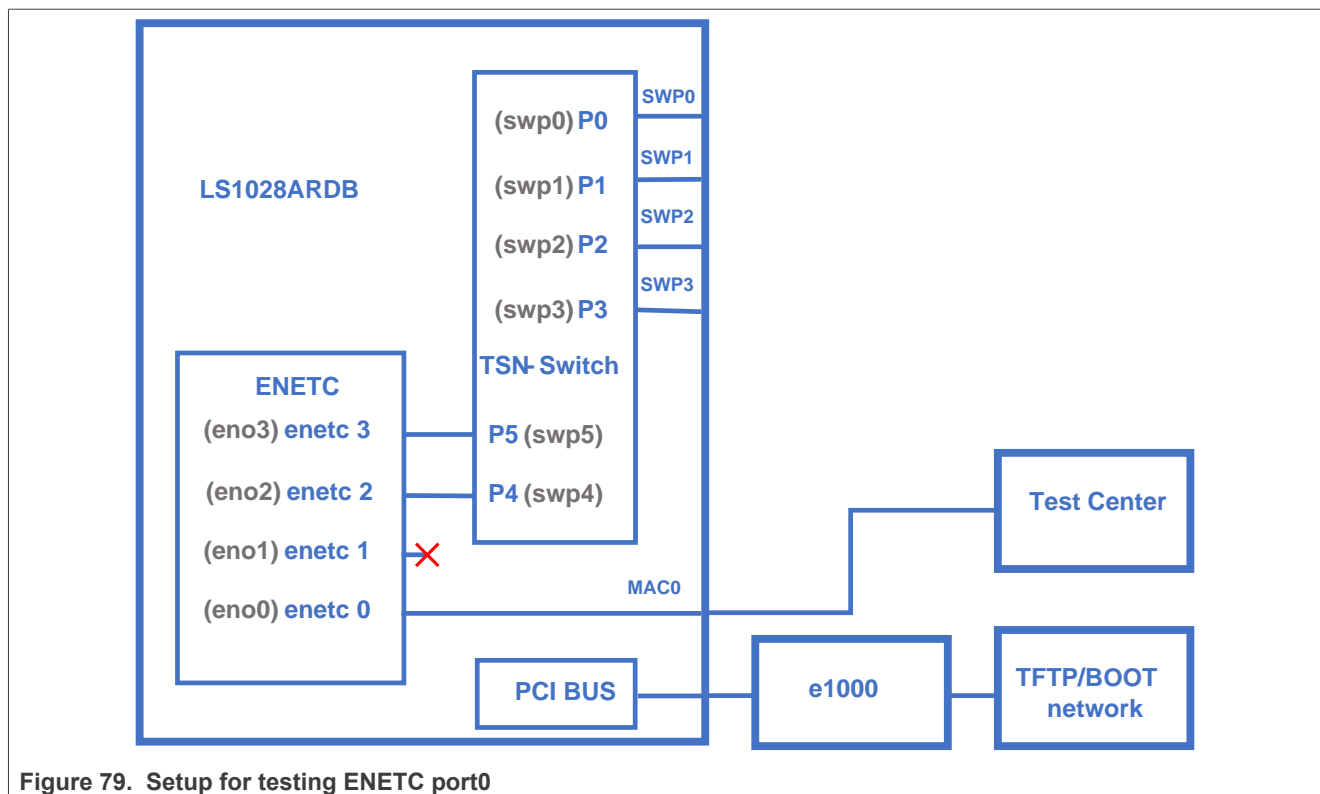


Figure 79. Setup for testing ENETC port0

Note: TestCenter is a device used to capture streams from *enetc0* of LS1028ARDB board. Users can use another board to capture streams by using *tcpdump* command and then use Wireshark to analyze it.

```
cat > qbv5.txt << EOF
t0 1111111b 1000000 t1 00000000b 1000000
EOF
qbvset --device eno0 --entryfile qbv5.txt
/home/root/samples/pktgen/pktgen_twoqueue.sh -i eno0 -q 3 -n 0
#The stream would get about half line rate
```

5.1.4.1.2.4 Using taprio Qdisc Setup Qbv

LS1028ARDB support the taprio qdisc to setup Qbv either. Below is an example setup.

```
#Qbv test do not require the mqprio setting.
# If mqprio is enabled, try to disable it by below command:
tc qdisc del dev eno0 root handle 1: mqprio
# Enable the Qbv for ENETC eno0 port
# Below command set eno0 with gate 0x01, means queue 0 open, the other queues
gate close.
tc qdisc replace dev eno0 parent root handle 100 taprio num_tc 8 map 0 1 2 3 4 5
6 7 queues 100 101 102 103 104 105 106 107 base-time 0 sched-entry S 01 300000
flags 0x2
# Ping through eno0 port should be ok
# Then close the gate queue 0. Open gate queue 1. The other queues gate close.
tc qdisc replace dev eno0 parent root handle 100 taprio num_tc 8 map 0 1 2 3 4 5
6 7 queues 100 101 102 103 104 105 106 107 base-time 0 sched-entry S 02 300000
flags 0x2
# Ping through eno0 port should be dropped
#Disable the Qbv for ENETC eno0 port as below
```



```
tc qdisc del dev eno0 parent root handle 100 taprio
```

5.1.4.1.3 Qbu

- If user has two LS1028ARDB boards, then link the two eno0 ports back to back. In this case, the test does not need the switch to be set up. Users can omit the steps 2, 3, and 4 and just perform steps 1, 5, and 6.
- If user has only one board, user can set the frame path from eno0 to switch by linking enetc ports MAC0 - SWP0. The setup enables the switch SWP0 port-merging capability. Then enetc eno0 can show the preemption capability. Use the setup as shown in the for the Qbu test.

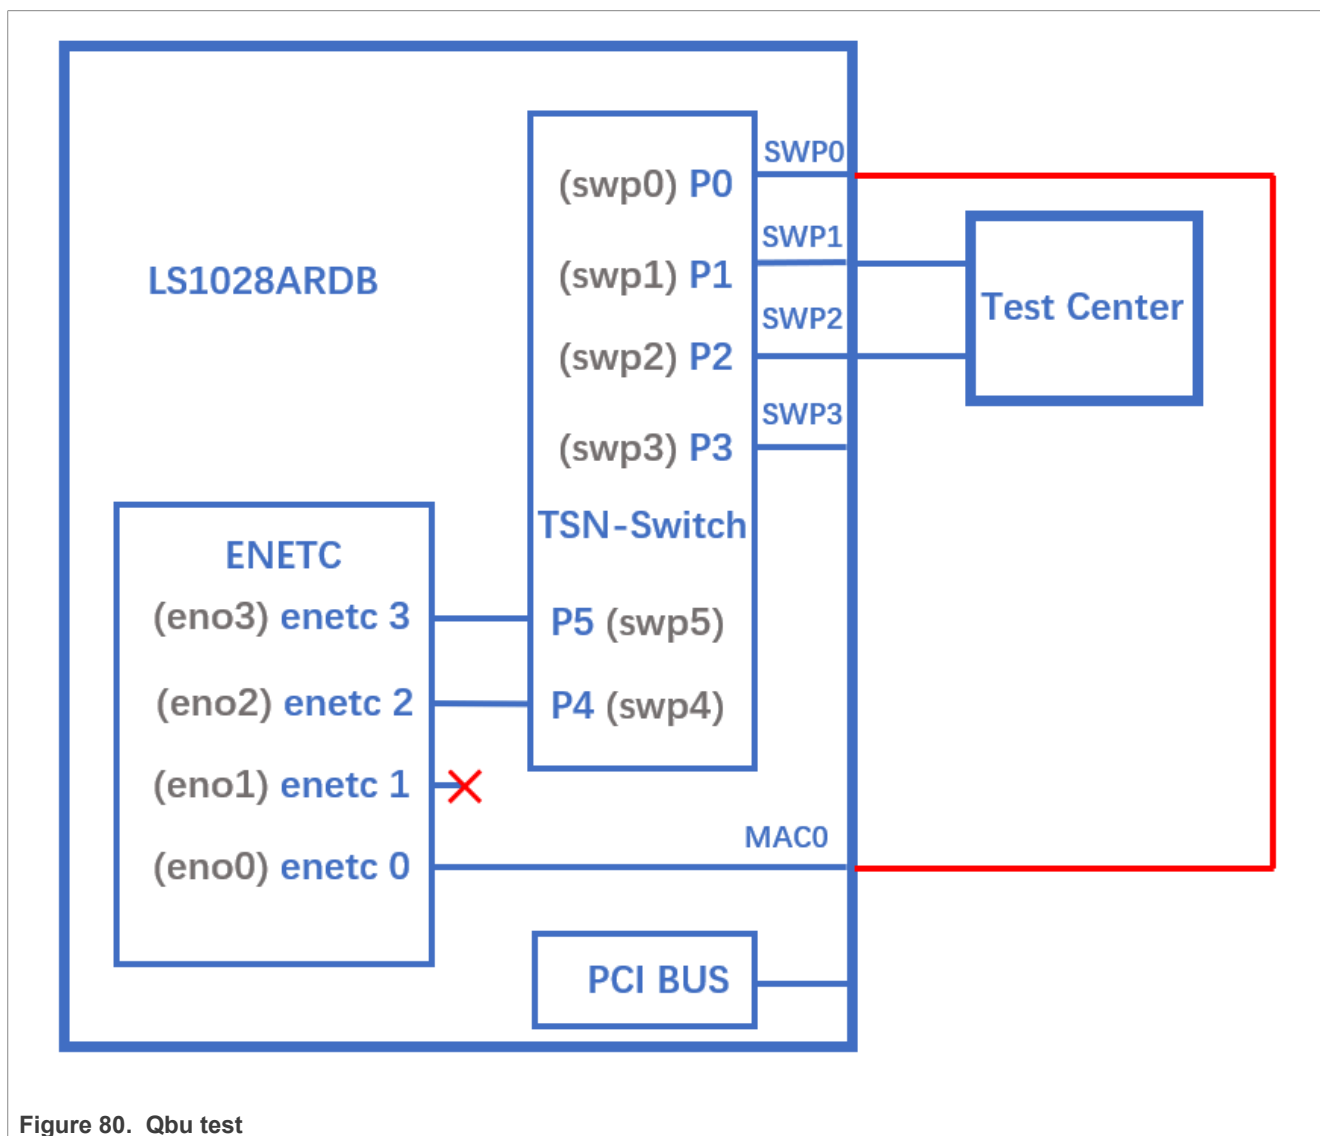


Figure 80. Qbu test

Before linking the cable between ENETC port0 to SWP0, set up the switch up (refer the [Switch configuration](#)) and set IP for ENETC port0. To make sure the ENETC port0 is linked to SWP0, use the steps below:

1. Ensure to enable the priority for each traffic class:

```
tc qdisc add dev eno0 root handle 1: mqprio num_tc 8 map 0 1 2 3 4 5 6 7
queues 1@0 1@1 1@2 1@3 1@4 1@5 1@6 1@7 hw 1
```


2. Make sure that the link speed is 1 Gbit/s by using the command:

```
ethtool eno0
```

3. If it is not 1 Gbit/s, set it to 1 Gbit/s by using the command:

```
ethtool -s swp0 speed 1000 duplex full autoneg on
```

4. Set the switch to enable merge (or user can link to another merge capability port in another board):

```
ethtool --set-mm swp0 tx-enabled on pmac-enabled on
```

5. ENETC port setting set and frame preemption test:

```
ip link set eno0 address 90:e2:ba:ff:ff:ff  
tsntool qbuset --device eno0 --preemptable 0xfe  
/home/root/samples/pktgen/pktgen_twoqueue.sh -i eno0 -q 0 -s 100 -n 20000 -m  
90:e2:ba:ff:ff:ff
```

pktgen would flood frames on TC0 and TC1.

6. Check the TX merge counter, if it has a non-zero value, it indicates that the Qbu is working.

```
tsntool regtool 0 0x11f18
```

Note: 0x11f18 counting the merge frame count:

0x11f18 Port MAC Merge Fragment Count TX Register (MAC_MERGE_MMFCTXR)

LS1028ARDB also supports ethtool setup for preemption as in the example below:

```
ethtool --set-frame-preemption eno0 preemptible-queues-mask 0xfe
```

This implies that we can get same result by using TC0 to pass express MAC and TC1~TC7 to pass preemptable MAC.

The ENETC also supports preemption verify. Use two boards to test preemption verification on eno0. Refer to [Section 5.1.3.4.1](#).

5.1.4.1.4 Qci

Use the following as the background setting:

- Set eno0 MAC address

```
ip link set eno0 address 10:00:80:00:00:00
```

Opposite port MAC address **99:aa:bb:cc:dd:ee** as frame provider as example.

- Use the figure below as the hardware setup.

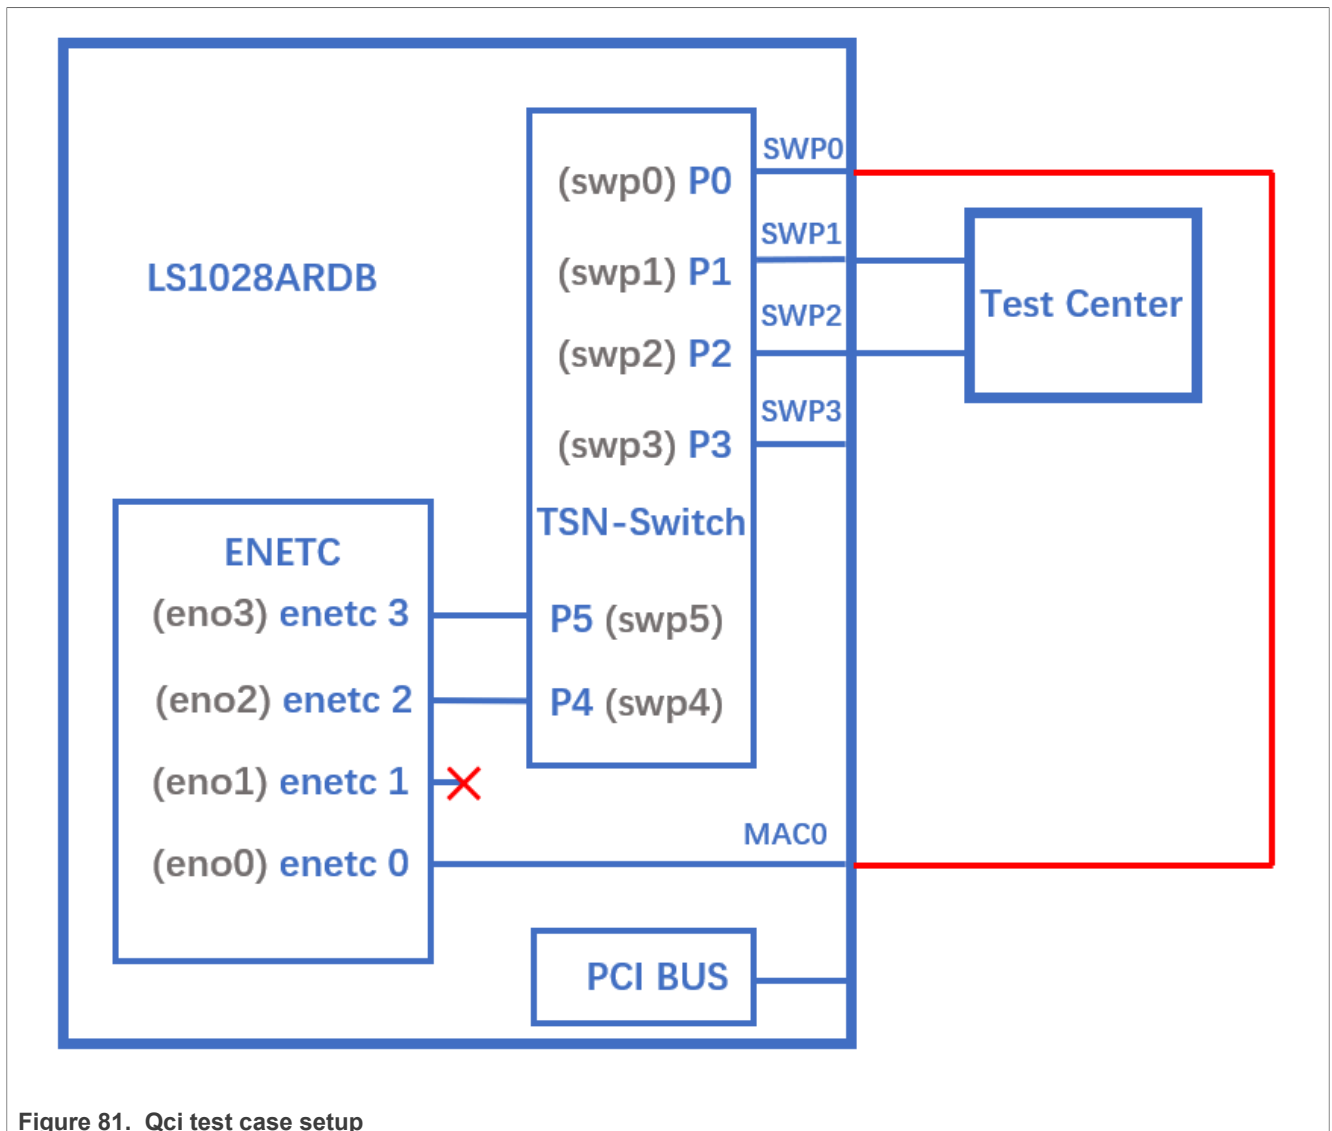


Figure 81. Qci test case setup

Note: TestCenter is a device to send streams to enetc0 of LS1028ardb board. User also can use another board to send streams.

5.1.4.1.4.1 Test SFI No Streamhandle

Qci PSFP can work for the streams without stream identify module, which are the streams without MAC address and vid filter. Such kind of filter setting always sets a larger index number stream for filter entry. The frames that are not filtered then flow into this stream filter entry.

The below example tests no streamhandle in a stream filter, set on stream filter entry index 2 with a gate stream entry id 2. Then none stream identifies frames would flow into the stream filter entry index 2 then pass the gate entry index 2, as shown in the following example:

```
tsntool> qcisfiset --device eno0 --index 2 --gateid 2
```

- Streams no streamhandle should pass this filter.

```
tsntool> qcisfiget --device eno0 --index 2
```

- Send a frame from the opposite device port (ping for example).

```
tsntool> qcisfiget --device eno0 --index 2
```

- Set Stream Gate entry 2

```
tsntool> qcisgiset --device eno0 --index 2 --initgate 1
```

- Send a frame from the opposite device port.

```
tsntool> qcisfiget --device eno0 --index 2
```

- Set Stream Gate entry 2, gate closes permanently.

```
tsntool> qcisgiset --device eno0 --index 2 --initgate 0
```

- Send a frame from the opposite device port.

```
tsntool> qcisfiget --device eno0 --index 2
#The result should look like below:
match pass gate_drop sdu_pass sdu_drop red
  1   0   1       1       0       0
```

5.1.4.1.4.2 Testing null stream identify entry

Null stream identify in stream identify module means trying to filter using destination MAC address and vlan id.

Following steps show the stream identify entry index 1 set with filtering destination mac address as 10:00:80:00:00:00 and vlan id ignored (with or without vland id). Then stream filter is set on the entry index 1 with stream gate index entry id 1.

1. Set main stream by closing gate.
2. Set Stream identify Null stream identify entry 1.

```
tsntool> cbstreamidset --device eno0 --index 1 --nullstreamid --nullldmac
0x0000000800010 --nulltagged 3 --nullvid 10 --streamhandle 100
```

3. Get stream identify entry index 1.

```
tsntool> cbstreamidget --device eno0 --index 1
```

4. Set Stream filter entry 1 with stream gate entry id 1.

```
tsntool> qcisfiset --device eno0 --streamhandle 100 --index 1 --gateid 1
```

5. Set Stream Gate entry 1, keep gate state close (all frames dropped. return directly if ask user for editing gate list).

```
tsntool> qcisgiset --device eno0 --index 1 --initgate 0
```

6. Send one frame from the opposite device port should pass to the close gate entry id 1.

```
tsntool> qcisfiget --device eno0 --index 1
```

7. The result should look like the output below:

```
match pass gate_drop sdu_pass sdu_drop red
  1  0  1  1  0  0
```

5.1.4.1.4.3 Testing source stream identify entry

Source stream identify means stream identify the frames by the source mac address and vlan id.

Use the following steps for this test:

1. Keep Stream Filter entry 1 and Stream gate entry 1.
2. Add stream2 in opposite device port: SMAC is 66:55:44:33:22:11 DMAC:20:00:80:00:00:00 (Not with destination mac address 10:00:80:00:00:00 which stream identify entry index 1 is filtering that dmac address)
3. Set Stream identify Source stream identify entry 3

```
tsntool> cbstreamidset --device eno0 --index 3 --sourcemacvid --sourcemac 0x112233445566 --sourcetagged 3 --sourcevid 20 --streamhandle 100
```

4. Send frame from opposite device port. The frame passes to stream filter index 1.

```
tsntool> qcisfiget --device eno0 --index 1
```

5.1.4.1.4.4 SGI stream gate list

Use the command below for this test:

```
cat > sgil.txt << EOF
t0 0b -1 1000000000 0
t1 1b -1 1000000000 0
EOF
tsntool> qcisfiset --device eno0 --index 2 --gateid 2
tsntool> qcisgiset --device eno0 --index 2 --initgate 1 --gatelistfile sgil.txt
#flooding frame size 64bytes from opposite device port.(iperf or netperf as
example)
tsntool> qcisfiget --device eno0 --index 2
```

Check the frames dropped and passed, they should be the same since stream gate list is setting 100ms open and 100ms close periodically.

5.1.4.1.4.5 FMI test

Only send green color frames (normally it is the TCI bit value in 802.1Q tag). Flooding the stream against the eno0 port speed to 10000 kbps/s:

```
tsntool> qcisfiset --device eno0 --index 2 --gateid 2 --flowmeterid 2
tsntool> qcifmiset --device eno0 --index 2 --cm --cf --cbs 1500 --cir 5000 --ebs 1500 --eir 5000
```

The 'cm' parameter set color mode enable means frames to separate green frames and yellow frames judged by the TCI bit in frame. Or else, any frames are green frames.

The 'cf' parameter sets the coupling flag enable. When CF is set to 0, the frames that are declared yellow are bound by EIR. When CF is set to 1, the frames that are declared Yellow are bound by CIR + EIR, depending on volume of the offered frames that are declared Green.

After the above commands are setup, since green frames are not larger than EIR + CIR 10 Mbit/s. So the green frame would not be dropped.

The below setting shows the dropped frames:

```
tsntool> qcifmiset --device eno0 --index 2 --cm --cf --cbs 1500 --cir 5000 --ebs 1500 --eir 2000
```

This case makes the green frames pass 5 Mbit/s in CIR, then it pass to the EIR space. However, EIR is 2 Mbit/s, so total EIR + CIR 7 Mbit/s still do not qualify the total 10 Mbit/s bandwidth. So green frame would be dropped part.

To get information of color frame counters showing at application layer, use the code as in the below example:

```
tsntool> qcifmiget --device eno0 --index 2
=====
bytecount drop dr0_green dr1_green dr2_yellow remark_yellow dr3_red remark_red
1c89 0 4c 0 0 0 0 0
=====
index = 2
cir = c34c
cbs = 5dc
eir = 4c4b3c
ebs = 5dc
couple flag
color mode
```

5.1.4.1.5 Qav

5.1.4.1.5.1 Using tsntool

[Figure 82](#) illustrates the hardware setup diagram for the Qav test.




Figure 82. Qav test setup

Note: TestCenter is a device to capture streams from enetc0 of LS1028ARDB board. Users can also use another board to capture streams by using `tcpdump` command, and use Wireshark network protocol analyzer to analyze results.

0. Ensure to enable the priority for each traffic class:

```
tc qdisc add dev eno0 root handle 1: mqprio num_tc 8 map 0 1 2 3 4 5 6 7 queues
1@0 1@1 1@2 1@3 1@4 1@5 1@6 1@7 hw 1
```

1. Run the following commands:

```
tsntool cbsset --device eno0 --tc 7 --percentage 60
tsntool cbsset --device eno0 --tc 6 --percentage 20
```

2. Check each queue bandwidth (pktgen requires enabling NET_PKTGEN in kernel)

```
/home/root/samples/pktgen/pktgen_sample01_simple.sh -i eno0 -q 7 -s 500 -n
30000
```

Wait a few seconds later to check the result. It should get about 60% percentage line rate.

```
/home/root/samples/pktgen/pktgen_sample01_simple.sh -i eno0 -q 6 -s 500 -n 30000
```

Wait a few seconds later to check the result. It should get about 20% percentage line rate.

5.1.4.1.5.2 Using CBS Qdisc to setup Qav

LS1028a supports the CBS qdisc to setup Credit-based Shaper. Below commands set CBS with 100 Mbit/s for queue 7 and 300 Mbit/s for queue 6.

```
tc qdisc add dev eno0 root handle 1: mqprio num_tc 8 map 0 1 2 3 4 5 6 7 queues 1@0 1@1 1@2 1@3 1@4 1@5 1@6 1@7 hw 1
tc qdisc replace dev eno0 parent 1:8 cbs locredit -1350 hicredit 150 sendslope -900000 idleslope 100000 offload 1
tc qdisc replace dev eno0 parent 1:7 cbs locredit -1050 hicredit 950 sendslope -700000 idleslope 300000 offload 1
```

```
# Try to flood stream here (require kernel enable NET_PKTGEN)
/home/root/samples/pktgen/pktgen_sample01_simple.sh -i eno0 -q 7 -s 500 -n 20000
/home/root/samples/pktgen/pktgen_sample01_simple.sh -i eno0 -q 6 -s 500 -n 20000
tc qdisc del dev eno0 parent 1:7 cbs
tc qdisc del dev eno0 parent 1:8 cbs
```

5.1.4.2 TSN configuration on Felix switch

The following sections describe examples for the basic configuration of TSN switch.

5.1.4.2.1 Switch configuration

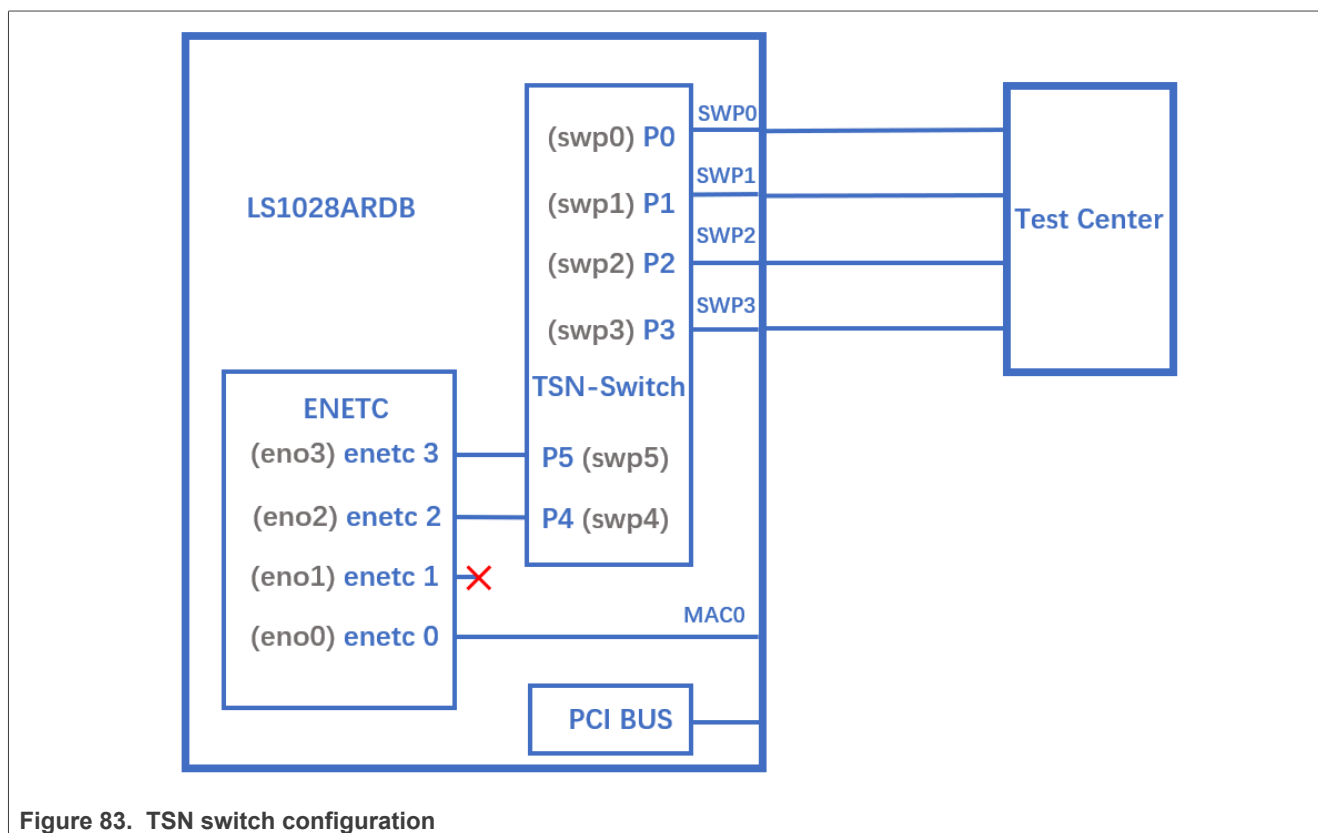


Figure 83. TSN switch configuration

Use the following commands to configure bridge on LS1028ARDB:

```
ls /sys/bus/pci/devices/0000:00:00.5/net/
```

Get switch device interfaces for swp0, swp1, swp2 and swp3 as shown below:

```
ip link set eno2 up
ip link add name switch type bridge vlan_filtering 1
ip link set switch up
ip link set swp0 master switch && ip link set swp0 up
ip link set swp1 master switch && ip link set swp1 up
ip link set swp2 master switch && ip link set swp2 up
ip link set swp3 master switch && ip link set swp3 up
```

5.1.4.2.2 Linuxptp test

To run PTP clock synchronization cases on TSN switch in:

- 4.3.5 Quick Start for IEEE 1588
- 4.3.6 Quick Start for IEEE 802.1AS

There are additional configurations of PTP packets trapping besides basic L2 switch mode configuration of “4.1.4.2.1 Switch configuration”. An available IP should be configured on bridge, and don’t configure IP on swpX interfaces.

```
$ ./switch-ptp-trap.sh
```



```
$ ifconfig switch <ip_address>
```

switch-ptp-trap.sh

```
# swp0, trap ptp
tc qdisc add dev swp0 clsact
tc filter add dev swp0 ingress chain 0 pref 49152 flower skip_sw action goto
chain 10000
tc filter add dev swp0 ingress chain 10000 pref 49152 flower skip_sw action goto
chain 11000
tc filter add dev swp0 ingress chain 11000 pref 49152 flower skip_sw action goto
chain 12000
tc filter add dev swp0 ingress chain 12000 pref 49152 flower skip_sw action goto
chain 20000
tc filter add dev swp0 ingress chain 20000 pref 49152 flower skip_sw action goto
chain 21000
tc filter add dev swp0 ingress chain 21000 pref 49152 flower skip_sw action goto
chain 30000
tc filter add dev swp0 ingress chain 20000 protocol 0x88f7 flower skip_sw action
trap action goto chain 21000
tc filter add dev tc filter add dev swp0 ingress chain 20000 protocol ip flower
skip_sw dst_ip 224.0.1.129 action trap action goto chain 21000
tc filter add dev tc filter add dev swp0 ingress chain 20000 protocol ip flower
skip_sw dst_ip 224.0.0.107 action trap action goto chain 21000

# swp1, trap ptp
tc qdisc add dev swp1 clsact
tc filter add dev swp1 ingress chain 0 pref 49152 flower skip_sw action goto
chain 10000
tc filter add dev swp1 ingress chain 10000 pref 49152 flower skip_sw action goto
chain 11000
tc filter add dev swp1 ingress chain 11000 pref 49152 flower skip_sw action goto
chain 12000
tc filter add dev swp1 ingress chain 12000 pref 49152 flower skip_sw action goto
chain 20000
tc filter add dev swp1 ingress chain 20000 pref 49152 flower skip_sw action goto
chain 21000
tc filter add dev swp1 ingress chain 21000 pref 49152 flower skip_sw action goto
chain 30000
tc filter add dev swp1 ingress chain 20000 protocol 0x88f7 flower skip_sw action
trap action goto chain 21000
tc filter add dev swp1 ingress chain 20000 protocol ip flower skip_sw dst_ip
224.0.1.129 action trap action goto chain 21000
tc filter add dev swp1 ingress chain 20000 protocol ip flower skip_sw dst_ip
224.0.0.107 action trap action goto chain 21000

# swp2, trap ptp
tc qdisc add dev swp2 clsact
tc filter add dev swp2 ingress chain 0 pref 49152 flower skip_sw action goto
chain 10000
tc filter add dev swp2 ingress chain 10000 pref 49152 flower skip_sw action goto
chain 11000
tc filter add dev swp2 ingress chain 11000 pref 49152 flower skip_sw action goto
chain 12000
tc filter add dev swp2 ingress chain 12000 pref 49152 flower skip_sw action goto
chain 20000
tc filter add dev swp2 ingress chain 20000 pref 49152 flower skip_sw action goto
chain 21000
```

```
tc filter add dev swp2 ingress chain 21000 pref 49152 flower skip_sw action goto
chain 30000
tc filter add dev swp2 ingress chain 20000 protocol 0x88f7 flower skip_sw action
trap action goto chain 21000
tc filter add dev swp2 ingress chain 20000 protocol ip flower skip_sw dst_ip
224.0.1.129 action trap action goto chain 21000
tc filter add dev swp2 ingress chain 20000 protocol ip flower skip_sw dst_ip
224.0.0.107 action trap action goto chain 21000

# swp3, trap ptp
tc qdisc add dev swp3 clsact
tc filter add dev swp3 ingress chain 0 pref 49152 flower skip_sw action goto
chain 10000
tc filter add dev swp3 ingress chain 10000 pref 49152 flower skip_sw action goto
chain 11000
tc filter add dev swp3 ingress chain 11000 pref 49152 flower skip_sw action goto
chain 12000
tc filter add dev swp3 ingress chain 12000 pref 49152 flower skip_sw action goto
chain 20000
tc filter add dev swp3 ingress chain 20000 pref 49152 flower skip_sw action goto
chain 21000
tc filter add dev swp3 ingress chain 21000 pref 49152 flower skip_sw action goto
chain 30000
tc filter add dev swp3 ingress chain 20000 protocol 0x88f7 flower skip_sw action
trap action goto chain 21000
tc filter add dev swp3 ingress chain 20000 protocol ip flower skip_sw dst_ip
224.0.1.129 action trap action goto chain 21000
tc filter add dev swp3 ingress chain 20000 protocol ip flower skip_sw dst_ip
224.0.0.107 action trap action goto chain 21000

# ebttables, route ptp, not bridge
ebtables --table broute --append BROUTING --protocol 0x88F7 --jump DROP
ebtables --table broute --append BROUTING --protocol 0x0800 --ip-protocol udp --
ip-destination-port 320 --jump DROP
ebtables --table broute --append BROUTING --protocol 0x0800 --ip-protocol udp --
ip-destination-port 319 --jump DROP
```

5.1.4.2.3 Qbv test setup for LS1028ARDB

The following figure describes the setup for Qbv test on LS1028ARDB.

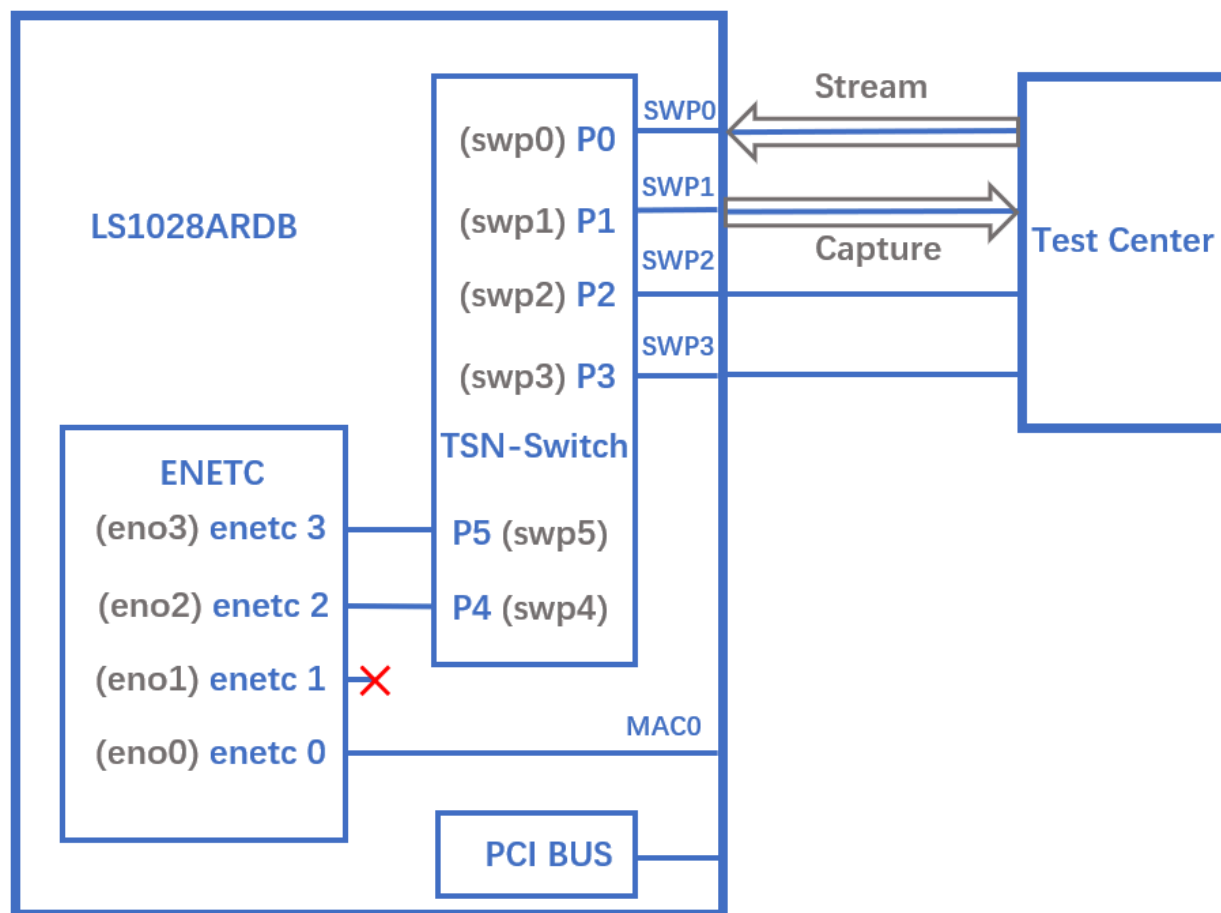


Figure 84. Qbv test

Reserve buffer for each queue on ingress and egress port to avoid resource depletion when Qbv gate is closed.

```

ingressport=0
egressport=1
for tc in {0..7}; do {
    devlink sb tc bind set pci/0000:00:00.5/$ingressport sb 0 tc $tc type
    ingress pool 0 th 3000
    devlink sb tc bind set pci/0000:00:00.5/$ingressport sb 1 tc $tc type
    ingress pool 0 th 10
    devlink sb tc bind set pci/0000:00:00.5/$egressport sb 0 tc $tc type egress
    pool 1 th 3000
    devlink sb tc bind set pci/0000:00:00.5/$egressport sb 1 tc $tc type egress
    pool 1 th 10
}
done

```

5.1.4.2.3.1 Using tsntool

Closing basic gates

Use the set of commands below for basic gate closing.

```
echo "t0 00000000b 20000" > qbv0.txt
#Explanation:
# 'NUMBER'      : t0
# 'GATE_VALUE'   : 00000000b
# 'TIME_LONG'    : 20000 ns
./tsntool
tsntool> verbose
tsntool> qbvset --device swp1 --entryfile ./qbv0.txt
#Send one broadcast frame to swp0 from TestCenter.
ethtool -S swp1
#Should not get any frame from swp1 on TestCenter.
echo "t0 11111111b 20000" > qbv0.txt
tsntool> qbvset --device swp1 --entryfile ./qbv0.txt
#Send one broadcast frame to swp0 on TestCenter.
ethtool -S swp1
#Should get one frame from swp1 on TestCenter.
```

Basetime test

For the basetime test, first get the current time in seconds:

```
#Get current time:
tsntool> ptptool -g -d /dev/ptp1

#add some seconds, for example user gets 200.666 time clock, then set 260.666 as
result

tsntool> qbvset --device swp1 --entryfile ./qbv0.txt --basetime 260.666

#Send one broadcast frame to swp0 on the TestCenter.
#Frame could not pass swp1 until time offset.
```

Qbv performance test

Use the following commands for the Qbv performance test:

```
cat > qbv5.txt << EOF
t0 11111111b 1000000
t1 00000000b 1000000
EOF
qbvset --device swp1 --entryfile qbv5.txt
```

#Send 1G rate stream to swp0 on TestCenter.

#The stream would get about half line rate from swp1.

Note: Each entry time must be larger than guard band, the guard band is set by "--maxsdu". If this parameter is not set, use default 1518Bytes, the least entry time is $(1518 \times 8) / 1G \approx 12\mu s$.

5.1.4.2.3.2 Tc-taprio usage

LS1028ARDB supports the taprio qdisc to setup Qbv either. Below is an example setup.

1. Enable the Qbv for swp1 port, set queue 1 gate open, set circle time to be 300 μ s.

```
tc qdisc replace dev swp1 parent root handle 100 taprio num_tc 8 map 0 1 2 3 4 5
6 7 \
    queues 1@0 1@1 1@2 1@3 1@4 1@5 1@6 1@7 base-time 0 sched-entry S 02
300000 flags 0x2
```

Note: Since the hardware can only use PCP, DSCP or other methods to classify QoS, it cannot map QoS to different hardware queues. mqprio is not implemented in the felix driver, so "map 0 1 2 3 4 5 6 7" in the tc-taprio command is invalid.

Note: Tc-taprio uses default port max SDU(1518B) as guard band value. Each entry time must be larger than guard band($1518 \times 8 / 1G \approx 12\mu s$).

2. Send one frame with PCP=1 in vlan tag to swp0 from TestCenter, so as to capture the frame from swp1.
3. Send one frame with PCP=2 in vlan tag to swp0 from TestCenter, gate is closed and the frame from swp1 cannot be captured.
4. Disable the Qbv for swp1 port as below:

```
tc qdisc del dev swp1 parent root handle 100 taprio
```

5.1.4.2.4 Qbu

[Figure 85](#) illustrates the setup for performing the Qbu test using the TSN switch.

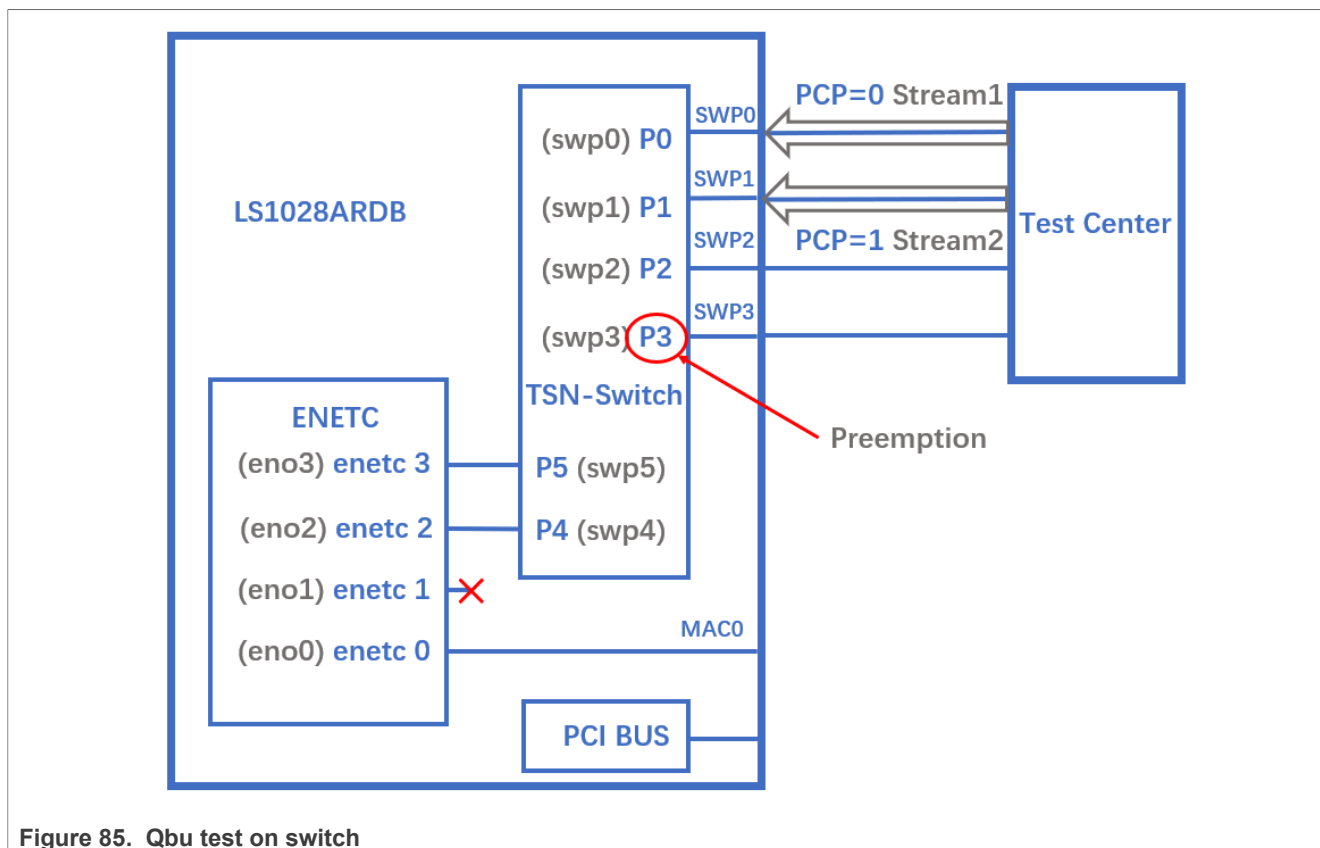


Figure 85. Qbu test on switch

5.1.4.2.4.1 Using tsntool for Qbu test

1. Disable the Cut-through mode before enabling preemption on switch ports.

```
# tsntool> ctset --device swp3 --queue_stat 0x0
```

2. Set queue 1 to be preemptable. There are two ways to set preemptable queues; users can choose either tsntool or ethtool to set it.

```
#tsntool command to set preemptable queues:
tsntool> qbuset --device swp3 --preemptable 0x02
```

3. Send two streams from TestCenter, set packet size to be 1500 byte and bandwidth to be 1G. Now, check the number of additional mPackets transmitted by PMAC using the command below:

```
ethtool -I --show-mm swp3 | grep MACMergeFragCountTx
```

4. Follow the steps below to perform Qbu combined with Qbv test.
Set queue 0 gate open 20 μ s, queue 1 gate open to 20 μ s.

```
cat > qbv0.txt << EOF
t0 00000001b 200000
t1 00000010b 200000
EOF
qbvset --device swp3 --entryfile qbv0.txt
```

Send two streams from TestCenter. Observe that packets in queue 1 are preempted when gate 1 is closed.

5.1.4.2.4.2 Ethtool usage

1. Set queue 1 to be preemptable. There are two ways to set preemptable queues. Users can choose either **tsntool** or **ethtool** to set it.

```
#ethtool command to set preemptable queues:
ethtool --set-frame-preemption swp3 preemptible-queues-mask 0x02 min-frag-
size 124
```

Explanation:

- **preemptible-queues-mask**: An 8-bit vector that specifies preemptable queues within the 8 priorities (with bit-0 for priority-0 and bit-7 for priority-7).
 - **min-frag-size**: specifies the least frame bytes that have been transmitted in the fragment. The minimum non-final fragment size is 64, 128, 192, or 256 octets (include 4 Bytes fragment header).
2. Send two streams from TestCenter. Set the packet size to be 1500 bytes and bandwidth to be 1 G. Now, check the number of additional mPackets transmitted by the PMAC:

```
ethtool -I --show-mm swp3 | grep MACMergeFragCountTx
```

3. Qbu combined with Qbv test.

Set queue 0 gate open 20 μ s, queue 1 gate open 20 μ s.

```
tc qdisc replace dev swp3 parent root handle 100 taprio num_tc 8 map 0 1 2 3
  4 5 6 7 \
    queues 1@0 1@1 1@2 1@3 1@4 1@5 1@6 1@7 base-time 0 \
      sched-entry S 01 200000 \
      sched-entry S 02 200000 flags 0x2
```

Send two streams from TestCenter. Note that packets in queue 1 are preempted when gate 1 closed.

4. The Felix switch port also supports preemption verification. Use two boards to test preemption verification on swp0-3. Refer to ["4.1.3.4.1 Preemption verify"](#).

5.1.4.2.5 Qci

The figure below illustrates the Qci test case setup.

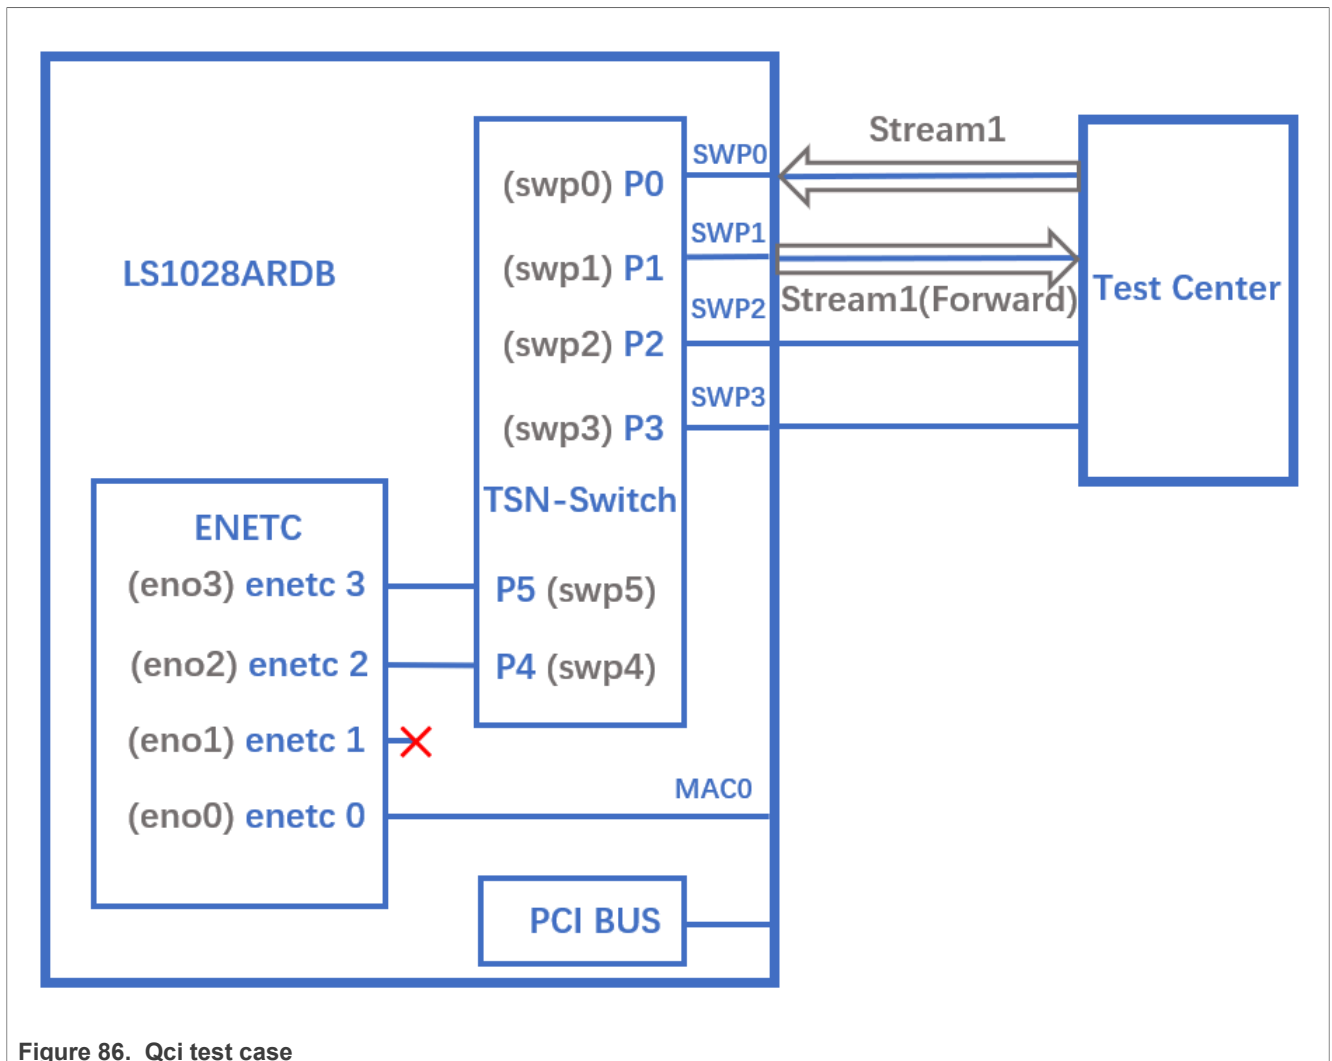


Figure 86. Qci test case

5.1.4.2.5.1 Using Tsntool for stream identification and other tests

The following section describes how to use tsntool for stream identification, stream gate control, SFI maxSDU test, FMI test, and Port-based SFI set.

Stream identification

Use the following commands for stream identification:

1. Set a stream to `swp0` on TestCenter. Edit the stream, set the destination MAC as: `00:01:83:fe:12:01`, Vlan ID : 1
2. Add the MAC to MAC table on LS1028a. (This step is not needed if the mac is already learned on port)

```
bridge fdb add 00:01:83:fe:12:01 dev swp1 vlan 1 master static
```

3. Use the destination MAC as: `00:01:83:fe:12:01`, vlan ID : 1 to set the stream identification on LS1028A.

```
tsntool> cbstreamidset --device swp1 --nullstreamid --index 1 --nullldmac  
0x000183fe1201 --nullvid 1 --streamhandle 1
```


Explanation:

- **device:** set the device port which is the stream forwarded to. If the {destmac, VID} is already learned by switch, switch will not care device port.
- **nulltagged:** switch only support nulltagged=1 mode, so there is no need to set it.
- **nullvid:** Use "bridge vlan show" to see the ingress VID of switch port.

```
tsntool> qcisfiset --device swp0 --index 1 --streamhandle 1 --gateid 1 --
priority 0 --flowmeterid 68
```

Explanation:

- **device:** can be any one of switch ports.
 - **index:** value is the same as streamhandle of cbstreamidset.
 - **streamhandle:** value is the same as streamhandle of cbstreamidset.
 - **flowmeterid:** PSFP Policer id, ranges from 63 to 383.
4. Send one frame, then check the frames.

```
ethtool -S swp1
ethtool -S swp2
```

Only swp1 can get the frame.

5. Use the following command to check and debug the stream identification status.

```
qcisfiget --device swp0 --index 1
```

Note: The parameter *streamhandle* is the same as *index* in stream filter set, we use *streamhandle* as SFID to identify the stream, and use *index* to set stream filter table entry.

Stream gate control

1. Use the following commands for stream gate control:

```
echo "t0 1b 3 50000 200" > sgi.txt
tsntool> qcisgiset --device swp0 --enable --index 1 --initgate 1 --initipv 0
--gatelistfile sgi.txt --basetime 0x0
```

Explanation:

- 'device': can be any one of switch ports.
 - 'index': gateid
 - 'basetime': It is the same as Qbv set.
2. Send one frame on TestCenter.

```
ethtool -S swp1
```

Note that the frame could pass, and green_prio_3 has increased.

3. Now run the following commands:

```
echo "t0 0b 3 50000 200" > sgi.txtx
tsntool> qcisgiset --device swp0 --enable --index 1 --initgate 1 --initipv 0
--gatelistfile sgi.txt --basetime 0x0
```

4. Next, send one frame on TestCenter.

```
ethtool -S swp1
```

Note that the frame could not pass.

SFI maxSDU test

Disable the cut-through mode on swp0 and swp1:

```
tsntool> ctset --device swp0 --queue_stat 0
tsntool> ctset --device swp1 --queue_stat 0
```

Use the following command to run this test:

```
tsntool> qcisfiset --device swp0 --index 1 --gateid 1 --priority 0 --flowmeterid
68 --maxsdu 200
```

Now, send one frame (frame size > 200) on TestCenter.

```
ethtool -S swp1
```

Users can observe that the frame could not pass.

FMI test

Use the following set of commands for the FMI test.

1. Reserve buffer for each queue on ingress port to receive yellow frames(dp=1) in switch.

```
ingressport=0
for tc in {0..7}; do {
    devlink sb tc bind set pci/0000:00:00.5/$ingressport sb 0 tc $tc type
    ingress pool 0 th 3000
    devlink sb tc bind set pci/0000:00:00.5/$ingressport sb 1 tc $tc type
    ingress pool 0 th 10
}
done
```

2. Run the command:

```
tsntool> qcifmisset --device swp0 --index 68 --cir 100000 --cbs 4000 --ebs
4000 --eir 100000
```

Note:

- The 'device' in above command can be any one of the switch ports.
- The index of *qcifmisset* must be the same as flowmeterid of *qcisfiset*.

3. Now, send one stream (rate = 100M) on TestCenter.

```
ethtool -S swp0
```

Note that all frames pass and get all green frames.

4. Now, send one stream (rate = 200M) on TestCenter.

```
ethtool -S swp0
```

Observe that all frames pass and get green and yellow frames.

5. Send one stream (rate = 300M) on TestCenter.

```
ethtool -S swp0
```

Note that not all frames could pass and get green, yellow, and red frames.

6. Send one yellow stream (rate = 100M) on TestCenter.

```
ethtool -S swp0
```

All frames pass and get all yellow frames.

7. Send one yellow stream (rate = 200M) on TestCenter.

```
ethtool -S swp0
```

Note that not all frames could pass and get yellow and red frames.

8. Test cf mode.

```
tsntool> qcifmiset --device swp0 --index 68 --cir 100000 --cbs 4000 --ebs 4000 --eir 100000 --cf
```

9. Send one yellow stream (rate = 200M) on TestCenter.

```
ethtool -S swp0
```

All frames pass and get all yellow frames (use CIR as well as EIR).

10. Send one yellow stream (rate = 300M) on TestCenter.

```
ethtool -S swp0
```

Note that not all frames could pass and get yellow and red frames.

Port-based SFI set

LS1028A switch can work on port-based PSFP set. This implies that when a null-identified stream is received on an ingress port, switch will use the port, default SFI.

Below example tests no streamhandle in qcisfiset to set a port, default SFI.

1. Use SFID 2 to set swp0 port as default SFI.

```
tsntool> qcisfiset --device swp0 --index 2 --gateid 1 --flowmeterid 68
```

After the port default SFI set, any stream sent from swp0 port will do the gate 1 and flowmeter 68 policy.

2. Set stream gate control.

```
echo "t0 1b 4 50000 200" > sgi.txt
tsntool> qcisgiset --device swp0 --enable --index 1 --initgate 1 --initipv 0 --
gatelistfile sgi.txt
```

3. Send any stream to swp0.

```
ethtool -S swp1
```

Note that the frame could pass, and green_prio_4 has increased.

5.1.4.2.5.2 Tc-flower usage

[Figure 87](#) illustrates the TC-flower-based Qci test case setup.

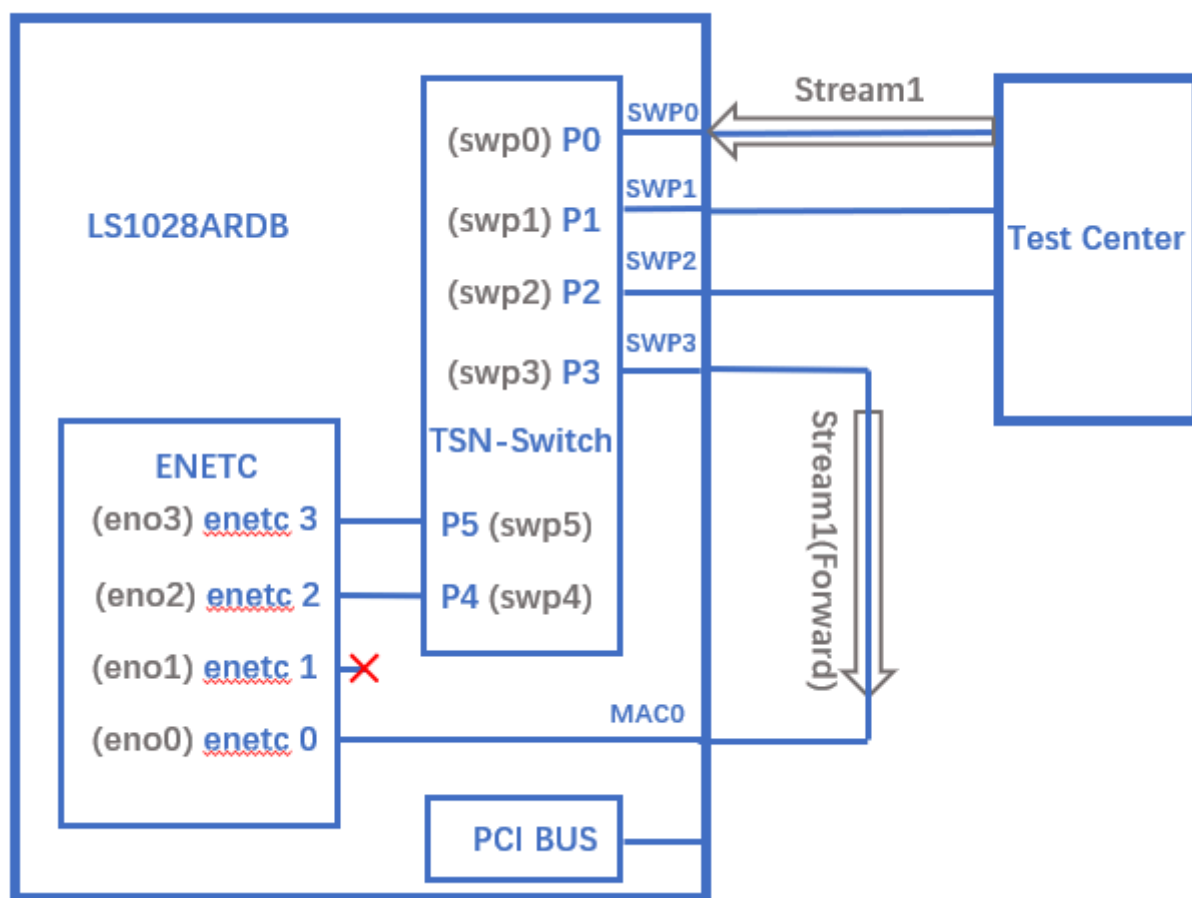


Figure 87. TC-flower based Qci test case

1. Add the MAC address 'CA:9C:00:BC:6D:68' in the MAC table by using `bridge fdb` command if it is not learned.

```
bridge fdb add dev swp3 CA:9C:00:BC:6D:68 vlan 1 master static
```

2. Register chains on ingress port swp0. Refer to [Section 5.4.2](#).

```
tc qdisc add dev swp0 clsact
tc filter add dev swp0 ingress chain 0 pref 49152 flower skip_sw action goto
chain 10000
tc filter add dev swp0 ingress chain 10000 pref 49152 flower skip_sw action goto
chain 11000
tc filter add dev swp0 ingress chain 11000 pref 49152 flower skip_sw action goto
chain 12000
tc filter add dev swp0 ingress chain 12000 pref 49152 flower skip_sw action goto
chain 20000
tc filter add dev swp0 ingress chain 20000 pref 49152 flower skip_sw action goto
chain 21000
tc filter add dev swp0 ingress chain 21000 pref 49152 flower skip_sw action goto
chain 30000
```

3. Set Qci on ingress port swp0.

a) Use the following commands to set Qci gate.

```
tc filter add dev swp0 ingress chain 30000 protocol 802.1Q flower skip_sw
dst_mac CA:9C:00:BC:6D:68 vlan_id 1 action gate index 1 base-time 0 sched-entry
CLOSE 6000 -1 -1
```

b). Use the following commands to set Qci flow meter.

```
tc filter add dev swp0 ingress chain 30000 protocol 802.1Q flower skip_sw
dst_mac CA:9C:00:BC:6D:68 vlan_id 1 action police index 1 rate 10Mbit burst
10000 conform-exceed drop/ok
```

c). Use the following commands to set Qci SFI priority.

```
tc filter add dev swp0 ingress chain 30000 protocol 802.1Q flower skip_sw
dst_mac CA:9C:00:BC:6D:68 vlan_id 1 vlan_prio 1 action gate index 1 base-time 0
sched-entry CLOSE 6000 -1 -1
```

d). Use the following commands to set both gate and flow meter.

```
tc filter add dev swp0 ingress chain 30000 protocol 802.1Q flower skip_sw
dst_mac CA:9C:00:BC:6D:68 vlan_id 1 action gate index 1 base-time 0 sched-entry
OPEN 6000 2 -1 action police index 1 rate 10Mbit burst 10000 conform-exceed
drop/ok
```

3. Send a stream from TestCenter, set the stream destination mac as CA:9C:00:BC:6D:68, set vid=1 and vlan_prio=1 in the vlan tag.

4. Using "tcpdump -i eno0 -w eno0.pcap" to receive the stream on eno0, check if packets are received.

5. Use the following commands to delete a stream rule.

```
tc -s filter show dev swp0 ingress chain 30000
tc filter del dev swp0 ingress chain 30000 pref 49152
```

Note:

- Each stream can only be added only once. If a user wants to update it, delete the rule and add a new one.
- MAC and VID of stream must have been learned in switch MAC table if the stream is required to be added.
- Qci gate cycle time is expected to be more than 5 μ s.
- Qci flow meter can only set *cir* and *cbs* now, and the policers are shared with ACL VCAPs.

5.1.4.2.6 Qav

[Figure 88](#) illustrates the Qav test case setup.

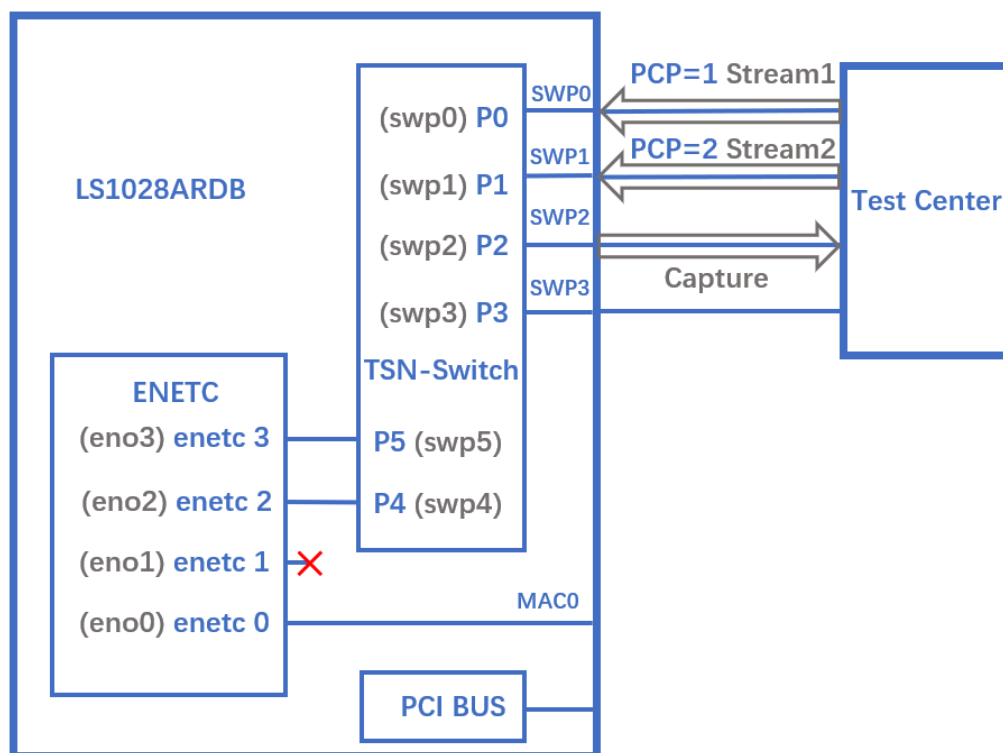


Figure 88. Qav test case

5.1.4.2.6.1 Tsntool usage

1. Set the percentage of two traffic classes:

```
tsntool> ctset --device swp0 --queue_stat 0x0
tsntool> ctset --device swp1 --queue_stat 0x0
tsntool> ctset --device swp2 --queue_stat 0x0
tsntool> cbsset --device swp2 --tc 1 --percentage 20
tsntool> cbsset --device swp2 --tc 2 --percentage 40
```

2. Send two streams from TestCenter, then check the frames count.

```
ethtool -S swp2
```

Note that the frame count of queue1 is half of queue2.

Note: Stream rate must larger than bandwidth limited of queue.

3. Capture frames on swp2 on TestCenter.

The Get Frame sequence is: (PCP=1), (PCP=2), (PCP=2), (PCP=1), (PCP=2), (PCP=2), ...

5.1.4.2.6.2 Tc-cbs usage

LS1028A supports the CBS qdisc to setup Credit-based Shaper. The below commands set CBS with 20 Mbit/s for queue 1 and 40 Mbit/s for queue 2.

1. Set the cbs of two traffic classes:

```
tc qdisc add dev swp2 root handle 1: mqprio num_tc 8 map 0 1 2 3 4 5 6 7 \
  queues 1@0 1@1 1@2 1@3 1@4 1@5 1@6 1@7 hw 0
tc qdisc replace dev swp2 parent 1:2 cbs locredit -1470 hicredit 30 \
  sendslope -980000 idleslope 20000 offload 1
tc qdisc replace dev swp2 parent 1:3 cbs locredit -1440 hicredit 60 \
  sendslope -960000 idleslope 40000 offload 1
```

2. Send one stream with PCP=1 from TestCenter, we can get the stream bandwidth is 20 Mbit/s from swp2.
3. Send two streams from TestCenter, then check the frames count.

```
ethtool -S swp2
```

Note: The frame count of queue1 is half of queue2.

4. Delete the cbs rules.

```
tc qdisc del dev swp2 parent 1:2 cbs
tc qdisc del dev swp2 parent 1:3 cbs
```

5.1.4.2.7 802.1CB

The [Figure 89](#) describes the test setup for the seamless redundancy test case.

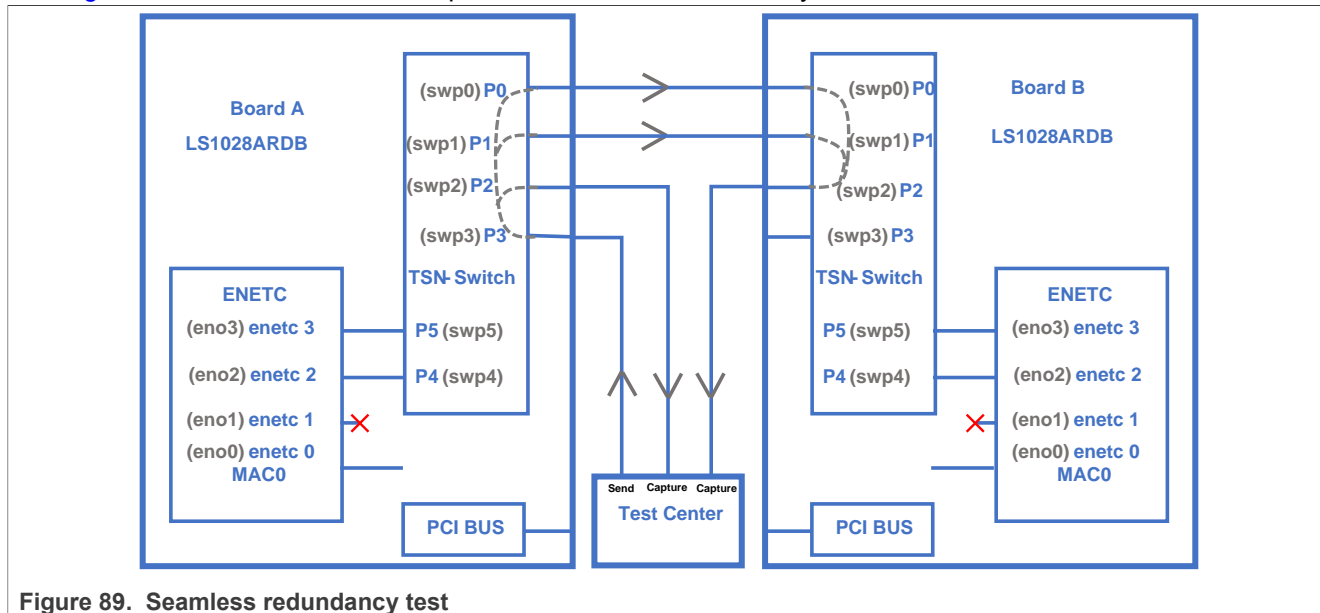


Figure 89. Seamless redundancy test

5.1.4.2.7.1 Sequence Generator test

Use the following set of commands for the 'Sequence Generator' test.

1. Configure switch ports to be forward mode.

On board A:

```
ifconfig eno2 up
ip link add name switch type bridge vlan_filtering 1
ip link set switch up
ip link set swp0 master switch && ip link set swp0 up
```

```
ip link set swp1 up
ip link set swp2 master switch && ip link set swp2 up
ip link set swp3 master switch && ip link set swp3 up
bridge vlan add dev swp0 vid 1 pvid
bridge vlan add dev swp2 vid 1 pvid
bridge vlan add dev swp3 vid 1 pvid
```

On board B

```
ifconfig eno2 up
ip link add name switch type bridge vlan_filtering 1
ip link set switch up
ip link set swp0 master switch && ip link set swp0 up
ip link set swp1 master switch && ip link set swp1 up
ip link set swp2 master switch && ip link set swp2 up
ip link set swp3 master switch && ip link set swp3 up
bridge vlan add dev swp0 vid 1 pvid
bridge vlan add dev swp1 vid 1 pvid
bridge vlan add dev swp2 vid 1 pvid
bridge vlan add dev swp3 vid 1 pvid
```

2. On board A, run the commands:

```
bridge fdb add 7E:A8:8C:9B:41:DD dev swp2 vlan 1 master static
tsntool> cbstreamidset --device swp0 --index 1 --nullstreamid --nulldmac
0x7EA88C9B41DD --nullvid 1 --streamhandle 1
tsntool> cbgen --device swp3 --index 1 --iport_mask 0x08 --split_mask 0x07 --
seq_len 16 --seq_num 2048
```

In the command above,

- **device:** can be any one of switch ports.
 - **index:** value is the same as streamhandle of cbstreamidset.
3. Send a stream from TestCenter to swp3 of board A, set destination mac as 7E:A8:8C:9B:41:DD.
 4. Capture frames on swp2 on TestCenter.
We can get frames from swp2 on TestCenter, each frame adds the sequence number: 23450801, 23450802, 23450803...
 5. Capture frames from swp2 of board B on TestCenter, we can get the same frames.

5.1.4.2.7.2 Sequence Recover test

Use the following steps for the **Sequence Recover** test:

1. On board B, run the following commands:

```
bridge fdb add 7E:A8:8C:9B:41:DD dev swp2 vlan 1 master static
tsntool> cbstreamidset --device swp2 --index 1 --nullstreamid --nulldmac
0x7EA88C9B41DD --nullvid 1 --streamhandle 1
tsntool> cbrec --device swp0 --index 1 --seq_len 16 --his_len 31 --
rtag_pop_en
```

In the cbrec command mentioned above:

- **device:** can be any one of switch ports.
 - **index:** value is the same as streamhandle of cbstreamidset.
2. Send a frame from TestCenter to swp3 of board A, set the destination MAC address to be 7E:A8:8C:9B:41:DD.
 3. Capture frames from swp2 of board B on the TestCenter, we can get only one frame without sequence tag.

5.1.4.2.8 TSN stream identification

TSN module uses QoS class to identify and control streams. There are three ways to identify the stream to different QoS class. These are explained in the following sections.

5.1.4.2.8.1 Stream identification based on PCP value of Vlan tag

The default QoS class is based on PCP of Vlan tag for a frame. If there is no Vlan tag for a frame, the default QoS class is 0.

Set the PCP value on TestCenter.

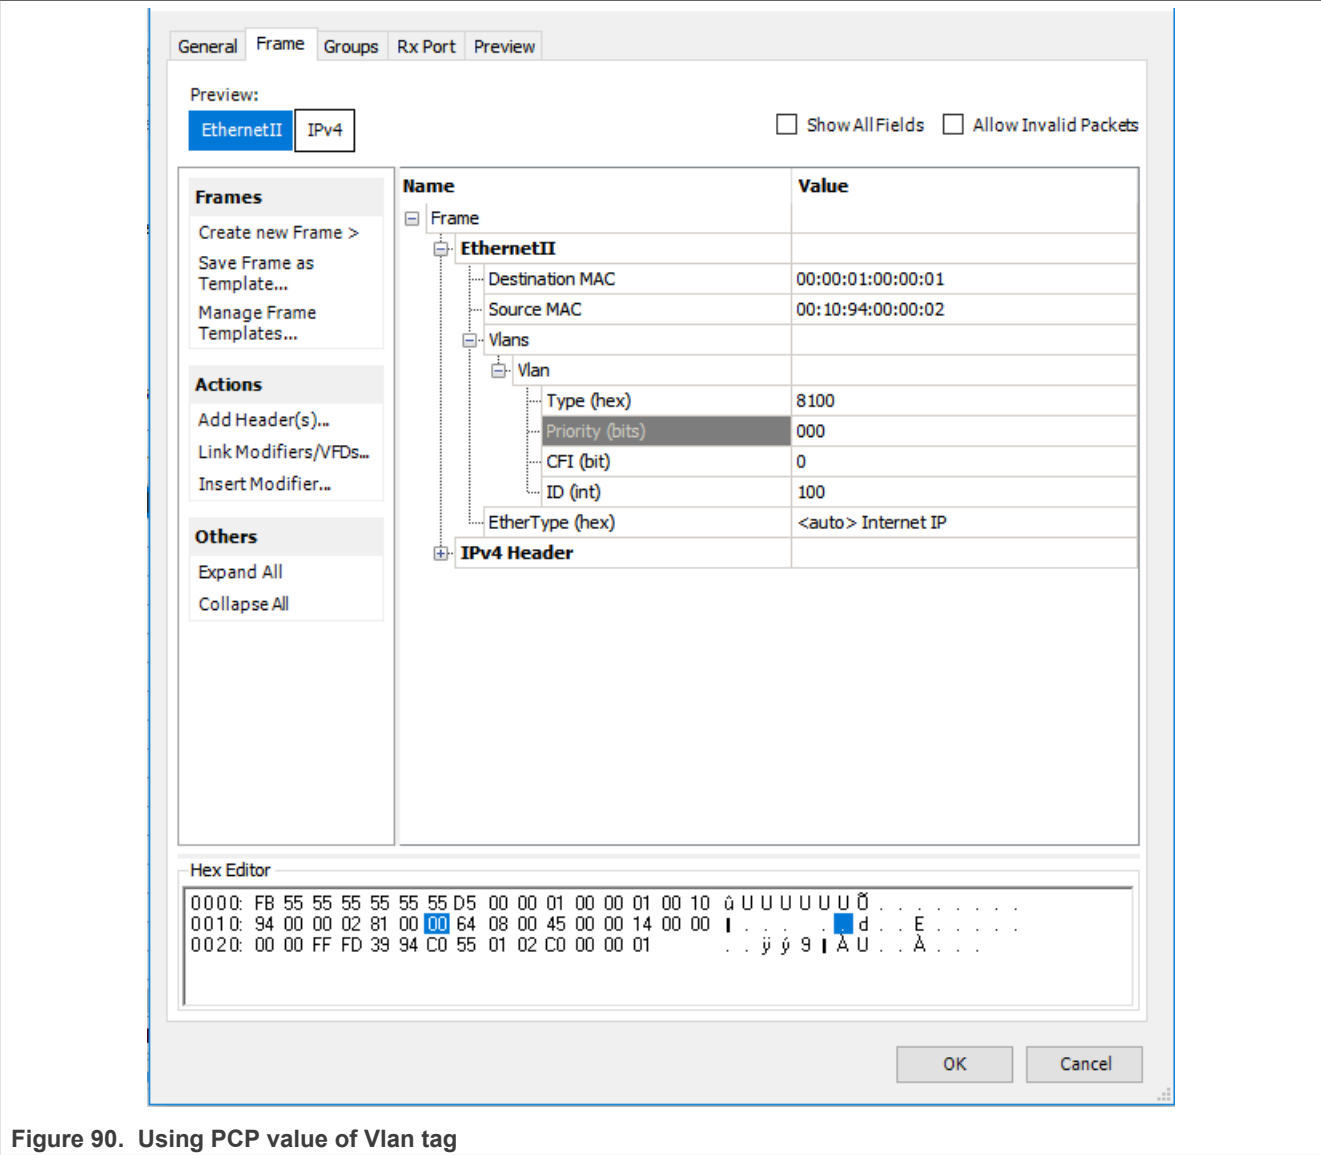


Figure 90. Using PCP value of Vlan tag

5.1.4.2.8.2 Based on DSCP of ToS tag

Use the below steps to identify stream based on DSCP value of ToS tag.

1. Map the DSCP value to a specific QoS class using the command below:

```
tsntool> dscpset --device swp0 --index 1 --cos 1 --dpl 0
```

Explanation:

- `index`: DSCP value of stream, 0-63.
 - `cos`: QoS class which is mapped to.
 - `dpl`: Drop level which is mapped to.
2. Set the DSCP value on TestCenter. DSCP value is the upper six bits of ToS in IP header, set the DSCP value on TestCenter as shown in [Figure 91](#).




Figure 91. Setting DSCP value on TestCenter

5.1.4.2.8.3 Based on qci stream identification

The following steps describe how to use qci to identify the stream and set it to a QoS class.

1. Identify a stream.

```
tsntool> cbstreamidset --device swp1 --nullstreamid --nullldmac 0x000183fe1201
--nullvid 1 --streamhandle 1
```

```
tsntool> qcisfiset --device swp0 --index 1 --gateid 1 --flowmeterid 68
```

- Set to Qos class 3 by using stream gate control.

```
echo "t0 1b 3 50000 200" > sgi.txt
tsntool> qcisgiset --device swp0 --enable --index 1 --initgate 1 --initip 0
--gatelistfile sgi.txt
```

Note: The Qci-based identity stream can only be used on both the ingress and egress are bridge ports. The flow injected or extracted through the CPU port cannot be configured for Qci.

5.2 GenAVB/TSN stack

This section describes the steps for configuring the GenAVB/TSN stack.

5.2.1 Introduction

The GenAVB/TSN Stack provides advanced implementation for Audio Video Bridging (AVB) and Time-Sensitive Networking (TSN) functionalities on NXP SoCs and hardware platforms, for both Endpoints and Bridges.

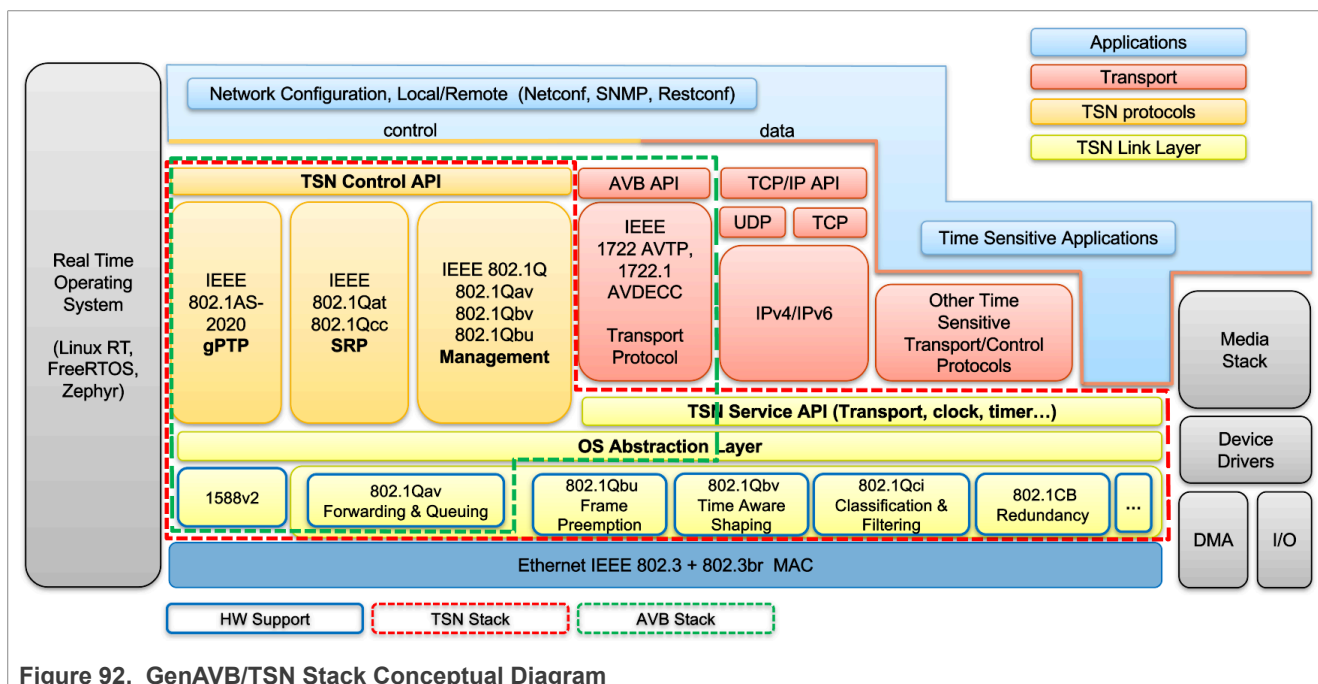


Figure 92. GenAVB/TSN Stack Conceptual Diagram

This section provides information on how to set up and evaluate the GenAVB/TSN Stack. In that context, it provides information on supported SoCs and boards, compile time software package configuration, and runtime configuration settings.

5.2.1.1 GenAVB/TSN Roles

The GenAVB/TSN stack supports the following roles:

- **AVB Endpoint**

AVB Endpoint functionality is provided in i.MX 93, i.MX 8M Plus, i.MX 8DXL, i.MX 8M Mini, and i.MX 6ULL SoCs (using hardware support if available).

- **AVB Bridge**

AVB Bridge functionality requires AVB hardware support, available in LS1028A SoC and SJA1105Q-EVB (paired with i.MX 93 14x14 EVK or i.MX 8DXL EVK).

- **AVB Hybrid**

AVB Hybrid functionality requires AVB hardware support, available in i.MX 93 14x14 EVK SoC paired with a SJA1105Q-EVB.

- **TSN Endpoint**

TSN Endpoint functionality requires TSN hardware support, available on i.MX 95, i.MX 93, i.MX 8M Plus, and i.MX 8DXL SoCs.

- **TSN Bridge**

TSN Bridge functionality requires TSN hardware support, available in LS1028A SoC.

5.2.1.2 GenAVB/TSN Protocols

The GenAVB/TSN stack includes following protocol stacks running in standalone userspaces processes.

5.2.1.2.1 gPTP Stack

The gPTP stack implements IEEE 802.1AS-2020 standard, and supports both time-aware Endpoint and Bridge systems. The stack runs fully in userspace, using Linux socket APIs for packet transmit, receive, and timestamping. Linux clock APIs are used for clock adjustment. Configuration files are used to configure the stack at initialization time and extensive logging is available at runtime.

5.2.1.2.2 SRP stack

The SRP stack implements MRP, MVRP, and MSRP defined in IEEE 802.1Q-2022, sections 10, 11, and 35, and supports both Endpoint and Bridge systems. The stack runs fully in userspace, using Linux socket APIs for packet transmit and receive. Linux tc and bridge netlink APIs are used to update Multicast FDB entries and FQTTSS Credit Based Shaper (CBS) configuration. Configuration files are used to configure the stack at initialization time and extensive logging is available at runtime.

5.2.1.2.3 AVTP Stack

The AVTP stack implements IEEE 1722-2016 standard, supporting both AVB Talker/Listener end stations and multiple Audio/Video stream formats. The stack runs in userspace, in combination with a Linux AVB kernel module, providing low latency network packet processing and AVTP packet encapsulation/decapsulation. The stack provides an API for external media applications through a run time library. The API allows external applications to act as sources of AVTP Talker streams/sinks of AVTP Listener streams.

5.2.1.2.4 AVDECC/Milan Stack

The AVDECC stack implements IEEE 1722.1-2013 standard, and supports Talker, Listener and Controller entities. The stack also implements "Milan Specification v1.2", which can be enabled at initialization time. AVDECC entity definitions are loaded from the filesystem and can be created based on a C header file definition. The stack provides an API for external media applications through a run time library.

5.2.1.2.5 MAAP Stack

The MAAP stack implements IEEE 1722-2016, Annex B. The stack provides an API for external media applications through a run time library, but it mainly serves the AVDECC/Milan stack.

5.2.1.3 User Library

The GenAVB/TSN stack provides a public C API through a user space dynamic library that can be linked into customer applications. These API's cover initialization of the stack, run time configuration as well as AVB and TSN endpoint functionality (AVTP Listener/Talker streaming, Layer 2 network sockets, ...).

5.2.1.4 AVB Kernel Module

Integration of AVB functionality in the Linux kernel is provided by a dedicated kernel module. The main role of this module is to add real time guarantees to the processing of AVTP packets.

5.2.1.5 Supported configurations

GenAVB/TSN stack currently supports the following boards and the associated roles:

- LS1028ARDB: gPTP Time-aware Bridge and SRP Bridge
- i.MX 95 19x19 LPDDR5 EVK: gPTP Time-aware Endpoint station and TSN Endpoint.
- i.MX 93 EVK: gPTP Time-aware Endpoint station, TSN Endpoint, and AVB Endpoint stack/applications.
- i.MX 93 9x9 LPDDR4 QSB: gPTP Time-aware Endpoint station and TSN Endpoint.
- i.MX 93 14x14 LPDDR4X EVK: gPTP Time-aware Endpoint station, TSN Endpoint and AVB Endpoint stack/applications. Along with AVB bridge and AVB hybrid stacks when connecting SJA1105Q-EVB Hardware.
- i.MX 8M Plus LPDDR4 EVK: gPTP Time-aware Endpoint station, TSN Endpoint, and AVB Endpoint stack/applications.
- i.MX 8M Mini LPDDR4 EVK: gPTP Time-aware Endpoint station and AVB Endpoint stack/applications.
- i.MX 8DXL LPDDR4 EVK: gPTP Time-aware Endpoint station, TSN Endpoint and AVB Endpoint stack/applications. Along with AVB bridge stack when connecting SJA1105Q-EVB Hardware.
- i.MX 6ULL 14x14 EVK: gPTP Time-aware Endpoint station and AVB Endpoint stack/applications.

The TSN stack supports and is enabled in the following Yocto Real-time Edge machines:

- ls1028ardb
- imx8mp-lpddr4-evk
- imx8dxlb0-lpddr4-evk
- imx93evk
- imx93-9x9-lpddr4-qsb
- imx93-14x14-lpddr4x-evk
- imx95-19x19-lpddr5-evk

The AVB Endpoint stack supports and is enabled in the following Yocto Real-time Edge machines:

- imx6ull14x14evk
- imx8mm-lpddr4-evk
- imx8mp-lpddr4-evk
- imx8dxlb0-lpddr4-evk
- imx93evk
- imx93-14x14-lpddr4x-evk

Follow Real-time Edge Software Yocto Project to get the code and build images for these platforms.

5.2.1.6 AVB endpoint example applications

The stack provides extensive example applications for media playback/capture and server. Please refer to “GenAVB/TSN Stack Evaluation User Guide” for detailed information. Refer [Section 1.5](#).

5.2.1.7 TSN endpoint example application

The TSN example application provides example code and re-usable middleware exercising the GenAVB/TSN API. It can be used to exercise and verify the real time behavior of the local system as well as TSN properties of the network between endpoints. [Figure 93](#) shows the TSN application cycle.

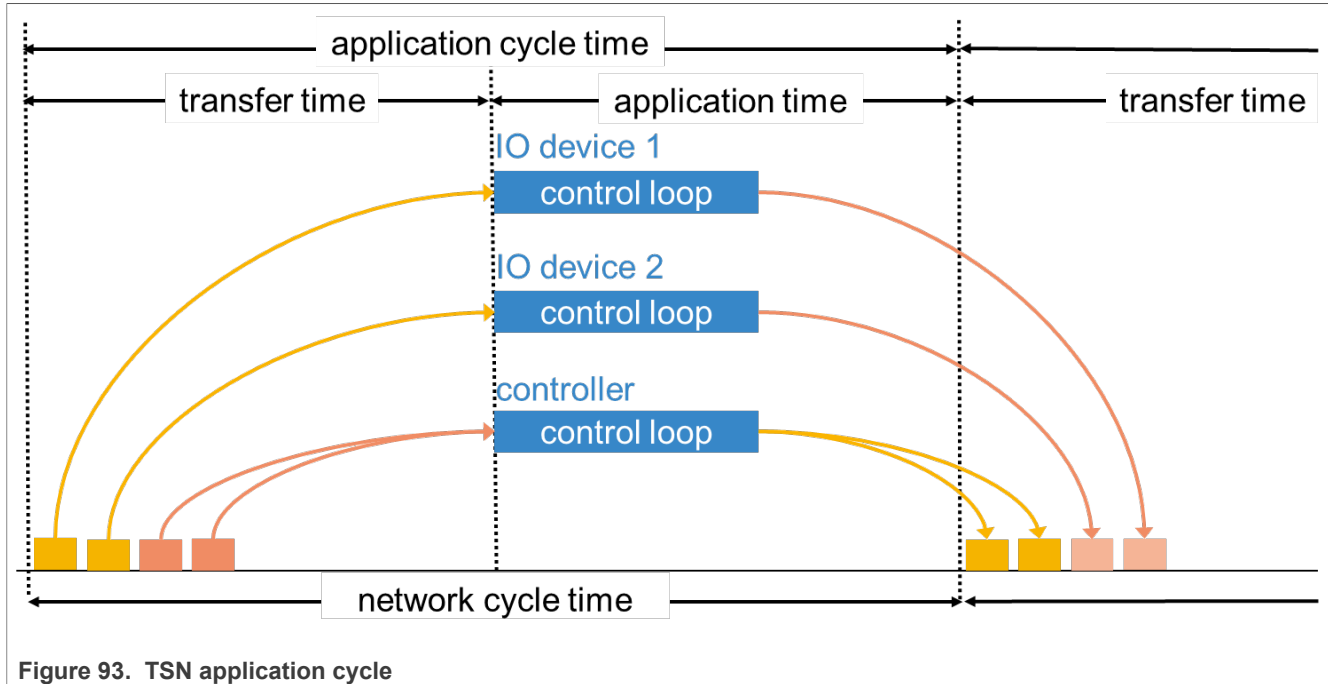


Figure 93. TSN application cycle

The TSN example application implements a control loop similar to industrial use cases requiring cyclic isochronous exchanges over the network.

The TSN endpoints run their application synchronized to a common time grid in the same gPTP domain so that they can send and receive network traffic in a cyclic isochronous pattern (the application cycle time is equal and synchronous to the network cycle time as shown in [Figure 93](#)). Currently the cycle is configured with a period of 2 ms, and periods as low as 1 ms have been confirmed to work as well. When the application is scheduled, frames from other endpoints are ready to be read and at the end of the application time frames are sent to other endpoints.




Figure 94. TSN application scheduling

As shown in [Figure 94](#), the controller and the IO devices are scheduled with a half cycle offset in order to reduce the processing latency.

The time sensitive traffic is layer 2 multicast with VLAN header and proprietary EtherType. Its priority is defined using the PCP field of the VLAN header.

In addition, the TSN application provides detailed logs and time sensitive traffic timing statistics (based on hardware timestamping of packets), which allow characterization of an entire real time distributed system.

Finally, a OPCUA server is implemented and offer the possibility to browse and retrieve the TSN application statistics exposed as OPCUA objects. The OPCUA server runs over TCP and allows access to any OPCUA client.

5.2.2 GenAVB/TSN stack start/stop

GenAVB/TSN stack can be manually started/stopped at runtime by using the commands listed below.

1. To start the TSN stack (if not already started) and start/stop the demo applications:

```
# avb.sh <start|stop>
```

2. To just start/stop the TSN stack (gPTP and SRP) use:

```
# tsnn.sh <start|stop>
```

3. To restart/stop all GenAVB/TSN processes, TSN stack, and demo applications:

```
avb.sh restart_all/stop_all
```

4. Real-time Edge also provides a `systemd` service to run `genavb-tsn` stack as a system service.

```
# systemctl enable genavb-tsn
# systemctl start genavb-tsn
```

5. The below commands can be used to stop or disable this service.

```
# systemctl stop genavb-tsn
# systemctl disable genavb-tsn
```


5.2.3 Use cases description

5.2.3.1 AVB endpoint

AVB endpoint use cases and example applications are described in the *GenAVB/TSN Stack AVB Endpoint User Guide* located in [Real Time Edge Documentation](#).

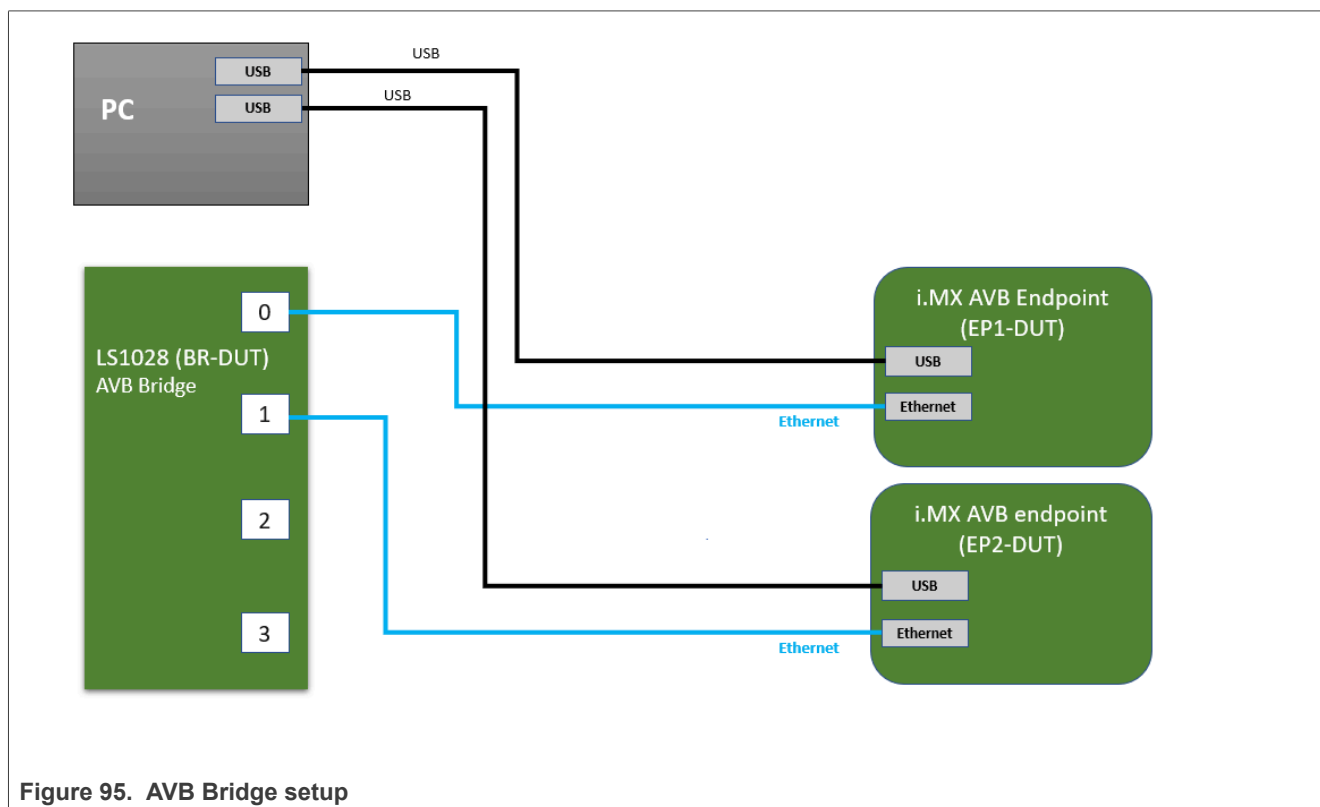
5.2.3.2 AVB Bridge

This use case illustrates an AVB Bridge (mixing gPTP and SRP stack) with other AVB Endpoints

5.2.3.2.1 Requirements

5.2.3.2.1.1 Using LS1028ARDB as AVB bridge

- [Two AVB endpoints](#)
- One AVB bridge (LS1028ARDB)



5.2.3.2.1.2 Using i.MX8DXL EVK with SJA1105Q EVB as AVB bridge

- [Two AVB endpoints](#) including:
 - an i.MX 8D XL with an IMX-RMII-PHY TJA1100
 - an i.MX 8M Plus with an Ethernet to BoardR-Reach converter
- One AVB bridge (i.MX8DXL with a SJA1105Q-EVB)

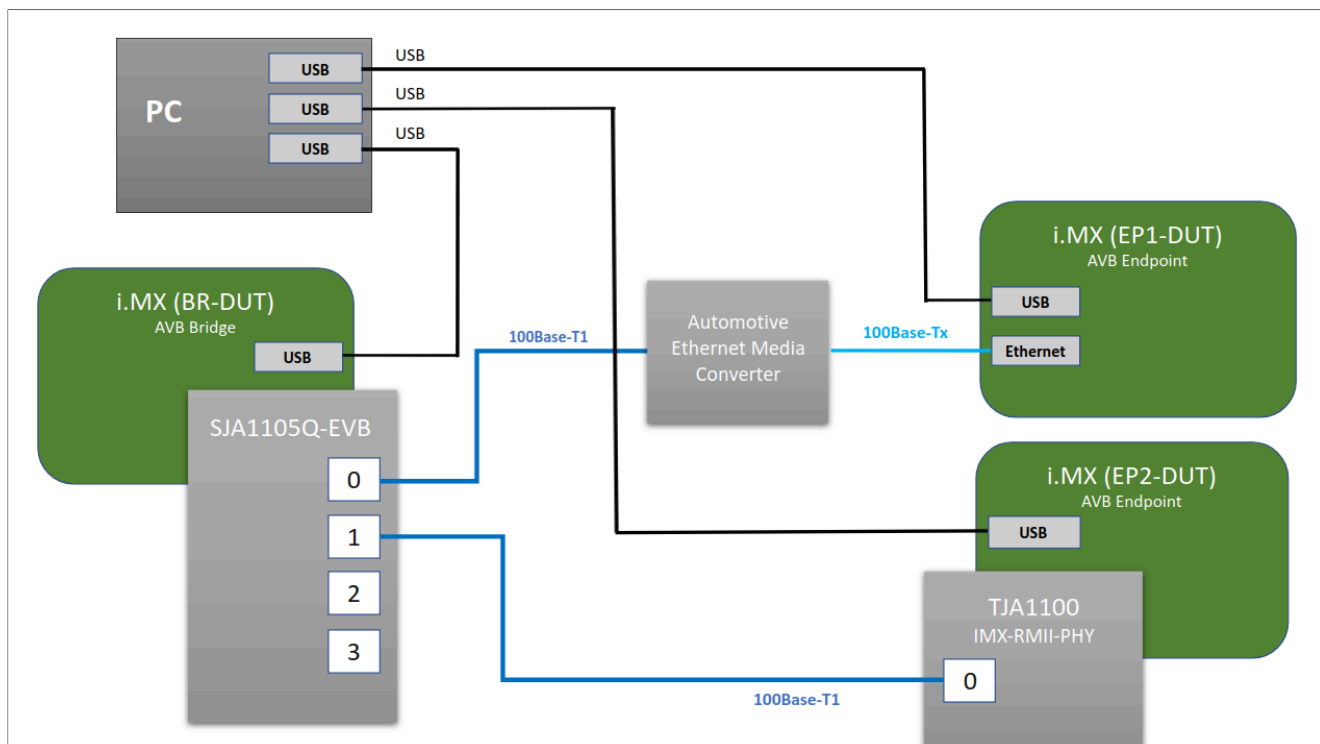


Figure 96. AVB Bridge setup

Setup the device tree:

- On the i.MX 8D XL Bridge with a SJA1105Q-EVB. Use the following device tree in U-boot:

```
setenv fdt_file imx8dxl-evk-enet0-sja1105.dtb
```

- On the i.MX 8DXL Endpoint with a TJA1100. Use the following device tree in U-boot:

```
setenv fdt_file imx8dxl-evk-enet0-tja1100-avb.dtb
```

Configure MASTER/SLAVE role of the PHYs:

- On the Bridge with SJA1105Q-EVB, check the default PHY role for the connected ports:

```
ethtool swpX
```

Note: The default role is of Master (ethtool reports forced-master under master-slave cfg).

- On the endpoint with TJA1100 PHY Daughter card, set the PHY role accordingly (master if SJA1105 port reports slave and vice versa):

```
ethtool -s eth0 master-slave forced-slave
```

- On the automotive media converter, set the role using the same method

5.2.3.2.1.3 Using i.MX 93 14x14 EVK with SJA1105Q EVB as AVB bridge

- Two AVB endpoints including the following:
 - an i.MX 8DXL board with an IMX-RMII-PHY TJA1100
 - an i.MX 8M Plus board with an Ethernet to BoardR-Reach converter
- One AVB bridge (i.MX 93 14x14 with a SJA1105Q-EVB connected on ENET2)

Note: SJA1105Q-EVB is using RGMII interface and the ENET connector on i.MX 93 14x14 EVK is usually set to RMII by default. A Hardware rework may be needed to set the interface to RGMII (mount R267, R272, R275 and R278 and unmount R288 and R293). Refer to board's schematics for more details.

See [Figure 96](#)

Setup the device tree:

- On the i.MX 93 14x14 Bridge with a SJA1105Q-EVB. Use the following device tree in U-boot:

```
setenv fdtfile imx93-14x14-evk-sja1105.dtb
```

- On the i.MX 8DXL Endpoint with a TJA1100. Use the following device tree in U-boot:

```
setenv fdt_file imx8dxl-evk-enet0-tja1100-avb.dtb
```

Configure Master/ Slave role of the PHYs:

- On the Bridge with SJA1105Q-EVB, check the default PHY role for the connected ports:

```
ethtool swpX
```

Note: Default role is Master (ethtool will report forced-master under master-slave cfg)

- On the endpoint with TJA1100 PHY Daughter card, set the PHY role accordingly (master if SJA1105 port reports slave and vice versa):

```
ethtool -s eth0 master-slave forced-slave
```

- On the automotive media converter, set the role using the same method

5.2.3.2.2 AVB network configuration

This topic describes AVB configuration.

5.2.3.2.2.1 Priority to traffic class mapping

The priority to traffic class mapping used for the bridge comes directly from the recommended mapping for two SR classes in IEEE Std 802.1Q-2018 Table 34-1:

Table 71. Priority to traffic class mapping

Priority	0	1	2	3	4	5	6	7
Traffic Class	1	0	6	7	2	3	4	5

The Bridge should be configured to forward VLAN tagged packets based on their PCP values according to this mapping, and should configure credit-based shapers on the two highest traffic classes (traffic class 6 and traffic class 7) for SR class A (priority 3) and SR class B (priority 2) traffic.

Refer to [Section 5.2.3.2.3](#) for the bridge PCP mapping configuration.

5.2.3.2.2.2 FQTSS Credit Based Shapers configuration

The SRP bridge stack relies on preconfigured qdiscs with specific handles to configure the hardware's credit-based shapers, on the two hardware queues with the two highest traffic classes, for every port. Thus, an mqprio qdisc with 8 traffic classes should be configured with the above priority to traffic class mapping and credit-based shapers qdiscs with the following handles: 0x9006 for CBS on traffic class 6 and 0x9007 for CBS on traffic class 7.

Refer to [Section 5.2.3.2.3](#) for the bridge qdisc configuration.

5.2.3.2.2.3 Linux Best Effort Traffic classification

Linux classifies egress packets, for assignment to traffic classes, based on skb priorities. To avoid assigning egress best effort traffic to traffic classes with configured credit-based shapers, the skb priorities should be rewritten so no packets with skb priorities 2 and 3 are present on egress. Furthermore, the bridge code is using the skb priority as the traffic class for packets injected from the CPU port, making packets with skb priorities 6 and 7 end up in the hardware's traffic classes 6 and 7 on the external ports which in turn harms traffic shaping. Again, forcing a remapping of these skb priorities avoids this scenario.

Refer to [Section 5.2.3.2.3](#) for the skb priorities remapping configuration.

5.2.3.2.2.4 Bridge VLAN awareness

A proper AVB bridge functioning requires that the switch forward AVB streams (with multicast destination MAC addresses and specific VLAN ID) only to ports configured in the Forwarding DataBase (FDB). For that, we should enable VLAN filtering on bridge level, add the desired VLAN ID to all ports and disable the default multicast flooding configuration (at least for the two highest priority queues) on all the external ports.

Refer to [Section 5.2.3.2.3](#) for the bridge vlan configuration.

5.2.3.2.3 Setup preparation

This has two steps described in the following sections.

5.2.3.2.3.1 GenAVB/TSN stack configuration

This configuration needs to be done once and is saved accross reboots.

Edit the GenAVB/TSN configuration file using the following command at the Linux prompt:

```
# vi /etc/genavb/config
```

And set the GENAVB_TSN_CONFIG correctly as "Bridge AVB".

Note: *It's the default configuration on LS1028A while the Bridge with SJA1105Q needs to be changed accordingly*

For a proper gPTP operation with AVB endpoints, the gPTP stack needs to compensate for PHY delay in PTP timestamps:

- AVB Bridge on LS1028A

In the `/etc/genavb/fgptp-br.cfg`, apply the settings (rxDelayCompensation and txDelayCompensation) described in [Table 85](#) on all bridge ports.

Attention: *The PHY Delay Compensation Values in [Table 85](#) are calibrated for 1 Gbps links. The i.MX AVB endpoints are configured to run by default with 100 Mbps links. These compensation values should be enough to keep pDelay values under 800 ns (propagation time threshold), and therefore the port would still be declared as Capable. If, with these values, the calculated propagation delay is still above 800 ns (or too close to it), adapt them accordingly (increase rxDelayCompensation and/or txDelayCompensation).*

Future releases shall have proper compensation values for each supported link speed.

- AVB Bridge on SJA1105Q

On AVB bridge setup with SJA1105Q-EVB, the calculated propagation delay is exceeding default propagation delay threshold (800ns) and the link would not be asCapable. To bypass that, increase the neighborPropDelayThreshold on all devices, including the bridge:

```
vi /etc/genvb/fgptp-br.cfg
```

Set the neighborPropDelayThreshold to 8000ns

The same should be done on endpoints connected to the bridge: /etc/genavb/fgptp.cfg

5.2.3.2.3.2 Bridge configuration

This configuration should be done after each boot. The user can either enter these commands manually or **execute a ready to use script provided by GenAVB/TSN stack**.

1. If you're using an i.MX8DXL EVK or i.MX 93 14x14 EVK Bridge, edit the GenAVB/TSN configuration file using the following command at the Linux prompt:

```
# vi /etc/genavb/config
```

And set the configuration GENAVB_TSN_CONFIG to use the config_avb_bridge file:

```
GENAVB_TSN_CONFIG=3
```

2. Execute the **automated configuration script** and start the AVB bridge stack:

```
# avb-bridge.sh
# avb.sh start
```

3. Alternatively, configure the bridge manually using the following commands:

- a. Setup bridge forwarding:

```
# ip link add name br0 type bridge
# ip link set br0 up
# ip link set master br0 swp0 up
# ip link set master br0 swp1 up
# ip link set master br0 swp2 up
# ip link set master br0 swp3 up
```

- b. Establish the PCP to QoS mapping for every port on the bridge using a custom tool (supported only on LS1028A):

```
# pcp_to_qos_map=([0]="1" [1]="0" [2]="6" [3]="7" [4]="2" [5]="3" [6]="4" [7]="5"); \
avb_ports="swp0 swp1 swp2 swp3"; \
for port in $avb_ports; do \
    for (( pcp=0; pcp < 8; ++pcp )); do \
        tsntool pcpmap -d $port -p $pcp -e 0 -c ${pcp_to_qos_map[$pcp]} -l 0; \
        tsntool pcpmap -d $port -p $pcp -e 1 -c ${pcp_to_qos_map[$pcp]} -l 1; \
    done ; \
done
```

- c. Configure the qdiscs and shapers, with the correct handles, and optionally offload the pcp to qos mapping for every external port

```
# pcp_to_qos_map=([0]="1" [1]="0" [2]="6" [3]="7" [4]="2" [5]="3" [6]="4" [7]="5"); \
avb_ports="swp0 swp1 swp2 swp3"; \
for port in $avb_ports; do \
    tc qdisc add dev $port root handle 100: mqprio num_tc 8 map ${pcp_to_qos_map[@]} queues 1@0 \
    1@1 1@2 1@3 1@4 1@5 1@6 1@7 hw <HW Offload> ; \
    tc qdisc replace dev $port handle 0x9007 parent 100:8 \
    cbs locredit -2147483646 hicredit 2147483647 sendslope -1000000 idleslope 0 offload \
    0 ; \
    tc qdisc replace dev $port handle 0x9006 parent 100:7 \
    cbs locredit -2147483646 hicredit 2147483647 sendslope -1000000 idleslope 0 offload \
    0 ; \
done
```

Note:

<HW offload>: set to 1 if the driver supports MQPrio PCP mapping offload (SJA1105Q), otherwise 0 (LS1028A).

The two most important CBS parameters for every port device are:

- the **parent**, which should match the traffic class 6 and 7,
- the **handle**, which should be 0x9006 and 0x9007.

The other parameters are initialization values and are overridden by the stack at runtime stream configuration:

- **offload** is set to 1 to offload the operation to hardware,
- **idleslope** and **sendslope** are set depending on stream,
- **port bit rates** and the **credit values** are kept at their minimum and maximum values as they do not directly affect the hardware shaping operation.

d. Setup skb priorities remapping for every external port:

```
# avb_ports="swp0 swp1 swp2 swp3"; \
for port in $avb_ports; do \
    tc qdisc add dev $port clsact; \
    tc filter add dev $port egress basic match 'meta(priority eq 2)' or \
    'meta(priority eq 3)' action skbedit priority 0; \
done
```

4. Enable Vlan filtering, set the correct Vlan IDs and disable multicast flooding, for every external port:

```
# ip link set br0 type bridge vlan_filtering 1; \
avb_ports="swp0 swp1 swp2 swp3"; \
for port in $avb_ports; do \
    bridge vlan add dev $port vid 2 master; \
    bridge link set dev $port mcast_flood off; \
done
```

5. Start the AVB and gPTP stacks:

```
# avb.sh start
```

- Since multicast traffic flooding is now disabled, adding MDB entries for AVDECC (ACMP/ADP) and MAAP protocols multicast addresses is needed. The following commands should be executed for every port facing an AVB endpoint.

```
# bridge mdb add dev br0 port <port> grp 91:e0:f0:01:00:00 permanent
# bridge mdb add dev br0 port <port> grp 91:e0:f0:00:ff:00 permanent
```

5.2.3.2.4 Evaluation instructions

1. Reset all endpoints and the bridge.
2. Using the procedures described above, configure the bridge and start the stack on all connected devices (bridge and endpoints)
3. After a few seconds, AVB endpoints should be synchronized through gPTP
4. Connect an SR class A (or SR class B) stream from EP-DUT2 as talker to EP-DUT1 as listener: the stream should be forwarded correctly to the listener endpoint

5.2.3.2.4.1 gPTP operation

If the gPTP protocol is running correctly on all devices, the following line should appear in the bridge gptp log file for every port connected to a gPTP capable device:

```
gptp_stats_dump: Port(0) domain(0, 0): Role: Master    Link: Up    asCapable: Yes
neighborGptpCapable: Yes DelayMechanism: P2P
...
gptp_stats_dump: Port(1) domain(0, 0): Role: Master    Link: Up    asCapable: Yes
neighborGptpCapable: Yes DelayMechanism: COMMON_P2P
```

Refer to [Section 5.2.3.5](#), for more details on gPTP Bridge operation.

5.2.3.2.4.2 SRP Operation

A detailed view on the SRP protocol communications (such as Domain declaration, SRP port boundary, Talker/Listener declarations and registration) can be followed by displaying the SRP specific logs from the TSN bridge stack log file `/var/log/tsn-br`:

```
# tail -f /var/log/tsn-br | grep srp
```

On stream connection, the FQTSS and FDB operation should be visible in the TSN bridge stack log file:

- Stack log shows the FQTSS configuration for the port facing the AVB listener:

```
fqtss_set_oper_idle_slope : logical_port(2) port (swp0, ifindex 5) tc(7)
cbs_qdisc_handle(9007:0): set idle_slope 7872000
```

- Stack log shows the FDB configuration for the port facing the AVB listener:

```
bridge_rtnetlink : add MDB: bridge (br0, ifindex 9) logical_port(2) port (swp0,
ifindex 5) mac_addr(91:e0:f0:00:fe:11) vlan_id(2)
```

Also, the same configuration can be checked using the Linux standard tools (`tc` and `bridge`)

- TC tool shows the FQTSS configuration for the port facing the AVB listener:

```
# tc qdisc show dev swp0
qdisc mqprio 100: root   tc 8 map 1 0 6 7 2 3 4 5 0 0 0 0 0 0 0 0
      queues:(0:0) (1:1) (2:2) (3:3) (4:4) (5:5) (6:6) (7:7)
qdisc pfifo 0: parent 9006: limit 1000p
qdisc pfifo 0: parent 9007: limit 1000p
qdisc pfifo_fast 0: parent 100:6 bands 3 priomap 1 2 2 2 1 2 0 0 1 1 1 1 1 1 1
1
qdisc pfifo_fast 0: parent 100:5 bands 3 priomap 1 2 2 2 1 2 0 0 1 1 1 1 1 1 1
1
qdisc pfifo_fast 0: parent 100:4 bands 3 priomap 1 2 2 2 1 2 0 0 1 1 1 1 1 1 1
1
qdisc pfifo_fast 0: parent 100:3 bands 3 priomap 1 2 2 2 1 2 0 0 1 1 1 1 1 1 1
1
qdisc pfifo_fast 0: parent 100:2 bands 3 priomap 1 2 2 2 1 2 0 0 1 1 1 1 1 1 1
1
qdisc pfifo_fast 0: parent 100:1 bands 3 priomap 1 2 2 2 1 2 0 0 1 1 1 1 1 1 1
1
qdisc cbs 9006: parent 100:7 hicredit 2147483647 locredit -2147483646 sendslope
-1000000 idleslope 0 offload 0
qdisc cbs 9007: parent 100:8 hicredit 2147483647 locredit -2147483648 sendslope
-992128 idleslope 7872 offload 1
```

- Bridge tool shows the FDB configuration for the port facing the AVB listener:

```
# bridge mdb show
dev br0 port swp0 grp 91:e0:f0:00:fe:11 permanent offload vid 2
```

5.2.3.3 AVB Hybrid

AVB Hybrid use cases and example applications are described in the *GenAVB/TSN Stack AVB Endpoint User Guide* located in [Real Time Edge Documentation](#).

5.2.3.4 gPTP Endpoint

Once the gPTP stack is started, logs can be displayed using the following command:

```
# tail -f /var/log/fgptp
```

In this log file, one can observe the role of the port in the time-aware system.

If the port of the endpoint is connected to another port capable of communicating a synchronized time, the following log appears for each gPTP domain:

```
gptp_stats_dump : Port(0) domain(0,0): Role: Slave Link: Up AS_Capable: Yes
neighborGptpCapable: Yes DelayMechanism: P2P
...
gptp_stats_dump : Port(0) domain(1,20): Role: Slave Link: Up AS_Capable: Yes
neighborGptpCapable: Yes DelayMechanism: COMMON_P2P
```

Role status can also take the value *Master* depending on the *Grandmaster Parameters* described in section [Section 5.2.4.2.2](#).

If a port is not connected, *Link* status takes the value *Down*.

If a port is not capable of communicating a synchronized time, *AS_Capable* status takes the value *No*.

If a port is using the Common Mean Link Delay Service (CMLDS) the Delay Mechanism takes the value COMMON_P2P, else the value P2P.

For further details about gPTP logs, refer to section [Section 5.2.5.1](#).

5.2.3.5 gPTP Bridge

LS1028ARDB, i.MX8DXL EVK or i.MX 93 14x14 EVK paired with a SJA1105Q-EVB can be used as a generic time-aware bridge, connected to other time-aware end stations or bridges.

By default, they don't forward packets if no bridge interface is configured under Linux. Enabling bridge interface is dependent on the board used. For example, the configuration of bridge interface is described in section [Section 5.1.4.2.1](#).

Note: The i.MX8DXL EVK uses *eth0* instead of *eno2*.

Once gPTP stack is started, logs can be displayed with the following command:

```
# tail -f /var/log/fgptp-br
```

In this log file, one can observe which ports are connected, which ports are currently communicating a synchronized time and what is the role of the port in the time-aware system.

If a port of the bridge is connected to another port capable of communicating a synchronized time, the following log should appear for each enabled gPTP domain:

```
gptp_stats_dump : Port(1) domain(0,0): Role: Master Link: Up asCapable: Yes
neighborGptpCapable: Yes DelayMechanism: P2P
...
gptp_stats_dump : Port(1) domain(1,20): Role: Master Link: Up asCapable: Yes
neighborGptpCapable: Yes DelayMechanism: COMMON_P2P
```

Role status can also take the value *Slave* depending on *parameters* described in section [Section 5.2.4.2.2](#).

If a port is not connected, *Link* status takes the value *Down*.

If a port is not capable of communicating a synchronized time, *AS_Capable* status takes the value *No*.

If a port is using the Common Mean Link Delay Service (CMLDS) the *DelayMechanism* takes the value *COMMON_P2P*, else the value *P2P*.

For further details about gPTP logs, refer to section [Section 5.2.5.2](#).

5.2.3.6 gPTP multiple domains

This use case illustrates two gPTP domains co-existing independently on a TSN network, over different 802.1AS-2020 Time-aware systems.

The first domain uses the PTP timescale whereas the second domain uses the ARB (arbitrary) timescale.

5.2.3.6.1 Requirements

The reference setup for gPTP multiple domains is made of:

- Two gPTP endpoints EP1-DUT and EP2-DUT:
 - (i.MX 8M Plus LPDDR4 EVK or i.MX 93 EVK connected through TSN interface `eth1`: `dwmac0` **Or**
 - i.MX 95 19x19 LPDDR5 EVK, i.MX 8DXL LPDDR4 EVK or i.MX 93 9x9 LPDDR4 QSB connected through TSN interface `eth0`);
- One gPTP bridge (LS1028ARDB): BR-DUT

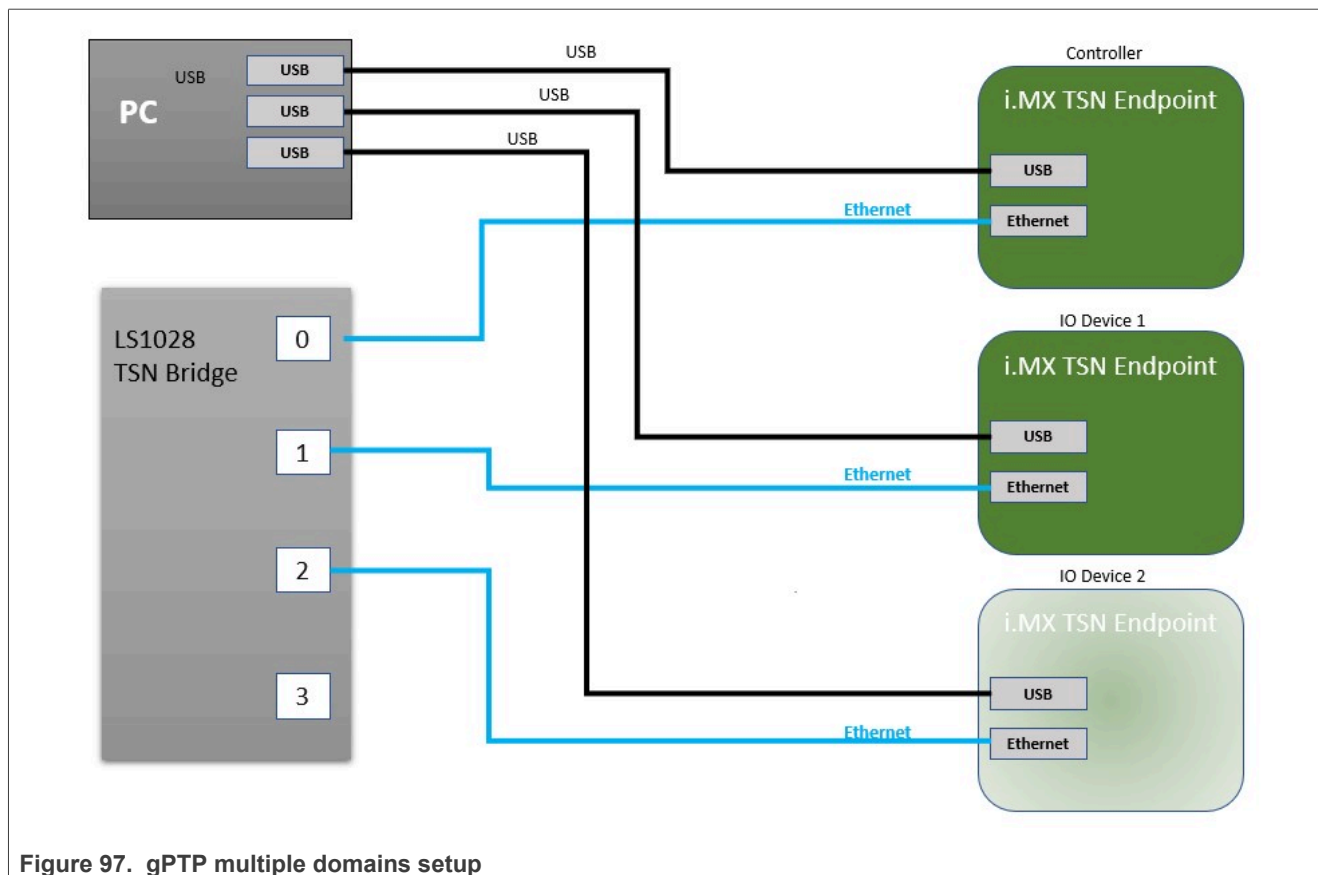


Figure 97. gPTP multiple domains setup

5.2.3.6.2 gPTP Stack Configuration

5.2.3.6.2.1 Domains configuration

The gPTP stack can enable or disable each domain independently through a configuration file.

The default configuration file (for example: `/etc/genavb/fgptp.cfg`) is for general gPTP parameters as well as domain 0 parameters. To enable other domains, new files must be created with '-N' appended to the filename (for example: `/etc/genavb/fgptp.cfg-1` for domain 1).

For gPTP multiple domains, all devices configuration should be changed to support two domains. The first domain (domain 0) must be assigned domain number 0. The second domain (domain 1) is assigned domain number 20.

BR-DUT is defined as the GrandMaster for the first domain (domain 0). EP1-DUT is defined as the GrandMaster for the second domain (domain 1).

On EP1-DUT, edit the file `/etc/genavb/fgptp.cfg-1` and change `domain_number` and `priority1` parameters as follows:

```
domain_number = 20
priority1 = 245
```

On EP2-DUT, edit the file `/etc/genavb/fgptp.cfg-1` and change `domain_number` parameter as follows:

```
domain_number = 20
```

On BR-DUT, edit the file `/etc/genavb/fgptp-br.cfg-1` and change `domain_number` parameter as follows:

```
domain_number = 20
```

Note:

- On Domain 0, BR-DUT is the GrandMaster with the highest priority (lowest value) among all devices in the domain (BR-DUT priority1=246, EP1-DUT and EP2-DUT priority1=248).
- On Domain 1, EP1-DUT is the GrandMaster with the highest priority (lowest value) among all devices in the domain (BR-DUT priority1=246, EP1-DUT priority1=245 and EP2-DUT priority1=248).
- By default,
 - All ports on Domain 0 are configured to use the per instance peer delay mechanism (DelayMechanism=P2P).
 - All ports on Domain 1 are configured to use the CMLDS (DelayMechanism=COMMON_P2P).

5.2.3.6.2.2 Clocks Configuration

The stack supports two types of virtual clocks as target gPTP clocks when enabling multi-domain support:

- Using internal software clocks
- Or using Linux Virtual PHC devices

Using internal software clocks

The GenAVB/TSN stack has support for userspace based software clocks that are used to track time in multiple domains.

By default, the stack is configured to use a software clock as target gPTP clock for the second domain (domain 20) while keeping the physical PHC device as target gPTP clock for the first domain (domain 0). Refer to `/etc/genavb/system.cfg` for full configuration details.

Using Linux Virtual PHCs on endpoints

Alternatively, the stack can be configured to use Virtual PHC devices on Linux to support gPTP multi domains, as follows:

1. Create two virtual clocks bounded to the PHC device, associated to the network interface, using the following command:

```
# echo 2 > /sys/class/ptp/ptpX/n_vclocks
```

Note:

Replace `ptpX` with the right PHC device associated with the used network interface:

- `ptp0`: on i.MX 95 19x19 LPDDR5 EVK, i.MX 8DXL LPDDR4 EVK and i.MX 93 9x9 LPDDR4 QSB.
- `ptp1`: on i.MX 8M Plus LPDDR4 EVK and i.MX 93 EVK.

2. The new virtual clocks devices (ptp) appear in the dmesg log. For example:

```
[ 245.771356] ptp ptp0: new virtual clock ptp2
[ 245.772108] ptp ptp0: new virtual clock ptp3
[ 245.772144] ptp ptp0: guarantee physical clock free running
```

3. Configure `/etc/genavb/config` to prevent the system configuration file to be overridden by startup scripts:

```
# vi /etc/genavb/config

CFG_SYSTEM_CFG_NO_OVERRIDE=1
```

4. Assign one virtual clock to each domain in the file `/etc/genavb/system.cfg.tsn`:

```
# vi /etc/genavb/system.cfg.tsn

[LOGICAL_PORT]
endpoint = ethX

[CLOCK]
endpoint_gptp_0 = /dev/ptp2
endpoint_gptp_1 = /dev/ptp3
endpoint_local = /dev/ptpX
```

Note:

Replace `ethX` and `ptpX` with the right network interface and PHC devices:

- `ptp0` and `eth0`: on *i.MX 95 19x19 LPDDR5 EVK*, *i.MX 8DXL LPDDR4 EVK*, and *i.MX 93 9x9 LPDDR4 QSB*.
- `ptp1` and `eth1`: on *i.MX 8M Plus LPDDR4 EVK* and *i.MX 93 EVK*.
- `endpoint_gptp_0` and `endpoint_gptp_1` must be set to each of the new virtual PHCs.

5.2.3.6.3 Evaluation instructions

Test Procedure

1. Start gPTP stack manually on all DUTs by issuing the command below:

```
# tsn.sh start
```

2. Wait for 30 s.
3. Check gPTP stack logs on BR-DUT (`/var/log/fgptp-br`), EP1-DUT and EP2-DUT (`/var/log/fgptp`)

Verification:

Check the following:

- After Step 3, the log on EP1-DUT reports Port 0 as synchronized on domain 0 only:

```
Port(0) domain(0, 0) SYNCHRONIZED - synchronization time (ms): 250
```

- After Step 3, the log on EP2-DUT reports Port 0 as synchronized on all domains :

```
Port(0) domain(0, 0) SYNCHRONIZED - synchronization time (ms): 250
Port(0) domain(1, 20) SYNCHRONIZED - synchronization time (ms): 250
```

- After Step 3, the log on BR-DUT reports Port 0 as synchronized on domain 1 only:

```
Port(0) domain(1, 20) SYNCHRONIZED - synchronization time (ms): 250
```

- The “*Initial adjustment*” message should be reported only once per synchronized domains (domain 0 for EP1-DUT and EP2-DUT, domain 1 for EP2-DUT and BR-DUT):

```
domain(0,0) Initial Adjustment, offset: 125486471315484 ns, freq_adj: 32764
```

```
domain(1,20) Initial Adjustment, offset: 125455671332661 ns, freq_adj: 16384
```

Once synchronization is achieved, all the reported clock offset average values should be stable within -50 to +50 ns range (domain 0 for EP1-DUT and EP2-DUT, domain 1 for EP2-DUT and BR-DUT):

```
domain(0,0) Offset between GM and local clock (ns): min -45 avg 0 max 35
```

```
domain(1,20) Offset between GM and local clock (ns): min -66 avg 0 max 15
```

5.2.3.7 TSN endpoint sample application

5.2.3.7.1 Requirements

- Two TSN endpoints among:
 - i.MX 8MPlus LPDDR4 EVK connected through TSN interface eth1: dwmac
 - i.MX 93 EVK connected through TSN interface eth1: dwmac
 - i.MX 93 14x14 EVK connected with IMXA12ETH-ATH daughter card through TSN interface eth1
 - i.MX 8DXL LPDDR4 EVK connected through TSN interface eth0
 - i.MX 93 9x9 LPDDR4 QSB connected through TSN interface eth0
 - i.MX 95 19x19 LPDDR5 EVK connected through TSN interface eth0
 - optionally an i.MX RT1170 EVK
- One TSN bridge (LS1028ARDB)

Note: The second IO Device is optional.

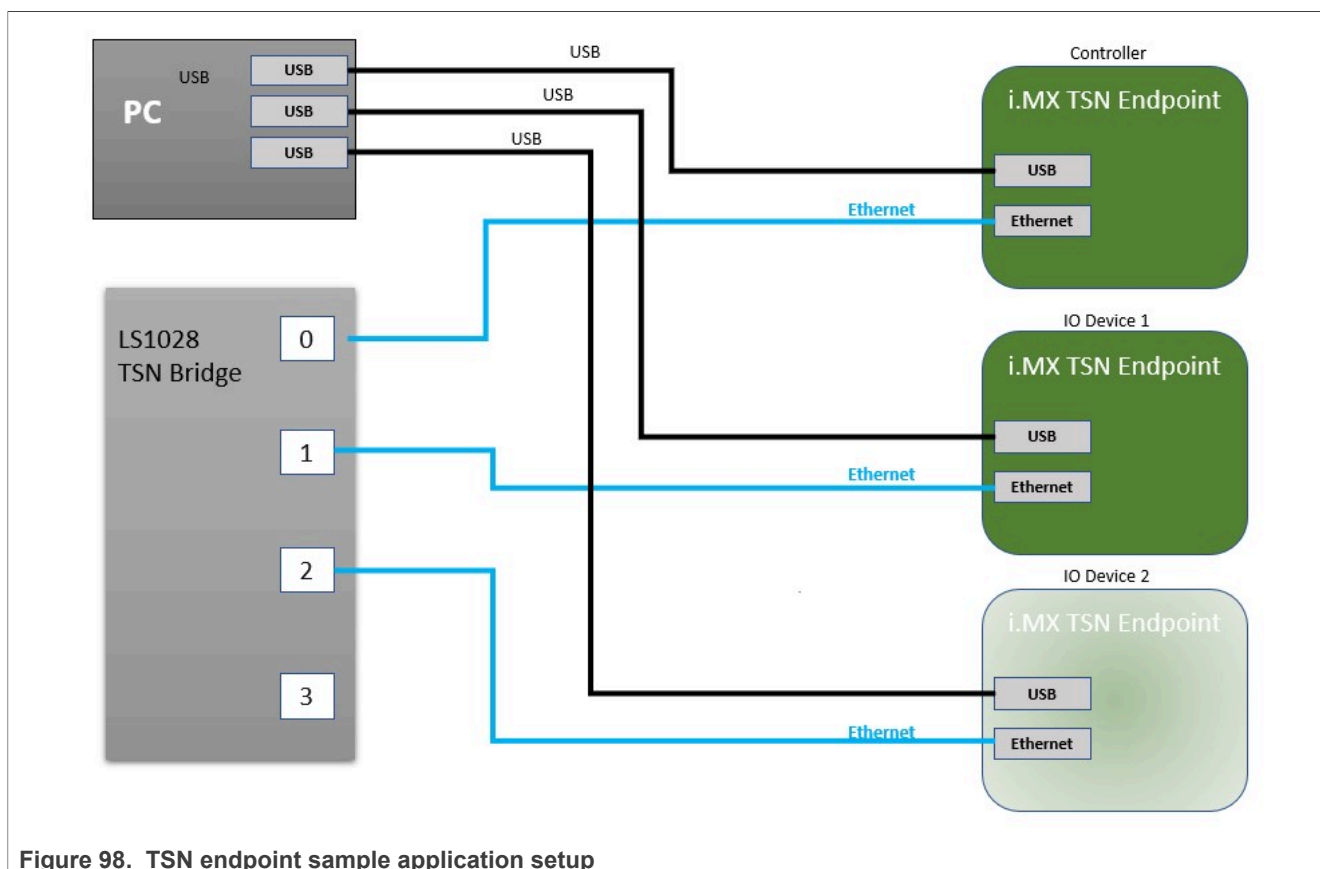


Figure 98. TSN endpoint sample application setup

5.2.3.7.1.1 Using i.MX 8MP EVK as Endpoint TSN

Make sure to update the boot parameter to use the device tree binary that includes the hardware description relative to the i.MX 8MP EVK board:

- Power on the board and stop the automatic boot process by pressing the space bar on the keyboard.
- Enter the following commands at the U-Boot prompt:

```
U-Boot > setenv fdtfile imx8mp-evk.dtb
U-Boot > saveenv
U-Boot > boot
```

The change is saved across reboots.

5.2.3.7.1.2 Using i.MX 8DXL EVK as Endpoint TSN

Make sure to update the boot parameter to use the device tree binary that includes the hardware description relative to the i.MX 8DXL EVK board:

- Power on the board and stop the automatic boot process by pressing the space bar on the keyboard.
- Enter the following commands at the U-Boot prompt:

```
U-Boot > setenv fdtfile imx8dxl-evk.dtb
U-Boot > saveenv
U-Boot > boot
```

The change is saved across reboots.

5.2.3.7.1.3 Using i.MX 93 EVK as Endpoint TSN

Make sure to update the boot parameter to use the device tree binary that includes the hardware description relative to the i.MX 93 EVK board:

- Power on the board and stop the automatic boot process by pressing the space bar on the keyboard.
- Enter the following commands at the U-Boot prompt:

```
U-Boot > setenv fdtfile imx93-11x11-evk.dtb
U-Boot > saveenv
U-Boot > boot
```

The change is saved across reboots.

5.2.3.7.1.4 Using i.MX 93 9x9 EVK as Endpoint TSN

Make sure to update the boot parameter to use the device tree binary that includes the hardware description relative to the i.MX 93 9x9 EVK board:

- Power on the board and stop the automatic boot process by pressing the space bar on the keyboard.
- Enter the following commands at the U-Boot prompt:

```
U-Boot > setenv fdtfile imx93-9x9-qsb.dtb
U-Boot > saveenv
U-Boot > boot
```

The change is saved across reboots.

5.2.3.7.1.5 Using i.MX 93 14x14 EVK as Endpoint TSN

To support TSN endpoint usecase with standard ethernet, an additional Ethernet PHY daughter card (IMXAI2ETH-ATH) is required. see [Figure 99](#).

Note: The IMXAI2ETH-ATH daughter card is using RGMII interface and the ENET connector on i.MX 93 14x14 EVK is usually set to RMII by default. A Hardware rework may be needed to set the interface to RGMII (mount R267, R272, R275 and R278 and unmount R288 and R293). Refer to board's schematics for more details.



Figure 99. Atheros Ethernet add on card (IMXAI2ETH-ATH)

Make sure to update the boot parameter to use the device tree binary that includes the hardware description relative to the i.MX 93 14x14 EVK board (Using ENET2):

- Power on the board and stop the automatic boot process by pressing the space bar on the keyboard.
- Enter the following commands at the U-Boot prompt:

```
U-Boot > setenv fdtfile imx93-14x14-evk-imxai2eth-ath.dtb
U-Boot > saveenv
U-Boot > boot
```

The change is saved across reboots.

5.2.3.7.1.6 Using i.MX 95 19x19 LPDDR5 EVK as Endpoint TSN

Make sure to update the boot parameter to use the device tree binary that includes the hardware description relative to the i.MX 95 19x19 LPDDR5 EVK board:

- Power on the board and stop the automatic boot process by pressing the space bar on the keyboard.
- Enter the following commands at the U-Boot prompt:

```
U-Boot > setenv fdtfile imx95-19x19-evk.dtb
U-Boot > saveenv
U-Boot > boot
```

The change is saved across reboots.

5.2.3.7.2 Configuring GenAVB/TSN stack and example applications

For some platforms, the GenAVB/TSN stack supports both modes: Endpoint TSN and Endpoint AVB.

By default, these platforms are configured as Endpoint TSN. The `GENAVB_TSN_CONFIG` parameter should be set to the right configuration using the file `/etc/genavb/config`:

```
# avb.sh stop_all
```

```
# vi /etc/genavb/config
```

Platforms that support both Endpoint AVB and Endpoint TSN (for example i.MX 8MP, i.MX 8DXL and i.MX 93), should have:

```
GENAVB_TSN_CONFIG=1
```

For Endpoint TSN mode, the change from one profile to another is made by modifying the `/etc/genavb/config_tsn` file. This file specifies the application configuration file. `APPS_CFG_FILE` (`apps-*.cfg`) points to a file containing a demo configuration (application to use, options...). It is parsed by the startup script `avb.sh`.

TSN configuration profile is made of the application configuration profile. The file `/etc/genavb/config_tsn` already lists the supported `cfg` files. Set the `PROFILE` variable to choose the desired configuration profile.

5.2.3.7.3 TSN network configuration

This topic describes TSN configuration.

5.2.3.7.3.1 Streams

The stream details can be used for analysis and also for computing scheduled traffic timings.

Table 72. TSN streams definition

Stream No	Source	Destination	Unicast / Multicast	Destination MAC Address	Vlan ID	Vlan PCP	Frame Length ^[1] (bytes)
Stream1	Controller	IO device(s)	Multicast	91:e0:f0:00:fe:70	2	5	84
Stream2	IO device 1	Controller	Multicast	91:e0:f0:00:fe:71	2	5	84
Stream3	IO device 2	Controller	Multicast	91:e0:f0:00:fe:80	2	5	84

[1] The frame length includes inter frame gap, preamble, start of frame and CRC (can be used as is for timing calculations)

5.2.3.7.3.2 Scheduled traffic

For deterministic packet transmission the use of scheduled traffic is required both on endpoints and bridges. The default scheduling configuration for the TSN endpoint example application, as shown in [Figure 94](#), leads to the following traffic schedules.

Endpoints

Endpoints are running a schedule with a 2000us period. The base offset of the schedule is aligned to gPTP time modulo 1 second.

Controller transmit gate (for Stream1) opens at 500us offset (relative to the period start).

IO device transmit gate (for Stream2/3) opens at 1000us + 500us offset (relative to the period start).

The gate open interval is around 4us (enough to accommodate the stream frame length plus some margin).

The 500us offset is related to the worst case application latency to send its frame to its peer(s). This value provides a good margin for a Linux PREEMPT-RT system but can be lowered on a well-tuned system.

Bridges

The schedule for all Bridges and all Bridge ports that transmit one of the streams above, must have a 2000 μ s period and a base offset aligned to gPTP time modulo 1 second.

One possible schedule is to open transmit gate (for the ports and queues transmitting Stream 1) at offset 500 μ s and use a gate open interval that accommodates the worst propagation delay.

It is also possible to use a fixed gate open interval but increase the transmit time offset at each hop along the stream path.

For ports and queues transmitting Stream 2 and 3, open the transmit gate at offset 1000 + 500 μ s.

5.2.3.7.4 Setup preparation

One of the TSN endpoint must be configured as the “controller” and the other one as an “IO device”. Both endpoints are connected to the TSN bridge.

Note:

1. On i.MX 8M Plus LPDDR4 EVK and i.MX 93 EVK: the TSN interface used is `eth1`.
2. On i.MX 95 19x19 LPDDR5 EVK, i.MX 8DXL LPDDR4 EVK and i.MX 93 9x9 LPDDR4 QSB: the TSN interface used is `eth0` rather than `eth1`.

5.2.3.7.4.1 Preparing the controller

To be done once:

1. Edit the GenAVB configuration file using the following command at the Linux prompt:

```
# vi /etc/genavb/config_tsn
```

2. Set the configuration profile to `PROFILE 1`:

```
PROFILE=1
```

3. Exit and save.
Perform the below steps at each boot:
4. Execute the `tsn.sh` script to make sure the gPTP stack is running/synchronized before configuring the Transmission Qdiscs (TC taprio) through the `tsn-app-setup.sh` below

```
tsn.sh start
```

Note: On i.MX 95 19x19 LPDDR5 EVK, make sure to re-setup the Taprio Qdisc (with `tsn-app-setup.sh`) after gPTP has synchronized (and after each gPTP synchronization loss)

5. The system configuration required for the `tsn-app` can be performed (after setting the correct profile) by using the following command (replace `ethX` with the right TSN network interface):
 - `eth1` on i.MX 8M Plus LPDDR4 EVK and i.MX 93 EVK
 - `eth0` on i.MX 95 19x19 LPDDR5 EVK, i.MX 8DXL LPDDR4 EVK and i.MX 93 9x9 LPDDR4 QSB

```
# tsn-app-setup.sh ethX
```

Note: This script sets many different settings to improve real time system behavior and to setup proper network configuration

- **VLAN configuration:** the script sets `vlan id 2` on the TSN interface as VLAN hardware filtering is enabled by default in kernel.
- **Low latency settings on network interface:** the script disables coalescing and flow control on the TSN interface.

- *Qdiscs and filters: the script sets taprio qdisc with proper parameters for TX and flower qdisc for RX classification.*
 - *Interrupts, network tasks, CPU affinities, and priorities: the scripts enables threaded NAPI in kernel and isolate tasks processing TSN traffic on a separate CPU core.*
6. Start the TSN demo application by using the following command:

```
avb.sh start
```

5.2.3.7.4.2 Preparing IO device(s)

To be done once:

1. Edit the GenAVB configuration file using the following command at the Linux prompt:

```
# vi /etc/genavb/config_tsn
```

2. Set the configuration profile to PROFILE 2:

```
PROFILE=2
```

3. Exit and save.

The below steps should be done at each boot:

4. Execute the `tsn.sh` script to make sure the gPTP stack is running/synchronized before configuring the Transmission Qdiscs (TC taprio) through the `tsn-app-setup.sh` below

```
tsn.sh start
```

Note: On i.MX 95 19x19 LPDDR5 EVK, make sure to re-setup the Taprio Qdisc (with `tsn-app-setup.sh`) after gPTP has synchronized (and after each gPTP synchronization loss)

5. The system configuration required for the tsn-app can be performed (after setting the correct PROFILE) by using the following command (replace `ethX` with the right TSN network interface):
- `eth1` on i.MX 8M Plus LPDDR4 EVK and i.MX 93 EVK
 - `eth0` on i.MX95 19x19 LPDDR5 EVK, i.MX 8DXL LPDDR4 EVK and i.MX 93 9x9 LPDDR4 QSB

```
# tsn-app-setup.sh ethX
```

Note: This script sets many different settings to improve real time system behavior and to setup proper network configuration

- *VLAN configuration: the script sets vlan id 2 on the TSN interface as VLAN hardware filtering is enabled by default in kernel.*
 - *Low latency settings on network interface: the script disables coalescing and flow control on TSN interface.*
 - *Qdiscs and filters: the script sets taprio qdisc with proper parameters for TX and flower qdisc for RX classification.*
 - *Interrupts, network tasks, CPU affinities, and priorities: the scripts enable threaded NAPI in kernel and isolate tasks processing TSN traffic on a separate CPU core.*
6. Start the TSN demo application using the following command:

```
avb.sh start
```

5.2.3.7.4.3 Preparing the Bridge

Refer to section [Section 5.1.2](#) and [Section 5.1.4.2.3.2](#) to configure scheduled traffic on the LS1028ARDB board.

Follow the schedule described in section [Section "Bridges"](#)

Follow the below steps at each boot:

1. Set up bridge forwarding:

```
# ip link add name br0 type bridge
# ip link set br0 up
# ip link set master br0 swp0 up
# ip link set master br0 swp1 up
# ip link set master br0 swp2 up
# ip link set master br0 swp3 up
```

2. Disable Pause frames:

```
# ethtool -A swp0 autoneg off rx off tx off
# ethtool -A swp1 autoneg off rx off tx off
# ethtool -A swp2 autoneg off rx off tx off
# ethtool -A swp3 autoneg off rx off tx off
```

3. Start the gPTP stack:

```
# tsn.sh start
```

4. Setup scheduled traffic (see above)

```
# tc qdisc del dev swp0 root
# tc qdisc del dev swp1 root
# tc qdisc del dev swp2 root
# tc qdisc replace dev swp0 root taprio \
    num_tc 8 \
    map_0 1 2 3 4 5 6 7 \
    queues 100 101 102 103 104 105 106 107 \
    base-time 1500000 \
    sched-entry S 0x20 20000 \
    sched-entry S 0xdf 1980000 \
    flags 0x2
# tc qdisc replace dev swp1 root taprio \
    num_tc 8 \
    map_0 1 2 3 4 5 6 7 \
    queues 100 101 102 103 104 105 106 107 \
    base-time 500000 \
    sched-entry S 0x20 20000 \
    sched-entry S 0xdf 1980000 \
    flags 0x2
# tc qdisc replace dev swp2 root taprio \
    num_tc 8 \
    map_0 1 2 3 4 5 6 7 \
    queues 100 101 102 103 104 105 106 107 \
    base-time 500000 \
    sched-entry S 0x20 20000 \
    sched-entry S 0xdf 1980000 \
    flags 0x2
```

5.2.3.7.4.4 Preparing the OPC UA client

In order to visualize the data exposed by the TSN endpoint application OPC UA server it is required to use an OPC UA client on a PC connected to the bridge.

1. Install an OPC UA client on a PC:

- a. FreeOpcUa: client with a Qt GUI interface.

Can be found here: <http://freeopcua.github.io/>

- b. opcua-commander: CLI alternative based nodejs node-opcua stack. Can be found here:

<https://github.com/node-opcua/opcua-commander>

2. Connect the PC to the bridge. If not already done, setup IP addresses on the endpoint running the TSN example application and also on the PC. Then, make sure you can successfully ping the endpoint using the PC.

5.2.3.7.5 Evaluation instructions

1. Reset all endpoints.
2. Using the procedures described above, start the gPTP stack on the bridge and the tsn-app application on the endpoints with the proper enabled scheduled traffic as configured above.
3. After a few seconds, it can be observed that the TSN endpoints are synchronized through gPTP and exchange packets at the rate of 500 packets per second (pps). Check the logs to observe this behavior.

5.2.3.7.5.1 gPTP operation

If the gPTP protocol is running correctly on an endpoint or on the bridge, the following line should appear in the gptp log file (refer to [Section 5.2.3.4](#) for more details):

```
gptp_stats_dump: Port(0) domain(0,0) : Role: Slave Link : Up AS_Capable: Yes  
DelayMechanism: P2P
```

If the device is grand master, the role field should be "Master" otherwise it should be "Slave". The line appears periodically, but the role should not change over time, except for significant events (such as a cable disconnection).

5.2.3.7.5.2 Baseline tsn-app operation

If the TSN endpoint sample application is running correctly and receiving valid packets, the following points may be verified in the tsn_app log file (refer to [Section 5.2.5.4](#) for more details).

The following line should appear at regular intervals:

```
socket_stats_print : link up
```

The "valid frames" counter should increment by 2500 (500 pps for 5 seconds) between two appearances of the following log:

```
socket_stats_print : valid frames : XXXXX
```

The various error counters should not increment (it is normal to have non-zero values, because of the startup period when gPTP and/or the remote tsn-app endpoint may not be running and stable):

- "sched early", "sched late", "sched missed", "sched timeout", "sched discount", "clock err"
- "err id", "err ts", "err underflow"
- "frames err" (for both RX and TX directions)

Note:

The checks above apply to all tsn-app endpoints, whether they be the controller or one of the IO devices.

5.2.3.7.5.3 Scheduled traffic evaluation with no concurrent traffic

The observations below assume an otherwise idle system receiving and sending traffic only through the tsn-app application, with a 802.1Qbv schedule in place on all devices (tsn-app endpoints, bridge).

Scheduling error statistics ("sched err") should respect the following:

- min around 8 μ s
- avg around 11 μ s
- max around 25 μ s

```
stats(0xaaab06ed74b0) sched err min 8817 mean 11120 max 22077 rms^2 125202075
stddev^2 1544829 absmin 7417 absmax 1882057
```

Processing time statistics ("processing time") should respect the following:

- min around 23 μ s
- avg around 29 μ s
- max around 70 μ s

```
stats(0xaaab06ed7910) processing time min 23400 mean 29185 max 59100 rms^2
857707540 stddev^2 5943315 absmin 19560 absmax 4143240
```

Traffic latency statistics should respect the following:

- min around 503 μ s
- avg around 503 μ s
- max around 503 μ s
- stddev^2 less than 3000

```
stats(0x419a28) traffic latency min 503417 mean 503503 max 503637 rms^2
253515981945 stddev^2 2004 absmin 503397 absmax 504337
```

5.2.3.7.5.4 Scheduled traffic evaluation with TX best-effort traffic

1. Connect a PC to the 4th port of the LS1028ARDB switch (swp3).
2. Run iperf3 in server mode on the PC (replace ethX by the PC interface connected to the LS1028):

```
# ifconfig ethX 192.168.1.10 up
# iperf3 -s &
# iperf3 -s -p 5202 &
# iperf3 -s -p 5203 &
# iperf3 -s -p 5204 &
```

3. Run iperf3 in client mode on the controller (replace ethX by the controller interface connected to the LS1028):

```
# ifconfig ethX 192.168.1.80
# taskset b iperf3 -c 192.168.1.10 -u -b 0 -i 2 -t 100 &
# taskset b iperf3 -p 5202 -c 192.168.1.10 -u -b 0 -l 64 -i 2 -t 100 &
# taskset b iperf3 -p 5203 -c 192.168.1.10 -u -b 0 -l 64 -i 2 -t 100 &
# taskset b iperf3 -p 5204 -c 192.168.1.10 -u -b 0 -l 64 -i 2 -t 100 &
```

4. Observe stats in the ts-n-app log files (a 2nd terminal may have to be opened through SSH). The values should match the table below (in μ s):

	min	mean	max	stddev^2
Sched err (controller)	21	29	41	
Processing time (controller)	47	80	260	

	min	mean	max	stddev^2
Traffic latency (controller and IO device)	503	503	503	<3000

5.2.3.7.5.5 Scheduled traffic evaluation with RX best-effort traffic

Note:

On SoCs using EnetQos/dwmac as ethernet controller, part of the TSN traffic (untagged gPTP traffic) is processed in the same queue as best-effort untagged traffic. To more easily validate tsn-app with best-effort traffic, we should add a VLAN tag with PCP=0 to best-effort packets so they are dispatched into a different queue on receive.

On SoCs using ENETC controller (e.g i.MX 95), the RX classification is able to steer gPTP traffic to different hardware queues than the best effort untagged traffic.

1. Connect a PC to the 4th port of the LS1028ARDB switch (swp3).
2. Run iperf3 in server mode on the controller (replace ethX by the controller interface connected to the LS1028):

```
# ip link add link ethX name ethX.5 type vlan id 5
# ifconfig ethX.5 192.168.5.80 up
# taskset b iperf3 -s &
# taskset b iperf3 -s -p 5202 &
# taskset b iperf3 -s -p 5203 &
# taskset b iperf3 -s -p 5204 &
```

3. Run iperf3 in client mode on the PC (replace ethX by the PC interface connected to the LS1028):

```
# ip link add link ethX name ethX.5 type vlan id 5
# ifconfig ethX.5 192.168.5.10 up
# iperf3 -c 192.168.5.80 -u -b 0 -i 2 -t 100 &
# iperf3 -p 5202 -c 192.168.5.80 -u -b 0 -i 2 -t 100 &
# iperf3 -p 5203 -c 192.168.5.80 -u -b 0 -i 2 -t 100 &
# iperf3 -p 5204 -c 192.168.5.80 -u -b 0 -i 2 -t 100 &
```

4. Observe stats in the tsn-app log file (a 2nd terminal may have to be opened through SSH). The values should match the table below (in μ s):

	min	mean	max	stddev^2
Sched err (controller)	9	13	26	
Processing time (controller)	25	33	70	
Traffic latency (controller and IO device)	503	503	503	<130000

5.2.3.7.5.6 Modifying the scheduling period of the TSN sample application

The default tsn-app period of 2 ms can be changed through a command-line option. The change has to be made on all endpoints (controller and devices). The 802.1 Qbv schedule must also be updated to reflect the new period. The example below shows how to modify the period from the default 2 ms down to 1 ms (this value has been confirmed to work on the latest builds).

On the controller:

1. Stop the application if it was already running:

```
# avb.sh stop
```

2. Edit the application configuration file:

```
# vi /etc/genavb/apps-tsn-network-controller.cfg
```

or for an IO device:

```
# vi /etc/genavb/apps-tsn-network-iodevice.cfg
```

3. Use the "-p" option to change the period. The below example sets the period to 1 ms (1000000 ns):

```
CFG_EXTERNAL_MEDIA_APP_OPT="-m network_only -r controller -p 1000000"
```

4. Update the traffic schedule using 'tc' command.

In the sample command below, replace `ethX` with the right TSN network interface:

- eth1 on i.MX 8M Plus LPDDR4 EVK and i.MX 93 EVK
- eth0 on i.MX 8DXL LPDDR4 EVK and i.MX 93 9x9 LPDDR4 QSB

```
# tc qdisc del dev ethX root
#tc qdisc replace dev ethX root taprio \
num_tc 5 \
map_0 0 1 1 2 2 3 4 0 0 0 0 0 0 0 \
queues 1@0 1@1 1@2 1@3 1@4 \
base-time 250000 \
sched-entry S 0x4 4000 \
sched-entry S 0x1b 996000 \
flags 0x2
```

5. Restart the tsan-app application:

```
# avb.sh start
```

On the IO device(s):

1. Stop the application if it was already running:

```
# avb.sh stop
```

2. Edit the application configuration file:

```
# vi /etc/genavb/apps-tsn-network-iodevice.cfg
```

3. Use the "-p" option to change the period. The below example sets the period to 1 ms (1000000 ns):

```
CFG_EXTERNAL_MEDIA_APP_OPT="-m network_only -r iodevice_N -p 1000000"
```

4. Update the traffic schedule using tc.

In the sample command below, replace `ethX` with the right TSN network interface:

- eth1 on i.MX 8M Plus LPDDR4 EVK and i.MX 93 EVK
- eth0 on i.MX 8DXL LPDDR4 EVK and i.MX 93 9x9 LPDDR4 QSB

```
# tc qdisc del dev ethX root
#tc qdisc replace dev ethX root taprio \
num_tc 5 \
map_0 0 1 1 2 2 3 4 0 0 0 0 0 0 0 \
queues 1@0 1@1 1@2 1@3 1@4 \
base-time 750000 \
sched-entry S 0x4 4000 \
sched-entry S 0x1b 996000 \
```

```
flags 0x2
```

- Restart the `tsn-app` application:

```
# avb.sh start
```

On the bridge, update the Qbv schedule on all ports:

```
# tc qdisc del dev swp0 root
# tc qdisc del dev swp1 root
# tc qdisc del dev swp2 root
# tc qdisc replace dev swp0 root taprio \
    num_tc 8 \
    map_0 1 2 3 4 5 6 7 \
    queues 1@0 1@1 1@2 1@3 1@4 1@5 1@6 1@7 \
    base-time 750000 \
    sched-entry S 0x20 20000 \
    sched-entry S 0xdf 980000 \
    flags 0x2
# tc qdisc replace dev swp1 root taprio \
    num_tc 8 \
    map_0 1 2 3 4 5 6 7 \
    queues 1@0 1@1 1@2 1@3 1@4 1@5 1@6 1@7 \
    base-time 250000 \
    sched-entry S 0x20 20000 \
    sched-entry S 0xdf 980000 \
    flags 0x2
# tc qdisc replace dev swp2 root taprio \
    num_tc 8 \
    map_0 1 2 3 4 5 6 7 \
    queues 1@0 1@1 1@2 1@3 1@4 1@5 1@6 1@7 \
    base-time 250000 \
    sched-entry S 0x20 20000 \
    sched-entry S 0xdf 980000 \
    flags 0x2
```

After that, the evaluation can follow the various use cases described previously with the default configuration: baseline operation, scheduled traffic evaluation with or without best-effort traffic.

Note:

An arbitrary low period might run into the scheduling limits of the systems, and result in errors in the `tsn-app` logs, as the systems may no longer be able to keep up with the requested pace.

5.2.3.7.5.7 Enabling AF_XDP sockets in TSN sample application

A new feature makes it possible to use AF_XDP sockets with the Linux `tsn-app` application, to take advantage of the lower latency offered by the AF_XDP path. The steps below describe how to reconfigure an i.MX 8M Plus LPDDR4 EVK, i.MX 8DXL LPDDR4 EVK, i.MX 93 EVK, or i.MX 93 9x9 LPDDR4 QSB board to use AF_XDP sockets.

Note: *i.MX 95 19x19 LPDDR5 EVK does not support AF_XDP use case for now.*

- Stop the application and TSN stack if they were already running:

```
# avb.sh stop_all
```

- Edit the application configuration file:

```
# vi /etc/genavb/apps-tsn-network-controller.cfg
```


3. To enable AF_XDP mode, replace the line:

```
CFG_EXTERNAL_MEDIA_APP_OPT="-m network_only -r controller"
```

With:

```
CFG_EXTERNAL_MEDIA_APP_OPT="-m network_only -r controller -x"
```

4. Attach the XDP program to the TSN interface. This step can be done at any time, even if the TSN sample application is still running with its default configuration, as long as it is done before restarting it in AF_XDP mode.

```
# ip l set dev ethX xdp obj /lib/firmware/genavb/genavb-xdp.bin
```

5. Restart the tsn-app application in AF_XDP mode:

```
# avb.sh start
```

After that, the evaluation can follow the various use cases described previously with the default configuration: baseline operation, scheduled traffic evaluation with or without best-effort traffic. Statistics should be similar to or better than the default configuration, except for traffic latencies: because AF_XDP currently cannot provide packet timestamps, traffic latencies display bogus values that should be ignored. The tables below summarize typical values (in μ s), on a setup using a 1 ms period.

Table 73. Timing statistics without any concurrent traffic

	min	mean	max
Sched err (controller)	6	7	16
Processing time (controller)	10	13	19
Total time (controller)	16	20	33

Table 74. Timing statistics with TX best-effort traffic

	min	mean	max
Sched err (controller)	18	25	51
Processing time (controller)	21	26	52
Total time (controller)	42	51	108

Table 75. Timing statistics with RX best-effort traffic

	min	mean	max
Sched err (controller)	7	9	34
Processing time (controller)	8	11	25
Total time (controller)	15	21	55

5.2.3.7.5.8 OPC UA server evaluation

The OPC UA server address is in this format : opc.tcp://<endpoint IP address>:4840/

Once connected, the server objects can be browsed and accessed. The same statistics described in [the TSN example application logs](#) are available as OPC UA objects. The OPC UA server traffic is classified as best effort and doesn't affect the time sensitive traffic.

See below screenshot using FreeOPCUA GUI client:

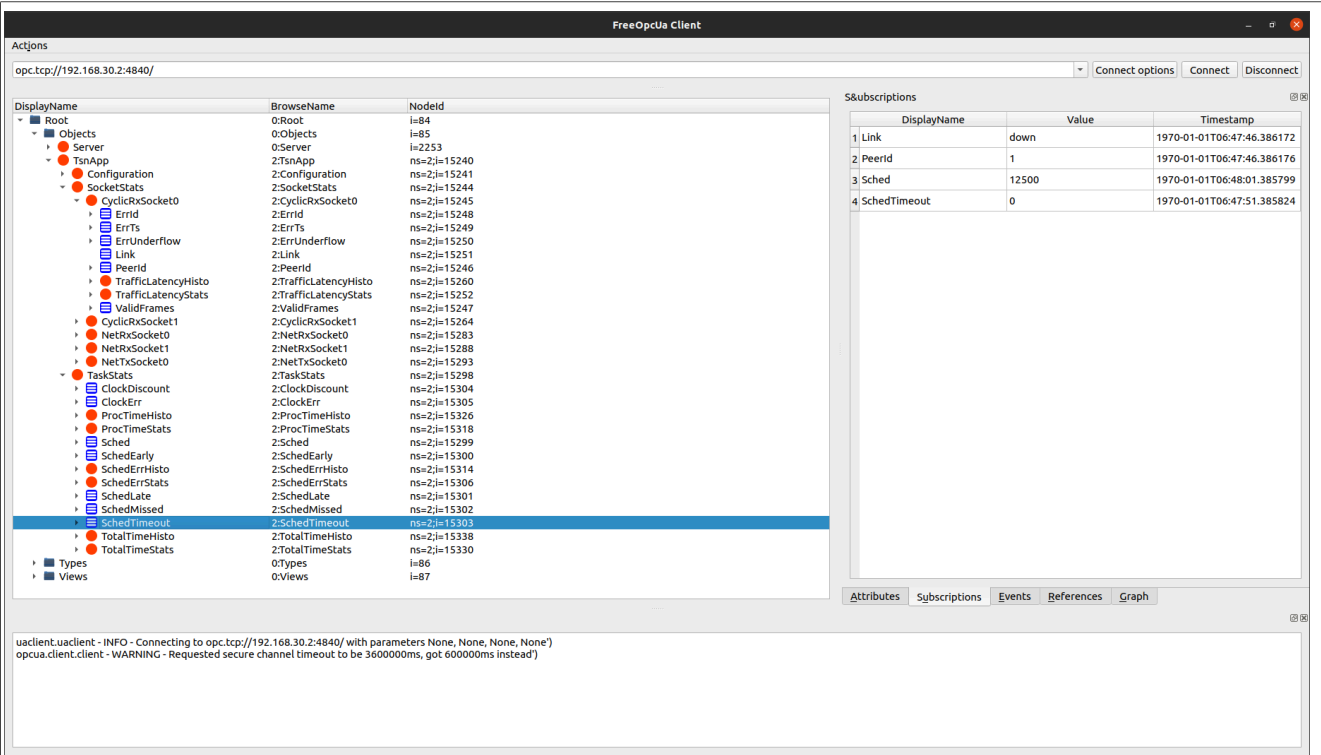


Figure 100. FreeOPCUA GUI client

5.2.4 Configuration files

5.2.4.1 System

The system configuration files (`system.cfg.<configuration>`) depending on configuration: Endpoint AVB, Endpoint TSN, Bridge, Hybrid AVB), located under `/etc/genavb/`, lists system network interface names, PTP hardware clock device names, XDP queues configuration and the Network protocols supported modes. On script startup, a symlink with the name `/etc/genavb/system.cfg` will be created to the right configuration file and will be used by the stack. The default values are used if the configuration file or the option key are missing. The values in the installed file may also required to be updated to match the system configuration.

Table 76. Logical ports

This section lists network interface names.
Currently endpoint package supports a single endpoint and bridge package a single bridge (with up to 5 ports).

Name	Key	Default value	Description
Endpoint Interface	endpoint	eth0	Endpoint network interface name. Only valid for endpoint package, otherwise should be set to “off”
Endpoint Hybrid Interface	endpoint_hybrid_port	off	Endpoint network interface name connected to the on-board bridge. Only valid for hybrid configuration,

Table 76. Logical ports...continued

This section lists network interface names.

Currently endpoint package supports a single endpoint and bridge package a single bridge (with up to 5 ports).

Name	Key	Default value	Description
			otherwise should be set to "off"
Bridge Name	bridge	br0	Bridge name. Only valid for bridge and hybrid package, otherwise should be set to "off"
Bridge 0 Interfaces	bridge_0	swp0, swp1, swp2, swp3, swp_cpu	Bridge 0 network interface names (comma separated). Only valid for bridge and hybrid package, otherwise should be set to "off"
Bridge Hybrid Interface	bridge_hybrid_port	off	Bridge network interface name connected to the on-board endpoint. Only valid for hybrid configuration, otherwise should be set to "off"

Table 77. Clock

This section lists clock device names.

Clocks names are either a PHC device name or a generic software clock (*sw_clock*). Local clock points to a PHC device, target clocks point to either:

- The same PHC device as local clock (gPTP time is reflected in the local clock)
- A generic software clock (in which case gPTP time is not reflected in the local clock).

Name	Key	Default value	Description
Endpoints gPTP domain 0 target clock	endpoint_gptp_0	/dev/ptp0	List of PTP devices for each endpoint's target clock for domain 0. Only valid for endpoint package.
Endpoints gPTP domain 1 target clock	endpoint_gptp_1	sw_clock	List of PTP devices for each endpoint's target clock for domain 1. Only valid for endpoint package.
Endpoint local clock	endpoint_local	/dev/ptp0	Endpoint clock for the local clock. Only valid for endpoint package.
Bridge gPTP domain 0 target clock	bridge_gptp_0	sw_clock	Bridge clock for gPTP domain 0 target clock. Only valid for bridge package.
Bridge gPTP domain 1 target clock	bridge_gptp_1	sw_clock	Bridge clock for gPTP domain 1 target clock. Only valid for bridge package.
Bridge local clock	bridge_local	/dev/ptp1	Bridge clock for the local clock. Only valid for bridge package.

Table 78. XDP

This section lists HW queues indexes for XDP packets.

Name	Key	Default value	Description
Endpoint XDP receive queues	endpoint_queue_rx	0, 0	Endpoint receive HW queues (one per port, comma seperated) for XDP packets.
Endpoint XDP transmit queues	endpoint_queue_tx	1, 1	Endpoint transmit HW queues (one per port, comma seperated) for XDP packets.

Table 79. Net Modes

This section lists the network modes used by the stack (TSN and AVB processes).

Network modes are configured per network protocol:

- AVB: using custom network sockets through the AVB kernel module.
- STD: using linux standard API sockets.
- XDP: using AF_XDP sockets (only available using sockets API in shared library).

Name	Key	Default value	Supported values	Description
AVTP net mode	avtp_net_mode	avb	avb, std	Net mode for the AVTP protocol.
AVDECC net mode	avdecc_net_mode	avb	avb, std	Net mode for the AVDECC protocol.
Endpoint SRP net mode	endpoint_srp_net_mode	avb	avb, std	Net mode for the SRP protocol on the Endpoint.
Bridge SRP net mode	bridge_srp_net_mode	std	avb, std	Net mode for the SRP protocol on the Bridge.
Endpoint GPTP net mode	endpoint_gptp_net_mode	avb	avb, std	Net mode for the GPTP protocol on the Endpoint.
Bridge GPTP net mode	bridge_gptp_net_mode	std	avb, std	Net mode for the GPTP protocol on the Bridge.

5.2.4.2 gPTP

The gPTP general parameters as well as default domain (domain 0) parameters are defined in the following configuration files depending on the package used:

- Endpoint package: /etc/genavb/fgptp.cfg
- Bridge package: /etc/genavb/fgptp-br.cfg

To enable other domains, new configuration files must be created with the associated domain instance appended to the configuration file name e.g.:

- Endpoint package, domain 1: /etc/genavb/fgptp.cfg-1
- Bridge package, domain 1: /etc/genavb/fgptp-br.cfg-1

Attention:

By default the GenAVB/TSN gPTP stack is packaged with the general parameters configuration file (fgptp.cfg or fgptp-br.cfg) and a reference configuration for domain 1 (fgptp.cfg-1 or fgptp-br.cfg-1)

5.2.4.2.1 General

Profile

The gPTP stack can operate in two different modes known as 'standard' or 'automotive' profiles.

When the 'standard' profile is selected, the gPTP stack operates following the specifications described in IEEE 802.1AS. When the 'automotive' profile is selected, the gPTP stack operates following the specifications described in the [AVnu AutoCDSFunctionalSpec_1.4](#) which is a subset of the IEEE 802.1AS standard optimized for automotive applications. IEEE 802.1AS-2020 features are not available in 'automotive' profile (e.g. Multiple domains).

The automotive environment is unique in that it is a closed system. Every network device is known prior to startup and devices do not enter or leave the network, except in the case of failures. Because of the closed nature of the automotive network, it is possible to simplify and improve gPTP startup performance. Specifically, functions like election of a grand master and calculations of wire delays are tasks that can be optimized for a closed system.

Reverse sync feature control

The Reverse Sync feature (Avnu specification) should be used for test/evaluation purpose only. Usually, to measure the accuracy of the clock synchronization, the traditional approach is to use a 1 Pulse Per Second (1PPS) physical output. While this is a good approach, there may be cases where using a 1PPS output is not feasible. More flexible and fully relying on software implementation the Reverse Sync feature serves the same objective using the standard gPTP Sync/Follow-Up messages to relay the timing information, from the Slave back to the GM.

Neighbor propagation delay threshold

The parameter `neighborPropDelayThresh` defines the propagation time threshold, above which a port is not considered capable of participating in the IEEE 802.1AS protocol (see IEEE 802.1AS-2020 - 11.2.2 Determination of `asCapable` and `asCapableAcrossDomains`). If a computed `neighborPropDelay` exceeds `neighborPropDelayThresh`, then `asCapable` is set to FALSE for the port. This setting does not apply to Automotive profile where a link is always considered to be capable or running IEEE 802.1AS.

IEEE 802.1AS-2011 Compatibility

The parameter `force_2011` defines if the gPTP Stack operates following the IEEE 802.1AS-2011 standard, i.e. disabling the IEEE 802.1AS-2020 specifics features such as Multiple Domain support. The use of this option may, in some cases, improve compatibility with gPTP equipment not supporting IEEE 802.1AS-2020 standard.

Table 80. General parameters
General configuration parameters^[1]

Name	Key	Default value	Range	Description
Profile	profile	"standard"	"standard" or "automotive"	Set fgptp main profile. "standard" - IEEE 802.1AS specs, "automotive" - AVnu automotive profile
Grandmaster ID	gm_id	"0x0001f2ffe0025fe"	64bits EUI format	Set static grandmaster ID in host order (used by automotive profile, ignored in case of standard profile)
Domains	domain_number	0: for default domain -1: for domains different from 0	-1 to 127	Disable (-1) or assign a gPTP domain number to a domain instance.

Table 80. General parameters...continued
General configuration parameters^[1]

Name	Key	Default value	Range	Description
802.1AS-2011 mode	force_2011	no	"no" or "yes"	Set to "yes" to force 802.1AS-2011 standard. "no" to enable 802.1AS-2020 full support.
Log output level	log_level	info	crit, err, init, info, or dbg	Set this configuration to dbg to enable debug mode
Reverse sync feature control	reverse_sync	0	0 or 1	Set to 1 to enable reverse sync feature.
Reverse sync feature interval	reverse_sync_interval	112	32 to 10000	Reverse sync transmit interval in ms units
Neighbor propagation delay threshold	neighborPropDelayThresh	800	32 to 10 000000	Neighbor propagation delay threshold expressed in ns
Statistics output interval	statsInterval	10	0 to 255	Statistics output interval expressed in seconds. Use 0 to disable statistics

[1] For domain instances other than 0, only domain_number is configurable in this section.

5.2.4.2.2 Grandmaster parameters

This section defines the native Grand Master capabilities of a time-aware system (see IEEE 802.1AS-2020 - 8.6.2 PTP Instance attributes). Grand Master capabilities parameters are defined in the main configuration file for gPTP domain 0 (e.g. fgptp.cfg) and in the additional per domain configuration files for other domains (e.g. fgptp.cfg-1).

gmCapable defines if the time-aware system is capable of being a grandmaster. By default gmCapable is set to 1 as in standard profile operation the Grand Master is elected dynamically by the BMCA. In case of automotive profile gmCapable must be set on each AED node to match the required network topology (that is, within a given gPTP domain only one node must have its gmCapable property set to 1).

priority1, priority2, clockClass, clockAccuracy and offsetScaledLogVariance are parameters used by the Best Master Clock algorithm to determine which of the Grand Master capable node within the gPTP domain has the highest priority/quality. Note that the lowest value for these parameters matches the highest priority/quality.

Table 81. Grandmaster parameters
Grandmaster capabilities parameters^[1]

Name	Key	Default value	Range	Description
Grandmaster capable setting	gmCapable	1	0 or 1	Set to 1 if the device has grandmaster capability. Ignored in automotive profile if the port is SLAVE.
Grandmaster priority1 value	priority1	248 for AED-E and 246 for AED-B	0 to 255	Set the priority1 value of this clock
Grandmaster priority2 value	priority2	248	0 to 255	Set the priority2 value of this clock
Grandmaster clock class value	clockClass	248	0 to 255	Set the class value of this clock

Table 81. Grandmaster parameters...continued
Grandmaster capabilities parameters^[1]

Name	Key	Default value	Range	Description
Grandmaster clock accuracy value	clockAccuracy	0xfe	0x0 to 0xff	Set the accuracy value of this clock
Grandmaster variance value	offsetScaledLogVariance	17258	0x0 to 0xffff	Set the offset scaled log variance value of this clock

[1] The parameters in this section are configurable for all supported domains.

5.2.4.2.3 Automotive parameters

The static pdelay feature is used only if the gPTP stack operates in automotive profile configuration.

At init time the gPTP stack's configuration file is parsed and based on neighborPropDelay_mode the specified initial_neighborPropDelay is applied to all ports and used for synchronization until a pdelay response from the peer is received. This is done only if no previously stored pdelay is available from the nvram database specified by nvram_file. As soon as a pdelay response from the peer is received the 'real' pdelay value is computed, and used for current synchronization. An indication may then be sent via callback up to the OS-dependent layer. Upon new indication the Host may update its nvram database and the stored value will be used at next restart for the corresponding port instead of the initial_neighborPropDelay. The granularity at which pdelay change indications are sent to the Host is defined by the neighborPropDelay_sensitivity parameter.

In the gPTP configuration file the neighborPropDelay_mode parameter is set to 'static' by default, meaning that a predefined propagation delay is used as described above while pdelay requests are still sent to the network.

The 'silent' mode behaves the same way as the 'static' mode except that pdelay requests are never sent at all to the network.

Optionally the neighborPropDelay_mode parameter can be set to standard forcing the stack to operate propagation delay measurements as specified in the 802.1AS specifications even if the automotive profile is selected.

(see AutoCDSFunctionalSpec-1_4 - 6.2.2 Persistent gPTP Values)

Table 82. Automotive parameters

Name	Key	Default value	Value & Range	Description
Pdelay mode	neighborPropDelay_mode	static	'static', 'silent' or 'standard'	Defines pdelay mechanism used
Static pdelay value	initial_neighborPropDelay	250	0 to 10000	Predefined pdelay value applied to all ports. Expressed in ns.
Static pdelay sensitivity	neighborPropDelay_sensitivity	10	0 to 1000	Amount of ns between two pdelay measurements required to trigger a change indication. Expressed in ns.
Nvram file name	nvram_file	/etc/genavb/fgptp.nvram		Path and nvram file name.

5.2.4.2.4 Timing

Pdelay requests and Sync messages sending intervals have a direct impact on the system synchronization performance. To reduce synchronization time while optimizing overall system load, two levels of intervals are defined. The first level called 'Initial', defines the messages intervals used until pdelay values have stabilized

and synchronization is achieved. The second level called 'Operational', defines the messages intervals used once the system is synchronized.

initialLogPdelayReqInterval and operLogPdelayReqInterval define the intervals between the sending of successive Pdelay_Req messages. initialLogSyncInterval and operLogSyncInterval define the intervals between the sending of successive Sync messages. initialLogAnnounceInterval defines the interval between the sending of successive Announce messages

(see AutoCDSFunctionalSpec-1_4 - 6.2.1 Static gPTP Values, IEC-60802 section 5, 802.1AS-2020 sections 10.7 and 11.5)

Table 83. Timing parameters

Name	Key	Default value	Value and Range	Description
Initial pdelay request interval value	initialLogPdelayReqInterval	0	0 to 3	Set pdelay request initial interval between the sending of successive Pdelay_Req messages. Expressed in log2 unit (default 0 -> 1s).
Initial sync interval value	initialLogSyncInterval	-3	-5 to 0	Set sync transmit initial interval between the sending of successive Sync messages. Expressed in log2 unit (default -3-> 125ms).
Initial announce interval value	initialLogAnnounceInterval	0	0 to 3	Set initial announce transmit interval between the sending of successive Announce messages. Expressed in log2 unit (default 0 -> 1s).
Operational pdelay request interval value	operLogPdelayReqInterval	0	0 to 3	Set pdelay request transmit interval used during normal operation state. Expressed in log2 unit (default 0 -> 1s).
Operational sync interval value	operLogSyncInterval	-3	-5 to 0	Set sync transmit interval used during normal operation state. Expressed in log2 unit (default -3 -> 125ms).

5.2.4.2.5 PORT_n

This section describes the settings per port where _n represents the port index starting at _n=1.

Table 84. Port related parameters

Name	Key	Default value	Value & Range	Description
Port role	portRole	disabled	'slave', 'master', 'disabled'	Static port role (ref. 802.1AS-2011, section 14.6.3, Table 10-1), applies to "automotive" profile only.
Ptp port enabled	ptpPortEnabled	1	0 or 1	Set to 1 if both time-synchronization and best master selection functions of the port should be used (ref. 802.1AS-2011, sections 14.6.4 and 10.2.4.12).

Table 84. Port related parameters...continued

Name	Key	Default value	Value & Range	Description
RX timestamp compensation	rxDelayCompensation	0	min=-100000 max=100000 (in ns units)	Compensation delay subtracted from receive timestamps.
TX timestamp compensation	txDelayCompensation	0	min=-100000 max=100000 (in ns units)	Compensation delay added to transmit timestamps.
Delay Mechanism	delayMechanism	P2P	'P2P' or 'COMMON_P2P'	Must be set to COMMON_P2P for all domains others than Domain 0. For Domain 0 the value can be either P2P or COMMON_P2P. In general, if a port is using the Common Mean Link Delay Service (CMLDS), the DelayMechanism should be set to COMMON_P2P.

The following table lists the recommended Rx and Tx compensation values to be applied to the supported NXP boards for optimized GPTP synchronization.

Table 85. PHY Delay Compensation Values

Board Type	PHY Name	Speed (Mbps)	Port	rxDelay Compensation	txDelay Compensation
LS1028ARDB	VSC8514	1000	ENETC2	274	349
I.MX 8M Plus EVK	RTL8211F	1000	ENET2	569	184
I.MX 8M Plus EVK	RTL8211F	100	ENET2	748	640

Note: The table indicates computed values but it is recommended to use a 50 ns propagation delay margin to avoid negative propagation delay measurements due to propagation delay jitter.

5.2.4.3 SRP

The SRP parameters are defined in the following configuration files, depending on the package used:

- Endpoint: /etc/genavb/srp.cfg
- Bridge: /etc/genavb/srp-br.cfg

The default values are used if the configuration file or the option key are missing. The values in the installed file may also required an update to match the system configuration.

Table 86. SRP General

This section lists general SRP stack component parameters.

Name	Key	Default value	Range	Description
Log output level	log_level	info	crit, err, init, info, or dbg	Log level for the SRP stack component.

Table 87. MSRP
This section lists MSRP parameters.

Name	Key	Default value	Range	Description
Enabled	enabled	1	0-disabled, 1-enabled	Enable/disable MSRP at runtime.

5.2.5 Log files

Several log files are available at runtime to monitor the different stack components.

5.2.5.1 gPTP Endpoint

Logs are stored in /var/log/fgptp.

- Linux command:

```
# tail -f /var/log/fgptp
```

- If the stack is configured in automotive mode, then the log contains:

```
Running fgptp in automotive profile on interface eth0
```

- Port Role, Port AS-capability and link Status are reported each time there is a change in the link state (link is 802.1AS capable or not) or upon Grand Master (GM) change. This information is also displayed regularly along with current synchronization and pdelay statistics for each of the enabled gPTP domain:

```
Port(0) domain(0,0): role changed from DISABLED to SLAVE
...
Port(0) domain(0,0): Slave - Link: Up - AS_Capable: Yes
```

- Selected Grand Master (GM) capabilities are reported upon new GM selection. *Root Identity* represents the clock ID of the currently selected GM. *Priority1*, *Priority2*, *Class* and *Accuracy* describe the clock quality of the selected GM. Finally, the *Source Port Identity* of the peer master port (e.g. the bridge port the local slave port is connected to). This information is displayed for each of the enabled gPTP domain:

```
domain(0,0) Grand master: root identity 00049ffffe039e35
domain(0,0) Grand master: priority1 245 priority2
domain(0,0) Grand master: class 248 accuracy 248
domain(0,0) Grand master: variance 17258
domain(0,0) Grand master: source port identity 0001f2fffe0025fe, port number 2
```

- Synchronization State* is reported upon GM selection (SYNCHRONIZED) or when no GM is detected (NOT SYNCHRONIZED). *Synchronization Time* expressed in ms represents the time it took for the local clock to reach synchronization threshold starting from the first SYNC message received. This information is displayed for each of the enabled domain.

```
Port(0) domain(0) SYNCHRONIZED - synchronization time (ms): 250
```

- Pdelay* (propagation delay) and local clock adjustments are printed out every 5 seconds. *PDelay* is expressed in ns units and represents the one-way delay from the endpoint and its peer master. *Correction* is expressed in parts per billion and represents the frequency adjustment performed to the local clock. *Offset* is expressed in ns represents the resulting difference between the locally adjusted clock and the reference gPTP GrandMaster's clock. (Min/Max/Avg and Variance are computed for both Correction and Offset statistics). *PDelay* is displayed only for Domain 0. *Correction* and *Offset* are displayed for each of the enabled domain.

```
Port 0 domain(0,0): Propagation delay (ns): 37.60 min 34 avg 36 max 45
variance 17
```

```

Port 0 domain(0,0): Correction applied to local clock (ppb): min -5603 avg 5572
max 5538 variance 148
Port 0 domain(0,0): Offset between GM and local clock (ns) min -12 avg 4 max 22
variance 111
...
Port 0 domain(1,20): Correction applied to local clock (ppb): min 32074 avg
32314 max 32574 variance 17695
Port 0 domain(1,20): Offset between GM and local clock (ns) min -61 avg 3 max
70 variance 1149

```

- The following per port per domain statistics (32 bits counters) are printed out every 15 seconds on slave and master entities:

Table 88. Port statistics displayed on slave and master entities

Receive counters	
PortStatRxPkts	Number of gPTP packets received (ether type 0x88F7)
PortStatRxSyncCount	Number of SYNC packets received
PortStatRxSyncReceiptTimeouts	Number of SYNC packets receive timeout
PortStatRxFollowUpCount	Number of FOLLOW-UP packets received
PortStatRxAnnounce	Number of ANNOUNCE packets received
PortStatAnnounceReceiptTimeouts	Number of ANNOUNCE packets timeout
PortStatAnnounceReceiptDropped	Number of ANNOUNCE packets dropped by the entity
PortStatRxSignaling	Number of SIGNALING packets received
PortStatRxPdelayRequest	Number of PDELAY REQUEST packets received
PortStatRxPdelayResponse	Number of PDELAY RESPONSE packets received
PortStatPdelayAllowedLostResponsesExceeded	Number of excess of allowed lost responses to PDELAY requests
PortStatRxPdelayResponseFollowUp	Number of PDELAY FOLLOW-UP packets received
PortStatRxErrEtype	Number of ether type errors (not 0x88F7)
PortStatRxErrPortId	Number of port ID errors
Transmit counters	
PortStatTxPkts	Number of gPTP packets transmitted
PortStatTxSyncCount	Number of SYNC packets transmitted
PortStatTxFollowUpCount	Number of FOLLOW-UP packets transmitted
PortStatTxAnnounce	Number of ANNOUNCE packets transmitted
PortStatTxSignaling	Number of SIGNALING packets transmitted
PortStatTxPdelayReques	Number of PDELAY REQUEST packets transmitted
PortStatTxPdelayResponse	Number of PDELAY RESPONSE packets transmitted
PortStatTxPdelayResponseFollowUp	Number of PDELAY FOLLOW-UP packets transmitted
PortStatTxErr	Number of transmit errors
PortStatTxErrAlloc	Number of transmit packets allocation errors
Miscellaneous counters	
PortStatAdjustOnSync	Number of adjustments performed upon SYNC received

Table 88. Port statistics displayed on slave and master entities...continued

PortStatMdPdelayReqSmReset	Number of reset of the PDELAY REQUEST state machine
PortStatMdSyncRcvSmReset	Number of reset of the SYNC RECEIVE state machine
PortStatHwTsRequest	Number of egress timestamp requests
PortStatHwTsHandler	Number of egress timestamp notification
PortStatNumSynchronizationLoss	Number or synchronization loss on the slave endpoint (e.g. GM change, GM reference clock discontinuity...)
PortStatNumNotAsCapable	Number of transitions from AS_Capable=TRUE to AS_Capable=FALSE

5.2.5.2 gPTP Bridge

Logs are stored in /var/log/fgptp-br.

- Linux command:

```
# tail -f /var/log/fgptp-br
```

- The bridge stack statistics are similar to the endpoint stack ones except that they are reported for each of the external ports of the switch (Port 0 to 3) and also for the internal port connected to the endpoint stack (Port 4) in case of Hybrid setup.
- Pdelay* (propagation delay) is printed only for Domain 0. *Link status*, *AS capability* and *Port Role* are printed out for each port and each gPTP domain.

```
Port 0 domain(0,0): Role: Disabled Link: Up AS_Capable: No neighborGptpCapable:
  No DelayMechanism: P2P
Port 1 domain(0,0): Role: Disabled Link: Up AS_Capable: No neighborGptpCapable:
  No DelayMechanism: P2P
Port 2 domain(0,0): Role: Disabled Link: Up AS_Capable: Yes
  neighborGptpCapable: Yes DelayMechanism: P2P
Port 2 domain(0,0): Propagation delay (ns): 433.98 min 425 avg 438 max 457
  variance 87
Port 3 domain(0,0): Role: Disabled Link: Up AS_Capable: No neighborGptpCapable:
  No DelayMechanism: P2P
Port 4 domain(0,0): Role Master Link: Up AS_Capable: Yes neighborGptpCapable:
  Yes DelayMechanism: P2P
Port 4 domain(0,0): Propagation delay (ns): 433.98 min 425 avg 438 max 457
  variance 87
...
Port 0 domain(1,20): Role: Disabled Link: Up AS_Capable: No
  neighborGptpCapable: No DelayMechanism: COMMON_P2P
Port 1 domain(1,20): Role: Disabled Link: Up AS_Capable: No
  neighborGptpCapable: No DelayMechanism: COMMON_P2P
Port 2 domain(1,20): Role: Disabled Link: Up AS_Capable: Yes
  neighborGptpCapable: Yes DelayMechanism: COMMON_P2P
Port 3 domain(1,20): Role: Disabled Link: Up AS_Capable: No
  neighborGptpCapable: No DelayMechanism: COMMON_P2P
Port 4 domain(1,20): Role Master Link: Up AS_Capable: Yes neighborGptpCapable:
  Yes DelayMechanism: COMMON_P2P
```

5.2.5.3 SRP Bridge

Logs are stored in /var/log/tsn-br.

- Linux command:

```
# tail -f /var/log/tsn-br | grep srp
```

- SRP protocol information is reported per port

```
INFO srp      msrp_vector_add_event      : port(0) domain(5, 2, 2)
MSRP_ATTR_TYPE_DOMAIN MRP_ATTR_EVT_JOINMT
INFO srp      msrp_vector_add_event      : port(0) domain(6, 3, 2)
MSRP_ATTR_TYPE_DOMAIN MRP_ATTR_EVT_JOINMT
INFO srp      msrp_vector_add_event      : port(1) domain(5, 2, 2)
MSRP_ATTR_TYPE_DOMAIN MRP_ATTR_EVT_JOINMT
INFO srp      msrp_vector_add_event      : port(1) domain(6, 3, 2)
MSRP_ATTR_TYPE_DOMAIN MRP_ATTR_EVT_JOINMT
INFO srp      msrp_vector_add_event      : port(2) domain(5, 2, 2)
MSRP_ATTR_TYPE_DOMAIN MRP_ATTR_EVT_JOINMT
INFO srp      msrp_vector_add_event      : port(2) domain(6, 3, 2)
MSRP_ATTR_TYPE_DOMAIN MRP_ATTR_EVT_JOINMT
INFO srp      msrp_vector_add_event      : port(4) domain(5, 2, 2)
MSRP_ATTR_TYPE_DOMAIN MRP_ATTR_EVT_JOINMT
INFO srp      msrp_vector_add_event      : port(4) domain(6, 3, 2)
MSRP_ATTR_TYPE_DOMAIN MRP_ATTR_EVT_JOINMT
INFO srp      msrp_vector_handler        : port(3) domain(5, 2, 2)
MRP_ATTR_EVT_MT
INFO srp      msrp_vector_handler        : port(3) domain(6, 3, 2)
MRP_ATTR_EVT_MT
INFO srp      msrp_vector_handler        : port(3) domain(5, 2, 2)
MRP_ATTR_EVT_JOINMT
INFO srp      msrp_vector_handler        : port(3) domain(6, 3, 2)
MRP_ATTR_EVT_JOINMT
INFO srp      msrp_vector_add_event      : port(3) domain(5, 2, 2)
MSRP_ATTR_TYPE_DOMAIN MRP_ATTR_EVT_JOININ
INFO srp      msrp_vector_add_event      : port(3) domain(6, 3, 2)
MSRP_ATTR_TYPE_DOMAIN MRP_ATTR_EVT_JOININ
INFO srp      msrp_vector_add_event      : port(3) domain(5, 2, 2)
MSRP_ATTR_TYPE_DOMAIN MRP_ATTR_EVT_JOININ
INFO srp      msrp_vector_add_event      : port(3) domain(6, 3, 2)
MSRP_ATTR_TYPE_DOMAIN MRP_ATTR_EVT_JOININ
INFO srp      msrp_vector_add_event      : port(0) domain(5, 2, 2)
MSRP_ATTR_TYPE_DOMAIN MRP_ATTR_EVT_MT
INFO srp      msrp_vector_add_event      : port(0) domain(6, 3, 2)
MSRP_ATTR_TYPE_DOMAIN MRP_ATTR_EVT_MT
INFO srp      msrp_vector_add_event      : port(1) domain(5, 2, 2)
MSRP_ATTR_TYPE_DOMAIN MRP_ATTR_EVT_MT
INFO srp      msrp_vector_add_event      : port(1) domain(6, 3, 2)
MSRP_ATTR_TYPE_DOMAIN MRP_ATTR_EVT_MT
INFO srp      msrp_vector_add_event      : port(2) domain(5, 2, 2)
MSRP_ATTR_TYPE_DOMAIN MRP_ATTR_EVT_MT
INFO srp      msrp_vector_add_event      : port(2) domain(6, 3, 2)
MSRP_ATTR_TYPE_DOMAIN MRP_ATTR_EVT_MT
INFO srp      msrp_vector_add_event      : port(4) domain(5, 2, 2)
MSRP_ATTR_TYPE_DOMAIN MRP_ATTR_EVT_MT
INFO srp      msrp_vector_add_event      : port(4) domain(6, 3, 2)
MSRP_ATTR_TYPE_DOMAIN MRP_ATTR_EVT_MT
```

5.2.5.4 TSN Endpoint example application

Logs are stored in /var/log/tsn_app.

The TSN application has various counters and statistics which help to validate:

- application scheduling and processing timing statistics
- network traffic correctness and latency statistics

Most of the information in the logs are either:

- Counters: single integer values counting specific events (frames received, transmitted, errors, etc)
- Statistics: composite data over a series of measurements: min (minimum during the last period), mean (average of measurements during the last period), max (maximum during the last period), rms[~](root mean square of measurements during the last period), stddev² (square of standard deviation during the last period), absmin(absolute minimum since the application start), absmax (absolute maximum since the application start)
- Histograms: number and size of slots of the histogram on the 1st line, array of counters for each slot on the 2nd line

5.2.5.4.1 Main TSN task

The main TSN task logs are described below:

- Scheduling counters (“sched” should increment of 500 per second):

```
INFO 1604531064    tsn_task_stats_print          tsn task(0x37d8d630)
INFO 1604531064    tsn_task_stats_print          sched                : 1700000
INFO 1604531064    tsn_task_stats_print          sched early         : 0
INFO 1604531064    tsn_task_stats_print          sched late          : 0
INFO 1604531064    tsn_task_stats_print          sched missed        : 0
INFO 1604531064    tsn_task_stats_print          sched timeout       : 0
INFO 1604531064    tsn_task_stats_print          clock discont       : 0
INFO 1604531064    tsn_task_stats_print          clock err           : 0
```

- Scheduling error statistics (difference between actual task scheduling time and programmed time, in ns):

```
stats(0x2ddd4560) sched err min 8062 mean 10662 max 18862 rms^2 116252492
stddev^2 2558383 absmin 2842 absmax 35082
```

- Scheduling error histogram (XXX ns bucket)

```
n_slot 101 slot_size 10000  
18081119 21722197 1676 7 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0  
0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0  
0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
```

- Processing time statistics (main task duration in ns)

```
stats(0x2ddd49c0) processing time min 21600 mean 27564 max 47460 rms^2
768483628 stddev^2 8664988 absmin 12540 absmax 152100
```

- Processing time histogram (XXX ns bucket)

```
n_slot 101 slot_size 1000  
0_0 0 0 0 0 0 0 0_0 0 0 0 9 13 7 259 302 187 14273 64041 91896 427011 1473790  
2154889 4239252 4681777 5682883 5516429 4915621 3789351 2688466 1709802  
1036649 587716 327653 170711 89620 47637  
25422 15211 9573 6908 4852 3694 3101 2910 3028 3036 3331 3091 2842 2375 1748  
1251 895 585 342 208 100 63 64 26 21 13 9 14 6 7 4 4 6 3 4 1 0 2 2 1 1 0 0 0 1  
0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
```

- Total time statistics (scheduling error + processing time, in ins)

```
stats(0x2ddd4e20) total time min 30082 mean 38227 max 55862 rms^2 1473506486
stddev^2 12160755 absmin 18962 absmax 170082
```

- Total time histogram (XXX ns bucket)

```
n_slot 101 slot_size 1000
0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 1 0 1 2 15 16 57 273 1674 7795 30165 80952
240707 651779 1497160 2290907 3606869 4077757 4761502 4647779 4239893 3631222
2831749 2209868 1631293 1173544
819740 536234 338535 201816 115114 65529 36480 21413 12892 8249 5613 4178 3310
3034 2890 2921 2615 2461 2132 1790 1427 1082 844 585 353 266 170 98 73 48 20
25 6 14 14 1 6 2 3 8 2 2 1 3 2 1 1 2 1 2 0 1 1 0 0 0 10 0
```

5.2.5.4.2 Network socket

Below is an example of the network socket logs:

- Low-level network socket. Only frames relevant to the network socket (Layer 2) are counted here:

```
INFO 1604531059 net_socket_stats_print net rx socket(0x37d8d660) 0
INFO 1604531059 net_socket_stats_print frames : 1697802
INFO 1604531059 net_socket_stats_print frames err : 0
INFO 1604531059 net_socket_stats_print net tx socket(0x37d8d6e0) 0
INFO 1604531059 net_socket_stats_print frames : 1697500
INFO 1604531059 net_socket_stats_print frames err : 0
```

5.2.5.4.3 Application socket

- Application header is checked at this level. Also, the timestamps from the remote peers are verified as well which guarantees that only expected and in sequence data is processed.

```
INFO 1604531069 socket_stats_print cyclic rx socket(0x419560)
net_sock(0x37d8d660) peer id: 1
INFO 1604531069 socket_stats_print valid frames : 1702497
INFO 1604531069 socket_stats_print err id : 0
INFO 1604531069 socket_stats_print err ts : 305
INFO 1604531069 socket_stats_print err underflow : 2
INFO 1604531069 socket_stats_print link up
```

- Traffic latency statistics (the difference between the theoretical scheduling time of the peer that sent the frame and the frame receive time (measured by the MAC), in ns)

```
stats(0x419a28) traffic latency min 503417 mean 503503 max 503637 rms^2
253515981945 stddev^2 2004 absmin 503397 absmax 504337
```

- Traffic latency histogram (XXX ns bucket)

[illegible]

5.3 IEEE 1588/802.1AS

IEEE 1588 is the IEEE standard for a precision clock synchronization protocol for networked measurement and control systems.

IEEE 802.1AS is the IEEE standard for local and metropolitan area networks – timing and synchronization for time-sensitive applications in bridged local area networks. It specifies the use of IEEE 1588 specifications where applicable in the context of IEEE Std 802.1D-2004 and IEEE Std 802.1Q-2005.

NXP's Layerscape platform provides hardware assist for 1588 compliant time stamping with the 1588 timer module to support applications of IEEE 1588/802.1AS.

5.3.1 Introduction

The i.MX and Layerscape platforms provided by NXP support hardware assist for 1588-compliant time stamping with the 1588 timer module. The software components require to run the IEEE 1588/802.1AS protocol utilizing hardware features that are listed below:

1. Linux PTP Hardware Clock (PHC) driver
2. Linux Ethernet controller driver with hardware timestamping support
3. A software stack application for IEEE 1588/802.1AS

Note: In this document, IEEE 1588 mentioned is IEEE 1588-2008 whereas the IEEE 802.1AS mentioned is IEEE 802.1AS-2011.

5.3.2 IEEE 1588 device types

There are five basic types of PTP devices in IEEE 1588.

1. **Ordinary clock:** A clock that has a single Precision Time Protocol (PTP) port in a domain and maintains the timescale used in the domain. It can serve as a source of time (if used as a master clock) or can synchronize with another clock (if used as a slave clock).
2. **Boundary clock**
A clock that has multiple Precision Time Protocol (PTP) ports in a domain and maintains the timescale used in the domain. It may serve as a source of time (be a master clock) or may synchronize to another clock (be a slave clock).
3. **End-to-end transparent clock**
A transparent clock that supports the use of the end-to-end delay measurement mechanism between slave clocks and the master clock.
4. **Peer-to-peer transparent clock**
A transparent clock that provides Precision Time Protocol (PTP) event transit time information. It also provides corrections for the propagation delay of the link connected to the port receiving the PTP event message. In the presence of peer-to-peer transparent clocks, delay measurements between slave clocks and the master clock are performed using the peer-to-peer delay measurement mechanism.
5. **Management node**
A device that configures and monitors clocks.

Note: Transparent clock is a device that measures the time taken for a PTP event message to transit the device. It provides this information to clocks receiving the PTP event message.

5.3.3 IEEE 802.1AS time-aware systems

In gPTP, there are only two types of time-aware systems: end stations and Bridges, while IEEE 1588 has ordinary clocks, boundary clocks, end-to-end transparent clocks, and P2P transparent clocks. A time-aware end station corresponds to an IEEE 1588 ordinary clock, and a time-aware Bridge is a type of IEEE 1588

boundary clock where its operation is very tightly defined, so much so that a time-aware Bridge with Ethernet ports can be shown to be mathematically equivalent to a P2P transparent clock in terms of how synchronization is performed.

1. Time-aware end station: An end station that is capable of acting as the source of synchronized time on the network, or destination of synchronized time using the IEEE 802.1AS protocol, or both.
2. Time-aware bridge: A bridge that is capable of communicating synchronized time received on one port to other ports, using the IEEE 802.1AS protocol.

5.3.4 Software stacks

5.3.4.1 linuxptp stack

Features of open source linuxptp

- Supports hardware and software time stamping via the Linux `SO_TIMESTAMPING` socket option.
- Supports the Linux PTP Hardware Clock (PHC) subsystem by using the `clock_gettime` family of calls, including the `clock_adjtimex` system call.
- Implements Boundary Clock (BC), Ordinary Clock (OC), Modular design allowing painless addition of new transports and clock servos.
- Implements unicast operation.
- Supports a number of profiles, including:
 - The automotive profile.
 - The default 1588 profile.
 - The enterprise profile.
 - The telecom profiles G.8265.1, G.8275.1, and G.8275.2.
 - Supports the NetSync Monitor protocol.
- Implements peer to peer one-step.
- Supports bonded, IPoIB, and vlan interfaces.

Note: The features listed are from *linuxptp website*. It does not mean all these features work on release boards. The hardware 1588 capability, driver support and *ptp4l* version needs to be considered. Refer to following user manual of this chapter for what had been verified.

Features added by Real-time Edge

- Supports IEEE 802.1AS-2011 in the role of time-aware bridge.
- Support dynamic direction in *ts2phc* to cooperate with *ptp4l*.

5.3.4.2 NXP GenAVB/TSN gPTP stack

Following are the features of the NXP GenAVB/TSN gPTP stack:

- Implements gPTP IEEE 802.1AS-2020, for both time-aware Endpoint and Bridge systems
- Implements gPTP BMCA
- Supports GrandMaster, Master, and Slave capabilities
- Supports multiple gPTP domains
- Supports Avnu Alliance Automotive profile
- Supports configuration profiles for the stack
- Supports hardware time stamping via the Linux `SO_TIMESTAMPING` socket option
- Supports the Linux PTP Hardware Clock (PHC) subsystem by using the `clock_gettime` family of calls, including the `clock_adjtimex` system call.

5.3.5 Quick Start for IEEE 1588

The following sections describe the clock verification process for ordinary, boundary, and transparent clocks for IEEE 1588.

5.3.5.1 Ordinary clock verification

Connect two network interfaces in a back-to-back manner for two boards. Make sure that there is no MAC address conflict on the boards, the IP addresses are set properly and ping the test network. Run `linuxptp` on each board. For example, `eth0` is used on each board.

```
$ ptp4l -i eth0 -m
```

On running the above command time, synchronization starts. Thereby, the slave `linuxptp` selected automatically synchronizes to the master. It also displays synchronization messages such as time offset, path delay and so on. For example, see the below log:

```
ptp4l[878.504]: master offset      -10 s2 freq    -2508 path delay    1826
ptp4l[878.629]: master offset      -5 s2 freq    -2502 path delay    1826
ptp4l[878.754]: master offset       0 s2 freq    -2495 path delay    1826
ptp4l[878.879]: master offset       9 s2 freq    -2482 path delay    1826
ptp4l[879.004]: master offset      -9 s2 freq    -2507 path delay    1826
ptp4l[879.129]: master offset     -24 s2 freq    -2530 path delay    1826
ptp4l[879.255]: master offset      -7 s2 freq    -2508 path delay    1826
ptp4l[879.380]: master offset      -2 s2 freq    -2502 path delay    1826
ptp4l[879.505]: master offset     -17 s2 freq    -2524 path delay    1827
ptp4l[879.630]: master offset       6 s2 freq    -2493 path delay    1827
ptp4l[879.755]: master offset       6 s2 freq    -2492 path delay    1827
ptp4l[879.880]: master offset       0 s2 freq    -2500 path delay    1827
```

Some other options of ptp4l

```
Delay Mechanism
-E          E2E, delay request-response (default)
-P          P2P, peer delay mechanism
Network Transport
-2          IEEE 802.3
-4          UDP IPV4 (default)
-6          UDP IPV6
```

Note: must keep same delay mechanism and network transport protocol used on two boards.

Configure master mode

By default, the BMC (Best Master Clock) algorithm selects the master clock. To appoint a specific clock as master, a lower "priority1" attribute value than the other clock can be set. Lower value takes precedence. For example, in the current case, specify one clock as master with the below option. (The other clock uses the default priority1 value 128.)

```
--priority1=127
```

One-step timestamping

Currently one-step timestamping is supported only on DPAA2. To use one-step timestamping, add the below option for `ptp4l` running.

```
--twoStepFlag=0
```

5.3.5.2 Boundary clock verification

At least three boards are needed for boundary clock verification. Below is an example for a three-boards network connection. Make sure that there is no MAC address conflict on the boards and the IP addresses are set properly. Then, ping the test network.




Figure 101. Hardware setup for boundary clock verification

Run linuxptp on Board1 (boundary clock).

```
$ ptp4l -i eth0 -i eth1 -m
```

Run linuxptp on Board2/Board3 (ordinary clock).

```
$ ptp4l -i eth0 -m
```

On running the above command, time synchronization starts, and the slaves linuxptp selected automatically synchronizes to the unique master with synchronization messages displayed such as time offset, path delay and so on. For example,

ptp4l[878.504]:	master offset	-10	s2	freq	-2508	path delay	1826
ptp4l[878.629]:	master offset	-5	s2	freq	-2502	path delay	1826
ptp4l[878.754]:	master offset	0	s2	freq	-2495	path delay	1826
ptp4l[878.879]:	master offset	9	s2	freq	-2482	path delay	1826
ptp4l[879.004]:	master offset	-9	s2	freq	-2507	path delay	1826
ptp4l[879.129]:	master offset	-24	s2	freq	-2530	path delay	1826
ptp4l[879.255]:	master offset	-7	s2	freq	-2508	path delay	1826
ptp4l[879.380]:	master offset	-2	s2	freq	-2502	path delay	1826
ptp4l[879.505]:	master offset	-17	s2	freq	-2524	path delay	1827
ptp4l[879.630]:	master offset	6	s2	freq	-2493	path delay	1827
ptp4l[879.755]:	master offset	6	s2	freq	-2492	path delay	1827
ptp4l[879.880]:	master offset	0	s2	freq	-2500	path delay	1827

Some other options of ptp4l

```
Delay Mechanism
-E      E2E, delay request-response (default)
-P      P2P, peer delay mechanism
Network Transport
-2      IEEE 802.3
-4      UDP IPV4 (default)
-6      UDP IPV6
```

Note: You must keep same delay mechanism and use same network transport protocol on these boards.

Configure master mode

By default, the BMC (Best Master Clock) algorithm selects the master clock. To appoint a specific clock as master, a lower "priority1" attribute value than the other clock can be set. Lower value takes precedence. For example, in current case, specify one clock as master using the below option. (The other clock is using the default priority1 value 128.)

```
--priority1=127
```

One-step timestamping

Currently one-step timestamping is supported only on DPAA2. To use one-step timestamping, add the below option for ptp4l running.

```
--twoStepFlag=0
```

5.3.5.3 Transparent clock verification

At least three boards are needed. Below is an example for three boards network connection. Make sure there is no MAC address conflict on the boards, the IP addresses are set properly, and ping the test network.

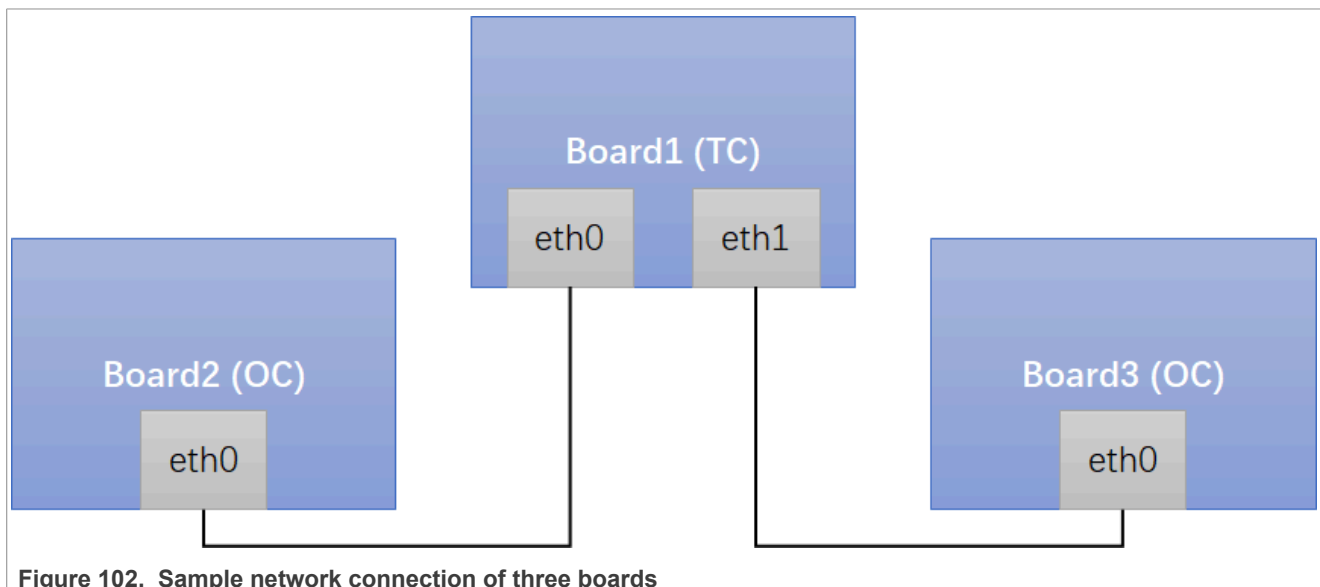


Figure 102. Sample network connection of three boards

Run `linuxptp` on Board1 (transparent clock). If you want Board1 works as E2E TC, use `E2E-TC.cfg`. If your want Board1 works as P2P TC, use `P2P-TC.cfg`.

```
$ ptp4l -i eth0 -i eth1 -f /etc/linuxptp/E2E-TC.cfg -m
```

Run `linuxptp` on Board2/Board3 (ordinary clock).

```
$ ptp4l -i eth0 -m -2
```

On running the above commands, time synchronization starts between ordinary clocks, and the slave `linuxptp` selected automatically synchronizes to the master with synchronization messages displayed such as time offset, path delay and so on.

5.3.6 Quick Start for IEEE 802.1AS

The following sections describe the steps for implementing IEEE 802.1AS on NXP boards. The following steps make use of `linuxptp` stack but similar commands can be executed with NXP GenAVB/TSN gPTP stack on supported boards, as described [here](#).

5.3.6.1 Time-aware end station verification

Connect two network interfaces in a back-to-back manner for two boards. Make sure that there is no MAC address conflict on the boards. Set the IP address properly and then ping the test network.

Remove the below option in the `/etc/linuxptp/gPTP.cfg` file to use by default the larger value, because the estimate path delay including PHY delay may exceed 800 ns since the hardware is using MAC timestamping.

```
neighborPropDelayThresh 800
```

Run `linuxptp` on each board. For example, `eth0` is used on each board.

```
$ ptp4l -i eth0 -f /etc/linuxptp/gPTP.cfg -m
```

On running the command above, time synchronization starts and the slave `linuxptp` selected automatically synchronizes to the master. It also displays synchronization messages such as time offset, path delay, and so on.

5.3.6.2 Time-aware bridge verification

At least three boards are needed for the time-aware bridge verification. Below is an example of the network connection amongst the three boards. Ensure that there is no MAC address conflict on the boards.




Figure 103. Setup for time-aware bridge verification

Remove the below option in `/etc/linuxptp/gPTP.cfg` file to use the default larger value, because estimated path delay including PHY delay may exceed 800 ns since hardware is using MAC timestamping.

```
neighborPropDelayThresh 800
```

Run linuxptp on Board1 (time-aware bridge) using the command below:

```
$ ptp4l -i eth0 -i eth1 -f /etc/linuxptp/gPTP.cfg -m
```

Run linuxptp on Board2/Board3 (time-aware end station) using the command:

```
$ ptp4l -i eth0 -f /etc/linuxptp/gPTP.cfg -m
```

Time synchronization will start between the three boards, and the linuxptp slaves selected will automatically synchronize to the unique master with synchronization messages displayed (such as time offset, path delay and so on).

5.3.7 Virtual clock use case

1. Connect the eth0 of two boards
2. Enable virtual clock on two boards. Ensure that the network interface eth0 uses ptp0. Create two virtual clocks for eth0:

```
$ echo 2 > /sys/class/ptp/ptp0/n_vclocks
```

3. Run linuxptp on the two boards:

```
$ ptp4l -i eth0 --phc_index 2 -f /etc/ptp4l_cfg/gPTP.cfg --domainNumber 2 -m
&
$ ptp4l -i eth0 --phc_index 3 -f /etc/ptp4l_cfg/gPTP.cfg --domainNumber 3 -m
&
```

5.3.8 Long term test

This section describes the long term test results for Linux PTP stack implementation.

5.3.8.1 Virtual clock use case

To run the virtual clock use case, follow the steps below:

1. Connect the eth0 of two boards.
2. Enable virtual clock on two boards
3. Ensure that the network interface eth0 uses ptp0. Then, create two virtual clocks for eth0 as shown below:

```
$ echo 2 > /sys/class/ptp/ptp0/n_vclocks
```

4. Run linuxptp on the two boards:

```
$ ptp4l -i eth0 --phc_index 2 -f /etc/ptp4l_cfg/gPTP.cfg --domainNumber 2 -m
&
$ ptp4l -i eth0 --phc_index 3 -f /etc/ptp4l_cfg/gPTP.cfg --domainNumber 3 -m
&
```

5.3.9 Known issues and limitations

1. When the LS1028A TSN switch in Linux is configured as an L2 switch, the interfaces should not be configured with IP addresses. Running linuxptp on these interfaces must use the Ethernet protocol instead of UDP/IP. The method is to add an option "-2" executing ptp4l command. For example:

```
$ ptp4l -i eth0 -2 -m
```

2. i.MX 8M Plus current dwmac driver (eth1) initializes some hardware functions during opening net device, including PTP initialization. Before that, the operations on it may not work, such as ethtool queries and PTP operations. So, the workaround is to do operations on the eth1 and PTP of dwmac only after running "ifconfig eth1 up".
3. If below error is reported during ptp4l running, just try to increase tx_timestamp_timeout. User space may need to wait longer for TX timestamp. For example, use option --tx_timestamp_timeout=20 while running ptp4l as shown below:

```
ptp4l[1560.726]: timed out while polling for tx timestamp
ptp4l[1560.726]: increasing tx_timestamp_timeout may correct this issue, but
it is likely caused by a driver bug
```

5.4 Networking

5.4.1 Q-in-Q on LS1028A Felix switch

1. Q-in-Q feature

Q-in-Q feature allows service providers to create a Layer 2 Ethernet connection between two user sites. Providers can segregate VLAN traffic of different users on a link or bundle using different user VLANs. When using Q-in-Q, the service VLAN tag (S-TAG: 0x88A8) prepend the 802.1Q VLAN tags (C-TAG: 0x8100) of the user.

2. Q-in-Q application scenario

In the following scenario, port **swp0** of the switch connects with Customer 1's LAN, **swp1** connects with the MAN of the ISP.

The traffic with VLAN tag is shown below:

uplink: Customer LAN (only C-TAG) -> swp0 -> swp1 (add S-TAG) -> ISP MAN (S-TAG + C-TAG)

downlink: ISP MAN (S-TAG + C-TAG) -> swp1 (pop S-TAG) -> swp0 (only C-TAG) -> Customer LAN




Figure 104. Q-in-Q application scenario

3. Q-in-Q configuration example

a. Enable swp1 Q-in-Q mode

```
devlink dev param set pci/0000:00:00.5 name qinq_port_bitmap value 2 cmode runtime
```

Note:

- 0000:00:00.5 is the PCIe bus and device number of ocelot switch.
- The value 2 is the bitmap for port 1. If port n is linked to ISP MAN, the related bit 'n' should be set to 1.

b. Create bridge and add ports:

```
ip link add dev br0 type bridge vlan_protocol 802.1ad
ip link set dev swp0 master br0
ip link set dev swp1 master br0
ip link set dev br0 type bridge vlan_filtering 1
```

c. Set swp0 pvid and set untagged for egress traffic:

```
bridge vlan del dev swp0 vid 1 pvid
bridge vlan add dev swp0 vid 100 pvid untagged
bridge vlan add dev swp1 vid 100
```

d. The result is displayed below:

```
Customer(tpid:8100 vid:111) -> swp0 -> swp1 -> ISP(STAG tpid:88A8 vid:100,
CTAG tpid:8100 vid:111)
ISP(tpid:88A8 vid:100 tpid:8100 vid:222) -> swp1 -> swp0 ->
Customer(tpid:8100 vid:222)
```


5.4.2 VCAP on LS1028A Felix switch

The VCAP is a content-aware packet processor for wire-speed packet inspection. It uses the 'tc flower' command to set the filter and actions. The following keys and actions are supported on LS1028A:

keys:

vlan_id
 vlan_prio
 dst_mac/src_mac for non IP frames
 dst_ip/src_ip
 dst_port/src_port

actions:

trap
 drop
 police
 vlan modify
 vlan push(Egress)

Use the following commands to set, get, and delete VCAP rules:

```
tc qdisc add dev swp0 clsact
tc filter add dev swp0 ingress chain [chain-id] protocol [ip/802.1Q] flower
  skip_sw [keys] action [actions]
tc -s filter show dev swp0 ingress chain [chain-id]
tc filter del dev swp0 ingress chain [chain-id] pref [pref_id]
tc qdisc add dev swp1 clsact
tc filter add dev swp1 egress protocol 802.1Q flower skip_sw [keys] action vlan
  push id [value] priority [value]
tc filter show dev swp1 egress
tc filter del dev swp1 egress pref [pref_id]
```

There are two ingress VCAPs and one egress VCAPs. The tc-flower chains are used on LS1028A ingress ports. Each action has a fixed chain. [Table 89](#) shows the chain allocation:

Table 89. Chain allocation

chain ID	Actions	Hardware module	keys
10000	skbedit priority	IS1 lookup 0	Source MAC address, source IP address (32 bits) outer VLAN, IP protocol, source TCP/UDP ports.
11000	vlan pop; vlan modify	IS1 lookup 1	Inner and outer VLAN, source and destination IP addresses (32 bits), IP protocol, source and destination TCP/UDP ports.
12000	goto chain [PAG]	IS1 lookup 2	Source MAC address, source IP address (32 bits) outer VLAN, IP protocol, source TCP/UDP ports.
20000-20255	police; trap	IS2 lookup 0	Source and destination MAC address, source and destination IP addresses (32 bits), IP protocol, source and destination TCP/UDP ports.
21000-21255	drop; redirect	IS2 lookup 1	Source and destination MAC address,

Table 89. Chain allocation...continued

chain ID	Actions	Hardware module	keys
			source and destination IP addresses (32 bits), IP protocol, source and destination TCP/UDP ports.
30000	gate; police	PSFP	destination MAC address and Vlan ID

Before using chains, users should register each chain and set chain pipeline order for a packet. The hardware ingress order is: **IS1->IS2->PSFP**.

```
tc qdisc add dev swp0 clsact
tc filter add dev swp0 ingress chain 0 pref 49152 flower skip_sw action goto chain 10000
tc filter add dev swp0 ingress chain 10000 pref 49152 flower skip_sw action goto chain 11000
tc filter add dev swp0 ingress chain 11000 pref 49152 flower skip_sw action goto chain 12000
tc filter add dev swp0 ingress chain 12000 pref 49152 flower skip_sw action goto chain 20000
tc filter add dev swp0 ingress chain 20000 pref 49152 flower skip_sw action goto chain 21000
tc filter add dev swp0 ingress chain 21000 pref 49152 flower skip_sw action goto chain 30000
```

After registering the chain, add rules to the corresponding chain. Following are the use cases for testing:




Figure 105. VCAP test

1. Drop all frames from source IP 192.168.2.1.

```
tc filter add dev swp0 ingress chain 21000 protocol ip flower skip_sw src_ip
192.168.2.1 action drop
```

Set source IP as 192.168.2.1 and send IP package from TestCenter, package will be dropped on swp0.

2. Limit bandwidth of HTTP streams to 10 Mbps.

```
tc filter add dev swp0 ingress chain 20000 protocol ip flower skip_sw
ip_proto tcp dst_port 80 action police rate 10Mbit burst 10000 conform-
exceed drop/pipe action goto chain 21000
```

Send TCP package and set destination port as 80 on TestCenter, set the stream bandwidth to 1Gbit/s, we can get a 10Mbps/s stream rate.

3. Filter frames that have a specific vlan tag (VID=1 and PCP=1). Then, modify the vlan tag (VID=2, PCP=2) and classified to QoS traffic class 2.

```
ip link set switch type bridge vlan_filtering 1
tc filter add dev swp0 ingress chain 11000 protocol 802.1Q flower skip_sw
vlan_id 1 vlan_prio 1 action vlan modify id 2 priority 2 action goto chain
12000
bridge vlan add dev swp0 vid 2
bridge vlan add dev swp1 vid 2
```

Set vid=1 and pcp=1 in vlan tag. Then, send IP package from TestCenter. Thus you can get a package with vid=2, pcp=2 from swp1 on TestCenter.

4. Push a specific vlan tag (vid=3, pcp=3) into frames (classified vid=2, pcp=2 in switch) egress from swp1.

```
tc qdisc add dev swp1 clsact
tc filter add dev swp1 egress protocol 802.1Q flower skip_sw vlan_id 2
vlan_prio 2 action vlan push id 3 priority 3
```

Set vid=1 and pcp=1 in vlan tag, then send IP package from TestCenter, the frame will hit rule in usecase 3 and retag the vlan (vid=2, pcp=2). Thus, users can get a frame with vid=3, pcp=3 from swp1 on TestCenter.

5. Push double vlan tag(Q-in-Q) into frames egress to swp1.

```
ip link add dev br0 type bridge
ip link set dev swp0 master br0
ip link set dev swp1 master br0
ip link set br0 type bridge vlan_filtering 1
bridge vlan add dev swp0 vid 222
bridge vlan add dev swp1 vid 222
tc qdisc add dev swp1 clsact
tc filter add dev swp1 egress protocol 802.1Q flower skip_sw \
vlan_id 222 vlan_prio 2 \
action vlan push id 200 priority 1 protocol 802.1AD \
action vlan push id 300 priority 3
```

Result: TX(tpid:8100 vid:222 pri:2) -> swp0 -> swp1 -> RX(S-TAG tpid:88A8 vid:200 pri:1, C-TAG tpid:8100 vid:300 pri:3)

6. Pop single or double vlan tag(Q-in-Q) from frames ingress from swp0.

```
ip link add dev br0 type bridge
ip link set dev swp0 master br0
ip link set dev swp1 master br0
tc filter add dev swp0 ingress chain 11000 \
protocol 802.1ad flower \
vlan_id 111 vlan_prio 1 vlan_ethertype 802.1q \
cvlan_id 222 cvlan_prio 2 cvlan_ethertype ipv4 \
action vlan pop action goto chain 12000
```

Result: TX(S-TAG tpid:88A8 vid:111 pri:1, C-TAG tpid:8100 vid:222 pri:2) -> swp0
-> swp1 -> RX(TAG tpid:8100 vid:222 pri:2)

```
tc filter add dev swp0 ingress chain 11000 \  
  protocol 802.1ad flower \  
  vlan_id 111 vlan_prio 1 vlan_ethtype 802.1q \  
  cvlan_id 223 cvlan_prio 2 cvlan_ethtype ipv4 \  
  action vlan pop \  
  action vlan pop action goto chain 12000
```

Result: TX(S-TAG tpid:88A8 vid:111 pri:1, C-TAG tpid:8100 vid:223 pri:2) -> swp0
-> swp1 -> RX(received packets without VLAN tag)

6 Protocols

This section describes the protocols supported.

6.1 EtherCAT MainDevice

Real-time Edge supports the usage of EtherCAT (Ethernet for Control Automation Technology) MainDevice and integrates the IGH EtherCAT MainDevice stack and SOEM. EtherCAT is verified on NXP platforms.

6.1.1 Introduction

EtherCAT is an Ethernet-based fieldbus system, invented by BECKHOFF Automation. The protocol is standardized in IEC 61158 and is suitable for both hard and soft real-time computing requirements in automation technology. EtherCAT was developed with an objective to apply Ethernet for automation applications requiring short data update times, low communication jitter, and reduced hardware costs.

Note:

- *Data update time is also called cycle time.*
- *Low communication jitter implies less than 100 μ s and for precise synchronization purposes, it is less than 1 μ s.*
- EtherCAT is Fast: 1000 digital I/O: 30 μ s, 100 SubDevice: 100 μ s.
- EtherCAT is Ethernet: Standard Ethernet at I/O level.
- EtherCAT is Flexible: Star, line, drop, with or without switch.
- EtherCAT is Inexpensive: Ethernet is mainstream technology, and therefore inexpensive.
- EtherCAT is Easy: everybody knows Ethernet and it is simple to use.

Currently, there are two open source EtherCAT MainDevices that Real-time Edge software supports. These are IGH for Cortex-A core and SOEM for Cortex-M core. Both IGH and SOEM are solutions that help users get rid of the low-level development directly on the EtherCAT protocol, and provide a common API for real-time applications. For more information, see <https://rt-labs.com/ethercat/> and <http://www.etherlab.org>.

IGH shows superior real-time characteristics since it is integrated into the Preempt-RT Linux kernel and uses native Ethernet drivers (only supported on a few specific NXP platforms). Also, IGH is more powerful and integrated by providing users with auxiliary functions such that it encourages users to focus on specific high-level programs. For example, IGH runs with an automatic MainDevice FSM (Finite State Machine). FSM is a EtherCAT SubDevice manager dynamically adapting the SubDevice configurations to the new topology and responding to requests from the application-layer. In addition, IGH provides a command-line tool in user space to display detailed information about MainDevice/SubDevice configurations and to list SDO dictionaries and reading/writing addresses.

SOEM (Simple Open EtherCAT MainDevice) is a C library that offers methods to send and receive EtherCAT frames. It is very light-weight but still powerful and stable. It can run on multiple operating systems, for example, Linux, Windows, or FreeRTOS. It can even run without an operating system (such as 'BareMetal'). Unlike IGH, users must configure PDOs and SDOs aligned with determined memory addresses in their own applications, which might cause uncertainties during development. On Real-time Edge software, SOEM is only supported on the Cortex-M core on i.MX 8M Mini LPDDR4 EVK and i.MX 8M Plus LPDDR4 EVK platforms and runs on FreeRTOS or BareMetal.

6.1.2 EtherCAT protocol

Following are the characteristics of the EtherCAT protocol:

- The EtherCAT protocol is optimized for process data and is transported directly within the standard IEEE 802.3 Ethernet frame using Ethertype 0x88a4.

- The data sequence is independent of the physical order of the nodes in the network; addressing can be in any order.
- Broadcast, multicast, and communication between SubDevice is possible, but must be initiated by the MainDevice.
- If IP routing is required, the EtherCAT protocol can be inserted into UDP/IP datagrams. This also enables any control with Ethernet protocol stack to address EtherCAT systems.
- It does not support shortened frames.

The [Figure 106](#) shows the EtherCAT frame structure.

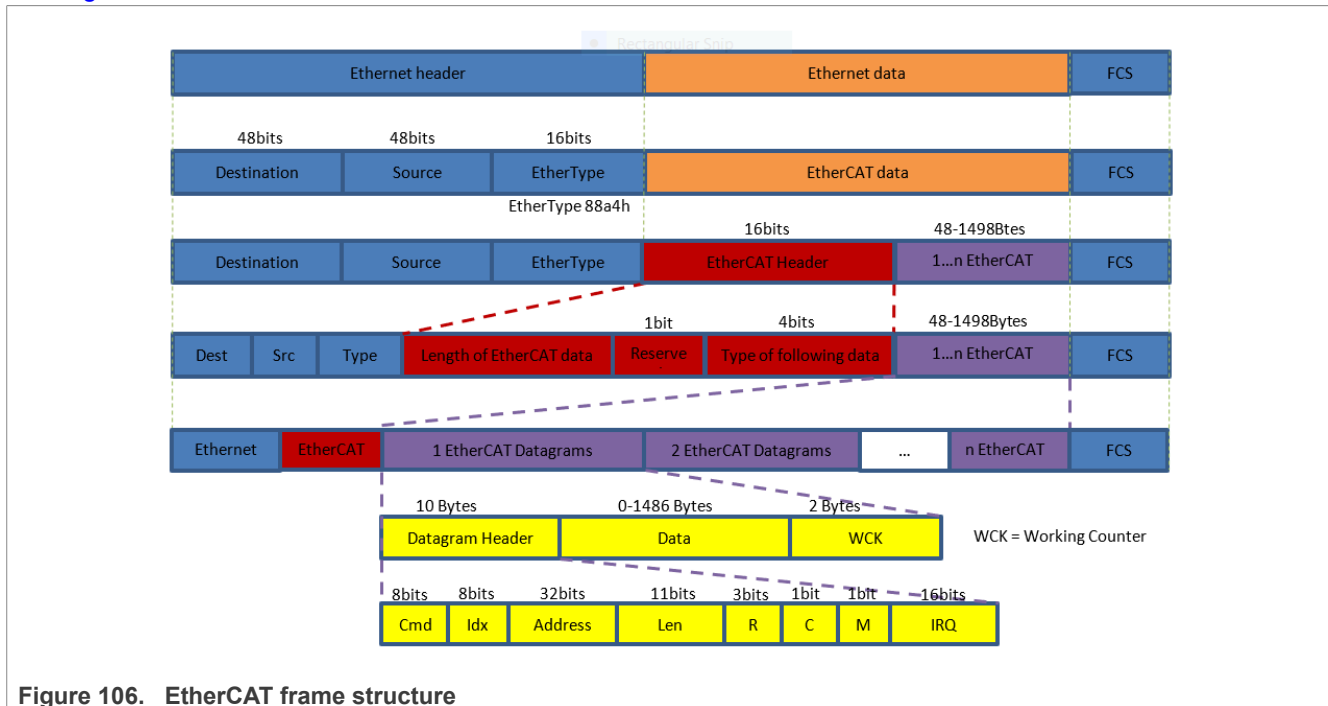


Figure 106. EtherCAT frame structure

6.1.3 IGH EtherCAT architecture

The components of the MainDevice environment are described below:

- **MainDevice module:** This is the kernel module containing one or more EtherCAT MainDevice instances, the 'Device Interface' and the 'Application Interface'.
- **Device modules:** These are EtherCAT-capable Ethernet device driver modules that offer their devices to the EtherCAT MainDevice via the device interface. These modified network drivers can handle network devices used for EtherCAT operation and 'normal' Ethernet devices in parallel. A MainDevice can accept a certain device and then, is able to send and receive EtherCAT frames. Ethernet devices declined by the MainDevice module are connected to the kernel's network stack, as usual.
- **Application:** A program that uses the EtherCAT MainDevice (usually for cyclic exchange of process data with EtherCAT SubDevice). These programs are not part of the EtherCAT MainDevice code, but require to be generated or written by the user. An application can request a MainDevice through the application interface. If this succeeds, it has the control over the MainDevice: It can provide a bus configuration and exchange process data. Applications can be kernel modules that use the kernel application interface directly. They also include user space programs, which use the application interface via the EtherCAT library or the RTDM library. The [Figure 107](#) shows the IGH EtherCAT MainDevice architecture.




Figure 107. IGH EtherCAT MainDevice architecture

6.1.3.1 IGH EtherCAT Device Drivers

The EtherCAT protocol is based on the Ethernet standard, so a MainDevice relies on standard Ethernet hardware to communicate with the bus. The term 'device' is used as a synonym for Ethernet network interface hardware. There are two kinds of device drivers modules:

1. Native Ethernet Device Drivers

Native Ethernet Device Drivers allow the EtherCAT MainDevice direct and exclusive access to the Ethernet hardware. This implies that the network device must not be connected to the kernel's stack as usual, which allows a high Real-time performance. In Real-time Edge software, there are three native Ethernet drivers:

- `ec_fec`: the native driver `ec_fec` can be used for the FEC MAC on i.MX 8DXL LPDDR4 EVK, i.MX 8M Mini LPDDR4 EVK, i.MX 8M Plus LPDDR4 EVK, and i.MX 93 EVK. This driver is not verified on other i.MX

platforms in this release. Please note that the original Ethernet fec driver must be recompiled to a module by reconfiguring Linux with command "make menuconfig" when using the `ec_fec` native driver.

- `ec_enetc`: the native driver "ec_enetc" is used for ENETC MAC on the LS1028ARDB platform.
- `ec_dpaa1`: the native driver "ec_dpaa1" is used for DPAA1 MAC on the LS1043ARDB and LS1046ARDB.

2. Generic Ethernet Device Driver

The Generic driver uses the lower layers of the Linux network stack to connect to the hardware, independently of the actual hardware driver. Therefore, it can be used by all the hardware platforms that Real-time Edge supports. However, the performance by using a generic Ethernet Device driver is a bit worse than the native driver, because the Ethernet frame data has to traverse the Linux network stack.

6.1.3.2 IGH EtherCAT setup

Before the IGH EtherCAT daemon starts, the Ethernet device and the Ethernet driver must be specified by setting the "MASTER0_DEVICE" and "DEVICE_MODULES" variables in the `/etc/ethercat.conf` file.

6.1.3.2.1 Specifying the Ethernet device

The Ethernet device is specified by setting `MASTER0_DEVICE` variable to the MAC address of the Ethernet device to use as the EtherCAT network interface as below:

```
MASTER0_DEVICE="00:04:9f:07:11:a6"
```

For LS1046ARDB or LS1043ARDB platforms, if multiple MainDevice are required, adding a non-empty variable `MASTER1_DEVICE` creates a second master, and so on.

6.1.3.2.2 Generic Ethernet driver

The generic Ethernet driver is enabled on Real-time Edge images by default for all platforms. It can be specified by setting "DEVICE_MODULES" variable to "generic" as shown below in the `/etc/ethercat.conf` file.

```
DEVICE_MODULES="generic"
```

6.1.3.2.3 Native Ethernet driver for i.MX 8M Mini LPDDR4 EVK, i.MX 8M Plus LPDDR4 EVK and i.MX 93 EVK

The native Ethernet driver "ec_fec" is enabled on Real-time Edge images by default for i.MX 8M Mini LPDDR4 EVK, i.MX 8M Plus LPDDR4 EVK and i.MX 93 EVK. And it can be specified by setting "DEVICE_MODULES" variable to "fec" as below on `/etc/ethercat.conf` file.

```
DEVICE_MODULES="fec"
```

But note that the original Ethernet fec driver must be compiled to modules by reconfiguring Linux menuconfig when the native Ethernet driver "ec_fec" is enabled. On Real-time Edge, the original fec Ethernet driver has been configured as a module by default.

6.1.3.2.4 Native Ethernet driver for LS1028ARDB

The native Ethernet driver `ec_enetc` is enabled on Real-time Edge images by default for LS1028ARDB. It can be specified by setting `DEVICE_MODULES` variable to `enetc` as shown in the below `/etc/ethercat.conf` file.

```
DEVICE_MODULES="enetc"
```

6.1.3.2.5 Native Ethernet driver for LS1043ARDB and LS1046ARDB

The native Ethernet driver '`ec_dpaa1`' is enabled by default for LS1046ARDB and LS1043ARDB platforms. There are up to 7 Ethernet ports on LS1046ARDB or LS1043ARDB. Hence, the multiple MainDevice and redundancy features can be supported on LS1046ARDB and LS1043ARDB platforms.

DPAA is a network co-processor that is more complex than ordinary network cards and is different from other native drivers mentioned above, . So the native driver "`ec_dpaa1`" must be configured by "`ethercat_port`" variable in U-Boot to notify the DPAA co-processor how to schedule data frames before EtherCAT stack starting. The format of "`ethercat_port`" variable is as below:

```
ethercat_port=master<x>_device, master<x>_backup, core_index;
```

The `master<x>_device` variable specifies the main Ethernet port for MainDevice with index 'x', while the `master<x>_backup` variable specifies the backup Ethernet port if redundancy is required. Every MainDevice must be bundled a specified CPU core. and `core_index` variable is used to specific which CPU core to be bundled for this MainDevice.

The relationship between Ethernet port name on LS1046ARDB chassis and the name in the Linux is in the [Table 90](#):

Table 90. LS1046ARDB chassis

Port name on chassis	Port name in Linux
RGMII1	fm1-mac3
RGMII2	fm1-mac4
SGMII1	fm1-mac5
SGMII2	fm1-mac6
10G Copper	fm1-mac9
10G SEP+	fm1-mac10




Figure 108. LS1046ARDB Ethernet chassis

The relationship between Ethernet port name on LS1043ARDB chassis and the name in the Linux is shown in [Table 91](#).

Table 91. LS1043ARDB chassis

Port name on chassis	Port name in Linux
QSGMII.P0	fm1-mac1
QSGMII.P1	fm1-mac2
QSGMII.P2	fm1-mac5
QSGMII.P3	fm1-mac6
RGMII1	fm1-mac3
RGMII2	fm1-mac4
10G Copper	fm1-mac9




Figure 109. LS1043ARDB Ethernet chassis

There are 3 instances of `ec_dpaa1` configuration for LS1046ARDB.

1. There is only one network topology, and no backup port for redundancy. The Ethernet port of this network is RGMII1, and the MainDevice is bundled to CPU core 3. For this case, the `ethernetat_port` variable is as below:

```
setenv ethernetat_port "fm1-mac3,,3;"
```

2. There are two EtherCAT network topologies, and each network topology has a backup port for redundancy. The main port of the first network is RGMII1, the backup port is RGMII2, and the MainDevice of this network is bundled to CPU core 3. The main port of the other network is SRGMII1, the backup port is SRGMII2, and the MainDevice of this network is bundled to CPU core 2. For this case, the `ethernetat_port` variable is as below:

```
setenv ethernetat_port "fm1-mac3, fm1-mac4, 3; fm1-mac5, fm1-mac6, 2;"
```

3. There are two EtherCAT network topologies, and only the first network topology has a backup port for redundancy. The main port of the first network is RGMII1, the backup port is RGMII2, and the MainDevice of this network is bundled to CPU core 3. The main port of the other network is SRGMII1, and the MainDevice of this network is bundled to CPU core 2. For this case, the `ethernetat_port` variable is as below:

```
setenv ethernetat_port "fm1-mac3, fm1-mac4, 3; fm1-mac5,, 2;"
```

As mentioned above, each MainDevice is bundled to a specific CPU core. So it is recommended to bundle the corresponding EtherCAT real-time application to the same core with the MainDevice.

6.1.3.2.6 Starting IGH EtherCAT

Check the mac address under uboot:

```
u-boot=> print ethaddr
```

Bind dtb file under uboot:

```
# for i.MX 6ULL EVK
u-boot=> setenv fdt_file imx6ull-14x14-evk-igh.dtb

# for i.MX 8M Mini LPDDR4 EVK
u-boot=> setenv fdtfile imx8mm-evk-igh.dtb

# for i.MX 8M Plus LPDDR4 EVK
u-boot=> setenv fdtfile imx8mp-evk-igh.dtb

# for i.MX 93 EVK
u-boot=> setenv fdtfile imx93-11x11-evk-igh.dtb

u-boot=> run bootcmd
```

After filling in the mac address found in uboot into “MASTER0_DEVICE” in file “/etc/ethercat.conf”, use the below command to start the IGH EtherCAT daemon:

```
$ ethercatctl start
```

Also, the below commands are used to stop or restart it.

```
# ethercatctl stop
# ethercatctl restart
```

Note: If the generic driver is used, make sure using “ifconfig <ethX> up” command to enable the network Card.

IGH provides a powerful auxiliary command-line tool, named “ethercat”. it is can be used to query the MainDevice and all SubDevice information and status. The usage is as below:

```
Usage: ethercat <COMMAND> [OPTIONS] [ARGUMENTS]
Commands (can be abbreviated):
alias      Write alias addresses.
config     Show slave configurations.
crc        CRC error register diagnosis.
cstruct    Generate slave PDO information in C language.
data       Output binary domain process data.
debug      Set the master's debug level.
domains    Show configured domains.
download   Write an SDO entry to a slave.
eoe        Display Ethernet over EtherCAT statistics.
foe_read   Read a file from a slave via FoE.
foe_write  Store a file on a slave via FoE.
graph      Output the bus topology as a graph.
master     Show master and Ethernet device information.
pdos       List Sync managers, PDO assignment and mapping.
reg_read   Output a slave's register contents.
reg_write  Write data to a slave's registers.
rescan     Rescan the bus.
sdos       List SDO dictionaries.
sii_read   Output a slave's SII contents.
sii_write  Write SII contents to a slave.
slaves     Display slaves on the bus.
soe_read   Read an SoE IDN from a slave.
soe_write  Write an SoE IDN to a slave.
states     Request application-layer states.
```

upload	Read an SDO entry from a slave.
version	Show version information.
xml	Generate slave information XML.

Real-time Edge also provides a systemd service to run IGH EtherCAT daemon as a system service.

```
# systemctl enable ethercat
# systemctl start ethercat
```

The below commands are used to stop or disable this service:

```
# systemctl stop ethercat
# systemctl disable ethercat
```

6.1.3.3 IGH EtherCAT User Space setup

IGH EtherCAT User Space supports i.MX 8M Mini LPDDR4 EVK, i.MX 8M Plus LPDDR4 EVK, and i.MX 93 EVK boards.

6.1.3.3.1 IGH EtherCAT User Space Introduction

IGH EtherCAT stack has both code in the kernel space and user space. In kernel space, it integrates multiple drivers for different networking devices. In the user space, it supplies a library (`libetherat.so/.a`) for user's application. Users can set/get the running parameters of the device using API of this library. The API reads or writes data via ioctl mechanism of Linux kernel.

ioctl invokes the routine implemented in MainDevice module (`ec_master.ko`). Then MainDevice module then calls the RX/TX function in different drivers. ioctl is a system call supplied by Linux. The updates for running parameters triggers system call of the kernel, though this introduces more system jitter.

On the other hand, there are mainly two important threads for using IGH EtherCAT stack. One is MainDevice operation thread running in kernel space, which is responsible for parsing the frames received from EtherCAT devices. The other thread is RT thread implemented by user, which triggers RX/TX functionality cyclically according to cycle time. Due to these two threads are running in different spaces, a context switch between user space and kernel space is required, which also increases system jitter and time latency.

The user space IGH EtherCAT stack runs wholly in user space. Therefore, it doesn't need ioctl. User can invoke the functions of MainDevice module directly, this can avoid system call, it can decrease jitter of system obviously.

From the perspective of threads, in user space stack, MainDevice operation thread is also running in user space, so we can refactor OP thread to one callback function, this callback function is called in the user's cyclic RT thread. In such a case, there is only one main thread in whole system, which avoids context switching. This also improves the performance, such as jitter and latency and other such parameters.

User space IGH EtherCAT stack has no code that runs in the kernel space, which makes it entirely independent of the kernel version.

6.1.3.3.2 Configure EtherCAT User Space dtb

The EtherCAT User Space dtb can be reconfigured by adding the below lines in the "`source/meta-real-time-edge/conf/distro/include/real-time-edge-base.inc`" file.

```
# dtb for EtherCAT userspace application
KERNEL_DEVICETREE:append:imx8mm-lpddr4-evk = " freescale/imx8mm-evk-ecat-
userspace.dtb"
```

```
KERNEL_DEVICETREE:append:imx8mp-lpddr4-evk = " freescale/imx8mp-evk-ecat-
userspace.dtb"
KERNEL_DEVICETREE:append:imx93evk = " freescale/imx93-11x11-evk-ecat-
userspace.dtb"
```

6.1.3.3.3 Starting IGH EtherCAT User Space

Use the steps below to start the IGH EtherCAT User Space.

1. Execute the following command to update the dtb file and boot into the kernel.

```
# on i.MX 8M Mini
uboot => setenv fdtfile imx8mm-evk-ecat-userspace.dtb
# on i.MX 8M Plus
uboot => setenv fdtfile imx8mp-evk-ecat-userspace.dtb
# on i.MX 93
uboot => setenv fdtfile imx93-11x11-evk-ecat-userspace.dtb
u-boot=> run bootcmd
```

2. The EtherCAT tool (named `ethercat_userspace`) invokes the `ioctl` command to scan and display EtherCAT devices. For the convenience of using the EtherCAT tool, one virtual `ioctl` interface is implemented to support this tool. One backend program (named `ethercat_master`) supplies the service for the tool. Firstly, run the backend program:

```
# mkdir -p /dev/hugepages
# mount -t hugetlbfs hugetlbfs /dev/hugepages
# echo 64 > /sys/kernel/mm/hugepages/hugepages-2048kB/nr_hugepages
# ethercat_master &
```

Then, run EtherCAT tool:

```
# ethercat_userspace slaves
```

3. Now, exit `ethercat_master` program before running the motor control example:

```
# fg
# Type 'ctrl + c' to exit IGH
// where "IGH" --> "ethercat_master program"
# cd /examples/igh-ethercat-userspace/
```

Note: If want to run the user application, exit and close `ethercat_master` program.

- For example, see the below `ec_motor_example` to control motor rotation:

```
./ec_motor_example 131072 1
In the above command, 131072 is the encoding rate of the motor, and 1 is the
number of motors.
```

Note: For better performance, user can isolate the working core from Linux, using `isolcpus=1` for U-Boot bootargs, user space IGH EtherCAT works on Core1 by default.

Note: About how to use user space IGH EtherCAT library, user can refer to the readme file of the below compile path:

```
{
$build_dir}/build-imx8mpevk-real-time-edge/tmp/work/imx8mp_lpddr4-evk-poky-
linux/igh-ethercat-userspace/1.5.2-r0/git/Readme_For_User_Space_IGH_EtherCAT.txt
```

6.1.3.4 The 'real-time-edge-servo' stack

The **real-time-edge-servo** is a CiA402 (also referred to as DS402) profile framework based on IgH CoE interface. (For details of an EtherCAT MainDevice stack, see [Section 6.1](#)). It abstracts the CiA 402 profile and provides an easily-usable API for the Application developer.

The real-time-edge-servo project consists of a basic library *libnservo* and several auxiliary tools.

The application developed with *libnservo* is flexible enough to adapt to the changing of CoE network by modifying the *xml* config file, which is loaded when the application starts. The *xml* config file describes the necessary information, which includes EtherCAT network topology, SubDevice and MainDevice configurations, and definitions of all the axles.

The stack has been tested on below CoE servo production: DELTA ASDA-B3, HCFA SV-X6EB, SV-X3EB, Just Motion Control 2HSS458-EC, and INOVANCE InoSV680N.

6.1.3.4.1 CoE network

The [Figure 110](#) illustrates a typical CoE network.

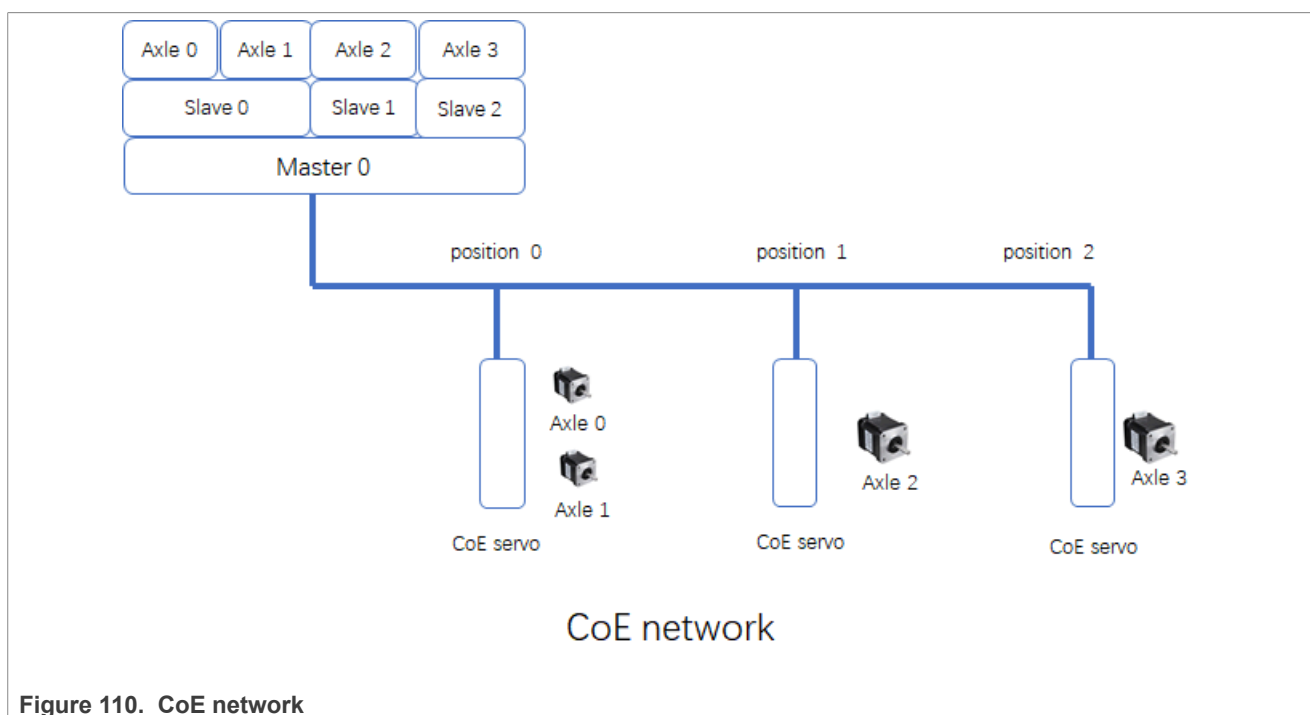


Figure 110. CoE network

There are three CoE servos on this network and we name them SubDevice *x* as the position they are. Each CoE servo can have more than one axle. The *libnservo* initiates the CoE network and encapsulates the details of network topology into axle nodes. Therefore, the developer can focus on the each axle operation without taking care of the network topology.

6.1.3.4.2 Libnservo architecture

real-time-edge-servo runs on top of the *Igh* EtherCAT stack. The *Igh* stack provides CoE communication mechanisms - Mailbox and Process Data. Using these mechanisms, real-time-edge-servo can access the CiA Object Dictionary located on CoE servo. The [Figure 111](#) shows the Libservo architecture.

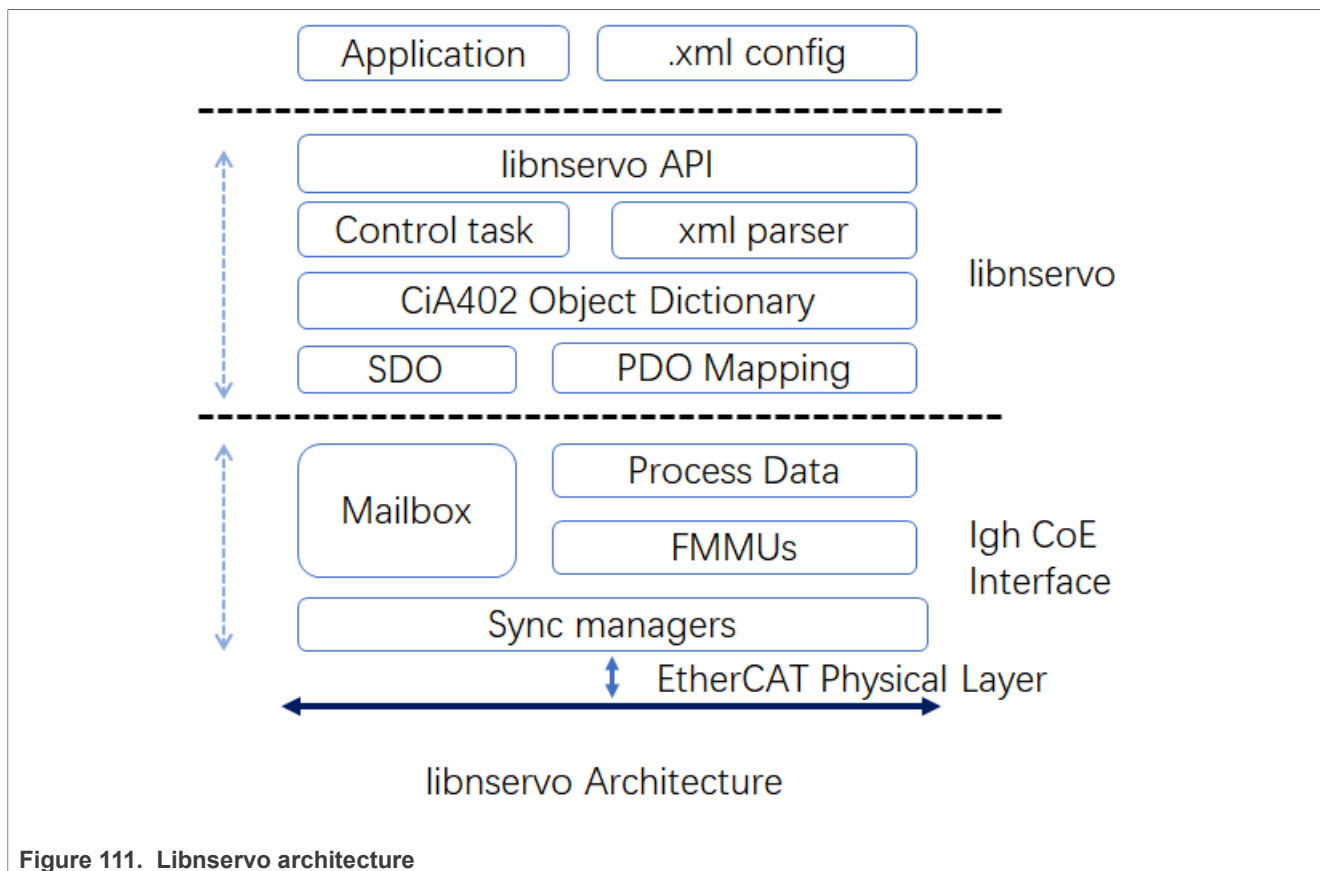


Figure 111. Libnservo architecture

The control task initiates the MainDevice and all SubDevice on the CoE network and registers all PDOs to Igh stack. It then constructs a data structure to describe each axle. Finally, the control task creates a task to run the user task periodically.

6.1.3.4.3 Xml configuration for real-time-edge-servo

This section describes how the xml config file that describes a CoE network should be configured.

The basic framework of the XML configuration is shown in the code below:

```
<?xml version="1.0" encoding="utf-8"?>
<Config Version="1.2">
  <PeriodTime>#100000000</PeriodTime>
  <MaxSafeStack>#8192</MaxSafeStack>
  <master_status_update_freq>#1</master_status_update_freq>
  <slave_status_update_freq>#1</slave_status_update_freq>
  <axle_status_update_freq>#1</axle_status_update_freq>
  <sync_ref_update_freq>#2</sync_ref_update_freq>
  <sched_priority>#90</sched_priority>
  <sched_policy>#SCHED_FIFO</sched_policy>
  <Masters>
    <Master>
      ...
    <\Master>
  <Master>
    ...
  <\Master>
  <Master>
    ...
  <\Master>
  <Master>
```



```

<Axles>
  <Axle>
    ...
  <\Axle>
  <Axle>
    ...
  <\Axle>
<\Axles>
</Config>

```

- All config elements must be within the <Config> element.
- All config elements shown above are mandatory.
- The numerical value started with # indicates that it is a decimal value.
- The numerical value started with #x indicates that it is a hexadecimal value.
- <PeriodTime> element indicates that the period of control task is 10 ms.
- <MaxSafeStack> indicates the stack size, and it is an estimated value. 8K is sufficient to meet the needs of most applications.
- <master_status_update_freq> element indicates the frequency of MainDevice status update. the value #1 means update the MainDevice status every task period.
- <slave_status_update_freq> element indicates the frequency of SubDevice status update. the value #1 signifies to update the SubDevice status every task period.
- <axle_status_update_freq> element indicates the frequency of axles status update. the value #1 signifies to update the axles status every task periods.
- <sync_ref_update_freq> element indicates the frequency of reference clock update. the value #2 signifies to update the axles status every two task periods.
- <sched_policy> element specifies which schedule policy is used for a user task.
- <sched_priority> element indicates the priority of the user task.
- <Masters> element can contain more than one MainDevice element. For most cases, there is only one MainDevice on a host.
- <Axles> element can contain more than one Axle element, which is an important feature for the developers.

6.1.3.4.3.1 MainDevice element

As shown in the [Section 6.1.3.4.1](#), the MainDevice can have many SubDevice, so the MainDevice element might consist of few *SubDevice* elements.

```

<Master>
  <Master_index>#0</Master_index>
  <Reference_clock>#0</Reference_clock>
  <Slave alias="#0" slave_position="#0">
    ....
  </Slave>
  <Slave alias="#1" slave_position="#1">
    ....
  </Slave>
</Master>

```

- <Master_index> element indicates the index of the MainDevice. As mentioned above, for many cases, there is only one MainDevice, so the value of this element is always #0.
- <Reference_clock> element indicates the SubDevice that is used the reference clock.
- <Slave> element indicates that there is a SubDevice on this MainDevice.

SubDevice element

```
<Slave alias="#0" slave_position="#0">
<VendorId>#x66668888</VendorId>
<ProductCode>#x20181302</ProductCode>
  <Name>2HSS458-EC</Name>
  <Emerg_size>#x08</Emerg_size>
<WatchDog>
  <Divider>#x0</Divider>
  <Intervals>#4000</Intervals>
</WatchDog>
<DC>
  <SYNC SubIndex='#0'>
    <Shift>#0</Shift>
  </SYNC>
</DC>
<SyncManagers force_pdo_assign="#1">
  <SyncManager SubIndex="#0">
    ...
  </SyncManager>
  <SyncManager SubIndex="#1">
    ...
  </SyncManager>
</SyncManagers>
<Sdos>
  <Sdo>
    ...
  </Sdo>
  <Sdo>
    ...
  </Sdo>
</Sdos>
</Slave>
```

- *alias* attribute means the alias name of this SubDevice.
- *slave_position* attribute means which position of the SubDevice is on this network.
- <Name>element is the name of the SubDevice.
- <Emerg_size> element is always 8 for all CoE device.
- <WatchDog> element is used to set the watch dog of this SubDevice.
- <DC> element is used to set the sync info.
- <SyncManagers> element should contain all syncManager channels.
- <Sdos> element contains the default value we want to initiate by SDO channel.

SyncManagers Element

For a CoE device, there are generally four syncManager channels.

- SM0: Mailbox output
- SM1: Mailbox input
- SM2: Process data outputs
- SM3: Process data inputs

```
<SyncManager SubIndex="#2">
  <Index>#x1c12</Index>
  <Name>Sync Manager 2</Name>
  <Dir>OUTPUT</Dir>
```

```

<Watchdog>ENABLE</Watchdog>
<PdoNum>#1</PdoNum>
<Pdo SubIndex="#1">
  <Index>#x1600</Index>
  <Name>RxPdo_1</Name>
  <Entry SubIndex="#1">
    ...
  </Entry>
  <Entry SubIndex="#2">
    ...
  </Entry>
</Pdo>
</SyncManager>

```

- <Index> element is the object address.
- <Name> is a name of this syncmanager channel.
- <Dir> element is the direction of this syncmanager channel.
- <Watchdog> is used to set watchdog of this syncmanager channel.
- <PdoNum> element means how many PDO we want to set.
- <Pdo SubIndex="#1"> element contains the object dictionary entry that is required to be mapped to.
 - <Index> PDO address.
 - <Name> PDO name
 - <Entry> the object dictionary we want to mapped.

The Entry element is used to describe an object dictionary that is required to be mapped to.

```

<Entry SubIndex="#1">
  <Index>#x6041</Index>
  <SubIndex>#x0</SubIndex>
  <DataType>UINT</DataType>
  <BitLen>#16</BitLen>
  <Name>statusword</Name>
</Entry>

```

Sdo element

The Sdo element is used to set the default value of a object dictionary.

```

<Sdo>
  <Index>#x6085</Index>
  <Subindex>#x0</Subindex>
  <value>#x1000</value>
  <BitLen>#32</BitLen>
  <DataType>DINT</DataType>
  <Name>Quick_stop_deceleration</Name>
</Sdo>

```

The element shown in figure above means set the Object Dictionary "6085" to 0x1000.

6.1.3.4.3.2 Axle element

```

<Axle master_index='#0' slave_position="#0" AxleIndex="#0" AxleOffset="#0">
  <Mode>pp</Mode>
  <Name>x-axle</Name>
  <reg_pdo>

```

```
...
    </reg_pdo>
    <reg_pdo>
...
    </reg_pdo>
</Axle>
```

- `master_index` attribute indicates which *master* this *axle* belongs to.
- `slave_position` attribute indicates which *slave* this *axle* belongs to.
- `AxleOffset` attribute indicates which *axle* this *axle* is on the SubDevice. As mentioned above, a CoE SubDevice could have more than one *axle* . If this axle is the second axle on the SubDevice, set `AxleOffset="#1"` .
- `<Mode>` indicates the mode this axle will work on.
- `<Name>` is the name of this axle.
- `<reg_pdo>` is the PDO entry we want to register.

`reg_pdo` element

```
<reg_pdo>
  <Index>#x606c</Index>
  <Subindex>#x0</Subindex>
  <Name></Name>
</reg_pdo>
```

6.1.3.4.4 Testing a CoE servo system

6.1.3.4.4.1 Hardware preparation

- A CoE servo system:
A CoE servo system includes a CoE servo and a motor. In this test, 'Delta ASDA-B3-E' or '2HSS458-EC' servo system shown as in the figure below is used.
- A board supported by Real-time Edge:
For this test, i.MX 8M Mini LPDDR4 EVK is used.



2HSS458-EC Servo System

Figure 112. '2HSS458-EC' servo system



Figure 113. ASDA-B3-E Servo System

6.1.3.4.4.2 Software preparation

Make sure the below config options are selected when configuring Real-time Edge.

- igh-ethercat
- libxml2
- real-time-edge-servo

6.1.3.4.4.3 CoE network detection

- Igh configuration
 - Configure the `MASTER0_DEVICE` field of the `/etc/ethercat.conf`
Set `MASTER0_DEVICE` to the MAC address to indicate which port the Igh uses.
 - Configure `DEVICE_MODULES="generic"` of the `/etc/ethercat.conf`

- Use the command below to start lgh service.

```
[root]# ethercatctl start
```

- Check CoE servo using below command.

```
[root]# ethercat slaves
0 0:0 PREOP + 2HSS458-EC
Or
[root]# ethercat slaves
0 0:0 PREOP + Delta ASDA-B3-E EtherCAT(CoE) Drive Rev0.00
```

6.1.3.4.4 Starting the test with 2HSS458-EC servo

Note:

The Position encoder resolution and Velocity encoder resolution of "2HSS458-EC" servo system are both 4000. It means the ratio of encoder increments per motor revolution.

Profile Position mode test

- Start the test service as below.

```
[root]# nservo_run -f /home/root/nservo_example/hss248_ec_config_pp.xml &
```

- Check whether the status of the SubDevice has been transferred from "PREOP" to "OP".

```
[root]# ethercat slaves
0 0:0 OP + 2HSS458-EC
```

- Check whether the phase of the MainDevice has been transferred from "Idle" to "Operation".

```
[root]# ethercat master | grep Phase
Phase: Operation
```

- Run below commands to test whether the motor works.

- Get current mode of axle 0.

```
[root]# nservo_client -a 0 -c get_mode
get_mode of the axle 0 : Profile Position Mode
```

- Get current position of axle 0.

```
[root]# nservo_client -a 0 -c get_current_position
get_current_position of the axle 0 : 0
```

- Get the profile speed of axle 0.

```
[root]# nservo_client -a 0 -c get_profile_velocity
get_profile_velocity of the axle 0 : 800000
```

The value 800000 indicates 200 revolutions per second.

- Set profile speed of axle 0.

```
[root]# nservo_client -a 0 -c set_profile_velocity:20000
set_profile_velocity of the axle 0 : 20000
```

Set the profile speed to 5 revolutions per second.

- Set target position of axle 0

```
[root]# nservo_client -c set_target_position:400000
set_target_position of the axle 0 : 400000
```

The value 400000 means that the motor will turn 100 rounds.

(target_position: 400000 - current_position:0) / 4000 = 100

- Get current speed of axle 0

```
[root]# nservo_client -a 0 -c get_current_velocity  
get_current_velocity of the axle 0 : 19999
```

- Get target position of axle 0

```
[root]# nservo_client -a 0 -c get_target_position  
get_target_position of the axle 0 : 400000
```

- Exit

```
[root]# nservo_client -c exit
```

Profile Velocity mode test

- Start the test service as below.

```
[root]# nservo_run -f /home/root/nservo_example/hss248_ec_config_pv.xml &
```

- Check whether the status of the SubDevice has been transferred from "PREOP" to "OP".

```
[root]# ethercat slaves  
0 0:0 OP + 2HSS458-EC
```

- Check whether the phase of the MainDevice has been transferred from "Idle" to "Operation".

```
[root]# ethercat master | grep Phase  
Phase: Operation
```

- Run below commands to test whether the motor works.

- Get current mode of axle 0.

```
[root]# nservo_client -a 0 -c get_mode  
get_mode of the axle 0 : Profile Velocity Mode
```

- Set target speed of axle 0.

```
[root]# nservo_client -a 0 -c set_target_velocity:40000  
set_target_velocity of the axle 0 : 40000
```

The value 40000 means that the motor will turn with 10 revolutions per second.

- Get current speed of axle 0.

```
[root]# nservo_client -a 0 -c get_current_velocity  
get_current_velocity of the axle 0 : 32000
```

- Get target speed of axle 0.

```
[root]# nservo_client -a 0 -c get_target_velocity  
get_target_velocity of the axle 0 : 40000
```

- Exit

```
[root]# nservo_client -c exit
```

6.1.3.4.4.5 Starting test with ASDA-B3-E servo system

Note: The position encoder resolution of "ASDA-B3-E" servo system is 16777216 (24 bits). It signifies that the ratio of encoder increments per motor revolution.

Profile Position mode test

1. Start the test service as below.

```
[root]# nservo_run -f /home/root/nservo_example/Delta-ASDA-B3-pp.xml &
```

Check whether the status of the SubDevice has been transferred from "PREOP" to "OP".

```
[root]# ethtercat slaves
0 0:0 OP + Delta ASDA-B3-E EtherCAT(CoE) Drive Rev0.00
```

2. Check whether the phase of the MainDevice has been transferred from "Idle" to "Operation".

```
[root]# ethtercat master | grep Phase Phase: Operation
```

3. Run the below commands to test whether the motor works.

- a. Get the current mode of axle 0.

```
[root]# nservo_client -a 0 -c get_mode
get_mode of the axle 0 : Profile Position Mode
```

- b. Get the current position of axle 0.

```
[root]# nservo_client -a 0 -c get_current_position
get_current_position of the axle 0 : 0
```

- c. Get the profile speed of axle 0.

```
[root]# nservo_client -a 0 -c get_profile_velocity
get_profile_velocity of the axle 0 : 0
```

- d. Set profile speed of axle 0.

```
[root]# nservo_client -a 0 -c set_profile_velocity:16777216
set_profile_velocity of the axle 0 : 16777216
```

- e. Set profile speed to 1 revolutions per second. Set target position of axle 0.

```
[root]# nservo_client -c set_target_position:167772160
set_target_position of the axle 0 : 167772160
```

The value 167772160 means that the motor turns 10 rounds. (target_position: 167772160 - current_position:0) / 16777216 = 10

- f. Get the current position of axle 0.

```
[root]# nservo_client -a 0 -c get_current_position
get_current_position of the axis 0 : 167772152
```

- g. Get target position of axle 0

```
[root]# nservo_client -a 0 -c get_target_position
get_target_position of the axle 0 : 167772160
```

4. Exit.

```
[root]# nservo_client -c exit
```

Profile Velocity mode test

1. Start the test service as below.

```
[root]# nservo_run -f /home/root/nservo_example/Delta-ASDA-B3-pv.xml &
```

2. Check whether the status of the SubDevice has been transferred from "PREOP" to "OP".

```
[root]# ethercat slaves
0 0:0 OP + Delta ASDA-B3-E EtherCAT(CoE) Drive Rev0.00
```

3. Check whether the phase of the MainDevice has been transferred from "Idle" to "Operation".

```
[root]# ethercat master | grep Phase Phase: Operation
```

4. Run the below commands to test whether the motor works.

- a. Get current mode of axle 0.

```
[root]# nservo_client -a 0 -c get_mode
get_mode of the axle 0 : Profile Velocity Mode
```

- b. Set target speed of axle 0.

```
[root]# nservo_client -a 0 -c set_target_velocity:600
set_target_velocity of the axle 0 : 600
```

The value 600 means that the motor turns with 60 revolutions per minute for ASDA-B3-E servo.

- c. Get current speed of axle 0.

```
[root]# nservo_client -a 0 -c get_current_velocity
get_current_velocity of the axle 0 : 600
```

- d. Get target speed of axle 0.

```
[root]# nservo_client -a 0 -c get_target_velocity
get_target_velocity of the axle 0 : 600
```

5. Exit

```
[root]# nservo_client -c exit
```

Cyclic Sync Position mode test

1. Start the test service as below.

```
[root]# nservo_run -f /home/root/nservo_example/Delta-ASDA-B3-csp.xml &
```

2. Check whether the status of the SubDevice has been transferred from "PREOP" to "OP".

```
[root]# ethercat slaves
0 0:0 OP + Delta ASDA-B3-E EtherCAT(CoE) Drive Rev0.00
```

3. Check whether the phase of the MainDevice has been transferred from "Idle" to "Operation".

```
[root]# ethercat master | grep Phase Phase: Operation
```

4. Run the below commands to test whether the motor works.

- a. Get current mode of axle 0.

```
[root]# nservo_client -a 0 -c get_mode
```



```
get_mode of the axle 0 : Cyclic sync Position Mode
```

- b. Load trajectory planning information from the command line for axle 0.

```
./nservo_client -a 0 -c set_tparrays:"Cyclic=1; Scale=46603; Bias=0;
Accel=8; Decel=8; Max_speed=3600; TpArrays=[(0:1000),(45:1000),(45:1000),
(90:1000)];"
```

The command parameters are described below:

- *Cyclic*: if this field is set, the motor will come back to the first point of *TpArrays* and begin running again
- *Scale*: this field is used to set the resolution of per unit. For example, the unit of position point in *TpArrays* is degree. So the Scale should be set to $167772160/360 = 46603$. 167772160 is the resolution per motor revolution.
- *Bias*: this field is used to set the bias value for the position point in *TpArrays*.
- *Accel*: the acceleration, unit^2 per second.
- *Decel*: the deceleration, unit^2 per second.
- *Max_speed*: The maximum speed, unit per second.
- *TpArrays*: is used to save *trajectory planning* information. Each element represents a position point and the time taken for the motor to rotate to this point from the last point. The unit of the time field is ms. Load trajectory planning information from a *tp* file for axle 0. Except loading the trajectory planning information from command line. This information could also be loaded from a *tp* file. The format of a sample *tp* file is shown below.

```
cat example/x6e_sv680_delta_tp_arrays
Axis=0; Cyclic=1; Scale=364; Bias=0; Accel=8; Decel=8; Max_speed=3600;
TpArrays=[(0:2000),(45:1000),(45:2000),(90:1000),(90:2000),(270:1000),
(270:2000),(0:1000)];
Axis=1; Cyclic=1; Scale=364; Bias=0; Accel=8; Decel=8; Max_speed=3600;
TpArrays=[(0:2000),(45:1000),(45:2000),(90:1000),(90:2000),(270:1000),
(270:2000),(0:1000)];
```

- *Axis*: the index of the axle.
- *Cyclic*: if this field is set, the motor will come back to the first point of *TpArrays* and run again.
- *Scale*: this field is used to set the resolution of per unit. For example, the unit of position point in *TpArrays* is degree. So the Scale should be set to $167772160/360 = 46603$. 167772160 is the resolution per motor revolution.
- *Bias*: this field is used to set bias value for the position point in *TpArrays*.
- *Accel*: the acceleration in unit^2 per second.
- *Decel*: the deceleration in unit^2 per second.
- *Max_speed*: The maximum speed, unit per second.
- *TpArrays*: are used to save *trajectory planning* information. Each element represents a position point and the time taken for the motor rotating to the this point from the last point. The unit of the time field is ms.

```
./nservo_client -c load_tp_file:"/home/root/nservo_example/Delta-ASDA-
B3-tp_arrays"
```

- *Cyclic*: if this field is set, the motor will come back to the first point of *TpArrays* and begin running again.
- *Scale*: this field is used to set the resolution of per unit. For example, the unit of the position point in *TpArrays* is degree. So the Scale should be set to $167772160/360 = 46603$. 167772160 is the resolution per motor revolution. /
- *Accel*: the acceleration, unit^2 per second.

- *Decel*: the deceleration, unit² per second.
 - *Max_speed*: The maximum speed, unit per second.
 - *TpArrays*: used to save trajectory planning information. Each element represents a position point and the time taken for the motor to rotate to this point from the last point. The unit of the time field is ms.
- Set the axle 0 to start running using the below command:

```
[root]# nservo_client -a 0 -c set_start  
set_start of the axis 0
```

- Get the *current_position* of axle 0 using the below command:

```
[root]# nservo_client -a 0 -c get_current_position  
get_current_position of the axis 0 : 2815922
```

- Set axle 0 to stop running using the below command:

```
[root]# nservo_client -a 0 -c set_stop  
set_stop of the axis 0
```

- Set all axes in CSP mode to start running using the below command:

```
[root]# nservo_client -c set_start_all
```

All CSP axis in ready status start to run.

- Set all axes in CSP mode to stop running using the below command:

```
[root]# nservo_client -c set_stop_all
```

All CSP axis in running status begin to stop.

- Exit

```
[root]# nservo_client -c exit
```

6.1.3.4.5 EtherCAT multiple axes control system

6.1.3.4.5.1 HCFA 60-axes servo system

Below is a 60-axes servo system built by HCFA, and this system consists of 60 X3E servo motors and also 60 pointers in the screen which could rotate 360 degrees under the corresponding servo motor control. As shown in the figure below, this system can render any character on the screen by rotating these pointers.

Any platforms supported on Real-time Edge could be used as the controller of this servo system, and the software is based on real-time-edge-servo stack.



Figure 114. 60-axes servo system built by HCFA

HCFA performance

The task period is 1 ms, and servo control mode is CSP.

- Native EtherCAT driver + IGH stack: 26 μ s
- Schedule latency: 200 μ s on i.MX 8MP, 220 μ s on i.MX 8M Mini
- Link propagation latency: 64 μ s
- Customer task: 690-700 μ s saved for app

Running this case (HCFA)

All software associated with this controller is integrated in Real-time Edge by default, and can be set up using the below steps:

- Start the nservo service as below:

```
[root]# nservo_run -f /home/root/nservo_example/x3e_csp_60_config.xml &
```

- Run the below commands to confirm whether the motor works:

```
[root]# nservo_client -a 59 -c get_mode
```

- Load the trajectory planning information:

```
[root]#nservo_client -c load_tp_file:"/home/root/nservo_example/
x3e_60_axis_nxp_logo"e file x3e_60_axis_nxp_logo includes all of trajectory
planning information for 60-axes to control this servo system to present NXP
logo.
```

- Start all servo motors:
After the trajectory planning information file is loaded, use the below command to start the system.

```
[root]# nservo_client -c set_start_all
```

6.1.4 SOEM EtherCAT MainDevice

Simple Open EtherCAT MainDevice (SOEM) is an open source EtherCAT MainDevice stack, designed to be easy to use and maintains a small footprint. The optimized native driver is implemented by building upon this stack and introduces various profiles and multi-axes use cases. These use cases are elaborated below:

- Native Ethernet driver for Ethernet controller of i.MX 8M Plus, i.MX 8M Mini, i.MX 93
- HAL layer for FreeRTOS and Baremetal
- CiA401 profile for digital IO
- CiA402 profile for motor control
- Multi-axes use case with market available servos and i.MXRT1180 as EtherCAT SubDevice with different control loop time

SOEM can run on FreeRTOS to scale down to MCU, Cortex-M core of MPU, and scale up to the Cortex-A core of MPU. The support on MCU is delivered through the MCUXpresso SDK.

Refer to MCUXpresso SDK for running SOEM on the supported NXP MCUs. On Real-time Edge software, SOEM examples can be run on Cortex-M Core and Cortex-A Core through the Heterogeneous Multicore Framework (see [Section 3.8](#)).

SOEM is supported on the below platforms:

- i.MX 8M Plus LPDDR4 EVK platform
- i.MX 8M Mini LPDDR4 EVK platform
- i.MX 93 EVK platform

For the i.MX 8M Plus as an example, it is equipped with four Cortex-A53 cores and one Cortex-M7 core. [Table 92](#) describes the use cases that can be executed on this MPU platform.

Table 92. SOEM use cases

SOEM examples	M7	A53	A53	A53	A53
digital_io	baremetal/FreeRTOS	FreeRTOS			
servo_motor	baremetal/FreeRTOS	FreeRTOS			
servo_motor_rt1180	baremetal/FreeRTOS	FreeRTOS			

Three examples are provided in this release:

- SOEM `digital_io` example for IO control using BECKHOFF EK1100 as EtherCAT SubDevice
- SOEM `servo_motor` example using Inovance SV680 servo as EtherCAT SubDevice
- SOEM `servo_motor_rt1180` example using i.MX RT1180-EVK as EtherCAT SubDevice

The support on MCU is delivered in the MCUXpresso SDK, refer to MCUXpresso SDK for running SOEM on MCU. On Real-time Edge software. SOEM examples can be run on Cortex-M Core and Cortex-A Core through the Heterogeneous Multicore Framework.

6.1.4.1 SOEM for i.MX 8M Plus LPDDR4 EVK platform

6.1.4.1.1 Setup hardware environment

For the **digital_io** example, the below hardware is required:

- i.MX 8M Plus LPDDR4 EVK platform
 - EK1100 - EtherCAT Coupler
 - EL2008 - EtherCAT Terminal, 8-channel digital output, 24V DC
 - EL1018 - EtherCAT Terminal, 8-channel digital input, 24 V DC
 - 24 V DC power supply




Figure 115. Running the digital_io example on i.MX 8M Plus LPDDR4 EVK

The channel-1 of the EL1018 is connected to a button to generate a 24 V pulse as a input signal to change the output of the channel-1 of the EL2008 labeled as "Dir". The channel-2 of the EL2008 outputs a square wave signal with a period of 250 μ s.

For **servo_motor** and **servo_motor_rt1180** examples, the below hardware is required:

- i.MX 8M Plus LPDDR4 EVK platform
- Inovance SV680 / i.MX RT1180 EVK platform
- Network Cable

There are two types of servos supported, Inovance SV680 and i.MX RT1180-EVK as EtherCAT SubDevice.

1. **Inovance SV680 as EtherCAT SubDevice:** Running SOEM Servo Motor example on i.MX 8M Plus EVK to control Inovance SV680.



Figure 116. Running SOEM Servo Motor example on i.MX 8M Plus EVK (Inovance SV680 as EtherCAT SubDevice)

2. **Running `ecat_examples/servo_motor` example on i.MX RT1180 EVK as EtherCAT SubDevice.**
Running SOEM Servo Motor example on i.MX 8M Plus EVK to control motor on i.MX RT1180 EVK.




Figure 117. Running `servo_motor` example on i.MX 8M Plus EVK (i.MX RT1180 EVK as EtherCAT SubDevice)

Note: Before running `servo_motor` example, you must download the MCU-SDK package for MIMXRT1180-EVK board from the NXP official website: [Welcome | MCUXpresso SDK Builder \(nxp.com\)](https://www.nxp.com/Welcome+MCUXpresso+SDK+Builder+nxp.com)

6.1.4.1.2 Building the demo images of Cortex-M

For the **digital_io** example, the demo images "freertos_soem_gpio_pulse" and "soem_gpio_pulse" are compiled with the i.MX 8M Plus LPDDR4 EVK target image compiling, and are installed in the "/example" directory of the target rootfs. For JTAG download, these images can also be found in the directory "<image-build-dir>/tmp/deploy/images/imx8mpevk/examples/" on the build host.

```
cd <image-build-dir>/
tree tmp/deploy/images/imx8mpevk/examples/mcux-sdk/soem-gpio-pulse
|— ddr_release
|   |— soem_gpio_pulse.elf
|   |— soem_gpio_pulse.bin
|— release
|   |— soem_gpio_pulse.elf
|   |— soem_gpio_pulse.bin
tree tmp/deploy/images/imx8mpevk/examples/mcux-sdk/freertos-soem-gpio-pulse
|— ddr_release
|   |— freertos_soem_gpio_pulse.elf
```



```
| └─ freertos_soem_gpio_pulse.bin
| └─ release
| └─ freertos_soem_gpio_pulse.elf
| └─ freertos_soem_gpio_pulse.bin
```

Also, the demo images can be compiled using the below command on i.MX 8M Plus LPDDR4 EVK build directory:

```
bitbake soem-gpio-pulse
```

Or

```
bitbake freertos-soem-gpio-pulse
```

For **servo_motor** and **servo_motor_rt1180** example, the demo images "freertos_soem_servo_motor", "soem_servo_motor", "freertos_soem_servo_motor_rt1180" and "soem_servo_motor_rt1180" are compiled with the i.MX 8M Plus LPDDR4 EVK target image. These demo images are installed into the "/" example" directory of the target rootfs. For JTAG download, these images can also be found in the "<image-build-dir>/tmp/deploy/images/imx8mpevk/examples/" directory on the build host.

Note: *freertos_soem_servo_motor* and *soem_servo_motor* are used for the **Inovance SV680** use case whereas "freertos_soem_servo_motor_rt1180" and "soem_servo_motor_rt1180" are used for the **i.MX RT1180 EVK** use case.

Inovance SV680 use case

```
cd <image-build-dir>/
tree tmp/deploy/images/imx8mpevk/examples/mcux-sdk/soem-servo-motor
└─ ddr_release
  └─ soem_servo_motor.elf
  └─ soem_servo_motor.bin
└─ release
  └─ soem_servo_motor.elf
  └─ soem_servo_motor.bin
tree tmp/deploy/images/imx8mpevk/examples/mcux-sdk/freertos-soem-servo_motor
└─ ddr_release
  └─ freertos_soem_servo_motor.elf
  └─ freertos_soem_servo_motor.bin
└─ release
  └─ freertos_soem_servo_motor.elf
  └─ freertos_soem_servo_motor.bin
```

i.MX RT1180 EVK use case

```
cd <image-build-dir>/
tree tmp/deploy/images/imx8mpevk/examples/mcux-sdk/soem-servo-motor-rt1180
└─ ddr_release
  └─ soem_servo_motor.elf
  └─ soem_servo_motor.bin
└─ release
  └─ soem_servo_motor.elf
  └─ soem_servo_motor.bin
tree tmp/deploy/images/imx8mpevk/examples/mcux-sdk/freertos-soem-servo_motor-rt1180
└─ ddr_release
  └─ freertos_soem_servo_motor.elf
  └─ freertos_soem_servo_motor.bin
└─ release
```

```
| └─ freertos_soem_servo_motor.elf
| └─ freertos_soem_servo_motor.bin
```

The demo images can also be compiled using the below commands on i.MX 8M Plus LPDDR4 EVK build directory:

```
bitbake soem-servo-motor
or
bitbake freertos-soem-servo-motor
or
bitbake soem-servo-motor-rt1180
or
bitbake freertos-soem-servo-motor-rt1180
```

6.1.4.1.3 Building the demo images of Cortex-A

For **digital_io** example, the demo images "soem_digital_io_ca53" are compiled with the i.MX 8M Plus LPDDR4 EVK target image compiling, and they are installed into the directory "/example" of the target rootfs. Images can also be found in the directory "<image-build-dir>/tmp/deploy/images/imx8mpevk/examples/" on building host.

```
cd <image-build-dir>/
tree tmp/deploy/images/imx8mpevk/examples/heterogeneous-multicore/soem-digital-
io-ca
|─ ddr_release
| └─ soem_digital_io_ca53.elf
| └─ soem_digital_io_ca53.bin
```

Also, the demo images can be compiled using the below command on i.MX 8M Plus LPDDR4 EVK build directory:

```
bitbake soem-digital-io-ca
```

For **servo_motor** and **servo_motor_rt1180** example, the demo images "soem_servo_ca53", "soem_servo_rt1180_ca53" are compiled with the i.MX 8M Plus LPDDR4 EVK target image compiling, and they are installed into the directory "/example" of the target rootfs. Images can also be found in the directory "<image-build-dir>/tmp/deploy/images/imx8mpevk/examples/" on building host.

Note: "soem_servo_ca53" are for "Inovance SV680" case, "soem_servo_rt1180_ca53" are for "i.MX RT1180 EVK" case.

```
cd <image-build-dir>/
tree tmp/deploy/images/imx8mpevk/examples/heterogeneous-multicore/soem-servo-ca
|─ ddr_release
| └─ soem_servo_ca53.elf
| └─ soem_servo_ca53.bin
tree tmp/deploy/images/imx8mpevk/examples/heterogeneous-multicore/soem-servo-
rt1180-ca
|─ ddr_release
| └─ soem_servo_rt1180_ca53.elf
| └─ soem_servo_rt1180_ca53.bin
```

Also, the demo images can be compiled using the below command on i.MX 8M Plus LPDDR4 EVK build directory:

```
bitbake soem-servo-ca
or
```



```
bitbake soem-servo-rtl180-ca
```

6.1.4.1.4 Running SOEM demo images using J-Link GDB Server

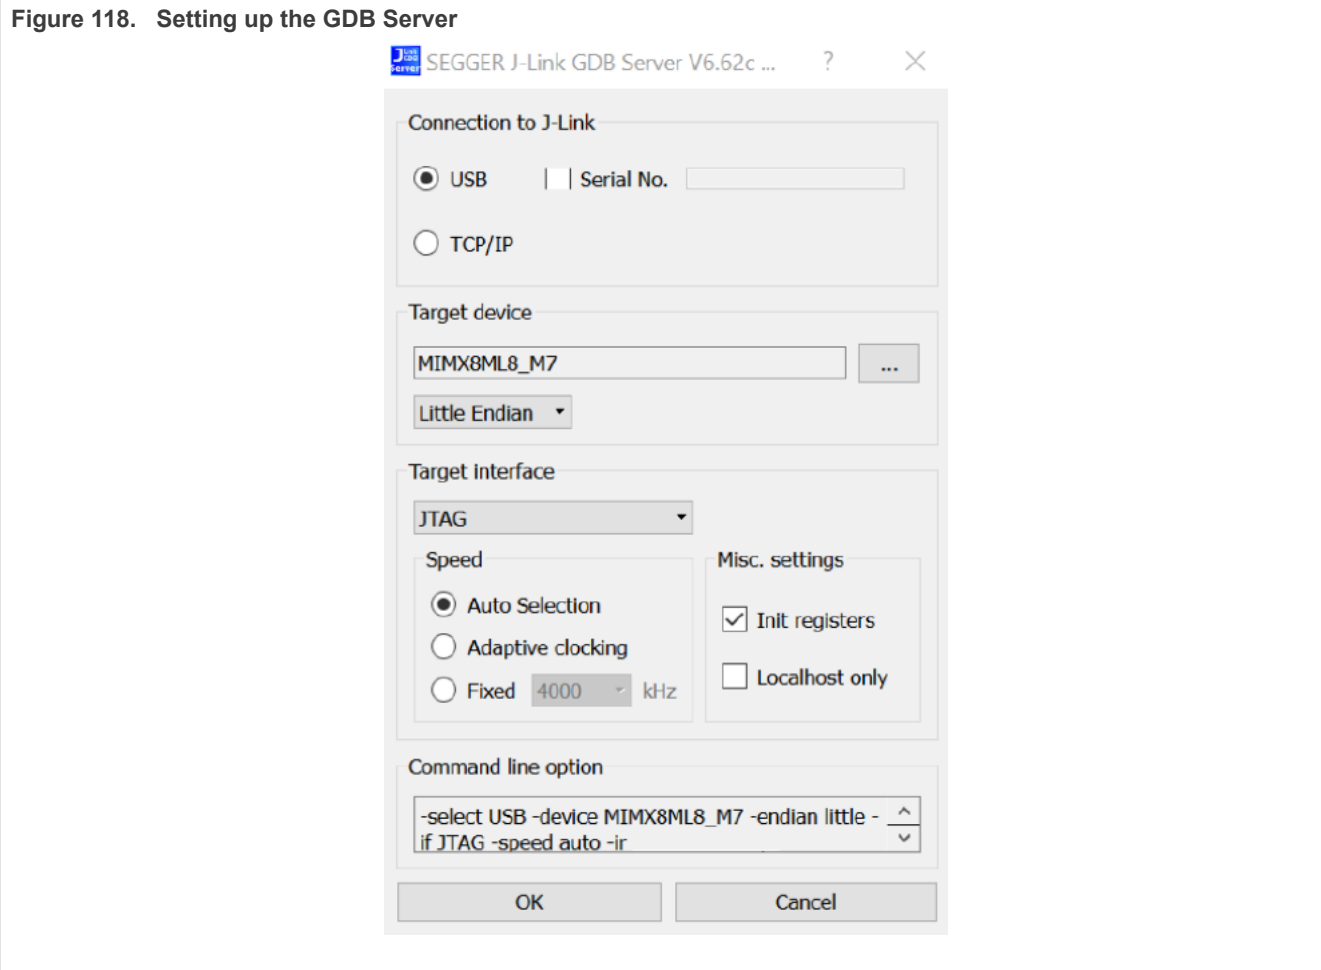
This section describes the steps to run a demo application using J-Link GDB Server application. After the J-Link interface is configured and connected, follow these steps to download and run the demo applications:

1. Connect the i.MX 8M Plus LPDDR4 EVK platform to your PC via USB cable between the USB-UART connector and the PC USB connector. If using a standalone J-Link debug pod, also connect it to the SWD/JTAG connector of the board.
2. Open the terminal application on the PC, such as PuTTY or TeraTerm, and connect to the debug serial port number. Configure the terminal with these settings:
 - 115200 baud rate
 - No parity
 - 8 data bits
 - 1 stop bit
3. Open the J-Link GDB Server application on Linux host. If the OS of your PC is Linux, using the below configuration to set up GDB Server. Assuming the J-Link software is installed, the application can be launched from a new terminal for the MIMX8ML_M7 device:

```
$ JLinkGDBServer -if JTAG -device
SEGGER J-Link GDB Server Command Line Version
JLinkARM.dll
Command line: -if JTAG -device MIMX8ML8_M7
-----GDB Server start settings-----
GDBInit file: none
GDB Server Listening port: 2331
SWO raw output listening port: 2332
Terminal I/O port: 2333
Accept remote connection: yes
< -- Skipping lines -->
Target connection timeout: 0 ms
-----J-Link related settings-----
J-Link Host interface: USB
J-Link script: none
J-Link settings file: none
-----Target related settings-----
Target device:
Target interface: JTAG
Target interface speed: 1000 kHz
Target endian: little
Connecting to J-Link...
J-Link is connected.
Firmware: J-Link V10 compiled Feb 2 2020 18:12:40
Hardware: V10.10
S/N: 600109545 Feature(s): RDI, FlashBP, FlashDL, JFlash, GDB
Checking target voltage...
Target voltage: 1.82 V
Listening on TCP/IP port 2331
Connecting to target...
J-Link found 1 JTAG device, Total IRLen = 4
JTAG ID: 0x5BA00477 (Cortex-M7)
Connected to target
Waiting for GDB connection...
```

4. Open the J-Link GDB Server application on Windows host. If the OS of your PC is Windows, use the below image to set up GDB Server. Assuming the J-Link software is installed, the application can be launched

by going to the Windows operating system “Start” menu and selecting “Programs -> SEGGER -> J-Link <version> J-Link GDB Server”. After server is launched, modify the settings as below. The target device selected for this demo is MIMX8ML8_M7 .



5. After GDB server is running, the screen looks like as below:

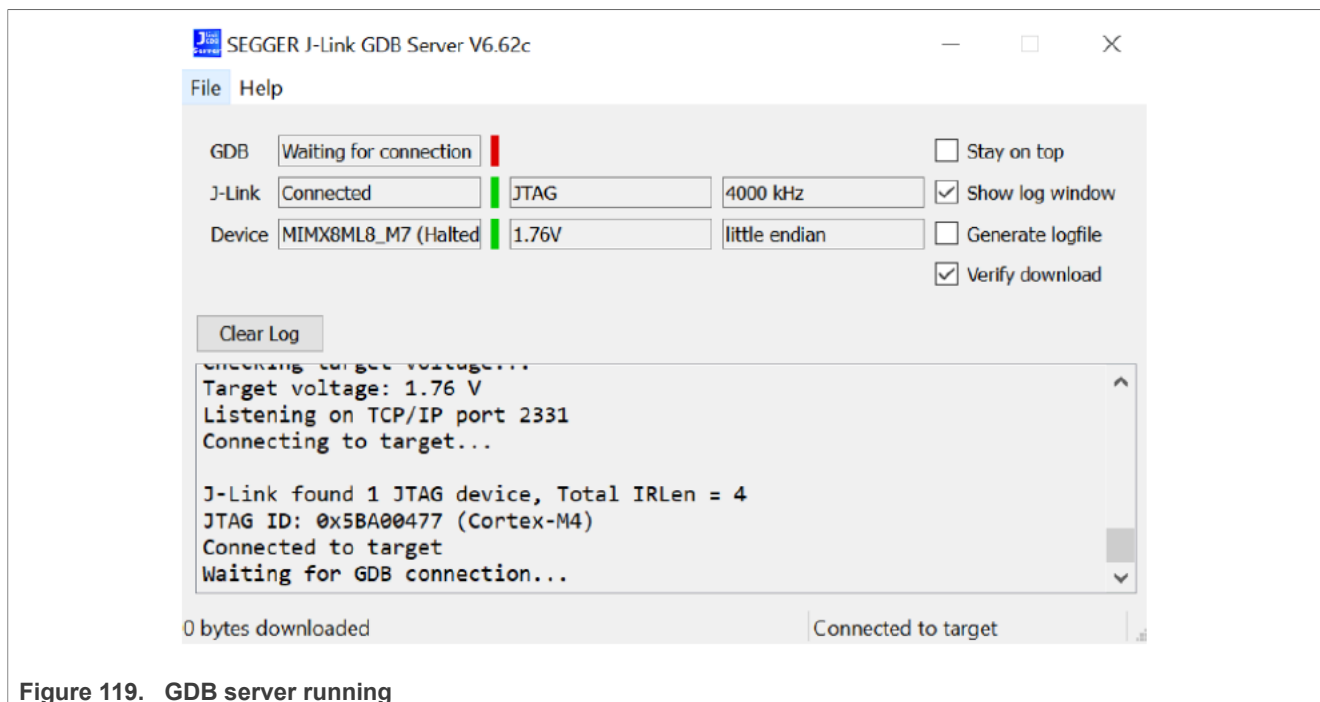


Figure 119. GDB server running

6. Start the GDB client:
7. Ensure that the `arm-none-eabi-gdb` is installed, change to the directory that contains the demo images output "`<build-dir>/tmp/deploy/images/imx8mpdevk/examples/soem-gpio-pulse`".

```
$ arm-none-eabi-gdb ./release/soem_gpio_pulse.elf
GNU gdb (GNU Tools for Arm Embedded Processors 9-2019-q4-major)
8.3.0.20190709-git
Copyright (C) 2019 Free Software Foundation, Inc.
License GPLv3+: GNU GPL version 3 or later < http://gnu.org/licenses/
gpl.html>
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.
Type "show copying" and "show warranty" for details.
This GDB was configured as "--host=i686-w64-mingw32 --target=arm-none-eabi".
Type "show configuration" for configuration details.
For bug reporting instructions, please see:
<http://www.gnu.org/software/gdb/bugs/>
Find the GDB manual and other documentation resources online at:
<http://www.gnu.org/software/gdb/documentation/>
For help, type "help". Type "apropos word" to search for commands related to
"word"...
Reading symbols from soem_gpio_pulse.elf...
(gdb)
```

8. Connect to the GDB server and load the binary by running the following commands:

```
a. target remote <GDB Server IP>:2331
b. monitor reset
c. monitor halt
d. load
```

9.


```
(gdb) target remote localhost:2331
Remote debugging using localhost:2331
0x00000008 in __isr_vector ()
(gdb) monitor reset Resetting target
```

```
(gdb) monitor halt
(gdb) load
Loading section .interrupts, size 0x240 lma 0x0
Loading section .text, size 0x3ab8 lma 0x240
Loading section .ARM, size size 0x8 lma 0x3cf8
Loading section .init_array, size 0x4 lma 0x3d00
Loading section .fini_array, size 0x4 lma 0x3d04
Loading section .data, size 0x64 lma 0x3d08
Start address 0x2f4, load size 15724
Transfer rate: 264 KB/sec, 2620 bytes/write.
(gdb)
```

The application is now downloaded and stopped at the reset vector. Execute the monitor go command to start the demo application.

```
(gdb) monitor go
```

6.1.4.1.5 Running SOEM demo images using U-Boot

This section describes the steps to write an SOEM demo image file to TCM or DRAM with the Real-time Edge image. The following steps describe how to use the U-Boot:

1. Connect the DEBUG UART slot on the board to your PC through the USB cable. The Windows OS installs the USB driver automatically, and the Ubuntu OS finds the serial devices as well.
2. On Windows OS, open the **Device Manager** and find USB serial Port in Ports (COM and LPT). Assume that the ports are COM9 and COM10. One port is for the debug message from the Cortex-A53 and the other is for the Cortex-M7. The port number is allocated randomly, so opening both is beneficial for development. On Ubuntu OS, identify the TTY device with name **/dev/ttyUSB*** to determine your debug port. Similar to Windows OS, opening both is beneficial for development.
3. Build and flash the `nxp-image-real-time-edge` image to a SD card and insert the SD card to the target board. Make sure to use the default boot SD slot and check the DIP boot switch configuration.
4. Open your preferred serial terminal for the serial devices. Set the speed to 115200 bps, 8 data bits, 1 stop bit (115200, 8N1), no parity. Then, power on the board.
5. Power on the board and hit any key to stop autoboot in the terminal window. Then, enter the U-Boot command line mode. You can then write the image and run it from TCM or DRAM with the following commands:

6. Running on Cortex-M

- a. If the `soem_gpio_pulse.bin` is made from the `<examples/...>/release` target, which means the binary file will run at TCM, use the following commands to boot:

```
=> ext4load mmc 1:2 0x48000000 /examples/mcux-sdk/soem-gpio-pulse/release/
soem_gpio_pulse.bin;
=> cp.b 0x48000000 0x7e0000 0x20000;
=> bootaux 0x7e0000
```

- b. If the `soem_gpio_pulse.bin` is made from the `<examples/...>/ddr_release` target, which means the binary file runs at DRAM, use the following commands:

```
=> ext4load mmc 1:2 0x80000000 /examples/mcux-sdk/soem-gpio-pulse/
ddr_release/soem_gpio_pulse.bin;
=> dcache flush
=> bootaux 0x80000000
```

- c. If the `soem_servo_motor.bin` is made from the `<examples/..>/release` target, which means the binary file will run at TCM, use the following commands to boot:

```
=> ext4load mmc 1:2 0x48000000 /examples/mcux-sdk/soem-servo-motor/
release/soem_servo_motor.bin;
=> cp.b 0x48000000 0x7e0000 0x20000;
=> bootaux 0x7e0000
```

Or

- d.

```
=> ext4load mmc 1:2 0x48000000 /examples/mcux-sdk/soem-servo-motor-rt1180/
release/soem_servo_motor.bin;
=> cp.b 0x48000000 0x7e0000 0x20000;
=> bootaux 0x7e0000
```

- e. If the `soem_servo_motor.bin` is made from the `<examples/..>/ddr_release` target, which means that the binary file runs at DRAM, use the following commands:

```
=> ext4load mmc 1:2 0x80000000 /examples/mcux-sdk/soem-servo-motor/
ddr_release/soem_servo_motor.bin;
=> dcache flush
=> bootaux 0x80000000
```

Or

```
=> ext4load mmc 1:2 0x80000000 /examples/mcux-sdk/soem-servo-motor-rt1180/
ddr_release/soem_servo_motor.bin;
=> dcache flush
=> bootaux 0x80000000
```

Note: If the Linux OS kernel runs together with M7, make sure the correct dtb file is used. This dtb file reserves resources used by M7 and avoids the Linux kernel from configuring them. Use the following command in U-Boot before running the kernel:

```
setenv fdtfile imx8mp-evk-rpmsg.dtb
```

Modification must be done in case you require that SOEM app still works when running Linux. M7 core needs exclusive access to Ethernet. So remove the Ethernet access from Linux Kernel by performing the following steps:

- a. In `kernel-source/arch/arm64/boot/dts/freescale/imx8mp-evk-rpmsg.dts` add at the end of the kernel device tree the following lines:

```
&fec {
    status = "disabled";
};

Recompile the kernel:
$ bitbake -f -c compile virtual/kernel
$ bitbake virtual/kernel
```

- b. Copy the new device tree and the kernel image to SD boot partition.
c. Now, SOEM still works while Linux is running.

7. Running on Cortex-A

```
=> ext4load mmc 1:2 0xC0000000 /examples/heterogeneous-multicore/soem-
digital-io-ca/ddr_release/soem_digital_io_ca53.bin;
=> dcache flush; icache flush;
=> cpu 2 release 0xC0000000
```

Or

```
=> ext4load mmc 1:2 0xC0000000 /examples/heterogeneous-multicore/soem-servo-  
ca/ddr_release/soem_servo_ca53.bin;  
=> dcache flush; icache flush;  
=> cpu 2 release 0xC0000000
```

Or

```
=> ext4load mmc 1:2 0xC0000000 /examples/heterogeneous-multicore/soem-servo-  
rt1180-ca/ddr_release/soem_servo_ca53.bin;  
=> dcache flush; icache flush;  
=> cpu 2 release 0xC0000000
```

6.1.4.2 SOEM for i.MX 8M Mini LPDDR4 EVK platform

6.1.4.2.1 Setting up the hardware environment

For the **digital_io** example, the below hardware is required:

- i.MX 8M Mini LPDDR4 EVK platform
- EK1100 - EtherCAT coupler
- EL2008 - EtherCAT Terminal, 8-channel digital output, 24 V DC
- EL1018 - EtherCAT Terminal, 8-channel digital input, 24 V DC
- 24 V DC power supply




Figure 120. Hardware setup for *digital_io* example on i.MX 8M Mini LPDDR4 EVK board

The channel-1 of the EL1018 is connected to a button to generate a 24 V pulse. This pulse is used as an input signal to change the output of the channel-1 of the EL2008 labeled as "Dir". The channel-2 of the EL2008 outputs a square wave signal with a period of 250 μ s.

For the **servo_motor** and the **servo_motor_rt1180** example, the below hardware is required:

- i.MX 8M Mini LPDDR4 EVK platform
- Inovance SV680 / i.MX RT1180 EVK platform
- Network cable

There are two ways to build the example:

1. Inovance SV680 as EtherCAT SubDevice. Running SOEM Servo Motor example on i.MX 8M Plus EVK to control Inovance SV680.




Figure 121. Running SOEM Servo Motor example on i.MX 8M Mini LPDDR4 EVK board (using Inovance SV680 as EtherCAT SubDevice)

Or

1. Running `ecat_examples/servo_motor` example on i.MX RT1180 EVK as EtherCAT SubDevice. Running SOEM Servo Motor example on i.MX 8M Plus EVK to control motor on i.MX RT1180 EVK.




Figure 122. Running SOEM Servo Motor example on i.MX 8M Mini LPDDR4 EVK board (using i.MX RT1180 EVK as EtherCAT SubDevice)

Note: Before running `servo_motor` example, users must download the MCU-SDK package for i.MX RT1180 EVK from the NXP official website:

[Welcome \[MCUXpresso SDK Builder \(nxp.com\)\]](https://www.nxp.com/developer/mcu-sdk-builder)

6.1.4.2.2 Building the demo images of Cortex-M

For **digital_io** example, the demo images "`freertos_soem_gpio_pulse`" and "`soem_gpio_pulse`" are compiled with the i.MX 8M Mini LPDDR4 EVK target image compiling, and are installed into the directory "/"

example" of the target rootfs. For JTAG download, these images can also be found in the directory "<image-build-dir>/tmp/deploy/images/imx8mmevk/examples/" on the build host.

```
cd <image-build-dir>/
tree tmp/deploy/images/imx8mmevk/examples/mcux-sdk/soem-gpio-pulse
|— ddr_release
|   |— soem_gpio_pulse.elf
|   |— soem_gpio_pulse.bin
|— release
|   |— soem_gpio_pulse.elf
|   |— soem_gpio_pulse.bin
tree tmp/deploy/images/imx8mmevk/examples/mcux-sdk/freertos-soem-gpio-pulse
|— ddr_release
|   |— freertos_soem_gpio_pulse.elf
|   |— freertos_soem_gpio_pulse.bin
|— release
|   |— freertos_soem_gpio_pulse.elf
|   |— freertos_soem_gpio_pulse.bin
```

Also, the demo images can be compiled using the below command on i.MX 8M Mini LPDDR4 EVK build directory:

```
bitbake soem-gpio-pulse
```

Or

```
bitbake freertos-soem-gpio-pulse
```

For **servo_motor** and **servo_motor_rt1180** example, the demo images "freertos_soem_servo_motor", "soem_servo_motor", "freertos_soem_servo_motor_rt1180" and "soem_servo_motor_rt1180" are compiled with the i.MX 8M Mini LPDDR4 EVK target image compiling, and they are installed into the directory "/example" of the target rootfs. For JTAG download, these images can also be found in the directory "<image-build-dir>/tmp/deploy/images/imx8mmevk/examples/" on building host.

Note: "freertos_soem_servo_motor", "soem_servo_motor" are for "Inovance SV680" case, "freertos_soem_servo_motor_rt1180" and "soem_servo_motor_rt1180" are for "i.MX RT1180 EVK" case.

```
cd <image-build-dir>/
tree tmp/deploy/images/imx8mmevk/examples/mcux-sdk/soem-servo-motor
|— ddr_release
|   |— soem_servo_motor.elf
|   |— soem_servo_motor.bin
|— release
|   |— soem_servo_motor.elf
|   |— soem_servo_motor.bin
tree tmp/deploy/images/imx8mmevk/examples/mcux-sdk/freertos-soem-servo_motor
|— ddr_release
|   |— freertos_soem_servo_motor.elf
|   |— freertos_soem_servo_motor.bin
|— release
|   |— freertos_soem_servo_motor.elf
|   |— freertos_soem_servo_motor.bin
```

```
cd <image-build-dir>/
tree tmp/deploy/images/imx8mmevk/examples/mcux-sdk/soem-servo-motor-rt1180
```



```

|— ddr_release
|   └─ soem_servo_motor.elf
|   └─ soem_servo_motor.bin
|— release
|   └─ soem_servo_motor.elf
|   └─ soem_servo_motor.bin
tree tmp/deploy/images/imx8mmevk/examples/mcux-sdk/freertos-soem-servo_motor-rt1180
|— ddr_release
|   └─ freertos_soem_servo_motor.elf
|   └─ freertos_soem_servo_motor.bin
|— release
|   └─ freertos_soem_servo_motor.elf
|   └─ freertos_soem_servo_motor.bin

```

Also, the demo images can be compiled using the below command on i.MX 8M Mini LPDDR4 EVK build directory:

```

bitbake soem-servo-motor
or
bitbake freertos-soem-servo-motor
or
bitbake soem-servo-motor-rt1180
or
bitbake freertos-soem-servo-motor-rt1180

```

6.1.4.2.3 Building the demo images of Cortex-A

For **digital_io** example, the demo images "soem_digital_io_ca53" are compiled along with the i.MX 8M Mini LPDDR4 EVK target image compilation, and are installed into the directory "/example" of the target rootfs. Images can also be found in the directory "<image-build-dir>/tmp/deploy/images/imx8mmevk/examples/" on the build host.

```

cd <image-build-dir>/
tree tmp/deploy/images/imx8mmevk/examples/heterogeneous-multicore/soem-digital-io-ca
|— ddr_release
|   └─ soem_digital_io_ca53.elf
|   └─ soem_digital_io_ca53.bin

```

Also, the demo images can be compiled using the below command on i.MX 8M Mini LPDDR4 EVK build directory:

```

bitbake soem-digital-io-ca

```

For **servo_motor** and **servo_motor_rt1180** example, the demo images "soem_servo_ca53", "soem_servo_rt1180_ca53" are compiled along with the i.MX 8M Mini LPDDR4 EVK target image, and are installed into the directory "/example" of the target rootfs. Images can also be found in the directory "<image-build-dir>/tmp/deploy/images/imx8mmevk/examples/" on building host.

Note: "soem_servo_ca53" file is for "Inovance SV680" case, whereas the "soem_servo_rt1180_ca53" file is for "i.MX RT1180 EVK" case.

```

cd <image-build-dir>/
tree tmp/deploy/images/imx8mmevk/examples/heterogeneous-multicore/soem-servo-ca
|— ddr_release
|   └─ soem_servo_ca53.elf

```

```
|      └─ soem_servo_ca53.bin
tree tmp/deploy/images/imx8mmevk/examples/heterogeneous-multicore/soem-servo-
rt1180-ca
|─ ddr_release
|   └─ soem_servo_rt1180_ca53.elf
|      └─ soem_servo_rt1180_ca53.bin
```

The demo images can also be compiled using the below command on i.MX 8M Mini LPDDR4 EVK build directory:

```
bitbake soem-servo-ca
or
bitbake soem-servo-rt1180-ca
```

6.1.4.2.4 Running SOEM demo images using J-Link GDB Server

This section describes the steps to run a demo application using J-Link GDB Server application. After the J-Link interface is configured and connected, follow these steps to download and run the demo applications:

1. Connect the i.MX 8M Mini LPDDR4 EVK platform to your PC via USB cable between the USB-UART connector and the PC USB connector. If using a standalone J-Link debug pod, also connect it to the SWD/JTAG connector of the board.
2. Open the terminal application on the PC, such as PuTTY or TeraTerm, and connect to the debug serial port number. Configure the terminal with these settings:
 - a. 115200 baud rate
 - b. No parity
 - c. 8 data bits
 - d. 1 stop bit
3. Open the J-Link GDB Server application on Linux host. If the OS of your PC is Linux, using the below configuration to set up GDB Server. Assuming the J-Link software is installed, the application can be launched from a new terminal for the MIMX8MM6_M4 device:

```
$ JLinkGDBServer -if JTAG -device
SEGGER J-Link GDB Server Command Line Version
JLinkARM.dll
Command line: -if JTAG -device MIMX8MM6_M4
-----GDB Server start settings-----
GDBInit file: none
GDB Server Listening port: 2331
SWO raw output listening port: 2332
Terminal I/O port: 2333
Accept remote connection: yes
< -- Skipping lines -->
Target connection timeout: 0 ms
-----J-Link related settings-----
J-Link Host interface: USB
J-Link script: none
J-Link settings file: none
-----Target related settings-----
Target device:
Target interface: JTAG
Target interface speed: 1000 kHz
Target endian: little
Connecting to J-Link...
J-Link is connected.
Firmware: J-Link V10 compiled Feb 2 2020 18:12:40
```

```
Hardware: V10.10
S/N: 600109545 Feature(s): RDI, FlashBP, FlashDL, JFlash, GDB
Checking target voltage...
Target voltage: 1.82 V
Listening on TCP/IP port 2331
Connecting to target...
J-Link found 1 JTAG device, Total IRLen = 4
JTAG ID: 0x5BA00477 (Cortex-M7)
Connected to target
Waiting for GDB connection...
```

4. Open the J-Link GDB Server application on Window host. If the OS of your PC is Windows, use the below image to set up GDB Server.
 - Ensure that the J-Link software is installed
 - Launch the application by going to the Windows operating system “Start” menu and selecting “Programs - > SEGGER -> J-Link <version> J-Link GDB Server”.
 - After server is launched, modify the settings as shown in [Figure 123](#). The target device selected for this demo is **MIMX8MM6_M4**.

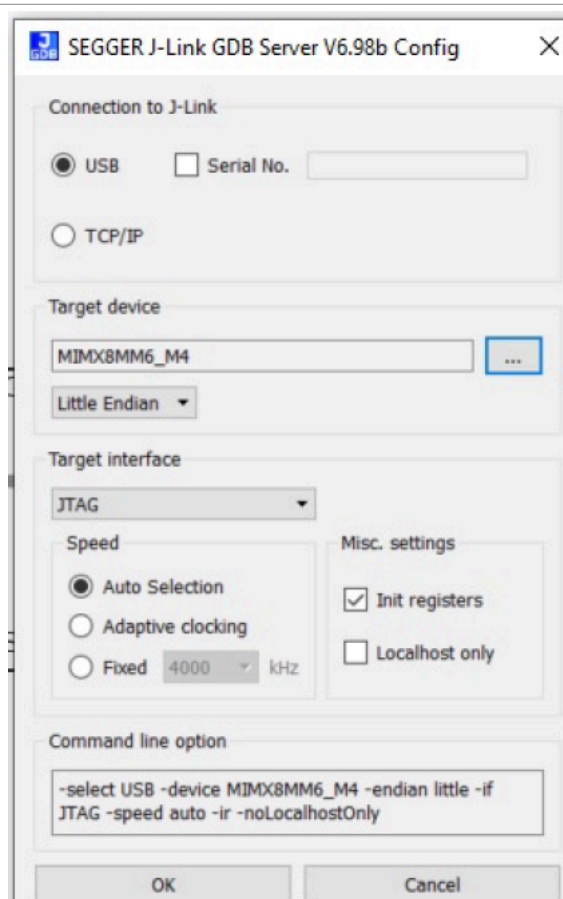


Figure 123. J-Link GDB Server application window

5. After GDB server begins to run, the screen looks as in [Figure 124](#):




Figure 124. J-Link connected and GDB server running

6. Start the GDB client

- Ensure that the `arm-none-eabi-gdb` is installed
- Change to the directory that contains the demo images output "`<build-dir>/tmp/deploy/images/imx8mmevk/examples/soem-gpio-pulse`".

```
$ arm-none-eabi-gdb ./release/soem_gpio_pulse.elf
GNU gdb (GNU Tools for Arm Embedded Processors 9-2019-q4-major)
8.3.0.20190709-git
Copyright (C) 2019 Free Software Foundation, Inc.
License GPLv3+: GNU GPL version 3 or later <http://gnu.org/licenses/gpl.html>
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.
Type "show copying" and "show warranty" for details.
This GDB was configured as "--host=i686-w64-mingw32 --target=arm-none-eabi".
Type "show configuration" for configuration details.
For bug reporting instructions, please see:
<http://www.gnu.org/software/gdb/bugs/>
Find the GDB manual and other documentation resources online at:
<http://www.gnu.org/software/gdb/documentation/>
For help, type "help". Type "apropos word" to search for commands related to
"word"...
Reading symbols from soem_gpio_pulse.elf...
(gdb)
```

7. Connect to the GDB server and load the binary by running the following commands:

```
a. target remote <GDB Server IP>:2331
b. monitor reset <p>
c. monitor halt
```

```
d. load
```

```
(gdb) target remote localhost:2331
Remote debugging using localhost:2331
0x00000008 in __isr_vector ()
(gdb) monitor reset Resetting target
(gdb) monitor halt
(gdb) load
Loading section .interrupts, size 0x240 lma 0x0
Loading section .text, size 0x3ab8 lma 0x240
Loading section .ARM, size size 0x8 lma 0x3cf8
Loading section .init_array, size 0x4 lma 0x3d00
Loading section .fini_array, size 0x4 lma 0x3d04
Loading section .data, size 0x64 lma 0x3d08
Start address 0x2f4, load size 15724
Transfer rate: 264 KB/sec, 2620 bytes/write.
(gdb)
```

8. The application now downloads and stops at the reset vector. Start the demo application by executing the `monitor go` command:

```
(gdb) monitor go
```

6.1.4.2.5 Running SOEM demo images by U-Boot

- **Running on Cortex-M**

The steps to write SOEM demo image file to TCM or DRAM using the Real-time Edge image on i.MX 8M Mini LPDDR4 EVK are same as followed for i.MX 8M Plus LPDDR4 EVK board. Refer Section 6.1.4.1.4 "Running SOEM demo images by U-Boot".

- **Running on Cortex-A**

```
=> ext4load mmc 1:2 0x93C00000 /examples/heterogeneous-multicore/soem-digital-
io-ca/ddr_release/soem_digital_io_ca53.bin;
=> dcache flush; icache flush;
=> cpu 2 release 0x93C00000
```

Or

```
=> ext4load mmc 1:2 0x93C00000 /examples/heterogeneous-multicore/soem-servo-ca/
ddr_release/soem_servo_ca53.bin;
=> dcache flush; icache flush;
=> cpu 2 release 0x93C00000
```

Or

```
=> ext4load mmc 1:2 0x93C00000 /examples/heterogeneous-multicore/soem-servo-
rt1180-ca/ddr_release/soem_servo_ca53.bin;
=> dcache flush; icache flush;
=> cpu 2 release 0x93C00000
```

6.1.4.3 SOEM for i.MX 93 EVK platform

6.1.4.3.1 Setting up the hardware environment

6.1.4.3.1.1 digital_io example

For **digital_io** example, the below hardware are required:

- i.MX 93 EVK platform
- EK1100 - EtherCAT Coupler
- EL2008 - EtherCAT Terminal, 8-channel digital output, 24V
- DCEL1018 - EtherCAT Terminal, 8-channel digital input, 24 V DC
- 24 V DC Power Supply



Figure 125. digital_io example for i.MX 93-EVK board

The channel-1 of the EL1018 is connected to a button to generate a 24 V pulse as a input signal to change the output of the channel-1 of the EL2008 labeled as "Dir". The channel-2 of the EL2008 outputs a square wave signal with a period of 250 μ s.

6.1.4.3.1.2 servo_motor and servo_motor_rt1180 examples

For **servo_motor** and **servo_motor_rt1180** example, the below hardware are required:

- i.MX 93 EVK platform
- Inovance SV680 / i.MX RT1180 EVK platform
- Network cable

There are two ways to build the example:

1. Inovance SV680 as EtherCAT SubDevice. Running SOEM Servo Motor example on i.MX 93 EVK to control Inovance SV680.




Figure 126. Inovance SV680 as EtherCAT SubDevice

Or

- Running `ecat_examples/servo_motor` example on i.MX RT1180 EVK as EtherCAT SubDevice. Running SOEM Servo Motor example on i.MX 93 EVK to control motor on i.MX RT1180 EVK.




Figure 127. `servo_motor` example on i.MX RT1180 EVK as EtherCAT SubDevice

Note: Before running the `servo_motor` example, you must download the MCU-SDK package for i.MX RT1180-EVK board from the NXP official website: [Welcome | MCUXpresso SDK Builder \(nxp.com\)](https://www.nxp.com/developers/boards-and-modules/boards-and-modules/rt1180-evk)

6.1.4.3.2 Building the demo images of Cortex-M

For `servo_motor` and `servo_motor_rt1180` examples, the demo images "soem_servo_motor" and "soem_servo_motor_rt1180" are compiled along with the i.MX 93 EVK target image. These images are installed in the "/example" directory of the target rootfs. For JTAG download, these images can also be found in the directory "<image-build-dir>/tmp/deploy/images/imx93evk/examples/" on building host.

Note: "soem_servo_motor" are for "Inovance SV680" case, "soem_servo_motor_rt1180" are for "i.MX RT1180 EVK" case.

```
cd <image-build-dir>/
tree tmp/deploy/images/imx93evk/examples/mcux-sdk/soem-servo-motor
|— release
|   |— soem_servo_motor.elf
|   |— soem_servo_motor.bin
```



```
cd <image-build-dir>/
tree tmp/deploy/images/imx93evk/examples/mcux-sdk/soem-servo-motor-rt1180
|— release
|   └─ soem_servo_motor.elf
|   └─ soem_servo_motor.bin
```

Also, the demo images can be compiled using the below command on i.MX 93 EVK build directory:

```
bitbake soem-servo-motor
or
bitbake soem-servo-motor-rt1180
```

6.1.4.3.3 Building the demo images of Cortex-A

6.1.4.3.3.1 digital_io example

For **digital_io** example, the demo images "soem_digital_io_ca55" are compiled with the i.MX 93 EVK target image compiling, and they are installed into the directory "/example" of the target rootfs. Images can also be found in the directory "<image-build-dir>/tmp/deploy/images/imx93evk/examples/" on building host.

```
cd <image-build-dir>/
tree tmp/deploy/images/imx93evk/examples/heterogeneous-multicore/soem-digital-
io-ca
|— ddr_release
|   └─ soem_digital_io_ca55.elf
|   └─ soem_digital_io_ca55.bin
```

Also, the demo images can be compiled using the below command on i.MX 93 EVK build directory:

```
bitbake soem-digital-io-ca
```

6.1.4.3.3.2 servo_motor and servo_motor_rt1180 examples

For **servo_motor** and **servo_motor_rt1180** example, the demo images "soem_servo_ca55", "soem_servo_rt1180_ca55" are compiled with the i.MX 93 EVK target image compiling, and they are installed into the directory "/example" of the target rootfs. Images can also be found in the directory "<image-build-dir>/tmp/deploy/images/imx93evk/examples/" on building host.

Note: "soem_servo_ca55" are for "Inovance SV680" case, "soem_servo_rt1180_ca55" are for "i.MX RT1180 EVK" case.

```
cd <image-build-dir>/
tree tmp/deploy/images/imx93evk/examples/heterogeneous-multicore/soem-servo-ca
|— ddr_release
|   └─ soem_servo_ca55.elf
|   └─ soem_servo_ca55.bin
tree tmp/deploy/images/imx93evk/examples/heterogeneous-multicore/soem-servo-
rt1180-ca
|— ddr_release
|   └─ soem_servo_rt1180_ca55.elf
|   └─ soem_servo_rt1180_ca55.bin
```


Also, the demo images can be compiled using the below command on i.MX 93 EVK build directory:

```
bitbake soem-servo-ca
or
bitbake soem-servo-rt1180-ca
```

6.1.4.3.4 Running SOEM demo images by using U-Boot

This section describes the steps to write an SOEM demo image file to the TCM or DRAM with the Real-time Edge image. The following steps describe how to use the U-Boot:

1. Connect the DEBUG UART slot on the board to your PC through the USB cable. The Windows OS installs the USB driver automatically, and the Ubuntu OS finds the serial devices as well.
 - a. On Windows OS, open the Device Manager, find USB serial Port in Ports (COM and LPT). Assume that the ports are COM9 and COM10. One port is for the debug message from the Cortex-A55 and the other is for the Cortex-M7. The port number is allocated randomly, so opening both is beneficial for development. On Ubuntu OS, find the TTY device with name /dev/ttyUSB* to determine your debug port. Similar to Windows OS, opening both is beneficial for development.
 - b. Build and flash the `nxp-image-real-time-edge` image to a SD card and insert the SD card to the target board. Make sure to use the default boot SD slot and check the DIP boot switch configuration.
 - c. Open your preferred serial terminal for the serial devices. Set the speed to 115200 bps, 8 data bits, 1 stop bit (115200, 8N1), no parity. Then, power on the board.
 - d. Power on the board and hit any key to stop autoboot in the terminal window. Then enter the U-Boot Command line mode. You can then write the image and run it from TCM or DRAM with the following commands:
2. **Running on Cortex-M** If the `soem_servo_motor.bin` is made from the `<examples/..>/release` target, which means the binary file runs at TCM, use the following commands to boot:

```
=> ext4load mmc 1:2 0xd0000000 /examples/mcux-sdk/soem-servo-motor/release/
soem_servo_motor.bin;
=> cp.b 0xd0000000 0x201e0000 20000;
=> bootaux 0x1ffe0000
```

Or

```
=> ext4load mmc 1:2 0xd0000000 /examples/mcux-sdk/soem-servo-motor-rt1180/
release/soem_servo_motor.bin;
=> cp.b 0xd0000000 0x201e0000 20000;
=> bootaux 0x1ffe0000
```

3. Running on Cortex-A

```
=> ext4load mmc 1:2 0xD0000000 /examples/heterogeneous-multicore/soem-
digital-io-ca/ddr_release/soem_digital_io_ca55.bin;
=> dcache flush; icache flush;
=> cpu 1 release 0xD0000000
```

Or

```
=> ext4load mmc 1:2 0xD0000000 /examples/heterogeneous-multicore/soem-servo-
ca/ddr_release/soem_servo_ca55.bin;
=> dcache flush; icache flush;
=> cpu 1 release 0xD0000000
```

Or

```
=> ext4load mmc 1:2 0xD0000000 /examples/heterogeneous-multicore/soem-servo-
rt1180-ca/ddr_release/soem_servo_ca55.bin;
```

```
=> dcache flush; icache flush;  
=> cpu 1 release 0xD0000000
```

6.1.5 CODESYS EtherCAT MainDevice

6.1.5.1 Overview

Based on Real-time Edge yocto project, a new yocto distro named "nxp-real-time-edge-plc" is added, which is specific to the PLC use case. This document describes how to build a Real-time Edge PLC image and how to setup a CODESYS test project as an example to drive motors.

- For more information about i.MX Yocto project, refer to: https://www.nxp.com.cn/docs/en/user-guide/IMX_YOCTO_PROJECT_USERS_GUIDE.pdf
- For additional information, refer to *Real Time Edge Yocto User Guide (RTEDGEYOCTOUG)* available on the URL: [REALTIME EDGE Documentation](#).

6.1.5.2 Features

Features are as follows:

- **Optimized native driver**

Typical industrial software, for example CODESYS or SOEM EtherCAT MainDevice stack, work on user space and communicate using Linux standard network interface. To reduce the latency when Ethernet raw packets pass from user space through Linux standard network, the network driver is optimized to avoid memory reallocation and copying, task rescheduling. The latency of the critical path is reduced for packet transmitting and receiving.

- **Light root filesystem**

Lighter root filesystem saves CPU cycles that could be used by the PLC user program.

- **Supported platforms**

- imx6ull14x14evk
- imx8mm-lpddr4-evk
- imx8mp-lpddr4-evk
- imx93evk

6.1.5.3 Building the image

This section provides detailed information along with the procedure for building an image.

6.1.5.3.1 Build configurations

A new yocto distro named "nxp-real-time-edge-plc" is added for the PLC use case, and below platforms are supported:

- imx6ull14x14evk
- imx8mm-lpddr4-evk
- imx8mp-lpddr4-evk
- imx93evk

Real-time Edge provides the script `real-time-edge-setup-env.sh` to simplify the setup for both i.MX and Layerscape boards. To use the script, the name of the specific machine to be built for and the desired distro must be specified. The script sets up a directory and the configuration files for the specified machine and distro.

The syntax for the `real-time-edge-setup-env.sh` script is shown below:

```
$ DISTRO=nxp-real-time-edge-plc MACHINE=<machine name> source real-time-edge-setup-env.sh -b <build dir>
```

where:

- **MACHINE=<machine name>** is the name of the supported platforms.
- **-b <build dir>** specifies the name of the build directory created by the `real-time-edge-setup-env.sh` script.

After the script runs, the working directory is the one just created by the script, specified with the `-b` option. A `conf` folder is created containing the files `bblayers.conf` and `local.conf`.

The `local.conf` file contains the machine and distro specifications. An example is shown below:

```
MACHINE ??= 'imx8mp-lpddr4-evk'
DISTRO ?= 'nxp-real-time-edge-plc'
ACCEPT_FSL_EULA = "1"
```

The MACHINE configuration can be changed by editing this file, if necessary.

6.1.5.3.2 Build scenarios

The following are build setup scenarios for various configurations. Set up the manifest and populate the Yocto Project layer sources using the commands below:

```
$ mkdir yocto-real-time-edge
$ cd yocto-real-time-edge
$ repo init -u https://github.com/nxp-real-time-edge-sw/yoctoreal-time-edge.git/
yocto-real-time-edge.git \
-b real-time-edge-scarthgap \
-m real-time-edge-3.1.0.xml
$ repo sync
```

The following sections give some specific examples.

• Real-time Edge PLC image on i.MX 6ULL EVK

```
$ DISTRO=nxp-real-time-edge-plc MACHINE=imx6ull14x14evk source real-time-edge-setup-env.sh -b build-imx-real-time-edge-plc
$ bitbake nxp-image-real-time-edge-plc
```

• Real-time Edge PLC image on i.MX 8M Mini EVK

```
$ DISTRO=nxp-real-time-edge-plc MACHINE=imx8mm-lpddr4-evk source real-time-edge-setup-env.sh -b build-imx-real-time-edge-plc
$ bitbake nxp-image-real-time-edge-plc
```

• Real-time Edge PLC image on i.MX 8M Plus EVK

```
$ DISTRO=nxp-real-time-edge-plc MACHINE=imx8mp-lpddr4-evk source real-time-edge-setup-env.sh -b build-imx-real-time-edge-plc
$ bitbake nxp-image-real-time-edge-plc
```

• Real-time Edge PLC image on i.MX 93 EVK

```
$ DISTRO=nxp-real-time-edge-plc MACHINE=imx93evk source real-time-edge-setup-env.sh -b build-imx-real-time-edge-plc
```

```
$ bitbake nxp-image-real-time-edge-plc
```

6.1.5.3.3 Image deployment

All filesystem images are deployed to the `<build_directory>/tmp/deploy/images` directory. Each image building creates a U-Boot image, a kernel image, and an image type based on the `IMAGE_FSTYPES` variable defined in the machine configuration file. Most machine configurations provide an SD card image (.wic) and a rootfs image (.tar). The SD card image contains a partitioned image (with U-Boot, kernel, rootfs, and other such files) suitable for booting the corresponding hardware.

Copy image to SD card:

The SD card image file, (`<image_name>.wic`) contains a partitioned image (with U-Boot, kernel, rootfs, and other files) suitable for booting the corresponding hardware. Run the following commands to copy the image to an SD card:

```
$ zstd -d <image_name>.wic.zst
$ sudo dd if=<image_name>.wic of=/dev/sd<disk> bs=1M conv=fsync
```

6.1.5.4 Running CODESYS on the boards

6.1.5.4.1 Running CODESYS on i.MX 6ULL

Execute the following command to update the dtb file and boot into kernel.

```
uboot ==> setenv fdt_file imx6ull-14x14-evk-ecat.dtb
uboot ==> run bootcmd
```

Set the CPU frequency governor to “*performance*” for **CPU** core to use the highest frequency.

```
# echo performance > /sys/devices/system/cpu/cpufreq/policy0/scaling_governor
```

Enable the EtherCAT port.

```
# ethercat_port=`ls /sys/bus/platform/devices/2188000.ethernet/net` && echo
$ethercat_port
# ifconfig $ethercat_port up
```

Assign IP address for the Ethernet port automatically if there is a DHCP server in the network.

```
# ethernet_port=`ls /sys/bus/platform/devices/20b4000.ethernet/net` && echo
$ethernet_port
# udhcpc -i $ethernet_port
```

Else assign IP address manually.

```
# ethernet_port=`ls /sys/bus/platform/devices/20b4000.ethernet/net` && echo
$ethernet_port
# ifconfig $ethernet_port 192.168.1.xx
```

Use **scp** command to copy **codesyscontrol.bin** into the board and start it:

```
# ./codesyscontrol.bin > CODESYS.log &
```

Note:

- Please refer to the next section to obtain `codesyscontrol.bin`.
- CODESYS logs can be viewed in the `CODESYS.log` file.

6.1.5.4.2 Running CODESYS on i.MX 8M Mini and i.MX 8M Plus

Execute the following command to isolate CPU1 (by default, the “EtherCAT_Task” thread created by the Realtime of CODESYS is bound into CPU1).

```
uboot ==> setenv mmcargs 'setenv bootargs ${jh_clk} ${mcore_clk} console=
${console} root=${mmccroot} isolcpus=1'
```

Execute the following command to update the dtb file and boot into kernel.

```
// on i.MX 8M Mini
uboot ==> setenv fdtfile imx8mm-evk-ecat.dtb
// on i.MX 8M Plus
uboot ==> setenv fdtfile imx8mp-evk-ecat.dtb

uboot ==> run bootcmd
```

Set the CPU frequency governor to “*performance*” for **CPU1** core to use the highest frequency.

```
echo performance > /sys/devices/system/cpu/cpu1/cpufreq/scaling_governor
```

Stop the “**udhcpc**” service which interface is connected to the EtherCAT servos.

```
# ethercat_net=`ls /sys/bus/platform/devices/30be0000.ethernet/net` && echo
$ethercat_net
# ps
// find "PID root 3468 S udhcpc -R -b -p /var/run/udhcpc.eth0.pid -i
$ethercat_net"
# kill -9 PID
```

Enable the EtherCAT port.

```
# ethercat_port=`ls /sys/bus/platform/devices/30be0000.ethernet/net` && echo
$ethercat_port
# ifconfig $ethercat_port up
```

Assign IP address for the Ethernet port automatically if there is a DHCP server in the network.

```
// assume the name of the USB network port is eth1
# udhcpc -i eth1
```

Else assign IP address manually.

```
// assume the name of the USB network port is eth1
# ifconfig eth1 192.168.1.xx
```

Use **scp** command to copy **codesyscontrol.bin** into the board and start it:

```
# ./codesyscontrol.bin > CODESYS.log &
```

Note:

- Please refer to the next section to obtain `codesyscontrol.bin`.
- CODESYS logs can be viewed in the `CODESYS.log` file.

6.1.5.4.3 Running CODESYS on i.MX 93

Execute the following command to isolate CPU1 (by default, the “EtherCAT_Task” thread created by the Realtime of CODESYS is bound into CPU1).

```
uboot ==> setenv mmccargs 'setenv bootargs ${jh_clk} ${mcore_clk} console=
${console} root=${mmccroot} isolcpus=1'
```

Execute the following command to update the dtb file and boot into kernel.

```
uboot ==> setenv fdtfile imx93-11x11-evk-ecat.dtb
uboot ==> run bootcmd
```

Stop the “udhcpc” service which interface is connected to the EtherCAT servos.

```
# ethercat_net=`ls /sys/bus/platform/devices/42890000.ethernet/net` && echo
$ethercat_net
# ps
// find "PID root 3468 S udhcpc -R -b -p /var/run/udhcpc.eth0.pid -i
$ethercat_net"
# kill -9 PID
```

Enable the EtherCAT port.

```
# ethercat_port=`ls /sys/bus/platform/devices/42890000.ethernet/net` && echo
$ethercat_port
# ifconfig $ethercat_port up
```

Assign IP address for the Ethernet port automatically if there is a DHCP server in the network.

```
# ethernet_port=`ls /sys/bus/platform/devices/428a0000.ethernet/net` && echo
$ethernet_port
# udhcpc -i $ethernet_port
```

Else assign IP address manually.

```
# ethernet_port=`ls /sys/bus/platform/devices/428a0000.ethernet/net` && echo
$ethernet_port
# ifconfig $ethernet_port 192.168.1.xx
```

Execute the following command to disable the "CPU1" idle state1.

```
# echo "1" > /sys/devices/system/cpu/cpu1/cpuidle/state1/disable
```

Execute the following command to disable the DDR auto clock gating.

```
# echo 0 > /sys/devices/platform/imx93-lpm/auto_clk_gating
```

Use **scp** command to copy **codesyscontrol.bin** into the board and start it:

```
# ./codesyscontrol.bin > CODESYS.log &
```

Note:

- Please refer to the next section to obtain `codesyscontrol.bin`.
- CODESYS logs can be viewed in the `CODESYS.log` file.

6.1.5.5 Setting up CODESYS

CODESYS is a powerful commercial PLC software programming tool. It supports IEC61131-3 standard IL, ST, FBD, LD, CFC, SFC, and six kinds of PLC programming languages. Users can choose different language editing subroutines in the same project, function module, and so on.

6.1.5.5.1 Board environment setup**i.MX 6ULL EVK Boards**

- **Hardware requirements**
 - Two network cables
 - Mini/micro USB cable
 - One servo motor (DELTA ASDA-B3)
 - Personal Computer on which the CODESYS has been installed
- **Preparing the example**
 - Connect the PC and the network port labeled **J1501 B** on the board.
 - Connect the servo motors to the network port labeled **J1501 A** on the board.
 - Connect a USB cable between the host PC and the OpenSDA USB port on the target board. Then, open a serial terminal with the following settings
 - 115200 baud rate
 - 8 data bits
 - No parity
 - One stop bit
 - No flow control

i.MX 8M Mini EVK Boards

- **Hardware requirements**
 - Two Network cables
 - Mini/micro USB cable
 - Network port converter with type-C interface

Note: *i.MX 8M Mini has only one network port, so it needs another network port connected with CODESYS IDE.*

 - One servo motor (DELTA ASDA-B3)
 - Personal Computer on which the CODESYS has been installed.
- **Preparing the example**
 - Connect the servo and the network port of the board.
 - Plug the Network port converter into the USB port (labeled PORT1) of the board, and then connect the Network port converter to the PC using a network cable.
 - Connect a USB cable between the host PC and the OpenSDA USB port on the target board, open a serial terminal with the following settings:
 - 115200 baud rate
 - 8 data bits
 - No parity
 - One stop bit

- No flow control

The connection diagram for i.MX 8M Mini board is as follows.



Figure 128. i.MX 8M Mini board connection

i.MX 8M Plus EVK Boards

- **Hardware requirements**

- Two Network cables
- Mini/micro USB cable
- One servo motor (DELTA ASDAB3)
- Personal Computer on which the CODESYS has been installed.

- **Preparing the example**

- Connect the PC to the network port labeled **ENET2** on the board.
- Connect the servo motors to the network port labeled **ENET1** on the board.
- Connect a USB cable between the host PC and the OpenSDA USB port on the target board, and open a serial terminal with the following settings:
 - 115200 baud rate
 - 8 data bits
 - No parity
 - One stop bit
 - No flow control

i.MX 93 EVK Boards

- **Hardware requirements**

- Two Network cables
- Mini/micro USB cable
- One servo motor(DELTA ASDAB3)

- Personal Computer on which the CODESYS has been installed
- **Preparing the example**
 - Connect the PC to the network port labeled **ENET1** on the board.
 - Connect the servo motors to the network port labeled **ENET2** on the board.
 - Connect a USB cable between the host PC and the OpenSDA USB port on the target board, and open a serial terminal with the following settings:
 - 115200 baud rate
 - 8 data bits
 - No parity
 - One stop bit
 - No flow control

6.1.5.5.2 CODESYS project setup

6.1.5.5.2.1 Downloading CODESYS runtime binary

Download the corresponding runtime binary file from CODESYS website and decompress it using the commands below:

```
# unzip CODESYS\ Control\ for\ Linux\ ARM64\ SL\ 4.7.0.0.package CODESYS.package
# cd Delivery/
# dpkg-deb -R codesyscontrol_linuxarm64_4.7.0.0-b.trunk.39_arm64.deb software
# cd software/opt/codesys/bin/
```

The runtime binary `codesyscontrol.bin` is located in the `software/opt/codesys/bin/` directory.

Note:

- **Download 32-bit version for the platform:**
 - i.MX 6ULL EVK
- **Download 64-bit version for the platforms:**
 - i.MX 8M Mini EVK
 - i.MX 8M Plus EVK
 - i.MX 93 EVK

6.1.5.5.2.2 Starting CODESYS runtime

Use the **scp** command to copy **codesyscontrol.bin** into the board and start it:

```
# ./codesyscontrol.bin > CODESYS.log &
```

CODESYS logs can be viewed in the CODESYS.log file.

6.1.5.5.2.3 Creating a new CODESYS project

This section describes the steps to create the CODESYS project in the machine on which the CODESYS tool has been installed.

- Open **CODESYS** and build a new project using the steps below:
 1. Click “**File**” -> “**New project**”.

2. Select **“Standard project”** as the template.
3. Rename this project.
4. Click **“OK”** to complete it.

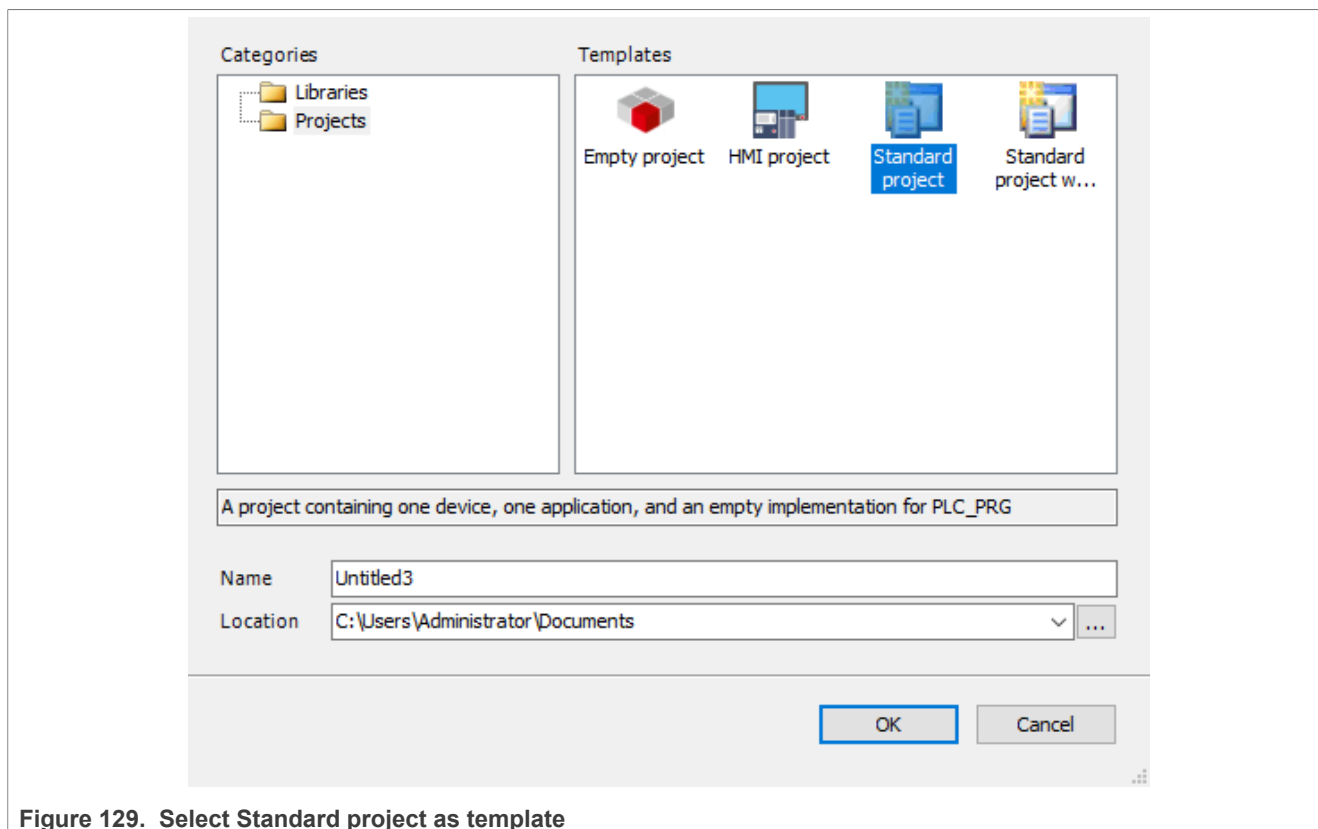


Figure 129. Select Standard project as template

• Set Device and PLC_PRG

1. Select **“CODESYS Control for Linux ARM64 SL”** or **“CODESYS Control for Linux ARM SL”** as the Device.
2. Select **“Structured Text(ST)”** as the PLC_PRG.
3. Click **“OK”** to complete it.

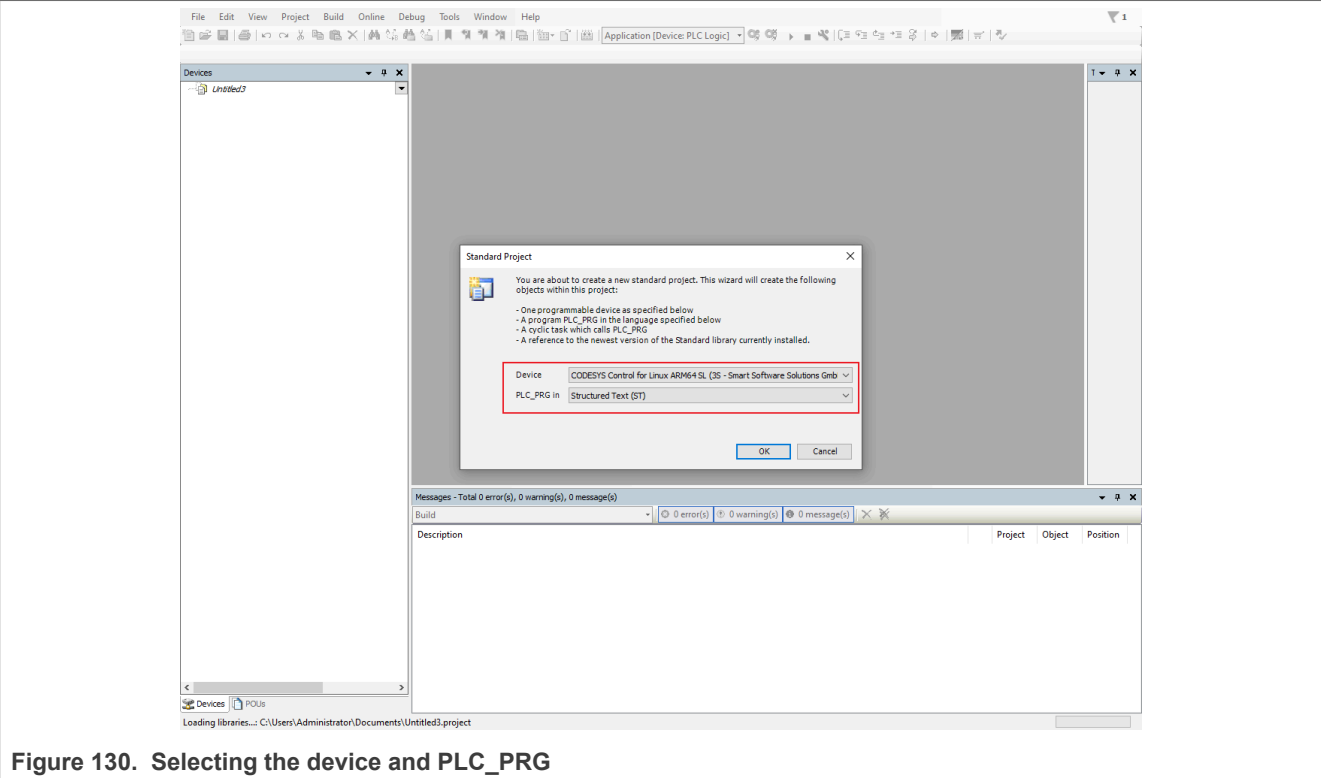


Figure 130. Selecting the device and PLC_PRG

6.1.5.5.2.4 Adding EtherCAT MainDevice and SubDevice

- Set EtherCAT MainDevice
 1. Right-click on the device -> **“Add Device”**
 2. Select **“EtherCAT Master SoftMotion”**
 3. Click **“Add Device”** to complete it.

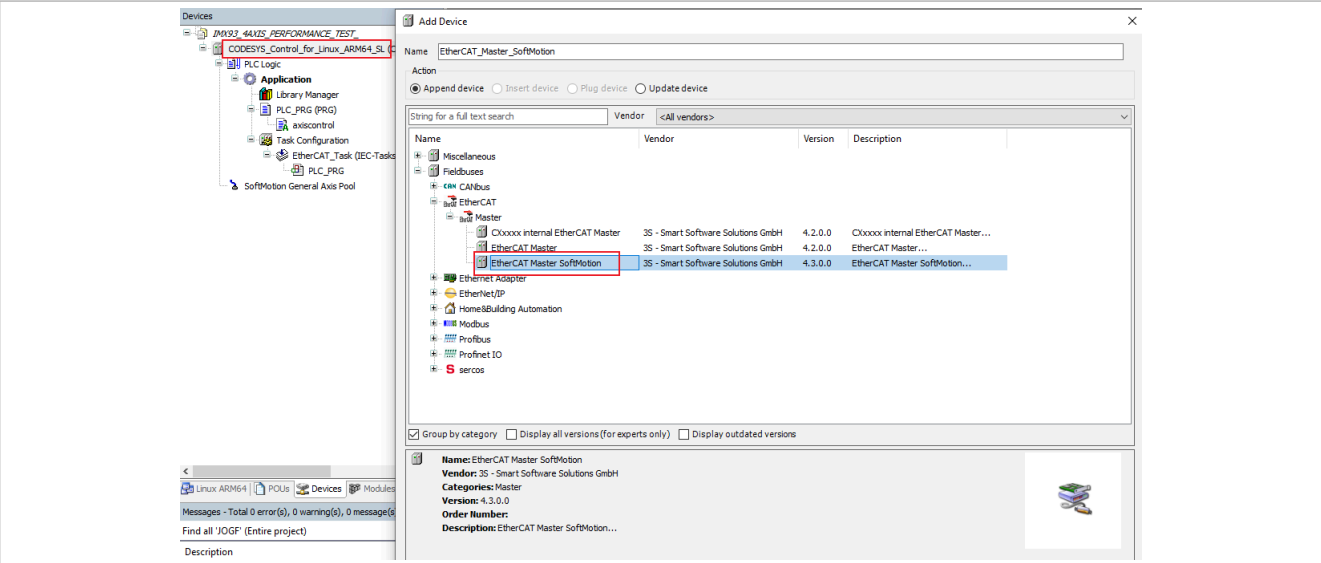


Figure 131. Selecting EtherCAT MainDevice

- Download DELTA ASDA-B3 device description file

https://downloadcenter.deltaww.com/en-US/DownloadCenter?v=1&q=ASDA-A3-E%2F%20B3-E%20XML%20File%20v0.04&sort_expr=cdate&sort_dir=DESC

- **Add DELTA device description file**

1. Click **“Tools”** -> **“Device Repository”**.
2. Click **“Install”**.
3. Select **“Delta ASDA-x3-E rev0.04_10EMC.XML”** file.
4. Click **“Close”** to complete it.

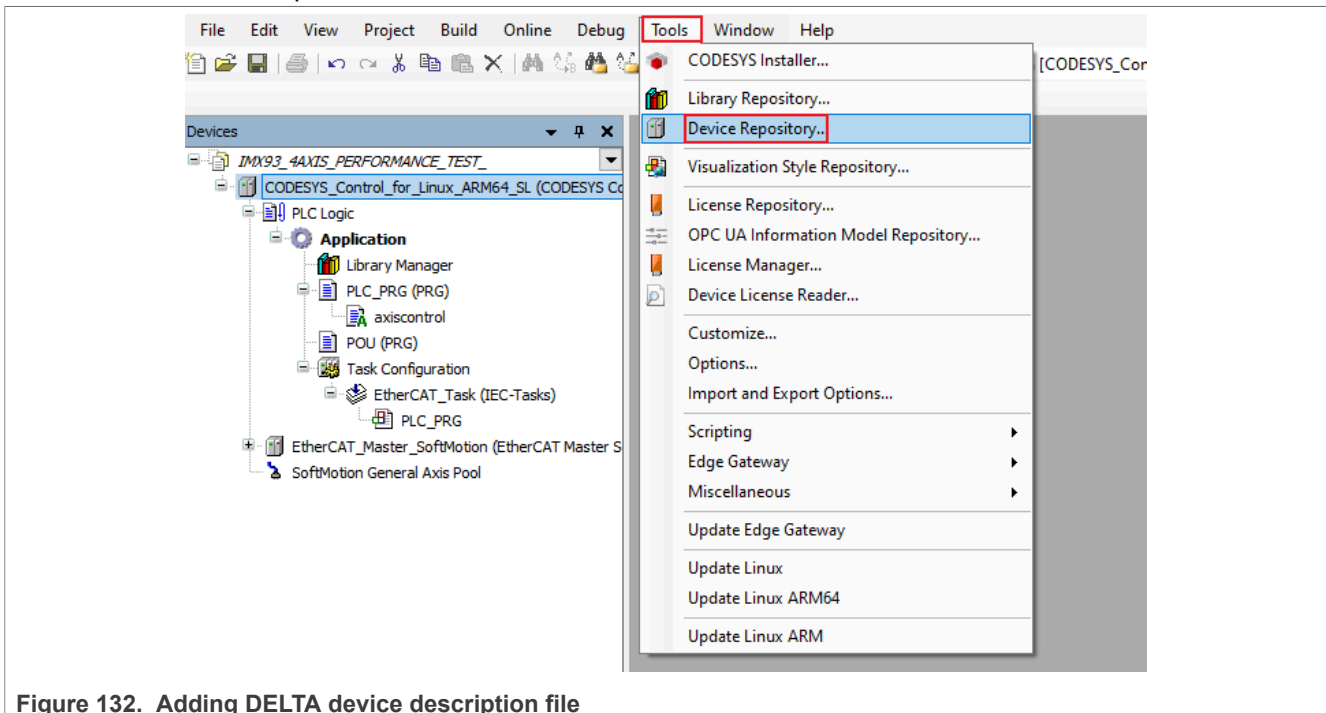


Figure 132. Adding DELTA device description file




Figure 133. View Installed devices

- **Add servo device**
 1. Right-click on the “**EtherCAT_Master_SoftMotion**” -> “**Add Device**”.
 2. Select the corresponding servo.
 3. Click “**Add Device**” to complete it.




Figure 134. Adding a Device




Figure 135. Select the device to be added

- Add CiA402 Axis device
1. Right-click on the servo -> “Add SoftMotion CiA402 Axis”.




Figure 136. Adding CiA402 Axis device

6.1.5.5.2.5 Adding PLC program

The below steps describe how to add the PLC program to the drive motor:

1. Right-click on the “**Application**” -> “**Add Object**” -> “**POU**” and name it “**PLC_PRG**”. (If it already exists, do not add it.) -> Pull the **PLC_PRG** under the **EtherCAT_Task**
2. Right-click on the “**PLC_PRG**” -> “**Add Object**” -> “**action**” and name it “**axiscontrol**”
3. Fill in the following PLC code:

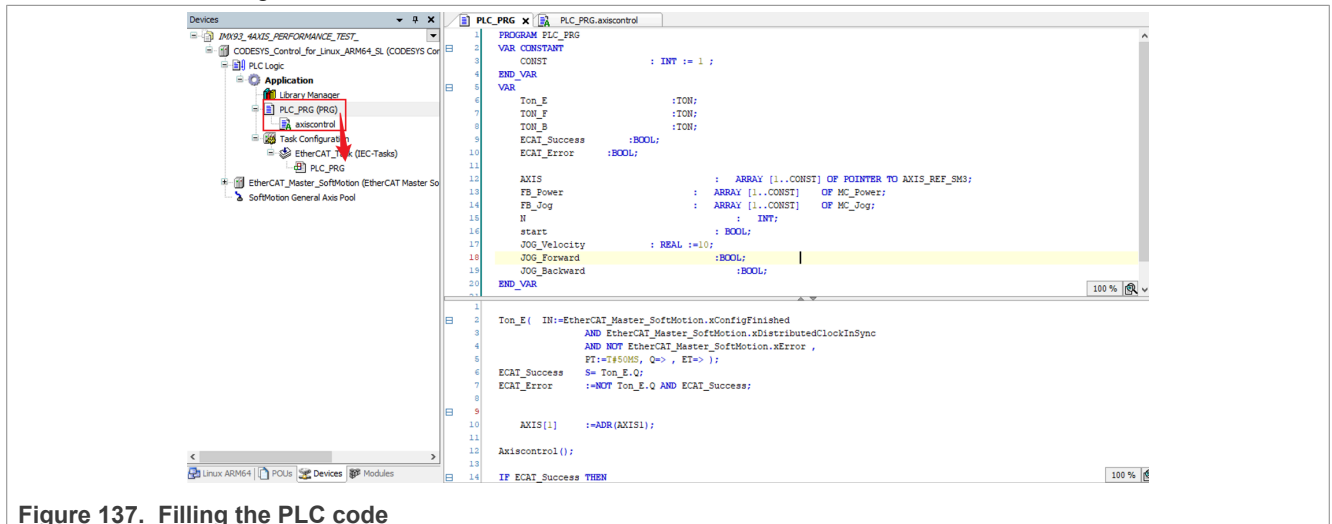


Figure 137. Filling the PLC code

The **PLC_PRG (PRG)** variable is as follows:

```
PROGRAM PLC_PRG
VAR CONSTANT
CONST : INT := 1; // The number of belt axes can be modified from here
END_VAR
VAR
Ton_E : TON;
TON_F : TON;
TON_B : TON;
ECAT_Success : BOOL;
ECAT_Error : BOOL;
AXIS : ARRAY [1..CONST] OF POINTER TO AXIS_REF_SM3;
FB_Power : ARRAY [1..CONST] OF MC_Power;
FB_Jog : ARRAY [1..CONST] OF MC_Jog;
N : INT;
start : BOOL;
JOG_Velocity : REAL := 10;
JOG_Forward : BOOL;
JOG_Backward : BOOL;
END_VAR
```

The **PLC_PRG (PRG)** code is as follows:

```
// MainDevice station operation judgment
Ton_E ( IN:=EtherCAT_Master_SoftMotion.xConfigFinished
AND EtherCAT_Master_SoftMotion.xDistributedClockInSync
AND NOT EtherCAT_Master_SoftMotion.xError ,
PT:=T#50MS, Q=> , ET=> );
ECAT_Success S= Ton_E.Q;
```

```

ECAT_Error :=NOT Ton_E.Q AND ECAT_Success;
// Axis pointer acquisition
AXIS[1] :=ADR(SM_Drive_GenericDSP402);
Axiscontrol();
IF ECAT_Success THEN
    start :=TRUE;
ELSE
    start :=FALSE;
END_IF
IF JOG_Forward=FALSE AND JOG_Backward =FALSE THEN
    JOG_Backward :=TRUE;
END_IF
TON_F (IN:= JOG_Forward, PT:= T#10S, Q=> , ET=> );
TON_B (IN:= JOG_Backward, PT:= T#10S, Q=> , ET=> );
IF TON_F.Q THEN
    JOG_Forward :=FALSE;
    JOG_Backward :=TRUE;
ELSIF TON_B.Q THEN
    JOG_Forward :=TRUE;
    JOG_Backward :=FALSE;
END_IF

```

The axis control code is as follows:

```

FOR N :=1 TO CONST BY 1 DO
    FB_Power[N] (
        Axis:= AXIS[N]^,
        Enable:= TRUE,
        bRegulatorOn:= start,
        bDriveStart:= TRUE,
        Status=> ,
        bRegulatorRealState=> ,
        bDriveStartRealState=> ,
        Busy=> ,
        Error=> ,
        ErrorID=> );
    FB_Jog[N] (
        Axis:= AXIS[N]^,
        JogForward:= FB_Power[N].Status AND JOG_Forward,
        JogBackward:= FB_Power[N].Status AND JOG_Backward,
        Velocity:= JOG_Velocity,
        Acceleration:= JOG_Velocity *10,
        Deceleration:= JOG_Velocity *10,
        Jerk:= ,
        Busy=> ,
        CommandAborted=> ,
        Error=> ,
        ErrorId=> );
END_FOR

```

6.1.5.5.2.6 Running CODESYS Project

1. In the Codesys_Control window, fill in the IP address of the board in respective textbox. Then, press **Enter** to connect to the board as shown in the figure below:

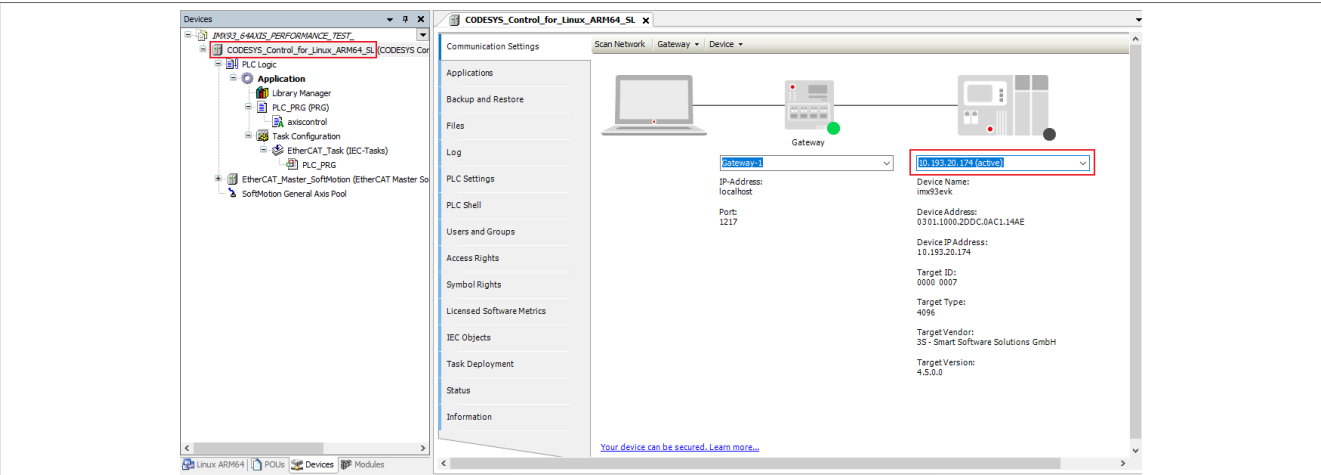


Figure 138. Adding IP address of the board

2. Select the network port on the board that is used to communicate with the servo.




Figure 139. Selecting the network port

Note: For network port selection, refer to [Section 6.1.5.5.1](#).

3. Click “Login” and “Start” button to run the PLC program:

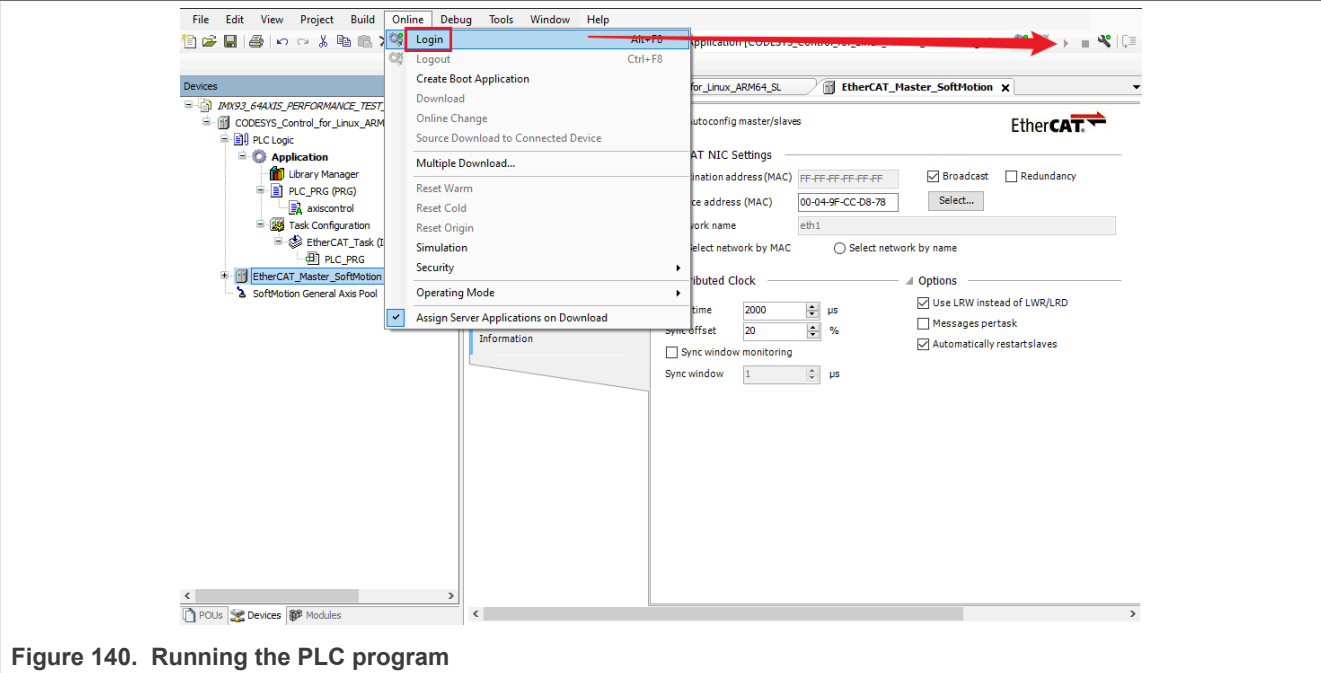


Figure 140. Running the PLC program

The motor will start to rotate after all steps above are done.

6.2 FlexCAN and CANopen

The following sections provide an introduction to the FlexCAN standard, details of the CAN bus, the CANopen communication system, details of how to integrate FlexCAN with Real-time Edge, and running a FlexCAN application.

6.2.1 Introduction

The LS1028ARDB board has the FlexCAN module. The FlexCAN module is a communication controller implementing the CAN protocol according to the CAN 2.0 B protocol specification. The main sub-blocks implemented in the FlexCAN module include an associated memory for storing message buffers, Receive (RX) Global Mask registers, Receive Individual Mask registers, Receive FIFO filters, and Receive FIFO ID filters. A general block diagram is shown in the following figure. The functions of these submodules are described in subsequent sections.

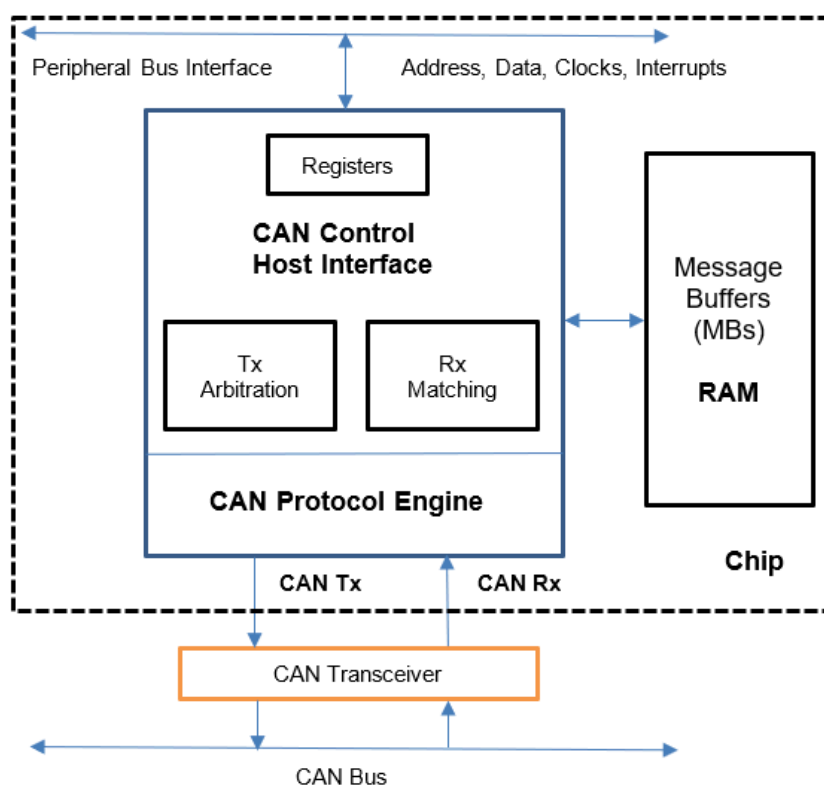


Figure 141. FlexCAN block diagram

6.2.1.1 CAN bus

CAN (Controller Area Network) is a serial bus system. A CAN bus is a robust [vehicle bus](#) standard designed to allow [microcontrollers](#) and devices to communicate with each other in applications without a [host computer](#). Bosch published several versions of the CAN specification and the latest is CAN 2.0 published in 1991. This specification has two parts; part A is for the standard format with an 11-bit identifier, and part B is for the extended format with a 29-bit identifier. A CAN device that uses 11-bit identifiers is commonly called CAN 2.0A and a CAN device that uses 29-bit identifiers is commonly called CAN 2.0B.

CAN is a [multi-master serial bus](#) standard for connecting Electronic Control Units [ECUs] also known as nodes. Two or more nodes are required on the CAN network to communicate. The complexity of the node can range from a simple I/O device up to an embedded computer with a CAN interface and sophisticated software. The node may also be a gateway allowing a standard computer to communicate over a USB or Ethernet port to the devices on a CAN network. All nodes are connected to each other through a two wire bus. The wires are a twisted pair with a 120 Ω (nominal) characteristic impedance.

High speed CAN signaling drives the CAN high wire towards 5 V and the CAN low wire towards 0 V when transmitting a dominant (0), and does not drive either wire when transmitting a recessive (1). The dominant differential voltage is a nominal 2 V. The termination resistor passively returns the two wires to a nominal differential voltage of 0 V. The dominant common mode voltage must be within 1.5 V to 3.5 V of common and the recessive common mode voltage must be within +/-12 V of common.

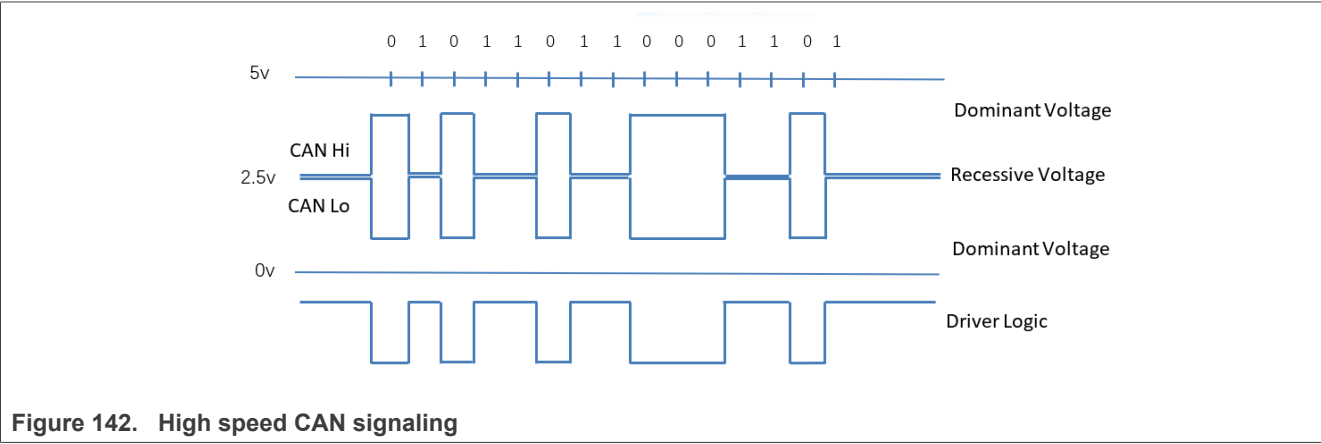


Figure 142. High speed CAN signaling

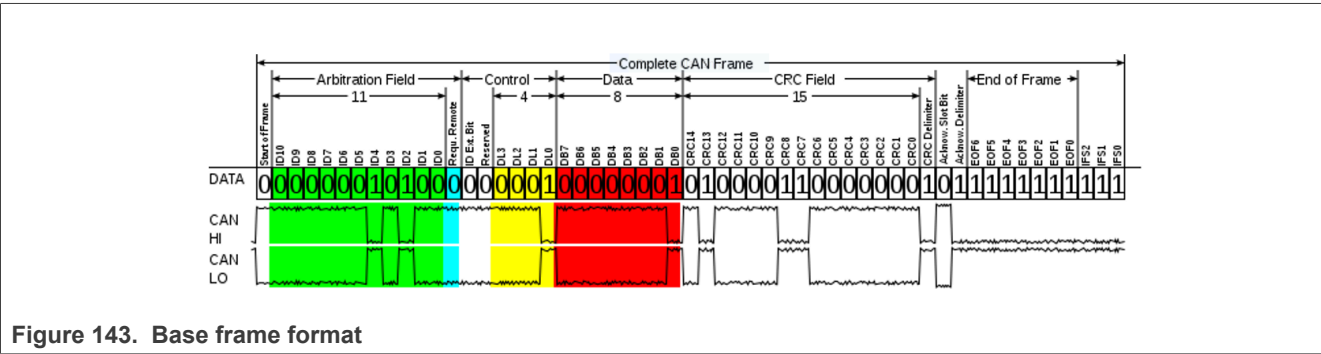


Figure 143. Base frame format

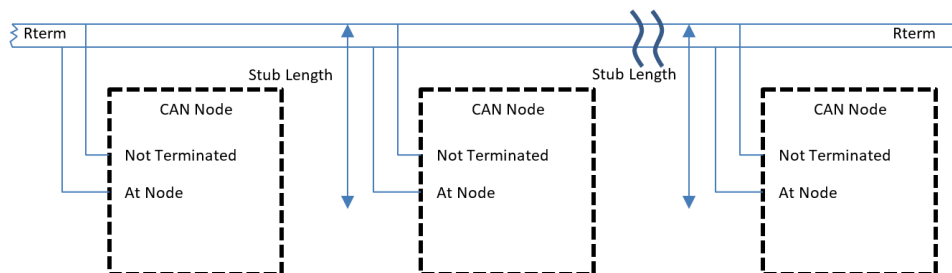


Figure 144. High speed CAN network

6.2.1.2 CANopen

CANopen is a CAN-based communication system. It comprises higher-layer protocols and profile specifications. CANopen has been developed as a standardized embedded network with highly flexible configuration capabilities. Today it is used in various application fields, such as medical equipment, off-road vehicles, maritime electronics, railway applications, and building automation.

CANopen provides several communication objects, which enable device designers to implement desired network behavior into a device. With these communication objects, device designers can offer devices that can communicate process data, indicate device-internal error conditions or influence and control the network behavior. As CANopen defines the internal device structure, the system designer knows exactly how to access a CANopen device and how to adjust the intended device behavior.

- **CANopen lower layers**

CANopen is based on a data link layer according to ISO 11898-1. The CANopen bit timing is specified in CiA 301 and allows the adjustment of data rates from 10 kbit/s to 1000 kbit/s. Although all specified CAN-ID addressing schemata are based on the 11-bit CAN-ID, CANopen supports the 29-bit CAN-ID as well. Nevertheless, CANopen does not exclude other physical layer options.

- **Internal device architecture**

A CANopen device consists of three logical parts. The CANopen protocol stack handles the communication via the CAN network. The application software provides the internal control functionality. The CANopen object dictionary interfaces the protocol as well as the application software. It contains indices for all used data types and stores all communication and application parameters. The CANopen object dictionary is most important for CANopen device configuration and diagnostics.

- **CANopen protocols**

- SDO protocol
- PDO protocol
- NMT protocol
- Special function protocols
- Error control protocols

The following figure shows the CANopen architecture.

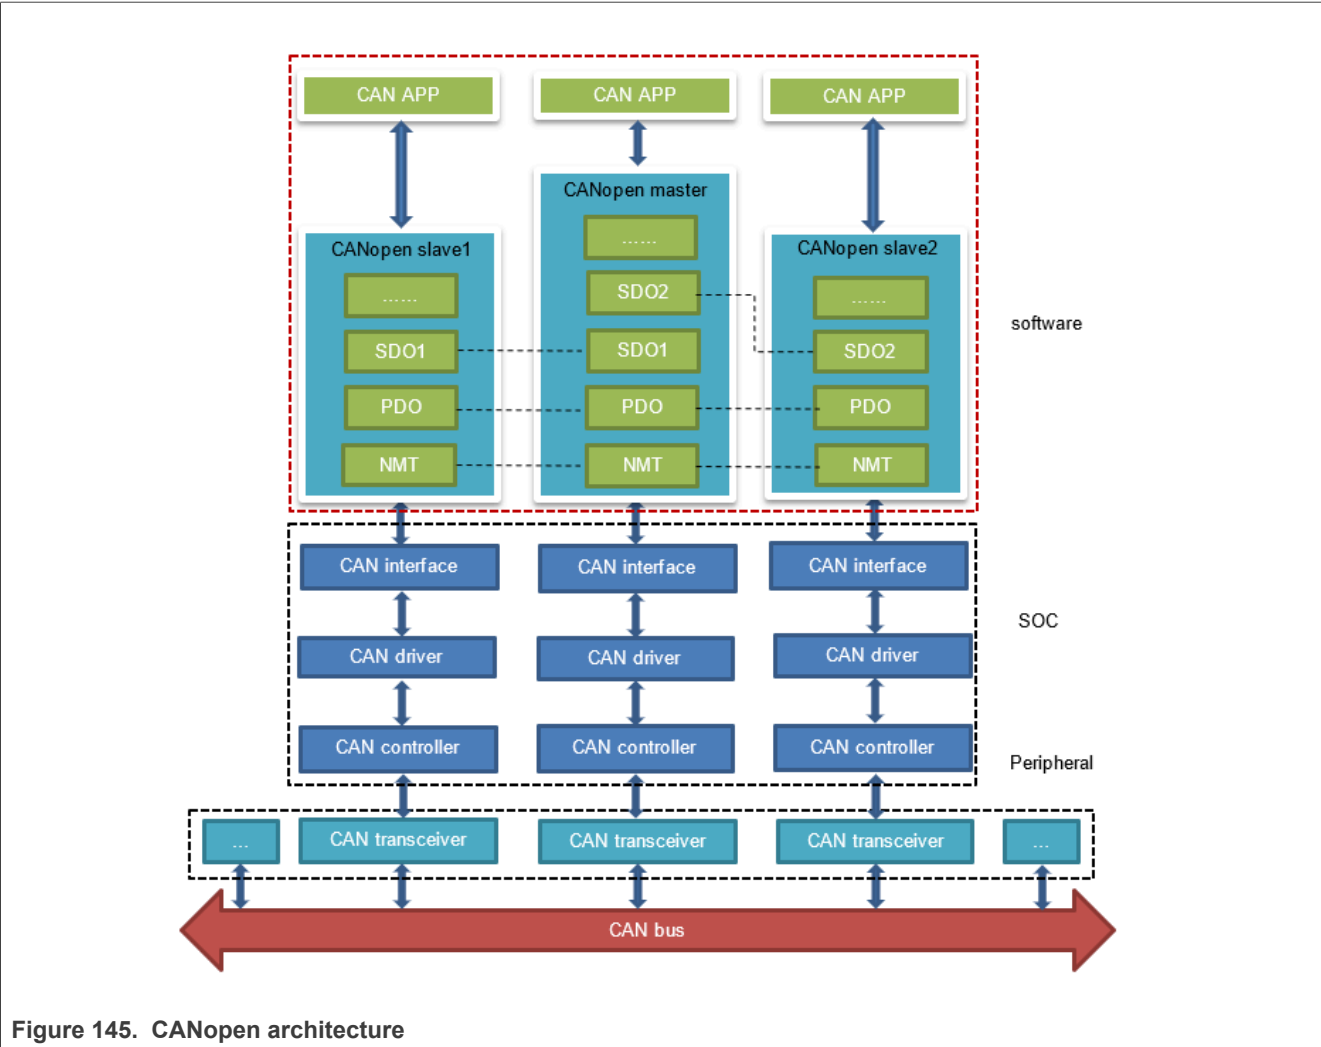


Figure 145. CANopen architecture

6.2.2 Introducing the function of CAN example code

CAN example code supports the CANopen protocol. It mainly implements three parts of functions: network manage function (NMT protocol), service data transmission function (SDO protocol), and process data transmission function (PDO protocol). NMT protocol can manage and monitor slave nodes, include heart beat message. SDO protocol can transmit single or block data. The PDO protocol can transmit process data that requires real time.

CAN example calls the CANopen interfaces, described in the table below:

Table 93. CAN Net APIs and their description

API name (type)	Description
UNS8 canReceive_driver (CAN_HANDLE fd0, Message * m)	SocketCAN receives CAN messages <ul style="list-style-type: none">fd0 – SocketCAN handlem – Receive buffer
UNS8 canSend_driver (CAN_HANDLE fd0, Message const * m)	SocketCAN sends CAN messages <ul style="list-style-type: none">fd0 – SocketCAN handlem – CAN message to be sent

Table 93. CAN Net APIs and their description...continued

API name (type)	Description
void setNodeId(CO_Data* d, UNS8 nodeId)	Set this node id value. <ul style="list-style-type: none"> d – object dictionary nodeId – id value (up to 127)
UNS8 setState(CO_Data* d, e_nodeState newState)	Set node state <ul style="list-style-type: none"> d – object dictionary newState – The state that needs to be set Returns 0 if OK, > 0 on error
void canDispatch(CO_Data* d, Message *m)	CANopen handles data frames that CAN receive. <ul style="list-style-type: none"> d – object dictionary m – Received CAN message
void timerForCan(void)	CANopen virtual clock counter.
UNS8 sendPDOrequest (CO_Data * d, UNS16 RPDOIndex)	Master node requests slave node to feedback specified data. <ul style="list-style-type: none"> d – object dictionary RPDOIndex – index value of specified data
UNS8 readNetworkDictCallback (CO_Data* d, UNS8 nodeId, UNS16 index, UNS8 subIndex, UNS8 dataType, SDOCallback_t Callback, UNS8 useBlockMode)	The master node gets the specified data from the slave node. <ul style="list-style-type: none"> d – object dictionary nodeId – the id value of slave node index – the index value of the specified data subIndex – the subindex value of the specified data dataType – the data type of the specified data Callback – callback function useBlockMode – specifies whether it is a block transmission
UNS8 writeNetworkDictCallBack (CO_Data* d, UNS8 nodeId, UNS16 index, UNS8 subIndex, UNS32 count, UNS8 dataType, void *data, SDOCallback_t Callback, UNS8 useBlockMode)	The master node sets the specified data to the slave node. <ul style="list-style-type: none"> d – object dictionary nodeId – the id value of slave node index – the index value of the specified data subIndex – the subindex value of the specified data count – the length of the specified data dataType – the data type of the specified data Callback – callback function useBlockMode – specifies whether it is a block transmission

6.2.3 Running a CAN application

The following sections describe the hardware and software preparation steps for running a CAN application.

6.2.3.1 Hardware preparation for LS1028ARDB

For LS1028ARDB, below hardware is required:

- LS1028ARDB board
- Two cables to connect CAN1 and CAN.

The hardware connection diagram is as shown in the following figure.



Figure 146. Physical connection for CAN using LS1028ARDB

6.2.3.2 Running the SocketCAN commands

This section describes the steps for running SocketCAN commands that can be performed on LS1028ARDB. These commands are executed on Linux. The standard SocketCAN commands are the following:

1. Open the can0 port.

```
$ ip link set can0 up
```

2. Close the can0 port.

```
$ ip link set can0 down
```

3. Set the baud rate to 500K for the can0 port

```
$ ip link set can0 type can bitrate 500000
```


4. Set can0 port to Loopback mode.

```
$ ip link set can0 type can loopback on
```

5. Send a message through can0. 002 (HEX) is node id, and this value must be 3 characters. 2288DD (HEX) is a message, and can take a value up to 8 bytes.

```
$ cansend can0 002#2288DD
```

6. Monitor can0 port and wait for receiving data.

```
$ candump can0
```

7. See can0 port details.

```
$ ip -details link show can0
```

Note: The third and fourth commands are valid when the state of can0 port is closed.

6.2.3.3 Testing CAN bus

Below is the sample code for testing the CAN bus on LS1028ARDB.

```
[root]# ip link set can0 down
[root]# ip link set can1 down
[root]# ip link set can0 type can loopback off
[root]# ip link set can1 type can loopback off
[root]# ip link set can0 type can bitrate 500000
[root]# ip link set can1 type can bitrate 500000
[root]# ip link set can0 up
[root]# ip link set can1 up
[root]# candump can0 &
[root]# candump can1 &
[root]# cansend can0 001#224466
can0 001 [3] 22 44 66
[root]# can1 001 [3] 22 44 66
[root]# cansend can1 001#224466
can0 001 [3] 22 44 66
can1 001 [3] 22 44 66
[root]# cansend can1 001#113355
can0 001 [3] 11 33 55
can1 001 [3] 11 33 55
[root]# cansend can0 000#224466
can0 000 [3] 22 44 66
```

6.3 OPC UA

OPC (originally known as “OLE for Process Control”, now “Open Platform Communications”) is a collection of multiple specifications, most common of which is OPC Data Access (OPC DA).

OPC Unified Architecture (OPC UA) was released in 2010 by the OPC Foundation as a backward incompatible standard to OPC Classic, under the name of IEC 62541.

OPC UA has turned away from the COM/DCOM (Microsoft proprietary technologies) communication model of OPC Classic, and switched to a TCP/IP based communication stack (asynchronous request/response), layered into the following:

- Raw connections
- Secure channels
- Sessions

6.3.1 OPC introduction

OPC UA defines:

- The transport protocol for communication (that can take place over HTTP, SOAP/XML or directly over TCP).
- A set of 37 'services' that run on the OPC server, and which clients call into, via an asynchronous request/response RPC mechanism.
- A basis for creating information models of data using object-oriented concepts and complex relationships.

The primary goal of OPC is to extract data from devices in the easiest way possible.

The *Information Model* provides a way for servers to not only provide data, but to do so in the most self-explanatory and intuitive way possible.

Note: Further references to 'OPC' in this document will imply OPC UA. OPC Classic is not discussed in this document.

Following are the typical scenarios for embedding an OPC-enabled device into a project:

- Manually investigate (“browse”) the server’s Address Space looking for the data user need using a generic, GUI client (such as UaExpert from Unified Automation, or the FreeOpcUa covered in this chapter).
- Using References and Attributes, understand the format it is in, and the steps that may be needed to convert the data.
- Have a custom OPC client (integrated into the application) subscribe directly to data changes of the node that contains the desired data.

In a typical use case:

- The OPC server runs near the source of information (in industrial contexts, this means near the physical process – for example, on a PLC on a plant floor).
- Clients consume the data at runtime (for example, logging into a database, or feeding it into another industrial process).

OPC-enabled applications can be composed: an industrial device may run an OPC client and feed the collected data into another physical process, while also exposing the latter by running an OPC server.

6.3.2 The node model

Data in an OPC server is structured in *Nodes*. The collection of all nodes that an OPC server exposes to its clients is known as an *Address Space*. Some nodes have a predefined meaning, while others have meaning that is unique to the *Information Model* of that specific OPC server.

Every Node has the following **Attributes**:

- an **ID** (unique)
- a **Class** (what type of node it is)
- a **BrowseName** (a string for machine use)
- a **DisplayName** (a string for human use)

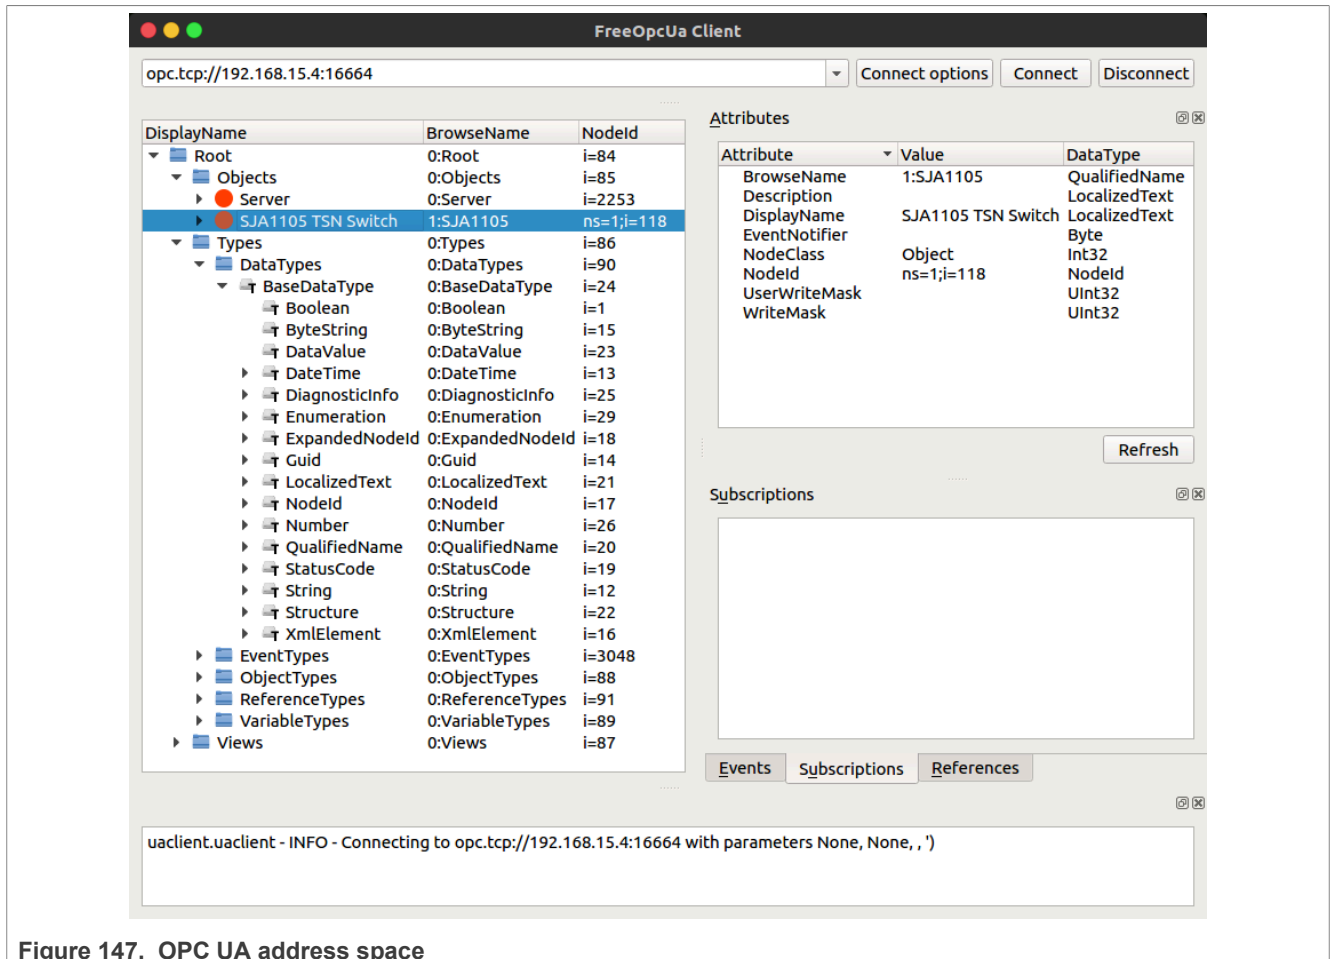


Figure 147. OPC UA address space

Shown on the left-hand side of the figure is the *Address Space* (collection of information that the server makes available to clients) of the OPC server found at `opc.tcp://192.168.15.4:16664`.

Selected is a node with NodeID `ns=1;i=118`, BrowseName=`1:SJA1105` and of NodeClass `Object`.

The full path of the selected node is `0:Root,0:Objects,1:SJA1105`.

6.3.3 Node Namespaces

Namespaces are the means for separating multiple Information Models present in the same Address Space of a server.

- Nodes that do not have the `ns=` prefix as part of the NodeID have an implicit `ns=0;` prefix (are part of the namespace zero).
- Nodes in *namespace * 0* have NodeID's pre-defined by the OPC UA standard. For example, the `0:Server` object, which holds self-describing information (capabilities, diagnostics, and vendor information), has a predefined NodeID of `ns=0;i=2253;`.

It is considered a good practice to not alter any of the nodes exposed in the *namespace * 0*.

6.3.4 Node classes

OPC nodes have an inheritance model, based on their *NodeClass*.

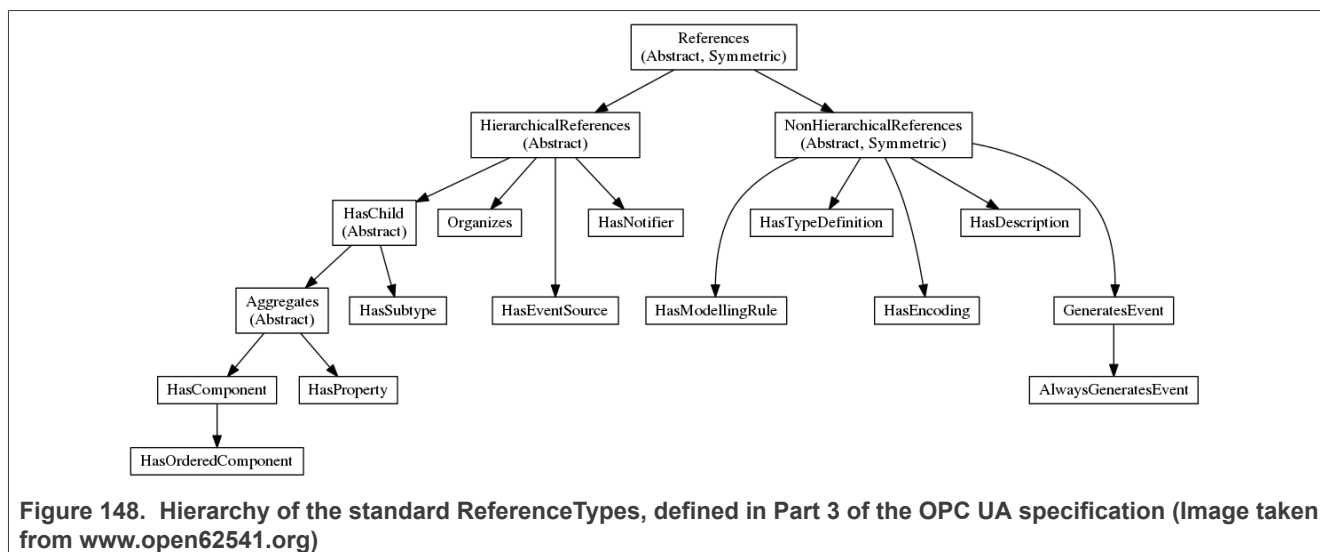
There are eight base node classes defined by the standard:

- Object
- Variable
- Method
- View
- ObjectType
- VariableType
- ReferenceType
- DataType

All nodes have the same base Attributes (inherited from the Node object), plus additional ones depending on their *NodeClass*.

6.3.5 Node graph and references

It may appear that nodes are only chained hierarchically, in a simple parent-child relationship. However, in reality nodes are chained in a complex directed graph, through *References* to other nodes.



In OPC, even ReferenceTypes are Nodes, and as such are structured hierarchically, as can be seen in the figure above.

The definitions of all OPC ReferenceTypes can be found under the `0:Root, 0:Types, 0:ReferenceTypes` path.

The semantics of OPC references can be enriched by creating custom ReferenceType nodes.




Figure 149. The 'Attributes' and 'References' views of the FreeOpcUa Client populated with details of the RGMII4 node

Selected in the Address Space is node `ns=1;i=197`. Conceptually, this represents one of the five Ethernet ports of the SJA1105 TSN switch.

Its NodeClass is Object, but it has a reference of type `HasTypeDefinition` to NodeID `ns=1;i=117` which is `1:EthPortType`. For this reason, the `1:RGMII4` node is of the custom ObjectType `EthPortType`.

6.3.6 Open62541

Real-time Edge integrates the Open62541 software stack (<https://open62541.org/>). This supports both server-side and client-side API for OPC UA applications. Only server-side capabilities of open62541 are being shown here.

Open62541 is distributed as a C-based dynamic library (`libopen62541.so`). The services run on pthreads, and the application code runs inside an event loop.

Enable open62541 in Real-time Edge file `./recipes-nxp/packagegroups/packagegroup-real-time-edge-industrial.bb`:

```
libopen62541 \
```

In order to install Open62541 example application, file `"meta-real-time-edge/conf/distro/include/libopen62541.inc"` has been included in distro configuration.

The following Open62541 example applications are included in the target image:

- `open62541_access_control_client`

- open62541_access_control_server
- open62541_client
- open62541_client_async
- open62541_client_connect
- open62541_client_connectivitycheck_loop
- open62541_client_connect_loop
- open62541_client_subscription_loop
- open62541_custom_datatype_client
- open62541_custom_datatype_server
- open62541_server_ctt
- open62541_server_inheritance
- open62541_server_instantiation
- open62541_server_loglevel
- open62541_server_mainloop
- open62541_server_nodeset
- open62541_server_repeated_job
- open62541_tutorial_client_events
- open62541_tutorial_client_firststeps
- open62541_tutorial_datatypes
- open62541_tutorial_server_datasource
- open62541_tutorial_server_firststeps
- open62541_tutorial_server_method
- open62541_tutorial_server_monitoreditems
- open62541_tutorial_server_object
- open62541_tutorial_server_variable
- open62541_tutorial_server_variabletype

6.3.7 OPC UA Pub/Sub over TSN

This section introduces OPC UA PubSub and demonstrates how TSN can be used to make deterministic and reliable transmission of OPC UA PubSub traffic as well as PTP traffic on a network co-existing with best effort traffic.

6.3.7.1 OPC UA Pub/Sub introduction

The 14th part of the OPC UA specification defines the OPC UA PubSub communication model. It provides an OPC UA Publish Subscribe model that complements the Client/Server communication model.

In PubSub, the participating OPC UA applications can assume the roles of Publishers and Subscribers. Publishers are the sources of data, while Subscribers consume that data. Communication in PubSub is message-based. Publishers send messages to a Message Oriented Middleware, without knowledge of what, if any, Subscribers there may be. Similarly, Subscribers express interest in specific types of data, and process messages that contain this data, without knowledge of what Publishers there are.

To cover a large number of use cases, OPC UA PubSub supports two largely different Message Oriented Middleware variants. These are:

1. A broker-based form, where the core component of the Message Oriented Middleware is a message Broker. Subscribers and Publishers use standard messaging protocols like AMQP or MQTT to communicate with the Broker.

2. A broker-less form, where the Message Oriented Middleware is the network infrastructure that is able to route datagram-based messages. Subscribers and Publishers use datagram protocols such as UDP or raw Ethernet as the transport protocol. In this form, the data sources (Publishers) and the data consumers (Subscribers) join a multicast group. Any data sent by a source to the group goes to all consumers subscribed to the same group. Joining is trivial in Ethernet (Layer 2): the network broadcasts multicast frames everywhere, leaving it to receivers to decide whether to pick up the frame based on the destination address. The OPC UA PubSub sample applications in this section will use this form.

Compared with client-server, the Publishers and Subscribers are decoupled. The number of Subscribers receiving data from a Publisher does not influence the Publisher. This makes PubSub suitable for applications where location independence and/or scalability are required.

One example use case for PubSub is generating logs to multiple systems. For example, sensors or actuators can write logs to a monitoring system, an HMI, an archive application for later querying, and so on. In this case, the data is sent cyclically.

6.3.7.2 OPC UA PubSub over TSN

In general, OPC UA operates at the upper layers of the OSI reference model for networking, whereas TSN is a Layer 2 protocol. TSN adds real-time capability to standard Ethernet. Operating at different layers, TSN and OPC UA PubSub complement each other, yielding a complete communication stack for the industrial Internet of Things. OPC UA standardizes the protocols by which applications exchange data and TSN enables this exchange to meet factories' timing requirements.

One of the key things is to define a mechanism for OPC UA nodes to tell the TSN layers how to prioritize data streams. This cross-layer control is essential to enabling operations technology (OT) using the OPC UA framework to get the data they need when they need it. It also enables time-sensitive OT to coexist on the same network as information technology (IT) functions. In this section, standard Linux tools (that is, tc) are used to map packets from different sources to different traffic classes in order to use TSN features like IEEE 802.1AS and IEEE 802.1Qbv.

6.3.7.3 OPC UA PubSub components

The following figure shows the different components of OPC UA PubSub and their relation to each other. The WriterGroup, DataSetWriter, and PublishedDataSet components define the data acquisition for DataSets, message generation, and transmission on the Publisher side. These parameters should be known on the Subscriber side to configure DataSetReaders to filter and process DataSetMessages.




Figure 150. OPC UA PubSub components

1. **PubSubConnection:** It represents settings needed for the transport protocol. One connection can have a number of writer groups and reader groups. A PubSub connection defines the used protocol and the network address for sending or receiving messages. In the case of using raw Ethernet as transport protocol, the network address can be an MAC multicast address. The Ethernet frame uses EtherType 0xB62C to encapsulate UADP (UA Datagram Protocol) NetworkMessages as payload without IP or UDP headers.
2. **PublishedDataSet:** It contains the collection of the published fields.
3. **WriterGroup:** Each writer group can have one or more DataSetWriters. A WriterGroup defines the timing (that is, publishing interval) and header settings for PubSub NetworkMessages sent by a Publisher.
4. **DataSetWriter:** It is the glue between the WriterGroup and the PublishedDataSet. Each DataSetWriter is bound to a single PublishedDataSet. A PublishedDataSet can have multiple DataSetWriters.
5. **ReaderGroup:** It is used to group a list of DataSetReaders and contains a few shared settings for them.
6. **DataSetReader:** It is the counterpart to a DataSetWriter on the Subscriber side. It defines the filter for the selection of the Publisher and DataSetWriter of interest. The parameters for the filter include the publisher identifier, WriterGroup identifier and DataSetWriter identifier.
7. **SubscribedDataSet:** Its parameters define the processing of the decoded DataSet in the Subscriber for one DataSetReader. The default processing is a mapping to target variables in the Subscriber address space.

6.3.7.4 OPC UA PubSub sample application

There are two sample applications for demonstrating OPC UA PubSub on NXP development boards. One acts as Publisher and the other acts as Subscriber.

On the Publisher:

1. A PubSubConnection is created with the required parameters passed in via command line arguments. This includes the network address URL (for example, *opc.eth://01-00-5E-00-00-01*) and the Ethernet interface (for example, eth1 for ENET2 on i.MX 8M Plus LPDDR4 EVK). Also the Publisher ID is hard coded to 2234.
2. A PublishedDataSet is added with several DataSetFields added. One of the DataSetFields is the CPU temperature measured by the thermal monitoring unit on i.MX 8M Plus. Another DataSetField is the TX HW timestamp of the published packet.

3. A WriteGroup is added with WriterGroup ID hard-coded to 100 and the publishing interval set to 1 second. The Publisher transmits one packet per second cyclically. Each cycle aligns with whole second using Linux system clock **CLOCK_REALTIME**.
4. A DataSetWriter is created with DataSetWriter ID hard-coded to 62541.

On the Subscriber:

1. A PubSubConnection is created with the required parameters passed in via command line arguments. These includes the network address URL (for example, *opc.eth://01-00-5E-00-00-01*) and the Ethernet interface (for example, eth1 for ENET2 on i.MX 8M Plus LPDDR4 EVK). Note that the Subscriber uses the same network address URL as the Publisher.
2. A ReaderGroup is added.
Note: *The Subscriber also runs cyclically with 1 second cycle time to receive packet. Each cycle aligns with whole second with 500 μ s offset to account for the application delay on the publisher and the path delay from publisher to subscriber.*
Linux system clock **CLOCK_REALTIME** is used.
3. A DataSetReader is added and configured with Publisher ID of 2234, WriterGroup ID of 100, and DataSetWriter ID of 62541. Note that all these parameters match the corresponding settings on the Publisher in order to filter the DataSetMessages to be processed by the DataSetReader.
4. A SubscribedDataSet is added with a list of targetVariables. The targetVariables corresponds to the DataSetFields in the PublishedDataSet on the Publisher.
5. Besides the above, the RX HW timestamp of the received packet is taken and the path delay is calculated by subtracting the RX HW timestamp taken on the Subscriber from the TX HW timestamp taken on the Publisher for the same packet. To achieve this, both the Publisher and the Subscriber must have synchronized time. This is achieved by running gPTP.

Both the Publisher and the Subscriber also run a OPC UA server. Users can use a OPC UA client running on a host PC to browse the server's Address Space on either the Publisher or the Subscriber.

6.3.7.5 OPC UA PubSub sample application over TSN

Hardware Requirements:

1. Two or three i.MX 8DXL LPDDR4 EVK, i.MX 8M Plus LPDDR4 EVK, or i.MX 93 EVK boards
2. One LS1028ARDB board

Software Requirements:

1. linuxptp package which provides tools such as `ptp4l`, `phc2sys`, `phc_ctl`, and `hwstamp_ctl`.
2. iproute2 package which provides tools such as `tc`.
3. Open source OPC UA stack open62541 compiled as shared library (`libopen62541.so`).
4. OPC UA PubSub sample application `opcua_pubsub_publisher` and `opcua_pubsub_subscriber` under `/home/root/open62541_example`.

All the above software tools and binaries are already in the rootfs.

The following sections use i.MX 8M Plus LPDDR4 EVK board as an example.

6.3.7.5.1 Case #1: two i.MX 8M Plus LPDDR4 EVK connected back-to-back

A simple setup could be made by connecting two i.MX 8M Plus LPDDR4 EVK boards back-to-back via **ENET2** as shown in [Figure 151](#). One i.MX 8M Plus LPDDR4EVK (Board A) acts as Publisher and the other (Board B) acts as Subscriber. Also, the **ENET1** interface on both boards is connected to the LAN (that is, the office network). The actual device name in Linux might be different.

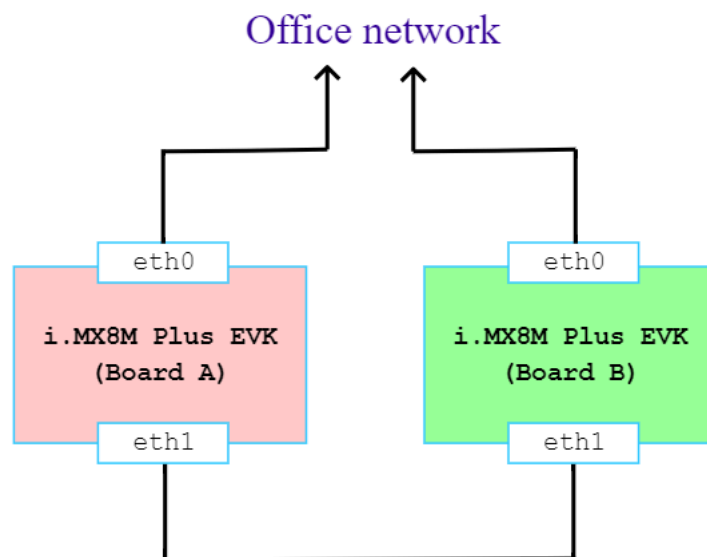


Figure 151. Two i.MX 8M Plus LPDDR4 EVK boards connected back-to-back

- On both boards, bring up ENET2 (that is, eth1):

```
# ip link set eth1 up
# ethtool eth1
```

Using the command `ethtool eth1` should display a message: `Link detected: yes`. Otherwise, check the hardware connection.

- On the Publisher (i.MX 8M Plus LPDDR4 EVK - Board A), add one tc filter rule to match OPC UA PubSub packet (EtherType 0xb62c) on ENET2 (that is, eth1) and modify SKB priority to 2.

```
# tc qdisc add dev eth1 clsact
# tc filter add dev eth1 egress prio 1 u32 match u16 0xb62c 0xffff at -2 action
  skbedit priority 2
# tc filter show dev eth1 egress
```

- On both boards, run `ptp4l` for PTP time synchronization and run `phc2sys` to synchronize PHC clock to Linux system clock (CLOCK_REALTIME). Also on the Subscriber (i.MX 8M Plus LPDDR4 EVK - Board B), use `hwstamp_ctl` to change the RX hardware timestamp setting to 'time stamp any incoming packet' to get the RX hardware timestamp of the packets transmitted by the Publisher.

On the Publisher (i.MX 8M Plus LPDDR4 EVK - Board A):

```
# cp /etc/ptp4l_cfg/gPTP.cfg .
# sed -i 's/priority1.*248/priority1\t\t246/g' ./gPTP.cfg
# ptp4l -i eth1 -p /dev/ptp1 -f ./gPTP.cfg -m > /var/log/ptp4l.log 2>&1 &
# phc2sys -s eth1 -O 0 -S 0.00002 -m > /var/log/phc2sys.log 2>&1 &
```

On the Subscriber (i.MX 8M Plus LPDDR4 EVK - Board B):

```
# ptp4l -i eth1 -p /dev/ptp1 -f /etc/ptp4l_cfg/gPTP.cfg -m > /var/log/ptp4l.log
  2>&1 &
# phc2sys -s eth1 -O 0 -S 0.00002 -m > /var/log/phc2sys.log 2>&1 &
# hwstamp_ctl -i eth1 -r 1
```

It is recommended to SSH to both boards to check the logs of ptp4l and phc2sys to continue to execute other commands on the serial console). On both boards, observe the logs of ptp4l and phc2sys to check the time synchronization progress by using the below commands:

```
# tail -f /var/log/ptp4l.log
# tail -f /var/log/phc2sys.log
```

On the Subscriber, the rms value reported by ptp4l shows the root mean square of the time offset between the PHC and the GM clock. If ptp4l consistently reports rms lower than 100 ns, the PHC is synchronized. See the example ptp4l log below:

```
root@imx8mpevk:~# tail -f /var/log/ptp4l.log
ptp4l[101.594]: rms    23 max    42 freq   -559 +/-   14 delay   779 +/-    0
ptp4l[102.596]: rms     7 max    15 freq   -545 +/-    9 delay   780 +/-    0
ptp4l[103.599]: rms   698 max  1282 freq  -1344 +/-  810 delay   781 +/-    0
ptp4l[104.600]: rms   242 max   491 freq  -1259 +/-  293 delay   782 +/-    0
ptp4l[105.601]: rms   289 max   310 freq   -583 +/-  107 delay   781 +/-    0
ptp4l[106.603]: rms   224 max   281 freq   -396 +/-   12 delay   782 +/-    0
ptp4l[107.604]: rms    87 max   134 freq   -432 +/-   18 delay   785 +/-    0
ptp4l[108.607]: rms    12 max    26 freq   -497 +/-   14 delay   785 +/-    0
ptp4l[109.608]: rms    15 max    26 freq   -531 +/-    8 delay   785 +/-    0
ptp4l[110.610]: rms    13 max    22 freq   -540 +/-    8 delay   785 +/-    0
ptp4l[111.612]: rms    10 max    14 freq   -547 +/-    6 delay   785 +/-    0
ptp4l[112.615]: rms     9 max    19 freq   -533 +/-   11 delay   784 +/-    0
ptp4l[113.616]: rms    11 max    20 freq   -521 +/-   10 delay   783 +/-    0
ptp4l[114.618]: rms     8 max    11 freq   -537 +/-    9 delay   782 +/-    0
ptp4l[115.620]: rms     6 max     8 freq   -533 +/-    8 delay   783 +/-    0
ptp4l[116.622]: rms     8 max    12 freq   -531 +/-   11 delay   783 +/-    0
ptp4l[117.624]: rms     8 max    13 freq   -534 +/-   11 delay   782 +/-    0
ptp4l[118.626]: rms     6 max     9 freq   -539 +/-    7 delay   782 +/-    0
ptp4l[119.628]: rms     8 max    17 freq   -529 +/-    9 delay   782 +/-    0
ptp4l[120.630]: rms    10 max    16 freq   -518 +/-   10 delay   783 +/-    0
ptp4l[121.633]: rms     5 max     8 freq   -527 +/-    7 delay   783 +/-    0
ptp4l[122.635]: rms     8 max    15 freq   -534 +/-   10 delay   784 +/-    0
ptp4l[123.636]: rms     7 max    13 freq   -536 +/-    9 delay   784 +/-    0
```

Figure 152. A sample ptp4l log

On both the Publisher and the Subscriber, the offset information reported by phc2sys shows the time offset between the PHC and the system clock (CLOCK_REALTIME). If phc2sys consistently reports offset lower than 100 ns, the System clock is synchronized. A sample phc2sys log is shown below:

```
root@imx8mpevk:~# tail -f /var/log/phc2sys.log
phc2sys[7349.412]: CLOCK_REALTIME phc offset      61 s2 freq    +58 delay   750
phc2sys[7350.413]: CLOCK_REALTIME phc offset     -81 s2 freq    -66 delay   750
phc2sys[7351.413]: CLOCK_REALTIME phc offset     -20 s2 freq    -29 delay   750
phc2sys[7352.413]: CLOCK_REALTIME phc offset      54 s2 freq    +39 delay   750
phc2sys[7353.414]: CLOCK_REALTIME phc offset      20 s2 freq    +21 delay   750
phc2sys[7354.414]: CLOCK_REALTIME phc offset     -31 s2 freq    -24 delay   750
phc2sys[7355.414]: CLOCK_REALTIME phc offset      48 s2 freq    +46 delay   750
phc2sys[7356.415]: CLOCK_REALTIME phc offset      -7 s2 freq     +5 delay   750
phc2sys[7357.415]: CLOCK_REALTIME phc offset     -57 s2 freq    -47 delay   750
phc2sys[7358.416]: CLOCK_REALTIME phc offset      20 s2 freq    +13 delay   750
phc2sys[7359.416]: CLOCK_REALTIME phc offset     -28 s2 freq    -29 delay   750
phc2sys[7360.416]: CLOCK_REALTIME phc offset      26 s2 freq    +16 delay   750
phc2sys[7361.417]: CLOCK_REALTIME phc offset      20 s2 freq    +18 delay   750
phc2sys[7362.417]: CLOCK_REALTIME phc offset     -14 s2 freq    -10 delay   750
```

Figure 153. A sample phc2sys log

After establishing the time synchronization successfully on both the Publisher and the Subscriber, we can configure TSN Qbv and run the OPC UA PubSub sample applications as in the following steps.

- On the Publisher (i.MX 8M Plus LPDDR4 EVK - Board A), configure TSN Qbv on ENET2 (that is, eth1) to map SKB priority to traffic class to hardware queue as below, set gate control list to have 2 entries and total cycle time of 1ms (queue 4 has 500 μ s for best effort traffic, queue 0 and queue 2 share 500us for OPC UA PubSub and PTP traffic as well as other traffic like ping), also set base time to 1ms so that the schedule is aligned to 1ms. This is just an example configuration for the schedule.

```
SKB priority 0 -> traffic class 0 -> queue 0
SKB priority 1 -> traffic class 1 -> queue 1
SKB priority 2 -> traffic class 2 -> queue 2
SKB priority 3 -> traffic class 3 -> queue 3
SKB priority 4 -> traffic class 4 -> queue 4
```

```
# tc qdisc replace dev eth1 parent root handle 100 taprio num_tc 5 map 0 1 2
  3 4 queues 1@0 1@1 1@2 1@3 1@4 base-time 001000000 sched-entry S 0x10 500000
  sched-entry S 0x05 500000 flags 2
# tc -g qdisc show dev eth1
```

Together with the tc filter rule configured previously, the above TSN qbv configuration on ENET2 distributes OPC UA PubSub traffic into TX hardware queue 2, PTP traffic into TX hardware queue 0. Also, we send best effort traffic to TX hardware queue 4. Other traffic such as pings can still go into TX hardware queue 0. Because the OPC UA PubSub and PTP traffic have different TX hardware queues and time slot than the best effort traffic, the latter cannot influence the former.

- On the Subscriber (i.MX 8M Plus LPDDR4 EVK - Board B), run the OPC UA PubSub Subscriber sample application. Run the Subscriber application before the Publisher application so that no packets sent by the Publisher are missed.

Note that octets in the MAC address must be separated by **hyphens (-)**.

```
# /home/root/open62541_example/opcua_pubsub_subscriber -u
opc.eth://01-00-5E-00-00-01 -d eth1
```

- On the Publisher (i.MX 8M Plus LPDDR4 EVK - Board A), run the OPC UA PubSub Publisher sample application.

Note that octets in the MAC address must be separated by **hyphens (-)**.

```
# /home/root/open62541_example/opcua_pubsub_publisher -u
opc.eth://01-00-5E-00-00-01 -d eth1
```

Example log on the Publisher:

```
root@imx8mpevk:~# /home/root/open62541_example/opcua_pubsub_publisher -u opc.eth://01-00-5E-00-00-01 -d eth1
[2020-09-20 11:02:20.030 (UTC+0000)] info/userland Transport Profile : http://opcfoundation.org/UA-Profile/Transport/pub
sub-eth-uadp
[2020-09-20 11:02:20.030 (UTC+0000)] info/userland Network Address URL : opc.eth://01-00-5E-00-00-01
[2020-09-20 11:02:20.030 (UTC+0000)] info/userland Ethernet Interface : eth1
[2020-09-20 11:02:20.482 (UTC+0000)] warn/server AccessControl: Unconfigured AccessControl. Users have all permissions.
[2020-09-20 11:02:20.482 (UTC+0000)] info/server AccessControl: Anonymous login is enabled
[2020-09-20 11:02:20.482 (UTC+0000)] warn/server Username/Password configured, but no encrypting SecurityPolicy. This ca
n leak credentials on the network.
[2020-09-20 11:02:20.482 (UTC+0000)] warn/userland AcceptAll Certificate Verification. Any remote certificate will be acce
pted.
[2020-09-20 11:02:20.483 (UTC+0000)] info/userland PubSub channel requested
[2020-09-20 11:02:20.483 (UTC+0000)] info/server Open PubSub ethernet connection.
[2020-09-20 11:02:20.483 (UTC+0000)] info/userland Publisher socket FD : 3
[2020-09-20 11:02:20.483 (UTC+0000)] info/userland Publisher Tx HW timestamp : Enabled
[2020-09-20 11:02:20.484 (UTC+0000)] info/userland Publisher cycle time : 1000.000000 ms
[2020-09-20 11:02:20.484 (UTC+0000)] info/userland Publisher thread priority : 78
[2020-09-20 11:02:20.484 (UTC+0000)] info/userland Publisher on CPU core : 3
[2020-09-20 11:02:20.484 (UTC+0000)] info/userland Publisher thread callback Id: 281472837534032
[2020-09-20 11:02:20.484 (UTC+0000)] info/network TCP network layer listening on opc.tcp://imx8mpevk:4840/
[2020-09-20 11:02:20.485 (UTC+0000)] info/userland Starting the publisher cycle at 1600599745.000000000
```

Figure 154. Example log on the Publisher

Example log on the Subscriber:

```
root@imx8mpvk:~# /home/root/open62541_example/opcu pubsub_subscriber -u opc.eth://01-00-5E-00-00-01 -d eth1
[2020-09-20 11:02:22.297 (UTC+0000)] info/userland Transport Profile : http://opcfoundation.org/UA-Profile/Transport/pub
sub-eth-uadp
[2020-09-20 11:02:22.297 (UTC+0000)] info/userland Network Address URL : opc.eth://01-00-5E-00-00-01
[2020-09-20 11:02:22.297 (UTC+0000)] info/userland Ethernet Interface : eth1
[2020-09-20 11:02:22.747 (UTC+0000)] warn/server AccessControl: Unconfigured AccessControl. Users have all permissions.
[2020-09-20 11:02:22.747 (UTC+0000)] info/server AccessControl: Anonymous login is enabled
[2020-09-20 11:02:22.747 (UTC+0000)] warn/server Username/Password configured, but no encrypting SecurityPolicy. This ca
n leak credentials on the network.
[2020-09-20 11:02:22.747 (UTC+0000)] warn/userland AcceptAll Certificate Verification. Any remote certificate will be acce
pted.
[2020-09-20 11:02:22.748 (UTC+0000)] info/userland PubSub channel requested
[2020-09-20 11:02:22.748 (UTC+0000)] info/server Open PubSub ethernet connection.
[2020-09-20 11:02:22.748 (UTC+0000)] info/userland Subscriber socket FD : 3
[2020-09-20 11:02:22.748 (UTC+0000)] info/userland Subscriber Rx HW timestamp : Enabled
[2020-09-20 11:02:22.748 (UTC+0000)] info/userland Subscriber thread priority : 81
[2020-09-20 11:02:22.749 (UTC+0000)] info/userland Publisher on CPU core : 3
[2020-09-20 11:02:22.749 (UTC+0000)] info/userland Subscriber thread callback Id: 281473844617552
[2020-09-20 11:02:22.749 (UTC+0000)] info/userland Starting the subscriber cycle at 1600599743.000500000
[2020-09-20 11:02:22.750 (UTC+0000)] info/network TCP network layer listening on opc.tcp://imx8mpvk:4801/
[2020-09-20 11:02:23.001 (UTC+0000)] info/userland Packet Sequence Number : 0
[2020-09-20 11:02:24.001 (UTC+0000)] info/userland Packet Sequence Number : 0
[2020-09-20 11:02:25.000 (UTC+0000)] info/userland Packet Sequence Number : 0
[2020-09-20 11:02:26.001 (UTC+0000)] info/userland Packet Sequence Number : 1
[2020-09-20 11:02:27.000 (UTC+0000)] info/userland Packet Sequence Number : 2
[2020-09-20 11:02:28.000 (UTC+0000)] info/userland Packet Sequence Number : 3
[2020-09-20 11:02:29.000 (UTC+0000)] info/userland Packet Sequence Number : 4
[2020-09-20 11:02:30.001 (UTC+0000)] info/userland Packet Sequence Number : 5
[2020-09-20 11:02:31.000 (UTC+0000)] info/userland Packet Sequence Number : 6
[2020-09-20 11:02:32.001 (UTC+0000)] info/userland Packet Sequence Number : 7
[2020-09-20 11:02:33.001 (UTC+0000)] info/userland Packet Sequence Number : 8
```

Figure 155. Example log on the Subscriber

- On a PC connected to office network and with OPC UA Client installed (that is, UaExpert as in below snapshots), we can browser either the OPC UA server's Address Space on either the Publisher or the Subscriber. (We assume that eth0 has obtained the IP address by DHCP automatically).
The URL of the OPC UA server on the Publisher is below:
opc.tcp://<IP_of_eth0_on_Publisher>:4840/
The URL of the OPC UA server on the Subscriber is below:
opc.tcp://<IP_of_eth0_on_Subscriber>:4801/
Example snapshot of UaExpert connected to the Publisher:




Figure 156. Sample snapshot of UaExpert connected to the Publisher

Example snapshot of UaExpert connected to the Subscriber:




Figure 157. Sample snapshot of UaExpert connected to the Subscriber

On the UaExpert client connected to the Subscriber, we can observe the CPU temperature published by the Publisher and the path delay from Publisher to Subscriber which is close to 800 ns.

- On the Publisher (i.MX 8M Plus LPDDR4 EVK - Board A), we can use pktgen to simulate high load best effort traffic which is sent to queue 4 of ENET2.

```
# /home/root/samples/pktgen/pktgen_sample01_simple.sh -i eth1 -q 4 -s 1000 -n 0
```

The OPC UA PubSub traffic and PTP traffic are protected by TSN Qbv by having different TX hardware queue and time slot than the best effort traffic. Hence, users can see consistent output on the console of the Publisher and the Subscriber, and the path delay from Publisher to Subscriber is still close to 800 ns.

In case TSN Qbv was not configured, after pktgen starts running, various issues may happen. First of all, the ptp4l application will show timeout issue as below.

Example error log of ptp4l on the Publisher:

Figure 158. Sample error log of ptp4l on the Publisher

```
ptp4l[436.575]: timed out while polling for tx timestamp
ptp4l[436.575]: increasing tx timestamp timeout may correct this issue, but it is likely caused by a driver bug
ptp4l[436.575]: port 1: send sync failed
ptp4l[436.575]: port 1: MASTER to FAULTY on FAULT DETECTED (FT_UNSPECIFIED)
ptp4l[436.600]: selected local clock 00049f.ffff.06fe8b as best master
ptp4l[436.600]: port 1: assuming the grand master role
```

Example error log of ptp4l on the Subscriber

Figure 159. Sample error log of ptp4l on the Subscriber

```
ptp4l[431.404]: rms 11 max 16 freq -152 +/- 14 delay 777 +/- 0
ptp4l[432.406]: rms 8 max 20 freq -155 +/- 10 delay 777 +/- 0
ptp4l[433.408]: rms 4 max 8 freq -146 +/- 4 delay 777 +/- 0
ptp4l[433.784]: port 1: SLAVE to MASTER on ANNOUNCE_RECEIPT_TIMEOUT_EXPIRES
ptp4l[433.784]: selected local clock 00049f.ffff.06fe83 as best master
ptp4l[433.784]: port 1: assuming the grand master role
ptp4l[452.679]: selected best master clock 00049f.ffff.06fe8b
ptp4l[452.680]: port 1: MASTER to SLAVE on RS_SLAVE
ptp4l[453.680]: rms 731 max 1281 freq +1135 +/- 357 delay 777 +/- 0
ptp4l[454.682]: rms 220 max 304 freq +157 +/- 200 delay 778 +/- 0
ptp4l[455.683]: rms 285 max 320 freq -252 +/- 54 delay 778 +/- 0
ptp4l[456.684]: rms 155 max 220 freq -299 +/- 18 delay 778 +/- 0
ptp4l[457.687]: rms 43 max 79 freq -231 +/- 27 delay 777 +/- 0
```

The OPC UA PubSub sample application may also be impacted by best-effort traffic without TSN, especially when the publish cycle time is very short (that is, 2 ms). Note that the publish cycle time is hard-coded to 1 second in the sample application to make observation easier on the Subscriber console as well as on the UaExpert client.

Several issues can be observed under high load best effort traffic without TSN. For example, the Publisher application might display a warning message *timed out while polling for tx timestamp!*. The Subscriber application might show that the packet sequence number stops incrementing. In addition, the path delay from Publisher to Subscriber displayed on UaExpert might show a large number due to packet transmission delay. When this issue happens, it can only be recovered by restarting both the Publisher and Subscriber applications.

6.3.7.5.2 Case #2: two i.MX 8M Plus LPDDR4 EVK boards connected to LS1028ARDB TSN switch

The setup could use one LS1028ARDB as TSN switch plus two or three i.MX 8M Plus LPDDR4 EVK boards. One i.MX 8M Plus LPDDR4 EVK (Board A) acts as Publisher and others act as Subscribers. [Figure 160](#) shows the block diagram of this setup. The ENET2 interface on each i.MX 8M Plus LPDDR4 EVK is connected to the switch port on LS1028ARDB. Also, the ENET1 interface on each i.MX 8M Plus LPDDR4 EVK is connected to LAN (that is, office network).

Note: The actual device name might be different in Linux.

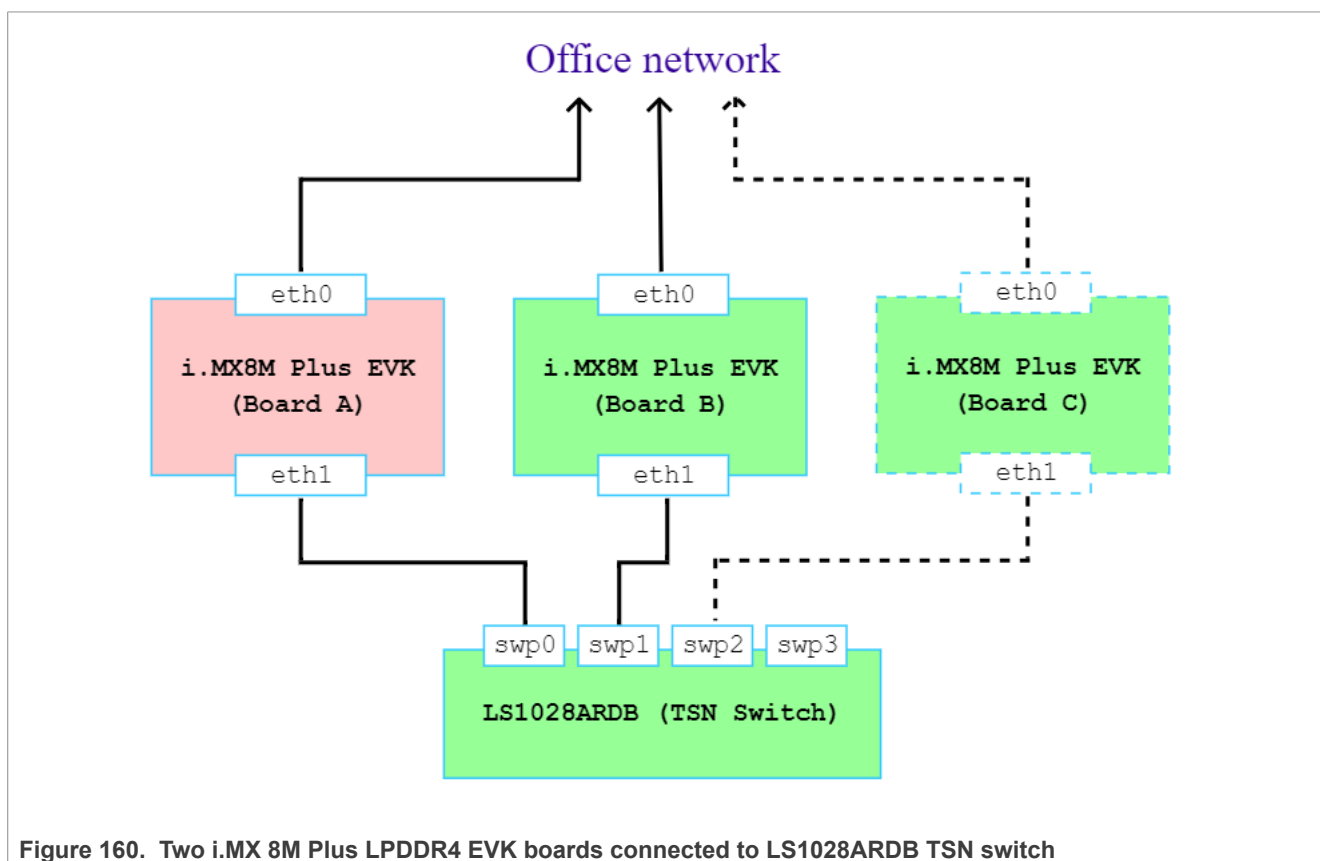


Figure 160. Two i.MX 8M Plus LPDDR4 EVK boards connected to LS1028ARDB TSN switch

The following steps assume that two i.MX 8M Plus LPDDR4 EVK boards are used. Board A acts as Publisher and Board B acts as Subscriber. In this setup, the switch port `swp0` is the ingress port for OPC UA PubSub traffic and best effort traffic. The switch port `swp1` is the egress ports for OPC UA PubSub traffic and best effort traffic.

Since the TSN switch on LS1028ARDB uses the value of VLAN PCP field to map traffic to different TX hardware queue on egress switch port (that is, `swp1`), we add VLAN header to the OPC UA PubSub packet and best effort packet. Note that the PTP packet is untagged without VLAN header.

- On LS1028ARDB, configure Ethernet bridge on TSN switch and enable VLAN filtering.

```
# ip link set eno2 up
# ip link set swp0 up
# ip link set swp1 up
# ip link add name br0 type bridge vlan_filtering 1
# ip link set dev swp0 master br0
# ip link set dev swp1 master br0
# ip link set dev br0 up
# bridge vlan add dev swp0 vid 100
# bridge vlan add dev swp1 vid 100
# bridge vlan show
```

- On both the Publisher (i.MX 8M Plus LPDDR4 EVK - Board A) and the Subscriber (i.MX 8M Plus LPDDR4 EVK - Board B), bring up ENET2 (that is, `eth1`):

```
# ip link set eth1 up
```

- On the Publisher (i.MX 8M Plus LPDDR4 EVK - Board A), add one tc filter to match OPC UA PubSub packet (EtherType 0xb62c) on ENET2 (that is, `eth1`) and modify SKB priority to 2.

```
# tc qdisc add dev eth1 clsact
# tc filter add dev eth1 egress prio 1 u32 match u16 0xb62c 0xffff at -2 action
  skbedit priority 2
# tc filter show dev eth1 egress
```

- On each board, run `ptp4l` for PTP time synchronization and run `phc2sys` to synchronize PHC clock to Linux system clock (CLOCK_REALTIME).

Also on the Subscriber (i.MX 8M Plus LPDDR4 EVK - Board B), use `hwstamp_ctl` to change the RX hardware timestamp setting to 'time stamp any incoming packet' in order to get the RX hardware timestamp of the packets transmitted by the Publisher.

On LS1028ARDB:

```
# ptp4l -i swp0 -i swp1 -p /dev/ptp1 -f /etc/ptp4l_cfg/gPTP.cfg -m > /var/log/
  ptp4l.log 2>&1 &
# phc2sys -s swp0 -O 0 -S 0.00002 -m > /var/log/phc2sys.log 2>&1 &
```

On the Publisher (i.MX 8M Plus LPDDR4 EVK - Board A):

```
# cp /etc/ptp4l_cfg/gPTP.cfg .
# sed -i 's/priority1.*248/priority1\t\t246/g' ./gPTP.cfg
# ptp4l -i eth1 -p /dev/ptp1 -f ./gPTP.cfg -m > /var/log/ptp4l.log 2>&1 &
# phc2sys -s eth1 -O 0 -S 0.00002 -m > /var/log/phc2sys.log 2>&1 &
```

On the Subscriber (i.MX 8M Plus LPDDR4 EVK - Board B):

```
# ptp4l -i eth1 -p /dev/ptp1 -f /etc/ptp4l_cfg/gPTP.cfg -m > /var/log/ptp4l.log
  2>&1 &
# phc2sys -s eth1 -O 0 -S 0.00002 -m > /var/log/phc2sys.log 2>&1 &
# hwstamp_ctl -i eth1 -r 1
```

On each board, observe the logs of `ptp4l` and `phc2sys` to check the time synchronization progress by using the below commands: (In order to continue to execute other commands on the serial console, performing an SSH to each board to check the logs of `ptp4l` and `phc2sys` is recommended).

```
# tail -f /var/log/ptp4l.log
# tail -f /var/log/phc2sys.log
```


On LS1028ARDB and the Subscriber, the rms value reported by ptp4l shows the root mean square of the time offset between the PHC and the GM clock. If ptp4l consistently reports rms lower than 100 ns, the PHC is synchronized. Refer to the example log of ptp4l in the back-to-back case.

On each board, the offset information reported by phc2sys shows the time offset between the PHC and the system clock (CLOCK_REALTIME). If phc2sys consistently reports offset lower than 100 ns, the System clock is synchronized. Refer to the example log of phc2sys in the back-to-back case.

After establishing the time synchronization successfully on each board, users can configure TSN Qbv and run the OPC UA PubSub sample applications using the following steps.

- On the Publisher (i.MX 8M Plus LPDDR4 EVK - Board A), configure TSN Qbv on ENET2 to map SKB priority to traffic class to hardware queue as below, set gate control list to have 2 entries and total cycle time of 1 ms (queue 4 has 500 μ s for best effort traffic, queue 0 and queue 2 share 500 μ s for OPC UA PubSub and PTP traffic as well as other traffic like ping), also set base time to 1 ms so that the schedule is aligned to 1 ms. This is just an example configuration for the schedule.

```
SKB priority 0 -> traffic class 0 -> queue 0
SKB priority 1 -> traffic class 1 -> queue 1
SKB priority 2 -> traffic class 2 -> queue 2
SKB priority 3 -> traffic class 3 -> queue 3
SKB priority 4 -> traffic class 4 -> queue 4
```

```
# tc qdisc replace dev eth1 parent root handle 100 taprio num_tc 5 map 0 1 2
  3 4 queues 1@0 1@1 1@2 1@3 1@4 base-time 001000000 sched-entry S 0x10 500000
  sched-entry S 0x05 500000 flags 2
# tc -g qdisc show dev eth1
```

Together with the tc filter rule configured previously, the above TSN Qbv configuration on ENET2 distributes OPC UA PubSub traffic into TX hardware queue 2, PTP traffic into TX hardware queue 0. Also, send best effort traffic to TX hardware queue 4. Other traffic like ping can still go into TX hardware queue 0. Because the OPC UA PubSub and PTP traffic have different TX hardware queues and time slot than the best effort traffic, the latter cannot influence the former.

- On LS1028ARDB, configure TSN Qbv on swp1, set gate control list to have 2 entries and total cycle time of 200 μ s. Ensure that queue 4 has 500 μ s for best effort traffic, queue 0 and queue 2 share 500 μ s for OPC UA PubSub and PTP traffic as well as other traffic such as ping. In addition, set the base time to 1 ms so that the schedule is aligned to 1 ms as 1 ms. Note that the TSN Qbv configuration on LS1028ARDB TSN switch is used to protect the OPC UA PubSub traffic from traffic which may enter the switch from other switch ports. In this use case, it is optional and is used for demonstration purpose only.

```
# tc qdisc replace dev swp1 root taprio num_tc 8 map 0 1 2 3 4 5 6 7 queues
  1@0 1@1 1@2 1@3 1@4 1@5 1@6 1@7 base-time 001000000 sched-entry S 0x10 500000
  sched-entry S 0x05 500000 flags 0x2
# tc -g qdisc show dev swp1
```

- With the above TSN Qbv configuration on egress switch port swp1,
 - OPC UA PubSub traffic goes into TX hardware queue 2 (Add VLAN header with PCP field set to 2 for OPC UA PubSub packet).
 - The best effort traffic goes into TX hardware queue 4 (Add VLAN header with PCP field set to 4 using pktgen for generating best effort traffic). Note that the PTP traffic is untagged without VLAN header and will use TX hardware queue 0 of swp1 to transmit to the Subscriber. Similar to the TSN Qbv configuration on Publisher, the OPC UA PubSub and PTP traffic have different TX hardware queues and time slot than the best effort traffic, the latter cannot influence the former.
- On the Subscriber (i.MX 8M Plus LPDDR4 EVK - Board B), run the OPC UA PubSub Subscriber sample application. Run the Subscriber application before the Publisher application so that no packet sent by the Publisher is missed.

Note that in the URL of below command 100.2 means VLAN ID 100 and PCP value 2 and it is separated from the MAC address using a colon.

```
# /home/root/open62541_example/opcua_pubsub_subscriber -u
opc.eth://01-00-5E-00-00-01:100.2 -d eth1
```

- On the Publisher (i.MX 8M Plus LPDDR4 EVK - Board A), run the OPC UA PubSub Publisher sample application.

Note that in the URL of below command 100.2 means VLAN ID 100 and PCP value 2 and it is separated from the MAC address using a colon.

```
# /home/root/open62541_example/opcua_pubsub_publisher -u
opc.eth://01-00-5E-00-00-01:100.2 -d eth1
```

Example log on the Publisher:

Figure 161. Example log on the Publisher

```
root@imx8mpevk:~# /home/root/open62541_example/opcua_pubsub_publisher -u opc.eth://01-00-5E-00-00-01:100.2 -d eth1
[2020-09-20 10:18:51.436 (UTC+0000)] info/userland Transport Profile : http://opcfoundation.org/UA-Profile/Transport/pubsub-eth-uadp
[2020-09-20 10:18:51.441 (UTC+0000)] info/userland Network Address URL : opc.eth://01-00-5E-00-00-01:100.2
[2020-09-20 10:18:51.441 (UTC+0000)] info/userland Ethernet Interface : eth1
[2020-09-20 10:18:51.920 (UTC+0000)] warn/server AccessControl: Unconfigured AccessControl. Users have all permissions.
[2020-09-20 10:18:51.920 (UTC+0000)] info/server AccessControl: Anonymous login is enabled
[2020-09-20 10:18:51.920 (UTC+0000)] warn/server Username/Password configured, but no encrypting SecurityPolicy. This can leak credentials on the network.
[2020-09-20 10:18:51.921 (UTC+0000)] warn/userland AcceptAll Certificate Verification. Any remote certificate will be accepted
.
[2020-09-20 10:18:51.921 (UTC+0000)] info/userland PubSub channel requested
[2020-09-20 10:18:51.921 (UTC+0000)] info/server Open PubSub ethernet connection.
[2020-09-20 10:18:51.922 (UTC+0000)] info/userland Publisher socket FD : 3
[2020-09-20 10:18:51.922 (UTC+0000)] info/userland Publisher Tx HW timestamp : Enabled
[2020-09-20 10:18:51.922 (UTC+0000)] info/userland Publisher cycle time : 1000.000000 ms
[2020-09-20 10:18:51.922 (UTC+0000)] info/userland Publisher thread priority : 78
[2020-09-20 10:18:51.922 (UTC+0000)] info/userland Publisher on CPU core : 3
[2020-09-20 10:18:51.922 (UTC+0000)] info/userland Publisher thread callback Id: 281473694241104
[2020-09-20 10:18:51.922 (UTC+0000)] info/userland Starting the publisher cycle at 1600597136.000000000
[2020-09-20 10:18:51.923 (UTC+0000)] info/network TCP network layer listening on opc.tcp://imx8mpevk:4840/
```

Example log on the Subscriber:

```
root@imx8mpevk:~# /home/root/open62541_example/opcua_pubsub_subscriber -u opc.eth://01-00-5E-00-00-01:100.2 -d eth1
[2020-09-20 10:19:43.151 (UTC+0000)] info/userland Transport Profile : http://opcfoundation.org/UA-Profile/Transport/pubsub-eth-uadp
[2020-09-20 10:19:43.151 (UTC+0000)] info/userland Network Address URL : opc.eth://01-00-5E-00-00-01:100.2
[2020-09-20 10:19:43.151 (UTC+0000)] info/userland Ethernet Interface : eth1
[2020-09-20 10:19:43.598 (UTC+0000)] warn/server AccessControl: Unconfigured AccessControl. Users have all permissions.
[2020-09-20 10:19:43.599 (UTC+0000)] info/server AccessControl: Anonymous login is enabled
[2020-09-20 10:19:43.599 (UTC+0000)] warn/server Username/Password configured, but no encrypting SecurityPolicy. This can leak credentials on the network.
[2020-09-20 10:19:43.599 (UTC+0000)] warn/userland AcceptAll Certificate Verification. Any remote certificate will be accepted
.
[2020-09-20 10:19:43.600 (UTC+0000)] info/userland PubSub channel requested
[2020-09-20 10:19:43.600 (UTC+0000)] info/server Open PubSub ethernet connection.
[2020-09-20 10:19:43.600 (UTC+0000)] info/userland Subscriber socket FD : 3
[2020-09-20 10:19:43.600 (UTC+0000)] info/userland Subscriber Rx HW timestamp : Enabled
[2020-09-20 10:19:43.600 (UTC+0000)] info/userland Subscriber thread priority : 81
[2020-09-20 10:19:43.600 (UTC+0000)] info/userland Publisher on CPU core : 3
[2020-09-20 10:19:43.600 (UTC+0000)] info/userland Subscriber thread callback Id: 281472940425552
[2020-09-20 10:19:43.600 (UTC+0000)] info/userland Starting the subscriber cycle at 1600597184.000500000
[2020-09-20 10:19:43.602 (UTC+0000)] info/network TCP network layer listening on opc.tcp://imx8mpevk:4801/
[2020-09-20 10:19:44.001 (UTC+0000)] info/userland Packet Sequence Number : 0
[2020-09-20 10:19:45.000 (UTC+0000)] info/userland Packet Sequence Number : 0
[2020-09-20 10:19:46.000 (UTC+0000)] info/userland Packet Sequence Number : 1
[2020-09-20 10:19:47.000 (UTC+0000)] info/userland Packet Sequence Number : 2
[2020-09-20 10:19:48.000 (UTC+0000)] info/userland Packet Sequence Number : 3
[2020-09-20 10:19:49.001 (UTC+0000)] info/userland Packet Sequence Number : 4
[2020-09-20 10:19:50.000 (UTC+0000)] info/userland Packet Sequence Number : 5
[2020-09-20 10:19:51.000 (UTC+0000)] info/userland Packet Sequence Number : 6
[2020-09-20 10:19:52.000 (UTC+0000)] info/userland Packet Sequence Number : 7
[2020-09-20 10:19:53.000 (UTC+0000)] info/userland Packet Sequence Number : 8
```

Figure 162. Example log on the Subscriber

- On a PC connected to office network and with OPC UA Client installed (that is, UaExpert as in below snapshots), we can browser either the OPC UA server's Address Space on either the Publisher or the Subscriber. (We assume that eth0 have got IP address by DHCP automatically.

The URL of the OPC UA server on the Publisher is below:

opc.tcp://<IP_of_eth0_on_Publisher>:4840/

The URL of the OPC UA server on the Subscriber is below:

opc.tcp://<IP_of_eth0_on_Subscriber>:4801/

Refer to the example snapshot of UaExpert in the back-to-back case. On the UaExpert client connected to the Subscriber, we can observe the CPU temperature published by the Publisher and the path delay from Publisher to Subscriber which is around 4 μ s. Compared to the 800 ns in the back-to-back case, the increased path delay is added by the bridge.

- On the Publisher (i.MX 8M Plus LPDDR4 EVK - Board A), we can use pktgen to simulate high load best effort traffic with VLAN ID set to 100 and VLAN PCP set to 4 in VLAN header.

```
# cp /home/root/samples/pktgen/pktgen_sample01_simple.sh /home/root/samples/
pktgen/pktgen_sample01_simple_vlan.sh
# sed -i '/^UDP_MAX=.*\n/a VLAN_ID=100\nVLAN_P=4' /home/root/samples/pktgen/
pktgen_sample01_simple_vlan.sh
# /home/root/samples/pktgen/pktgen_sample01_simple_vlan.sh -i eth1 -q 4 -s 1000
-n 0
```

Note: In order to protect the OPC UA PubSub traffic and PTP traffic, the TSN Qbv should be configured to have different TX hardware queue and time slot from the best effort traffic on both the Publisher and TSN switch. This ensures that users can get consistent output on the console of the Publisher and the Subscriber, and the path delay from Publisher to Subscriber is still around 4 μ s.

In case TSN Qbv was not configured, after pktgen starts running, various issues might occur. Refer to the issues detailed in the back-to-back case.

- On LS1028ARDB, it is possible to check the status of TX packets of swp1 by using the command below:

```
# ethtool -S swp1 | grep -i "tx_green_prio_"
```

Example log below: (tx_green_prio_0 mainly for PTP traffic, tx_green_prio_2 mainly for OPC UA PubSub traffic, tx_green_prio_4 mainly for best effort traffic generated by pktgen).

```
root@ls1028ardb:~# ethtool -S swp1 | grep -i "tx_green_prio_"
tx_green_prio_0: 6207
tx_green_prio_1: 0
tx_green_prio_2: 242
tx_green_prio_3: 0
tx_green_prio_4: 20039651
tx_green_prio_5: 0
tx_green_prio_6: 0
tx_green_prio_7: 16
root@ls1028ardb:~#
```

Figure 163. Sample log after checking the stats of TX packets of swp1

6.3.8 OPC UA client installation and usage

6.3.8.1 UaExpert

The UaExpert is an OPC UA Client developed by Unified Automation. It is free to download. Before downloading, you need to register on the following link to create a free account. Then login using your account, download the installation file and install it on a host PC. The UaExpert is available for both Windows and Linux.

<https://www.unified-automation.com/downloads/opc-ua-clients.html>

Below steps shows how to use UaExpert to connect to an OPC UA server on a Windows 10 PC.

- Open the UaExpert GUI. Click on the 'Add Server' button.




Figure 164. UaExpert GUI

- The 'Add Server' window will pop up. Select Custom Discovery and double click '< Double click to Add Server... >'. The 'Enter URL' window will pop up. Input IP address and port number of the OPC UA server separated by colon. For example, the complete URL is `opc.tcp://10.193.20.15:4840` in below snapshot. Click **OK**.

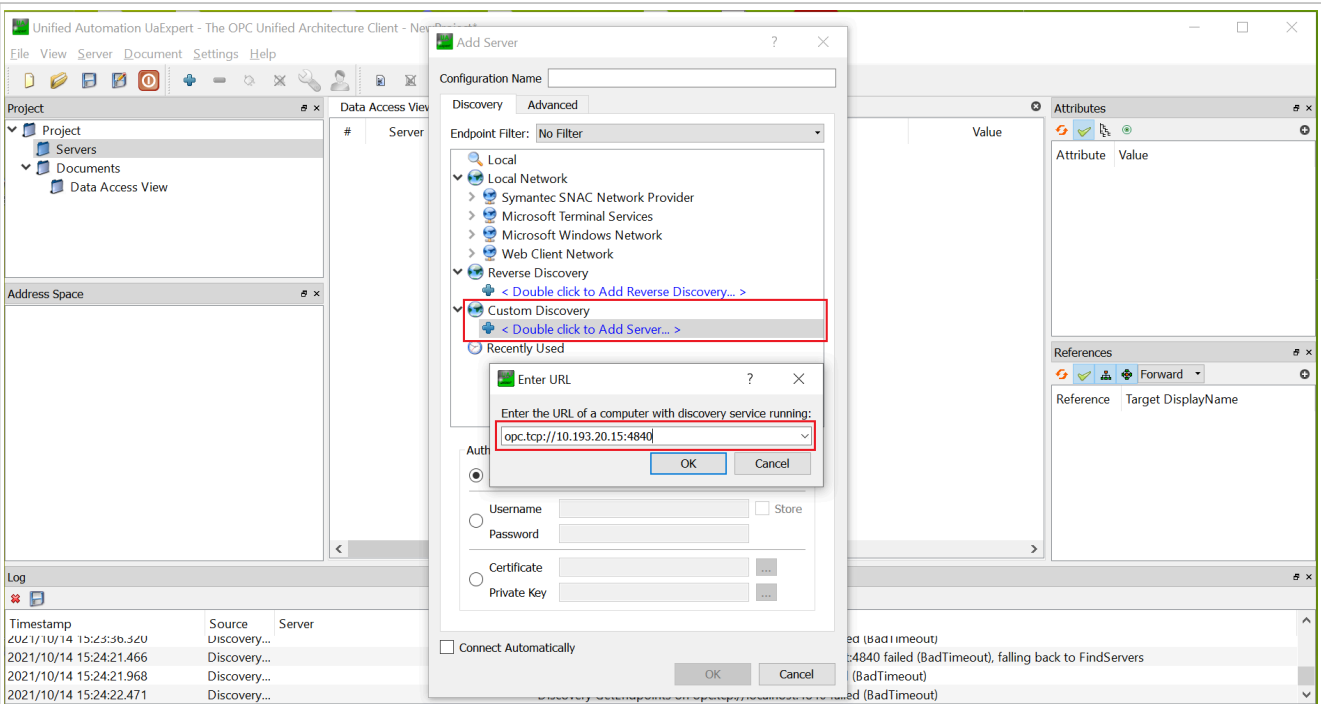


Figure 165. Adding the OPC UA server

- The new server (that is, *opc.tcp://10.193.20.15:4840*) will be listed under Custom Discovery. Click to expand it. Then click to expand 'open62541-based OPC UA Application (opc.tcp)'. A 'Replaced Hostname' window will pop up. Click 'Yes'.

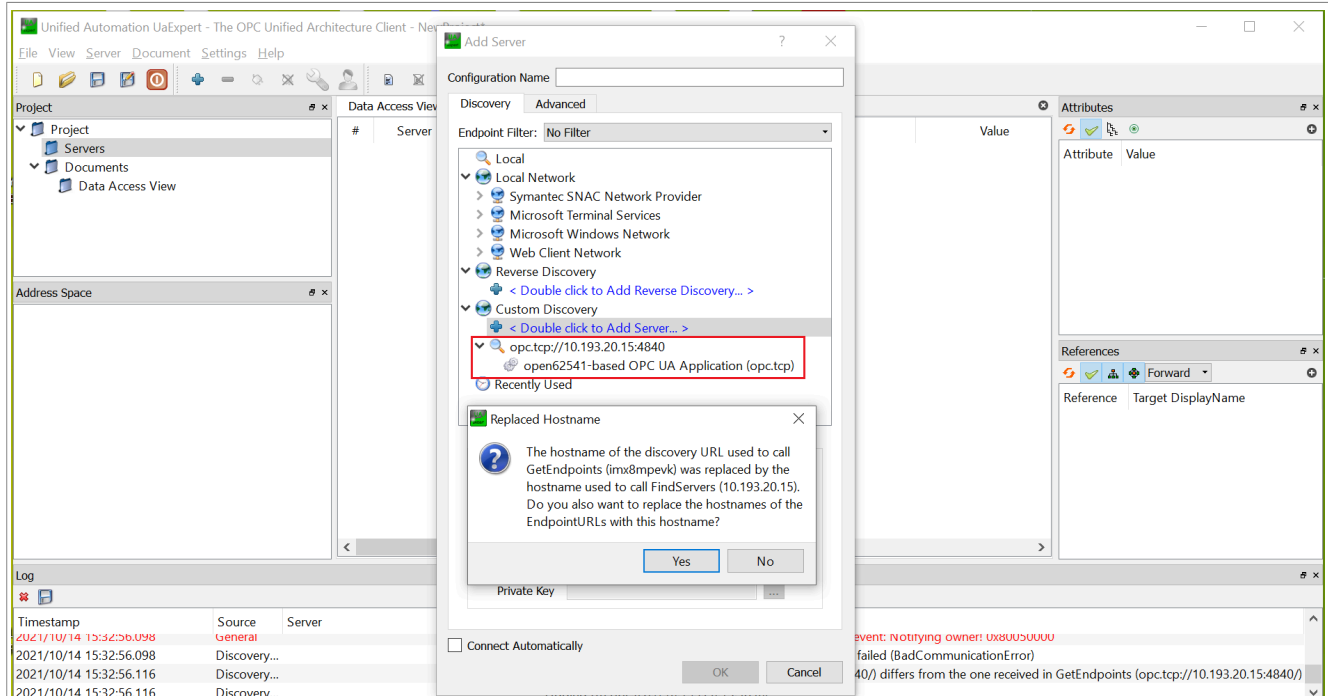


Figure 166. Server listed under Custom Discovery

- Click to select 'None – None (...)' and click OK.


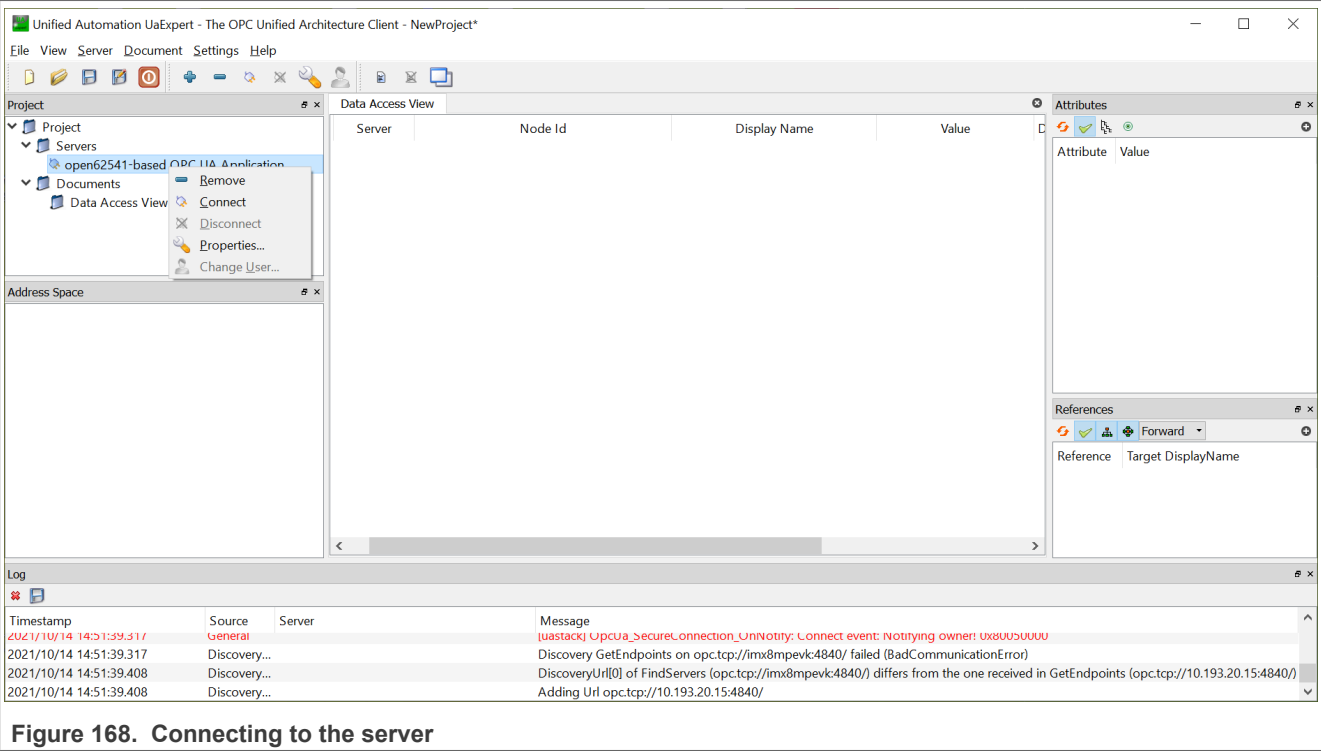



Figure 167. Selecting the hostname

- Right click on the server listed under 'Servers' and click 'Connect'.



- You are now connected to the OPC UA server and can browse or monitor its object. To monitor the value of an object, you can drag and drop the object to the 'Data Access View' area.



6.3.8.2 FreeOpcUa

FreeOpcUa is a project to implement an open-source (LGPL/GPL) OPC UA stack and associated tools. A GUI client from FreeOpcUa is available. It is written using freeopcua python api and pyqt. Use below command to install it on a Linux PC using pip3. Make sure python3 and python3-pip is installed.

```
$ sudo pip3 install opcua-client
```

For installation on Windows, please refer to the instructions available from below link:

<https://github.com/FreeOpcUa/opcua-client-gui>

Below steps shows how to use FreeOpcUa GUI client to connect to an OPC UA server on a Ubuntu 18.04 PC.

1) Launch the FreeOpcUa GUI client from the terminal on the Linux host PC:

```
$ opcua-client
```

In the FreeOpcUa GUI client, input the URL (that is, *opc.tcp://10.193.20.15:4840*) and click 'Connect'. You are now connected to the OPC UA server and can browse or monitor its object.

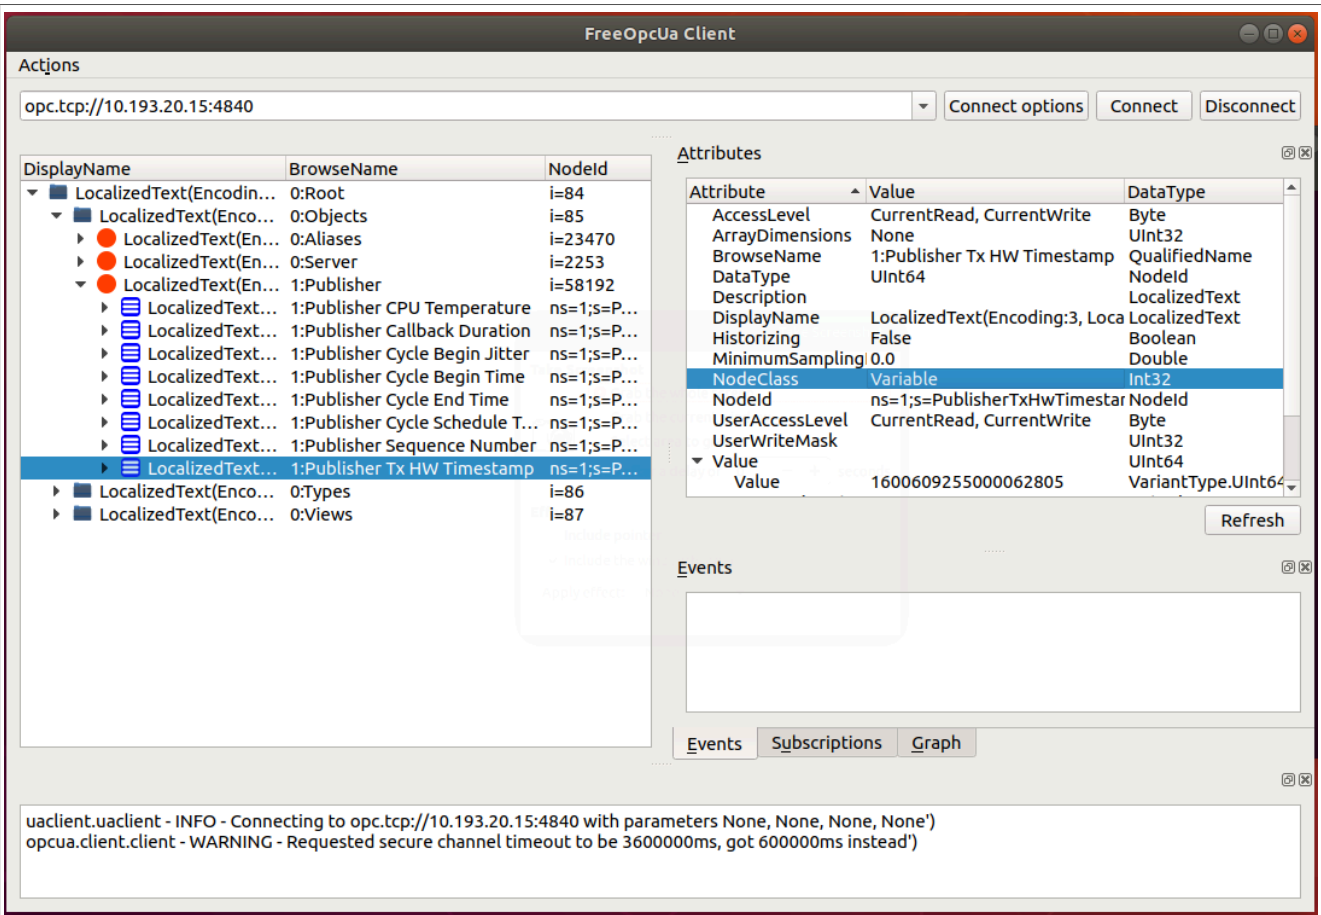


Figure 170. FreeOpcUa GUI client Connect options

6.4 NETCONF/YANG

This section provides an overview of the NETCONF protocol and Yang (a data modeling language for NETCONF). It describes the applications, installation and configuration steps, operation examples, Web UI demo, and troubleshooting aspects of NETCONF. It also describes how to enable the NETCONF feature in the Real-time Edge software.

6.4.1 Overview

The NETCONF protocol defines a mechanism for device management and configuration retrieval and modification. It uses a remote procedure call (RPC) paradigm and a system of exposing device (server) capabilities. These features enable a client to adjust to the specific features of any network equipment. NETCONF further distinguishes between state data (which is read only) and configuration data (which can be modified). Any NETCONF communication happens on four layers as shown in the table below. XML is used as the encoding format.

Table 94. The NETCONF layers

Layer	Purpose	Example
1	Content	Configuration data, Notification data
2	Operations	<edit-config>
3	Messages	<rpc>, <rpc-reply>, <notification>
4	Secure	Transport SSH

YANG is a standards-based, extensible, hierarchical data modeling language used to model the configuration and state data used by NETCONF operations, remote procedure calls (RPCs), and server event notifications. The device configuration data is stored in the form of an XML document. The specific nodes in the document and the allowed values are defined by a model, which is usually in YANG format or possibly transformed into YIN format with XML-based syntax. There are many such models created directly by IETF to support standardization and unification of the NETCONF interface of the common network devices.

For example, the general system settings of a standard computer are described in the IETF-system model ([rfc7317](#)). The configuration of its network interfaces defined by the IETF-interfaces model ([rfc7223](#)).

However, it is common for every system to have some specific parts exclusive to it. In such cases, there are mechanisms defined to enable extensions while keeping the support for the standardized core. This whole mechanism is designed in a liberal fashion, therefore, the configuration does not concern the network strictly. Even RPCs additional to those defined by NETCONF can be characterized. Therefore, it allows the client to request an explicit action from the server.

A YANG module defines a data model through its data, and the hierarchical organization of and constraints on that data. A module can be a complete, standalone entity, or it can reference definitions in other modules and sub-modules as well as augment other data models with additional nodes. The module dictates how the data is represented in XML.

A YANG module defines not only the syntax but also the semantics of the data. It explicitly defines relationships between and constraints on the data. This enables user to create syntactically correct configuration data that meets constraint requirements and enables user to validate the data against the model before uploading it and committing it on a device.

For information about NETCONF, see [RFC 6241](#), NETCONF Configuration Protocol.

For information about YANG, see [RFC 6020](#), YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF), and related RFCs.

6.4.2 Netopeer2

6.4.2.1 Overview

[Netopeer2](#) is a set of tools implementing network configuration tools based on the NETCONF protocol. This is the second generation of the toolset, originally available as the Netopeer project. It is based on the new generation of the NETCONF and YANG libraries - **libyang** and **libnetconf2**. The Netopeer2 server uses **sysrepo** as a NETCONF datastore implementation. It allows developers to control their devices via NETCONF and operators to connect to their NETCONF-enabled devices.

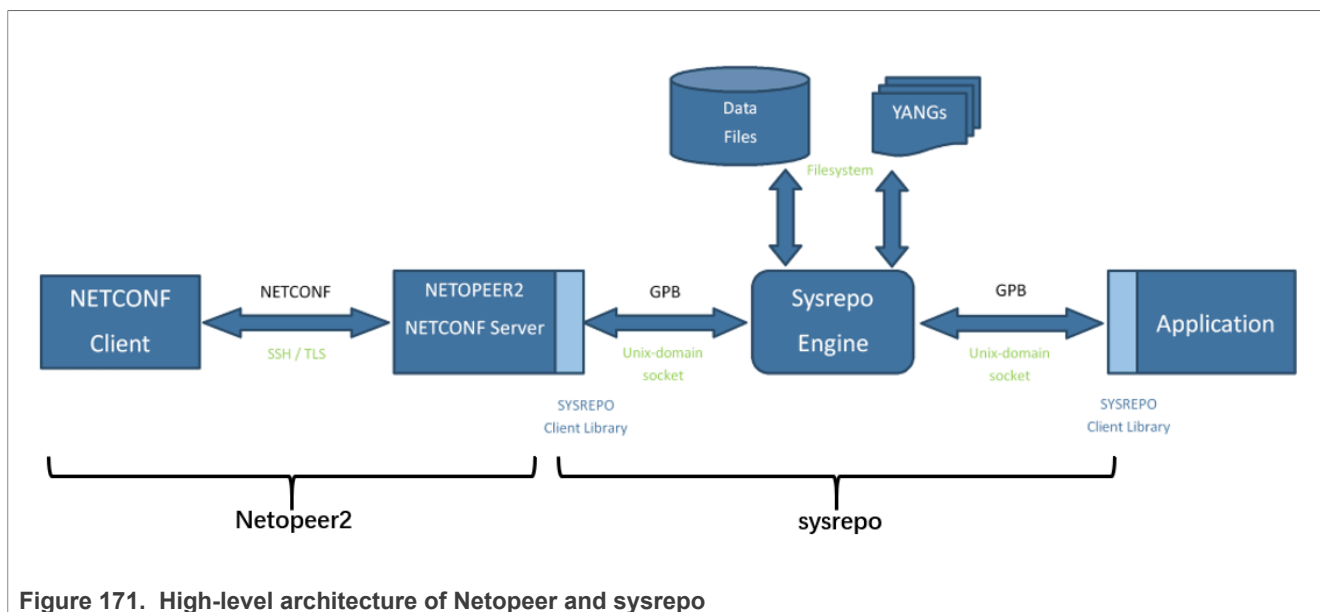


Figure 171. High-level architecture of Netopeer and sysrepo

6.4.2.2 Installing Netopeer2 on Ubuntu22.04

Use the following steps for installing Netopeer2-cli on Ubuntu22.04 operating systems.

1. Install the following packages:

```
$ sudo apt update
$ sudo apt install -y git cmake build-essential bison autoconf \
dh-autoreconf flex libavl-dev libprotobuf-c-dev protobuf-c-compiler \
zlib1g-dev libgcrypt20-dev libssh-dev libev-dev libpcrc3-dev
```

2. Install Netopeer2:

```
sudo apt install -y netopeer2 sysrepo-pluginind libyang2-tools
```

Attention: You may encounter an error as shown below:

```
/usr/share/netopeer2/merge_hostkey.sh: line 30: /dev/null: restricted: cannot
redirect output
dpkg: error processing package netopeer2 (--configure):
installed netopeer2 package post-installation script subprocess returned
error exit status 1
Errors were encountered while processing:
 netopeer2
E: Sub-process /usr/bin/dpkg returned an error code (1)
```

Use the below workaround for the above mentioned error:

```
sudo /usr/share/netopeer2/merge_hostkey.sh
sudo /usr/share/netopeer2/merge_config.sh
sudo apt install -f
```

6.4.2.3 Sysrepo

[Sysrepo](#) is an [YANG](#)-based configuration and operational state data store for Unix/Linux applications.

Applications can use sysrepo to store their configuration modeled by provided YANG model instead of using, for example, flat configuration files. Sysrepo ensures data consistency of the data stored in the datastore and enforces data constraints defined by the YANG model. Applications can currently use C language API of sysrepo Client Library to access the configuration in the datastore. However, the support for other programming languages is planned for later. Since sysrepo uses [Google Protocol Buffers \(GPB\)](#) as the interface between the datastore and client library. Therefore, writing of a native client library for any programming language that supports GPB is possible.

For information about sysrepo, refer to the URL: <https://netopeer.liberouter.org/doc/sysrepo/master/html/>.

6.4.2.4 Netopeer2 server

To implement network configuration management based on the NETCONF protocol, Netopeer2 includes a server (`netopeer2-server`). It is based on the new generation of the NETCONF and YANG libraries - libyang and libnetconf2. Netopeer2 uses sysrepo as a NETCONF datastore implementation.

6.4.2.5 Netopeer2 client

Netopeer2 includes a simple command-line NETCONF client, `netopeer2-cli`, which is built and installed by default. The Netopeer2 client is primarily used for Netopeer2 server testing. It allows the user to establish a NETCONF session with a NETCONF-enabled device to obtain or edit its configuration data.

6.4.2.6 Workflow in application practice

In practical application, we use the YANG language to model the device and generate the YANG model. The model is then instantiated to generate configuration files in XML format. The device was then configured using this configuration file as input via netopeer.




Figure 172. Workflow for netopeer

6.4.3 Configuration

6.4.3.1 Enabling NETCONF feature

This feature is enabled by default in Real-time Edge software.

The below packages are enabled by default in Real-time Edge software:

```
libyang libnetconf2 sysrepo netopeer2-server real-time-edge-sysrepo
```

sysrepo-tsn is daemon application to implement tsn configuration based on **sysrepo**. It is enabled for **LS1028ARDB**, **i.MX 8DXL LPDDR4 EVK**, **i.MX 8M Plus LPDDR4 EVK**, and **i.MX 93 EVK**.

Note: For **LS1028ARDB** board, *Qbv*, *Qbu*, *Qci*, stream identification in *CB*, *IP*, *MAC*, and *VLAN* are supported.

6.4.3.2 Netopeer2-server

The netopeer2-server is the NETCONF protocol server running as a system daemon. The netopeer2-server is based on sysrepo and libnetconf2 library.

6.4.3.3 Netopeer2-cli

6.4.3.3.1 Netopeer2 CLI commands

Following are the Netopeer2 CLI commands:

1. **help:** Displays a list of commands. The `--help` option is also accepted by all commands to show detailed information about the command.
2. **connect:** Connects to a NETCONF server.

```
connect [--help] [--ssh] [--host <hostname>] [--port <num>] [--login <username>]
```

The **connect** command has the following arguments:

- **--login** user name: Specifies the user to log in as on the NETCONF server. If not specified, current local user name is taken.
 - **--port** num
 - Port to connect to on the NETCONF server. By default, port 830 for SSH transport is used.
 - **host**
 - Hostname or ip-address of the target NETCONF server.
3. **disconnect**: disconnects from a NETCONF server.
 4. **commit**
 - Performs the NETCONF `commit` operation. For details, see RFC 6241, section 8.3.4.1.
 5. **copy-config**: Performs NETCONF `copy-config` operation. For details, see RFC 6241 section 7.3.

```
copy-config [--help] --target running|startup|candidate|url:<url> (--source
  running|startup|candidate|url:<url> | --src-config[=<file>])
  [--defaults report-all|report-all-tagged|trim|explicit]
```

Where, the arguments are the following:

- **--defaults** mode: Use: with the `--defaults` capability with specified retrieval mode. For details, refer to the RFC 6243 section 3 or *WITH-DEFAULTS* section of this manual.
 - **--target** datastore: Specifies the target datastore for the `copy-config` operation. For description of the datastore parameter, refer to [Section 6.4.3.3.2](#).
 - **--source** datastore: Specifies the source datastore for the `copy-config` operation. For description of the datastore parameter, refer to [Section 6.4.3.3.2](#).
6. **delete-config** Performs NETCONF `delete-config` operation. Refer to section 7.4 of the RFC 6241 specification for more details.

```
delete-config [--help] --target startup|url:<url>
```

Where

- **target** datastore: Specifies the target datastore for the `delete-config` operation.
7. **edit-config**

Performs NETCONF `edit-config` operation. For details, refer to RFC 6241 section 7.2.

```
edit-config [--help] --target running|candidate (--config[=<file>] | --url
  <url>)
  [--defop merge|replace|none] [--test set|test-only|test-then-set] [--
  error stop|continue|rollback]
```

Where

- **--defop** operation:

Specifies default operation for applying configuration data.

 - `merge`: Merges configuration data at the corresponding level. By default, the value is `merge`.
 - `replace`: Edits configuration data completely replaces the configuration in the target datastore.
 - `none`: The target datastore is unaffected by the edit configuration data, unless and until the edit configuration data contains the operation attribute to request a different operation. For more information, see the EDIT-CONFIG section of RFC 6241.

Note: To delete non-mandatory items, *nc:operation="delete"* should be added into the end of start tag of the item to be deleted. At the same time, the namespace *xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0"* should also be added to the start tag of the root node. Mandatory items cannot be deleted individually. They can only be deleted with their parent node.
- **--error** action

Sets reaction to an error.

 - `stop`: Aborts the operation on first error. This is the default value.

- **Continue:** Continues to process configuration data on error. The error is recorded and negative response is returned.
 - **Rollback:** Stops the operation processing on error and restore the configuration to its complete state at the start of this operation. This action is available only if the server supports rollback-on-error capability (see RFC 6241 section 8.5).
 - **--test option**
Performs validation of the modified configuration data. This option is available only if the server supports :validate:1.1 capability (see RFC 6241 section 8.6).
 - **set:** Does not perform validation test.
 - **test-only:** Does not apply the modified data, only performs the validation test.
 - **test-then-set:** Performs a validation test before attempting to apply modified configuration data. `test-then-set` is the default value.
 - **--config file**
 - Specifies path to a file containing edit configuration data. The content of the file is placed into the `config` element of the edit-config operation. Therefore, it does not have to be a well-formed XML document with only a single root element. If neither `--config` nor `--url` is specified, user is prompted to write edit configuration data manually. For examples, see the EDIT-CONFIG section of RFC 6241.
 - **--url URI**
 - Specifies remote location of the file containing the configuration data hierarchy to be modified, encoded in XML under the element `config` in the `urn:ietf:params:xml:ns:netconf:base:1.0` namespace. Note, that this differs from file parameter, where the `config` element is not expected.
 - **--target**
 - Target datastore to modify. For description of possible values, refer to [Section 6.4.3.3.2](#). Note that the url configuration datastore cannot be modified.
8. **get:** Performs NETCONF `get` operation. Receives both the status as well as configuration data from the current running datastore. Refer to section 7.7 of the RFC 6241 specification for more details. The command format is as follows:

```
get [--help] [--filter-subtree[=<file>]] | --filter-xpath <XPath>] [--defaults
report-all|report-all-tagged|trim|explicit] [--out <file>]
```

- **--defaults mode**
 - Use with the `-defaults` capability with specified retrieval mode. For further details, refer to the Section 3 or 'WITH-DEFAULTS' section of the RFC 6241 specification.
 - **--filter [file]**
 - Specifies if the request will contain subtree filter (RFC 6241 section 6). The option is able to accept path to the file containing the filter specification. If the path is not specified, user is prompted to write the filter specification manually.
9. **get-config** Performs NETCONF `get-config` operation. Retrieves only configuration data from the specified `target_datastore`. For details, refer to RFC 6241 section 7.1.

```
get-config [--help] --source running|startup|candidate [--filter-
subtree[=<file>]] | --filter-xpath <XPath>]
```

```
[--defaults report-all|report-all-tagged|trim|explicit] [--out <file>]
```

10. **--defaults** mode

- Use: with the `--defaults` capability with specified retrieval mode. For more details see RFC 6243 section 3 or WITH-DEFAULTS section of this manual.

11. **--filter** [file]

- Specifies if the request will contain subtree filter (RFC 6241 section 6). The option is able to accept path to the file containing the filter specification. If the path is not specified, user is prompted to write the filter specification manually.

12. **--target**

- Target datastore to retrieve. For description of possible values, refer to [Section 6.4.3.3.2](#). Note, that the url configuration datastore cannot be retrieved.

13. **lock**

Performs the NETCONF `lock` operation to lock the entire configuration datastore of a server. For details, see RFC 6241 section 7.5.

```
lock [--help] --target running|startup|candidate
```

Where the

- **--target:** specifies the target datastore to lock. For description of possible values, refer to [Section 6.4.3.3.2](#). Note, that the url configuration datastore cannot be locked.

14. **unlock:** Performs the NETCONF `unlock` operation to release a configuration lock, previously obtained with the `lock` operation. Refer to section 7.6 of the RFC 6241 specification for more details.

```
unlock [--help] --target running|startup|candidate
```

where

- **--target:** specifies the target datastore to unlock. For description of possible values, refer to [Section 6.4.3.3.2](#). Note, that the url configuration datastore cannot be unlocked.

15. **verb**

- Enables/disables verbose messages.

16. **quit**

- Quits the program.

6.4.3.3.2 Netopeer2 CLI datastore

Following are the netopeer2 CLI datastores:

• **running**

- This is the base NETCONF configuration datastore holding the complete configuration currently active on the device. This datastore always exists.

• **startup**

- The configuration datastore holding the configuration loaded by the device when it boots. This datastore is available only on servers that implement the `:startup` capability.

• **candidate**

- The configuration datastore that can be manipulated without impacting the device's current configuration and that can be committed to the running configuration datastore. This datastore is available only on servers that implement `:candidate` capability.

• **url:URI**

- Refers to a remote configuration datastore located at URI. The file that the URI refers to contains the configuration data hierarchy to be modified, encoded in XML under the element `config` in the `urn:ietf:params:xml:ns:netconf:base:1.0` namespace. This datastore is available only on servers that implement the `:url` capability.

6.4.3.4 Sysrepocfg

The **sysrepocfg** command-line tool allows editing, importing, and exporting configuration stored in the Sysrepo datastore. It also allows users to edit the startup or running configuration of a specified module in a preferred text editor. It propagates the performed changes into the datastore transparently for all subscribed applications. Moreover, the user can export the current configuration into a file or get it printed to the standard output. Similarly, an already prepared configuration can be imported from a file or read from the standard input.

In the background, **sysrepocfg** uses the Sysrepo client library for any data manipulation rather than directly accessing configuration data files. Therefore, it effectively inherits all main features of Sysrepo, such as YANG-based data validation, transaction, and concurrency support. Most importantly, subscribed applications are notified about the changes made using **sysrepocfg** and can immediately take the new configuration into account.

6.4.3.5 Sysrepoctl

The **sysrepoctl** provides functions to manage modules. It can be configured using the options and commands described below.

operation-operations

- **--help**: Prints the generic description and a list of commands. The detailed description and list of arguments for the specific command are displayed by using **--help** argument of the command.
- **--install**: Installs specified schema into sysrepo (**--yang** or **--yin** must be specified).
- **--uninstall**: Uninstalls specified schema from sysrepo (**--module** must be specified).
- **--list**: Lists YANG modules installed in sysrepo (note that Conformance Installed implies also Implemented).
- **--change**: Changes specified module in sysrepo (**--module** must be specified).

Other-options

- **--yang**: Path to the file with schema in YANG format (**--install** operation).
- **--yin**: Path to the file with schema in YIN format (**--install** operation).
- **--module**: Name of the module to be operated on (**--change**, **--feature-enable**, **--feature-disable** operations, **--uninstall** - several modules can be delimited with ',').
- **--permissions**: Access permissions of the module's data in chmod format (**--install**, **--change** operations).

Examples

- Installs a new module by specifying YANG file, ownership and access permissions:

```
sysrepoctl --install /home/user/ietf-interfaces.yang --owner root --group root --permissions 600
```

- Changes the ownership and permissions of an existing YANG module:

```
sysrepoctl --change ietf-interfaces --owner root --group root --permissions 644
```

- Enables a feature within a YANG module:

```
sysrepoctl --change ietf-interfaces --enable-feature if-mib
```

- Uninstall a module:

```
sysrepoctl --uninstall ietf-interfaces
```

6.4.3.6 List of yang models

Table 95. Revision of yang models

YANG Model Name	Revision	Enabled Features
ietf-interfaces	2018-02-20	
ietf-ip	2018-02-22	ipv4-non-contiguous-netmasks
ieee802-dot1q-bridge	2023-10-26	
ieee802-dot1q-sched	2023-10-22	scheduled-traffic
ieee802-dot1q-sched-bridge	2023-10-26	
ieee802-dot1q-preemption	2023-10-26	frame-preemption
ieee802-dot1q-preemption-bridge	2023-10-26	frame-preemption
ieee802-dot1q-stream-filters-gates	2023-07-03	closed-gate-state
ieee802-dot1q-stream-filters-gates-bridge	2023-07-03	
ieee802-dot1q-psfp	2023-10-26	psfp
ieee802-dot1q-psfp-bridge	2023-10-26	
ieee802-dot1cb-stream-identification	2021-12-08	
ieee802-dot1cb-stream-identification-types	2021-12-09	
ieee802-dot1cb-frer	2021-12-08	
ieee802-dot1cb-frer-types	2021-12-08	
ieee802-dot1q-pb	2023-10-22	
ieee802-dot1q-qci-augment	2019-05-20	
nxp-bridge-vlan-tc-flower	2020-04-02	

6.4.3.7 Operation examples

The following figure describes the steps to configure the device via netopeer2:

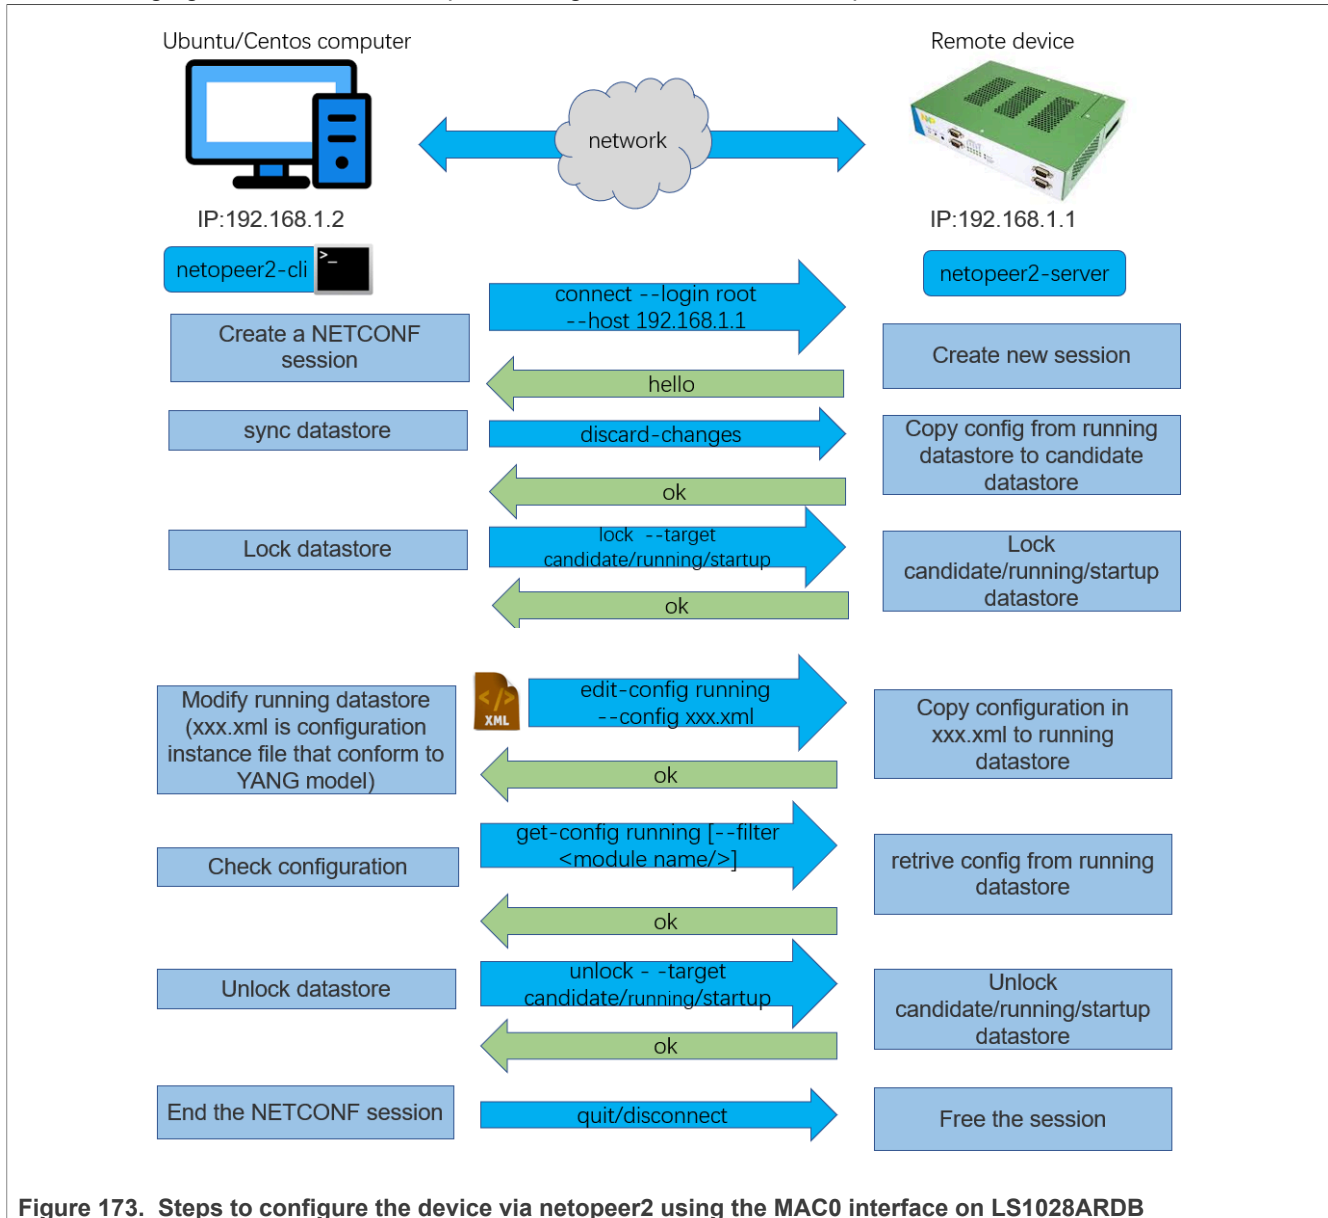


Figure 173. Steps to configure the device via netopeer2 using the MAC0 interface on LS1028ARDB

The [sysrepo-tsn](#) application includes few *instance files* to configure TSN features on the board:

- [Instance files for TSN configuration](#)

Users can configure the TSN functions using these instance files. Before starting, ensure that **sysrepo-pluginind**, **sysrepo-tsn**, and **netopeer2-server** are running on the board. Use the following steps to configure the TSN feature.

1. On a Linux command console, execute the netopeer2-cli command to start:

```
$ netopeer2-cli
```

2. Connect to netopeer2-server on the board (this example assumes that the IP address of the board is 10.193.20.53):

```
> connect --login root --host 10.193.20.53
```

3. Get the status data of the server:

```
> get
```

4. Get the configuration data in the running datastore:

```
> get-config --source running
```

5. Configure the QBV feature:

```
> edit-config --target running --config=qbv-swp0-enable.xml
```

6. Check the configuration data of QBV:

```
> get-config --source running --filter-xpath /ietf-interfaces:interfaces/  
interface[name='swp0']/ieee802-dot1q-bridge:bridge-port/ieee802-dot1q-  
preemption-bridge:frame-preemption-parameters
```

7. Copy configuration data in the **running** datastore to **startup** datastore:

```
> copy-config --source running --target startup
```

8. Disconnect the netopeer2-server:

```
> disconnect
```

6.4.3.8 Application scenarios on LS1028A

Note: The related xml file in the following cases can be obtained from the link: <https://github.com/nxp-real-time-edge-sw/real-time-edge-sysrepo/blob/master/Instances>.

Note: The interface name in the xml file should match with the actual interface name used on the board.

1. Prerequisites:

- a. Start netopeer2-cli on the computer with netopeer2-cli installed:

```
$ netopeer2-cli
```

- b. Connect to the netopeer2-server using the command below:

```
> connect --login root --host 10.193.20.53
```

2. Configure the IP address:

- a. Edit the configuration file, change Ethernet interface name and IP:

```
$ vim ietf-ip-cfg.xml
```

- b. Send the configuration file:

```
> edit-config --target running --config=ietf-ip-cfg.xml
```

3. Configure the MAC address for the bridge:

- a. Create a bridge named br1.

Note: The command only applies to the board with a switch, for example, LS1028ARDB.

```
$ ip link add name br1 type bridge
```

- b. Edit the configuration file, change bridge name and MAC:

```
$ vim ietf-mac-cfg.xml
```

- c. Send the configuration file:

```
$ edit-config --target running --config=ietf-mac-cfg.xml
```

4. Add VLAN for Ethernet interface:

- a. Create bridge named "br1" if not existing:

Note: Enter this command on the device console

```
$ ip link add name br1 type bridge
```

- b. Edit the configuration file to change the interface name and VLAN ID:

```
$ vim ietf-vlan-cfg.xml
```

- c. Send the configuration file:

```
> edit-config --target running --config=ietf-vlan-cfg.xml
```

5. Configure LS1028ARDB Qbv via tc.

- a. Edit the configuration file to change the interface name and VLAN ID:

```
$ vim qbv-swp0-enable.xml
```

- b. Send the configuration file:

```
> edit-config --target running --config=qbv-swp0-enable.xml
```

- c. Show the result.

Note: Enter this command on the device console

```
# tc qdisc show dev swp0
```

Note: If using tc or ethtool commands instead of libtsn, enable "real-time-edge-sysrepo-tc" in conf/distro/include/real-time-edge-base.inc as shown below:

```
REAL_TIME_EDGE_SYSREPO_ls1028ardb = "real-time-edge-sysrepo-tc"
```

Otherwise, disable "real-time-edge-sysrepo-tc":

```
REAL_TIME_EDGE_SYSREPO_ls1028ardb = ""
```

- For LS1028ARDB board, if real-time-edge-sysrepo-tc is enabled, you should set prerequisite for swpx (swp0 swp1 or swp2 ...) port using the following commands:

```
# tc qdisc add dev swpx ingress
# tc filter add dev swpx ingress chain 0 pref 49152 flower skip_sw action
  goto chain 10000
# tc filter add dev swpx ingress chain 10000 pref 49152 flower skip_sw
  action goto chain 11000
# tc filter add dev swpx ingress chain 11000 pref 49152 flower skip_sw
  action goto chain 12000
# tc filter add dev swpx ingress chain 12000 pref 49152 flower skip_sw
  action goto chain 20000
# tc filter add dev swpx ingress chain 20000 pref 49152 flower skip_sw
  action goto chain 21000
# tc filter add dev swpx ingress chain 21000 pref 49152 flower skip_sw
  action goto chain 30000
```

6. Configure LS1028ARDB Qci via tc using the steps below.

- a. Create a bridge named "switch" if not existing

Note: Enter this command on the device console.

```
# ip link add name switch type bridge
```

- b. Edit and send configuration file:

```
edit-config --target running --config=switch-qci-fm-gate-enable.xml
```

- c. Show the result.

Note: Enter this command on the device console.

```
# tc filter show dev swp0 ingress
```

- d. Disable the configuration.

```
> edit-config --target running --config=switch-qci-fm-gate-disable.xml
```

Note:

- The destination-address in instance file should be learned by switch.
- Users should send switch-qci-fm-gate-disable.xml after switch-qci-fm-gate-enable.xml

7. Configure LS1028ARDB Qbu via ethtool using the steps below.

Note: Disable cut through on the target board first by executing the command below:

```
# tsntool ctset --device swp0 --queue_stat 0x0
```

- a. Edit and send configuration file:

```
> edit-config --target running --config=qbu-swp0.xml
```

- b. Show the result:

```
# ethtool --show-frame-preemption swp0
```

8. Configure LS1028ARDB VLAN ID and priority filter via tc:

- a. Edit configuration file, change the interface name and action_spec:

```
$ vim ietf-br-vlan-cfg.xml
```

- b. Send the configuration file:

```
> edit-config --target running --config=ietf-br-vlan-cfg.xml
```

9. Configure i.MX 8DXL / i.MX 8M Plus / i.MX 93 Qbv via tc.

- a. Edit and send configuration file:

```
> edit-config --target running --config=qbv-eth1-enable.xml
```

- b. Display the result using the command below:

```
# tc qdisc show dev eth1
```

10. Configure i.MX 8DXL / i.MX 8M Plus / i.MX 93 Qbu via ethtool.

- a. Edit and send configuration file:

```
> edit-config --target running --config=qbu-eth1.xml
```

- b. Display the result using the command below:

```
# ethtool --show-frame-preemption eth1
```

6.4.4 Troubleshooting

1. Connection fails at client side:

```
nc ERROR: Remote host key changed, the connection will be terminated!
nc ERROR: Checking the host key failed.
cmd_connect: Connecting to the 10.193.20.4:830 as user "root" failed.
```

Fixing:

The reason is that the SSHD key changed at the server.

- First, users should get host list using the command `knownhosts`.
- Then, remove the related item. For example `knownhosts --del 19`.

2. Request could not be completed because the relevant data model content does not exist.

```
type:      application
tag:       data-missing
severity:  error
path:      /ietf-interfaces:interfaces/interface[name='eno0']/ieee802-dot1q-
sched:gate-parameters/admin-gate-states
message:   Request could not be completed because the relevant data model
           content does not exist.
```

Fixing:

The reason is that the configuration data in xpath does not exist in the datastore. Such as deleting a node that does not exist.

When encountering such an error, user can get configuration data in the board with `get-config` command, and check whether the operation type (`add/delete/modify`) of the node in the path is reasonable or not.

6.5 Wireless on LS1028A

6.5.1 NFC

NFC click board is a mikroBUS add-on board with a versatile near field communications controller from NXP — the [PN7120](#) NFC devices are used in contactless payment systems, electronic ticketing, smartcards. In retail and advertising, inexpensive NFC tags can be embedded into packaging labels, flyers, or posters.

This board is fully compliant with NFC Forum specifications. This implies that users can use the full potential of NFC and its three distinct operating modes listed below:

1. Card emulation
2. Read/Write
3. P2P

6.5.1.1 Introduction

The NXP's PN7120 IC integrates an Arm Cortex-M0 MCU, which enables easier integration into designs, because it requires fewer resources from the host MCU. The integrated firmware provides all NFC protocols for performing the contactless communication in charge of the modulation, data processing, and error detection.

The board communicates with the target board MCU through the mikroBUS I2C interface, in compliance with NCI (NFC controller interface) 1.0 host protocols. RST and INT pins provide additional functionality. The board uses a 3.3 V power supply.

6.5.1.2 PN7120 features

PN7120 IC embeds a new generation RF contactless front-end, supporting various transmission modes according to NFCIP-1 and NFCIP-2, ISO/IEC14443, ISO/IEC 15693, ISO/IEC 18000-3, MIFARE, and FeliCa specifications. It embeds an Arm Cortex-M0 microcontroller core loaded with the integrated firmware supporting the NCI 1.0 host communication.

6.5.1.3 Hardware preparation

Use the following hardware items for the NFC clickboard demo setup:

1. LS1028ARDB
2. NFC click board
3. NFC sample card (tag)

Note: Users should insert the NFC click board into the LS1028ARDB mikroBUS1 slot.

6.5.1.4 Software preparation

In order to support NFC click board, use the following steps:

1. In Real-time Edge, libnfc-nci is enabled by default.
2. In Linux kernel configuration, make sure the below options are enabled:

```
[*] Networking support --->
    <M>   NFC subsystem support    --->
        Near Field Communication (NFC) devices --->
            <M>   NXP PN5XX based driver
```

Note: The **NXP PN5XX based driver** only supports the Module mode.

3. Use the `make` command to create the images.

6.5.1.5 Testing the NFC click board

Use the following steps for testing the NFC Clickboard:

1. Install NFC driver module

```
[root]# modprobe pn5xx_i2c.ko
```

2. The following log appears at the console after the above command is successful. The error information can be ignored in this case.

```
[root@OpenIL:~]# insmod /lib/modules/4.14.47-ipipe/kernel/[ 195.547601] random: crng init
done
[ 195.551016] random: 5 urandom warning(s) missed due to ratelimiting
[root@OpenIL:~]# insmod /lib/modules/4.14.47-ipipe/kernel/drivers/misc/nxp-pn5xx
/pn5xx_i2c.ko
[ 777.503246] pn54x_dev_init
[ 777.506048] pn54x_probe
[ 777.508523] pn544 7-0028: FIRM GPIO <OPTIONAL> error getting from OF node
[ 777.515344] pn544 7-0028: CLKREQ GPIO <OPTIONAL> error getting from OF node
[ 777.522347] pn544 7-0028: 7-0028 supply nxp,pn54x-pvdd not found, using dummy regulator
[ 777.530424] pn544 7-0028: 7-0028 supply nxp,pn54x-vbat not found, using dummy regulator
[ 777.538490] pn544 7-0028: 7-0028 supply nxp,pn54x-pmuvcc not found, using dummy regulator
[ 777.546723] pn544 7-0028: 7-0028 supply nxp,pn54x-sevdd not found, using dummy regulator
```

3. Run the **nfcDemoApp** application:

```
[root]# nfcDemoApp poll
```

```
[root@OpenIL:~]# nfcDemoApp poll
#####
##                                NFC demo                                ##
#####
##                                Poll mode activated                        [ 1251.20807
1] pn54x_dev_open : 10,55
##
####[ 1251.212807] pn54x_dev_ioctl, cmd=1074063617, arg=1
#####[ 1251.219006] pn544_enable power on
#####
... press enter to quit ...

[ 1251.431597] pn54x_dev_ioctl, cmd=1074063617, arg=0
[ 1251.436416] pn544_disable power off
[ 1251.647586] pn54x_dev_ioctl, cmd=1074063617, arg=1
[ 1251.652401] pn544_enable power on
NfcHcpX:8103
NfcHcpR:8180
NfcHcpX:810103020304
NfcHcpR:8180
NfcHcpX:81010143da67663bda6766
NfcHcpR:8180
NfcHcpX:810204
NfcHcpR:818000
Waiting for a Tag/Device...
```

4. Put the NFC Sample Card (tag) on top of the NFC click board:

```
Waiting for a Tag/Device...

NFC Tag Found

Type :          'Type A - Mifare UL'
NFCID1 :        '04 67 66 D2 9C 39 81 '
Record Found :
                NDEF Content Max size :    '868 bytes'
                NDEF Actual Content size :   '29 bytes'
                ReadOnly :                  'FALSE'

Read NDEF URL Error

29 bytes of NDEF data received :
D1 01 19 55 01 6E 78 70 2E 63 6F 6D 2F 64 65 6D 6F 62 6F 61 72 64 2F 4F 4D
35 35 37 38

NFC Tag Lost

Waiting for a Tag/Device...
```

Display of the above information indicates successful card reading.

6.5.2 Bluetooth Low Energy

This chapter introduces the features of the Bluetooth Low Energy P click board and how to use it on NXP's LS1028A reference design board (RDB).

6.5.2.1 Introduction

The **BLE P Click** board carries the nRF8001 IC that allows user to add Bluetooth 4.0 to user's device. The click communicates with the target board MCU through mikroBUS SPI (CS, SCK, MISO, MOSI), RDY and ACT lines, and runs on 3.3 V power supply.

The **BLE P Click** board features a PCB trace antenna, designed for the 2400 MHz to 2483.5 MHz frequency band. The maximum device range is up to 40 meters in open space.

6.5.2.2 Bluetooth Low Energy

LS1028ARDB support Bluetooth Low Energy click board, Bluetooth Low Energy P click carries the nRF8001 IC that allows user to add Bluetooth 4.0 to the device.

6.5.2.3 Features

Following are the features provided by BLE P Click board:

- nRF8001 Bluetooth low energy RF transceiver
 - 16 MHz crystal oscillator
 - Ultra-low peak current consumption <14 mA
 - Low current for connection-oriented profiles, typically 2 µA
- PCB trace antenna (2400-2483.5 MHz, up to 40 meters)
- BLE Android app
- Interface: SPI (CS, SCK, MISO, MOSI), RDY, and ACT lines
- 3.3 V power supply

6.5.2.4 Hardware preparation

Use the following hardware items for the BLE P Click board demo setup:

1. LS1028ARDB
2. BLE P Click board
3. Android phone (option)

The figure below depicts the hardware setup required for the demo:



Figure 174. BLE P Click board hardware setup

6.5.2.5 Software preparation

Use these steps for the BLE P click board demo software setup:

- Download the JUMA UART (Android app) by using the link: <https://apkpure.com/juma-uart/com.juma.UART>
- Then, run the steps below in order to support BLE P click board:
 1. In Real-time Edge, libblep is enabled by default.
 2. In Linux kernel configuration, make sure the below options are enabled:

```
Device Drivers --->
  SPI support --->
    <*>   Freescale DSPI controller
    <*>   User mode SPI device driver support
```

3. Use the `make` command to create the images.

6.5.2.6 Testing the BLE P click board

Use the following steps for testing the BLE P click board:

1. Running the blep_demo application:

The following log is displayed to indicate that the BLE P click board is initialized. After this, users can scan from their mobile phone or computer's Bluetooth device for the BLE P click board. The name of the BLE P click board used is “MikroE”.

```
root@OpenIL-Ubuntu:~# blep_demo
spi mode: 0x0
bits per word: 8
max speed: 500000 Hz (500 KHz)
Please input a command!
Event device started: Setup
Error:no
Start setup command
.....
Setup complete
Event device started: Standby
Advertising started : Tap Connect on the nRF UART app
Error:no
Send broadcast command successfully
```

Figure 175. Initialization log

2. Connection log

Connect the BLE P click board via mobile app. On successful connection, the following log is displayed. Thereafter, the application can communicate with the BLE P click board.

```
Evt Connected
Evt Pipe Status
Evt link connection interval changed
ConnectionInterval:0x0006
SlaveLatency:0x0000
SupervisionTimeout:0x01F4
Evt Pipe Status
Evt link connection interval changed
ConnectionInterval:0x0027
SlaveLatency:0x0000
SupervisionTimeout:0x01F4
```

Figure 176. Connection log

3. Disconnection log

Click the **Disconnect** button of the Android APP to disconnect from the BLE P click board. The following log displays that the disconnection is successful:

```
Evt Disconnected
Advertising started : Tap Connect on the nRF UART app Send broadcast command successfully
```

Figure 177. Disconnection log

4. Command line introduction

The **blep_demo** application supports four command lines: **devaddr**, **name=**, **version**, and **echo**.

a. devaddr

This command is used to obtain the MAC address of the BLE P click board. User can run this command at any time.

```
devaddr
Please input a command!
Device address:DC:E2:6C:17:07:45
```

Figure 178. devaddr command log

b. name=

This command is used to set the Bluetooth name of the BLE P click board while broadcasting. No spaces are required after the equal sign "=", and the content after the 'equal to' sign is the set name. The maximum length is 16 characters.

```
name=ble_demo
Name set. New name: ble_demo, 8
Please input a command!
Unknow event:0x00
Set local data successfully
```

Figure 179. name log

c. **version**

This command is used to obtain the version of the BLE P click board. User can run this command at any time.

```
version
Please input a command!
Unknow event:0x00
Device version
Configuration ID:0x41
ACI protocol version:2
Current setup format:3
Setup ID:0x00
Configuration status:open(VM)
```

Figure 180. version log

d. **echo**

This command is used to send a string to the Android app. This command should be executed after the connection is established. The maximum length is 20 characters.

The below log displays the message displayed after user tries to send a string when no connection is established:

```
echo hi
Please input a command!
Unknow event:0x00
ACI Evt Pipe Error: Pipe #9
Pipe Error Code: 0x83
Pipe Error Data: 0x00
Please connect the device before sending data
```

Figure 181. User input log (no connection)

The below log is displayed when user sends a string after a connection is established:

```
echo hello,world!
Please input a command!
Unknow event:0x00
The number of data command buffer is 1
```

Figure 182. User input log (after a connection is established)

5. **Receiving data**

When the Android app sends a string:

```
DataReceivedEvent: hi.yugxdr
```

Figure 183. Received data log

6.5.3 BEE

This chapter introduces the features of the BEE Click Board and how to use it on LS1028ARDB.

6.5.3.1 BEE/ZigBEE

LS1028ARDB supports BEE click board, which can implement the MRF24J40MA 2.4 GHz IEEE 802.15.4 radio transceiver module from Microchip.

6.5.3.2 Features

The features of the BEE Click Board are listed below:

- PCB antenna
- MRF24J40MA module
- Low current consumption (TX 23 mA, RX 19 mA, Sleep 2 μ A)
- ZigBee stack
- MiWi™ stack
- SPI Interface
- 3.3 V power supply

6.5.3.3 Hardware preparation

Use the following hardware items for the BEE Click Board demo setup:

- Two LS1028ARDB boards
- Two BEE Click Boards

The [Figure 184](#) describes the hardware setup for the BEE Click Board.




Figure 184. BEE Click Board hardware setup

Note: The WA pin of BEE Click Board connects with the NC pin.

6.5.3.4 Software preparation

In order to support BEE click board, use the following steps:

1. In Real-time Edge, `libbee` is enabled by default.
2. In Linux kernel configuration, make sure the below options are enabled:

```
Device Drivers --->
  SPI support --->
    <*>   Freescale DSPI controller
    <*>   User mode SPI device driver support
  *- GPIO Support --->
    [*]   /sys/class/gpio/... (sysfs interface)
          Memory mapped GPIO drivers --->
            [*] MPC512x/MPC8xxx/QorIQ GPIO support
```

3. Use the `make` command to create the images.

6.5.3.5 Testing the BEE click board

The test application `bee_demo` is created by using the BEE Click Board library. This application can transfer the file between two BEE Click Boards.

1. Users should create a file in any path. For example, `./samples/test.txt`.
2. First, start a server node by running the command below:

```
bee_demo -s -f=XXX
```

The command parameters are as below:

- **-s:** This device node acts as a server.
- **-f=XXX:** This parameter is valid only on the server node. XXX is the file path (relative or absolute) to be transferred.

```
[root]# ls
samples
[root]# bee_demo -s -f=./samples/test.txt
spi mode: 0x0
bits per word: 8
max speed: 500000 Hz (500 KHz)
BEE Click Board Demo.
This node is a server node.
Waiting for a client
Reading the content of the file
```

3. Start a client node on another LS1028ARDB by running the command `bee_demo -c`. In the above command, the parameter `-c` implies that this device node acts as a client. After receiving the file, the client node automatically exits. The received file is saved in the current path.

```
[root]# ls
samples
[root]# bee_demo -c
spi mode: 0x0
bits per word: 8
max speed: 500000 Hz (500 KHz)
BEE Click Board Demo.
This node is a client node.
Starting to get a file
Send the SEQ_REQ command.
Send the SEQ_START command.
Send the SEQ_START command.

[root]# ls
samples  test.txt
[root]#
```

4. The following log indicates that the server node has finished sending a file.

```
Send the SEQ_INFO command.
Start to send the file
It's completed to send a file.
```

6.6 SAI on LS1028ARDB

SAI on LS1028ARDB is enabled. Due to the pins are shared between IEEE 1588 and SAI, therefore the default setting is to enable IEEE 1588 and disable SAI.

Following below steps to enable SAI in Real-time Edge Linux.

1. Enable SAI support in Real-time Edge software

```
$ cd yocto-real-time-edge/sources/meta-real-time-edge
# Open file "conf/distro/include/real-time-edge-base.inc", add "sai" to
"DISTRO_FEATURES:append:ls1028ardb" like this:
DISTRO_FEATURES:append:ls1028ardb = " jailhouse real-time-edge-libbee real-
time-edge-libblep libnfc-nci \
    wayland-protocols weston imx-gpu-viv libdrm kmscube \
    real-time-edge-sysrepo tsn-scripts wayland alsa sai"
```

2. Build the image

```
$ cd yocto-real-time-edge
$ DISTRO=nxp-real-time-edge MACHINE=ls1028ardb source real-time-edge-setup-
env.sh -b build-ls1028ardb
$ bitbake nxp-image-real-time-edge
```

3. Turn on "Lineout Playback Switch" and boot up LS1028ARDB with new image, enter Linux prompt:

```
$ amixer -c 0 cset name='Lineout Playback Switch' on
numid=11,iface=MIXER,name='Lineout Playback Switch'
; type=BOOLEAN,access=rw-----,values=1
: values=on
```

4. Run audio test (insert 3.5 mm headphone in AUDIO port on LS1028ARDB, some voice can be heard from left and right).

```
$ speaker-test -c 2 -l 10 -t wav
speaker-test 1.2.5.1
Playback device is default
Stream parameters are 48000Hz, S16_LE, 2 channels
WAV file(s)
Rate set to 48000Hz (requested 48000Hz)
Buffer size range from 96 to 1048576
Period size range from 32 to 349526
Using max buffer size 1048576
Periods = 4
was set period_size = 262144
was set buffer_size = 1048576
 0 - Front Left
 1 - Front Right
Time per period = 7.964220
 0 - Front Left
 1 - Front Right
Time per period = 2.976679
 0 - Front Left
 1 - Front Right
...
```

6.7 Wi-Fi on i.MX 8DXL EVK

6.7.1 Wi-Fi card information

The Wi-Fi card shipped with i.MX 8DXL EVK is called '**1XL M.2 Module**' which is co-developed by Embedded Artists and Murata. The table below describes this Wi-Fi card.

Table 96. Wi-Fi card

Feature	Description
Manufacturer	Embedded Artists
Part Number	EAR00387
Module	Murata LBEE5ZZ1XL (also known as 1XL)
Chipset	NXP 88W9098
Wi-Fi Standards	802.11 a/b/g/n/ac/ax 2x2 MU-MIMO, Wi-Fi 6
Frequency	2.4 GHz and 5 GHz
Wi-Fi Interface	PCIe 3.0 (in M.2 form factor)

The NXP 88W9098 wireless SoC is the industry's first Wi-Fi 6 solution based on the latest IEEE 802.11ax standard with an innovative Concurrent Dual Wi-Fi (CDW) architecture. CDW supports separate Wi-Fi networks simultaneously using both 2.4 GHz and 5 GHz frequency bands.

6.7.2 Hardware Setup

Install the Wi-Fi card '**1XL M.2 Module**' to the M.2 connector (J17) on i.MX 8DXL EVK.

6.7.3 Software Enablement

The following instructions show how to bring up the Wi-Fi module on i.MX 8DXL EVK and connect to Access Point (AP) in Station Mode.

1. Boot up Linux using DTB `imx8dxl-evk-rpmsg.dtb`.

```
=> setenv fdt_file imx8dxl-evk-rpmsg.dtb
=> boot
```

2. Load the Wi-Fi kernel driver module `moal.ko`.

The module parameters are in configuration file `/lib/firmware/nxp/wifi_mod_para.conf`.

The NXP Wi-Fi SoCs require a firmware image to be loaded on power-up/reset. The firmware image for NXP 88W9098 with PCIe interface is `/lib/firmware/nxp/pcieuart9098_combo_v1.bin`.

```
# modprobe moal mod_para=nxp/wifi_mod_para.conf
```

- Verify the kernel debug messages printed in serial console. Notice the bold texts.

```
[ 44.855282] mlan: loading out-of-tree module taints kernel.
[ 44.893129] wlan: Loading MWLAN driver
[ 44.897479] wlan: Register to Bus Driver...
[ 44.903192] wlan_pcie 0000:01:00.0: enabling device (0000 -> 0002)
[ 44.909628] Attach moal handle ops, card interface type: 0x206
```



```
[ 44.916925] PCIE9098: init module param from usr cfg
[ 44.922056] card_type: PCIE9098, config block: 0
[ 44.926746] cfg80211_wext=0xf
[ 44.929821] max_vir_bss=1
[ 44.932503] cal_data_cfg=none
[ 44.935612] ps_mode = 1
[ 44.938076] auto_ds = 1
[ 44.940610] host_mlme=enable
[ 44.947387] fw_name=nxp/pcieuart9098_combo_v1.bin
[ 44.952178] rx_work=1 cpu_num=2
[ 44.955403] Attach mlan adapter operations.card_type is 0x206.
[ 44.967149] Request firmware: nxp/pcieuart9098_combo_v1.bin
[ 45.247953] FW download over, size 682784 bytes
[ 45.510922] WLAN FW is active
[ 45.513922] on_time is 79996054375
[ 45.531982] VDLL image: len=160672
[ 45.535888] fw_cap_info=0xc8fcffa3, dev_cap_mask=0xffffffff
[ 45.541571] max_p2p_conn = 8, max_sta_conn = 64
[ 45.549005] wlan: mlan0 set max_mtu 2000
[ 45.569666] wlan: uap0 set max_mtu 2000
[ 45.584434] wlan: wfd0 set max_mtu 2000
[ 45.604604] wlan: version = PCIE9098--17.92.1.p116.1-MM5X17344.p3-GPL-
(FP92)
[ 45.617407] wlan_pcie 0000:01:00.1: enabling device (0000 -> 0002)
[ 45.623935] Attach moal handle ops, card interface type: 0x206
[ 45.630063] PCIE9098: init module param from usr cfg
[ 45.635290] card_type: PCIE9098, config block: 1
[ 45.640039] cfg80211_wext=0xf
[ 45.643271] max_vir_bss=1
[ 45.645925] cal_data_cfg=none
[ 45.650328] ps_mode = 1
[ 45.653099] auto_ds = 1
[ 45.655658] host_mlme=enable
[ 45.662576] fw_name=nxp/pcieuart9098_combo_v1.bin
[ 45.667565] rx_work=1 cpu_num=2
[ 45.670767] Attach mlan adapter operations.card_type is 0x206.
[ 45.687250] Request firmware: nxp/pcieuart9098_combo_v1.bin
[ 45.694421] WLAN FW already running! Skip FW download
[ 45.701756] WLAN FW is active
[ 45.704867] on_time is 80092507000
[ 45.710734] VDLL image: len=160672
[ 45.714827] fw_cap_info=0x68fcffa3, dev_cap_mask=0xffffffff
[ 45.720564] max_p2p_conn = 8, max_sta_conn = 64
[ 45.727958] wlan: mmmlan0 set max_mtu 2000
[ 45.741123] wlan: muap0 set max_mtu 2000
[ 45.747316] wlan: mwfd0 set max_mtu 2000
[ 45.754666] wlan: version = PCIE9098--17.92.1.p116.1-MM5X17344.p3-GPL-
(FP92)
[ 45.762585] wlan: Register to Bus Driver Done
[ 45.767017] wlan: Driver loaded successfully
```

3. Verify the Wi-Fi interface.

```
# ifconfig -a
```

Notice the bold texts in the output.

```
# ifconfig -a
```

```
.....
```

```

mlan0: flags=4098<BROADCAST,MULTICAST> mtu 1500
    ether 00:50:43:20:12:34 txqueuelen 1000  (Ethernet)
    RX packets 0  bytes 0 (0.0 B)
    RX errors 0  dropped 0  overruns 0  frame 0
    TX packets 0  bytes 0 (0.0 B)
    TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0
mmlan0: flags=4098<BROADCAST,MULTICAST> mtu 1500
    ether 00:50:43:20:52:56 txqueuelen 1000  (Ethernet)
    RX packets 0  bytes 0 (0.0 B)
    RX errors 0  dropped 0  overruns 0  frame 0
    TX packets 0  bytes 0 (0.0 B)
    TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0
muap0: flags=4098<BROADCAST,MULTICAST> mtu 1500
    ether 00:50:43:20:53:56 txqueuelen 1000  (Ethernet)
    RX packets 0  bytes 0 (0.0 B)
    RX errors 0  dropped 0  overruns 0  frame 0
    TX packets 0  bytes 0 (0.0 B)
    TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0
mwfd0: flags=4098<BROADCAST,MULTICAST> mtu 1500
    ether 02:50:43:20:52:56 txqueuelen 1000  (Ethernet)
    RX packets 0  bytes 0 (0.0 B)
    RX errors 0  dropped 0  overruns 0  frame 0
    TX packets 0  bytes 0 (0.0 B)
    TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0
uap0: flags=4098<BROADCAST,MULTICAST> mtu 1500
    ether 00:50:43:20:13:34 txqueuelen 1000  (Ethernet)
    RX packets 0  bytes 0 (0.0 B)
    RX errors 0  dropped 0  overruns 0  frame 0
    TX packets 0  bytes 0 (0.0 B)
    TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0
wfd0: flags=4098<BROADCAST,MULTICAST> mtu 1500
    ether 02:50:43:20:12:34 txqueuelen 1000  (Ethernet)
    RX packets 0  bytes 0 (0.0 B)
    RX errors 0  dropped 0  overruns 0  frame 0
    TX packets 0  bytes 0 (0.0 B)
    TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0

```

4. Enable Wi-Fi interface. Either `mlan0` or `mmlan0` can be used in the following steps.

```
# ifconfig mlan0 up
```

5. Scan the visible Wi-Fi access points.

```
# iw dev mlan0 scan
```

An sample output is below. Output for only one AP is shown for simplicity. Notice the bold text for the frequency and SSID.

```

# iw dev mlan0 scan
[ 88.895756] wlan: mlan0 START SCAN
[ 90.941865] wlan: SCAN COMPLETED: scanned AP count=11
.....
BSS 00:fe:c8:d3:00:6e(on mlan0)
    TSF: 90867505 usec (0d, 00:01:30)
    freq: 5320
    beacon interval: 102 TUs
    capability: ESS Privacy SpectrumMgmt RadioMeasure (0x1111)
    signal: -73.00 dBm

```

```
last seen: 4 ms ago
SSID: NXPOPEN
Supported rates: 18.0 24.0* 36.0 48.0 54.0
TIM: DTIM Count 0 DTIM Period 1 Bitmap Control 0x0 Bitmap[0] 0x0
Country: CN      Environment: Indoor/Outdoor
      Channels [36 - 48] @ 23 dBm
      Channels [52 - 64] @ 23 dBm
      Channels [149 - 165] @ 30 dBm
BSS Load:
  * station count: 8
  * channel utilisation: 34/255
  * available admission capacity: 23437 [*32us]
Power constraint: 0 dB
HT capabilities:
  Capabilities: 0x19ee
    HT20/HT40
    SM Power Save disabled
    RX HT20 SGI
    RX HT40 SGI
    TX STBC
    RX STBC 1-stream
    Max AMSDU length: 7935 bytes
    DSSS/CCK HT40
    Maximum RX AMPDU length 65535 bytes (exponent: 0x003)
    Minimum RX AMPDU time spacing: 8 usec (0x06)
    HT RX MCS rate indexes supported: 0-23
    HT TX MCS rate indexes are undefined
RSN:
  * Version: 1
  * Group cipher: CCMP
  * Pairwise ciphers: CCMP
  * Authentication suites: PSK
  * Capabilities: 4-PTKSA-RC 4-GTKSA-RC (0x0028)
HT operation:
  * primary channel: 64
  * secondary channel offset: below
  * STA channel width: any
  * RIFS: 0
  * HT protection: no
  * non-GF present: 1
  * OBSS non-GF present: 0
  * dual beacon: 0
  * dual CTS protection: 0
  * STBC beacon: 0
  * L-SIG TXOP Prot: 0
  * PCO active: 0
  * PCO phase: 0
RM enabled capabilities:
  Capabilities: 0x73 0xc0 0x00 0x00 0x00
    Link Measurement
    Neighbor Report
    Beacon Passive Measurement
    Beacon Active Measurement
    Beacon Table Measurement
    Transmit Stream/Category Measurement
    Triggered Transmit Stream/Category
    Nonoperating Channel Max Measurement Duration: 0
    Measurement Pilot Capability: 0
Extended capabilities:
  * Proxy ARP Service
  * BSS Transition
```

```

* DMS
* QoS Map
* WNM-Notification
* Operating Mode Notification
* Max Number Of MSDUs In A-MSDU is unlimited
VHT capabilities:
  VHT Capabilities (0x0f8379b2):
    Max MPDU length: 11454
    Supported Channel Width: neither 160 nor 80+80
    RX LDPC
    short GI (80 MHz)
    TX STBC
    SU Beamformer
    SU Beamformee
  VHT RX MCS set:
    1 streams: MCS 0-9
    2 streams: MCS 0-9
    3 streams: MCS 0-9
    4 streams: not supported
    5 streams: not supported
    6 streams: not supported
    7 streams: not supported
    8 streams: not supported
  VHT RX highest supported: 0 Mbps
  VHT TX MCS set:
    1 streams: MCS 0-9
    2 streams: MCS 0-9
    3 streams: MCS 0-9
    4 streams: not supported
    5 streams: not supported
    6 streams: not supported
    7 streams: not supported
    8 streams: not supported
  VHT TX highest supported: 0 Mbps
VHT operation:
* channel width: 0 (20 or 40 MHz)
* center freq segment 1: 0
* center freq segment 2: 0
* VHT basic MCS set: 0xffc0
Transmit Power Envelope:
* Local Maximum Transmit Power For 20 MHz: 1 dBm
* Local Maximum Transmit Power For 40 MHz: 1 dBm
WMM:
* Parameter version 1
* u-APSD
* BE: CW 15-1023, AIFSN 3
* BK: CW 15-1023, AIFSN 7
* VI: CW 7-15, AIFSN 2, TXOP 3008 usec
* VO: CW 3-7, AIFSN 2, TXOP 1504 usec
DS Parameter set: channel 64
.....

```

6. Check the current link status.

```
# iw dev wlan0 link
```

The output should be as below:

```
# iw dev wlan0 link
Not connected.
```

7. Configure the Wi-Fi network SSID and password in `/etc/wpa_supplicant.conf`. First delete lines from line 5 to end. Then run `'wpa_passphrase <ssid> <password>'` and append the output to `/etc/wpa_supplicant.conf`. This step only needs to be done once and is saved across reboots.

```
# sed -i '5,$ d' /etc/wpa_supplicant.conf
# wpa_passphrase <ssid> <password> >> /etc/wpa_supplicant.conf
```

A sample `/etc/wpa_supplicant.conf` file is shown below:

```
# cat /etc/wpa_supplicant.conf
ctrl_interface=/var/run/wpa_supplicant
ctrl_interface_group=0
update_config=1
network={
    ssid="NXPOPEN"
    psk=c5397a26ced00bcb3545de1e0b421f76c9b6ed2c72a36150b87d5951b755cf7c
}
```

8. Connect to the WLAN with the given SSID in `/etc/wpa_supplicant.conf`.

```
# wpa_supplicant -B -i wlan0 -c /etc/wpa_supplicant.conf -D nl80211
```

9. Check the link status again for the WLAN it is connected to.

```
# iw dev wlan0 link
```

A sample output is below:

```
# iw dev wlan0 link
Connected to fc:5b:39:5f:5d:ce (on wlan0)
    SSID: NXPOPEN
    freq: 5745
    RX: 52226 bytes (109 packets)
    TX: 2454 bytes (20 packets)
    signal: -75 dBm
    rx bitrate: 24.0 MBit/s
    tx bitrate: 6.5 MBit/s VHT-MCS 0 VHT-NSS 1
    bss flags:
    dtim period:    1
    beacon int:    102
```

10. Run DHCP client to get IP address.

```
# udhcpc -i wlan0
```

11. Check the IP address of AP and ping the AP to check connectivity.

```
# ip route
# ping <IP address of AP>
```

The sample output is shown below:

```
# ip route
default via 192.168.0.1 dev wlan0 metric 10
192.168.0.0/23 dev wlan0 proto kernel scope link src 192.168.0.158
# ping 192.168.0.1
PING 192.168.0.1 (192.168.0.1) 56(84) bytes of data.
```

```
64 bytes from 192.168.0.1: icmp_seq=1 ttl=255 time=5.39 ms
64 bytes from 192.168.0.1: icmp_seq=2 ttl=255 time=5.37 ms
64 bytes from 192.168.0.1: icmp_seq=3 ttl=255 time=6.13 ms
64 bytes from 192.168.0.1: icmp_seq=4 ttl=255 time=7.02 ms
.....
```

12. Check the connectivity to the public website.

```
# echo nameserver "8.8.8.8" >> /etc/resolv.conf
# ping www.nxp.com
PING e6860.h.akamaiedge.net (23.7.170.207) 56(84) bytes of data.
64 bytes from a23-7-170-207.deploy.static.akamaitechnologies.com
(23.7.170.207): icmp_seq=1 ttl=50 time=311 ms
64 bytes from a23-7-170-207.deploy.static.akamaitechnologies.com
(23.7.170.207): icmp_seq=2 ttl=50 time=356 ms
64 bytes from a23-7-170-207.deploy.static.akamaitechnologies.com
(23.7.170.207): icmp_seq=3 ttl=50 time=299 ms
64 bytes from a23-7-170-207.deploy.static.akamaitechnologies.com
(23.7.170.207): icmp_seq=4 ttl=50 time=340 ms
64 bytes from a23-7-170-207.deploy.static.akamaitechnologies.com
(23.7.170.207): icmp_seq=5 ttl=50 time=282 ms
.....
```

6.8 MODBUS

MODBUS is an application layer messaging protocol, positioned at level 7 of the OSI model. The protocol enables client/server communication between devices connected on different types of buses or network.

6.8.1 Libmodbus introduction

Real-time Edge integrates libmodbus library. Libmodbus is a free software library used to send or receive data with a device that conforms to the Modbus protocol. It contains various backends to communicate over different networks (for example, serial in RTU mode or Ethernet in TCP IPv4/IPv6).

All the i.MX series and Layerscape series of boards support modbus. It can be used to write both:

- Client applications that reads/writes data from various devices.
- Server applications that provide data to several clients.

The official website that contains the latest version of the documentation for libmodbus is <https://libmodbus.org/>.

6.8.2 Modbus-Simulator introduction

Modbus-Simulator is a Modbus tool based on libmodbus library. And it contains a modbus client and a modbus device simulator.

Modbus-device-Simulator supports both **TCP** and **RTU** modes, and each mode supports the following functions.

Features supported by TCP:

- Gets CPU temperature of a device
- Gets the status of the LED light of a device
- Modifies the status of the LED light of a device

Features supported by RTU:

- Gets CPU temperature of a device
- Gets the status of the LED light of a device
- Modifies the status of the LED light of a device
- Modifies the slave address of the device
- Modifies the baud rate of the device

Note: Only the i.MX 8M Mini and i.MX 8M Plus boards support LED light function and CPU temperature function, other boards define a dummy variable. By default Modbus-Simulator is disable except i.MX 8M Mini and i.MX 8M Plus, taking I.MX 93 as an example, add the following line to `meta-real-time-edge/conf/distro/include/real-time-edge-base.inc` to enable it.

```
DISTRO_FEATURES:append:mx93-nxp-bsp = " modbus-simulator"
```

6.8.3 Modbus-Simulator usage

SYNOPSIS

```
{modbus_device_simulator|modbus_client_simulator} [OPTION]... {RTU-PARAMS|TCP-PARAMS}... {SERIALPORT|HOST}... [WRITE-DATA]...
```

6.8.3.1 Parameter description

The parameter of [Option] is the general parameter of TCP and RTU startup.

Table 97. Parameter description

Parameter option	Description
--debug,	show debug information.
-m,	connection Type TCP/RTU, optional parameter: {tcp rtu}.
-a,	slave address.
-c,	read and write data number.
-t,	function codes, the following function codes are available. (0x01) Read Coils, (0x02) Read Discrete Inputs, (0x05) Write Single Coil (0x03) Read Holding Registers, (0x04) Read Input Registers, (0x06) Write Single Register
-r,	register start address.
-o,	response timeout(ms).

The parameter of RTU-PARAMS is the general parameter of RTU startup.

Table 98. Parameter description

Parameter option	Description
-b,	baud rate, optional parameter: {4800 9600 19200 115200}. NOTE: when running the modbus_device_simulator, directly select the above parameters, when running the modbus_client_simulator, select the index of the parameters: {0 1 2 3}.
-d,	data bits, optional parameter: {7 8}.
-s,	stop bits, optional parameter: {1 2}.
-p,	verify type, optional parameter: {none even odd}.

The parameter of TCP-PARAMS is the general parameter of TCP startup.

Table 99. Parameter description

Parameter option	Description
-p,	port.

6.8.3.2 Examples of TCP and RTU

Examples of TCP modbus_device_simulator and modbus_client_simulator startup are as follows:

1. Start modbus_device_simulator locally with port 1502.

```
# modbus_device_simulator --debug -m tcp -p 1502 0.0.0.0
```

2. Start modbus_client_simulator and connect to the modbus_device_simulator with 127.0.0.1 and port 1502, function code is Write Single Coil. Function: Change the status of the LED light to 1.

```
# modbus_client_simulator --debug -m tcp -t 0x05 -r 0 -p 1502 127.0.0.1  
0xFF00
```

Examples of RTU modbus_device_simulator and modbus_client_simulator startup are as follows:

1. Start modbus_device_simulator serial connection, slave address is 1, baud rate is 115200, verify type is none, device is /dev/ttymx2.

```
# modbus_device_simulator --debug -m rtu -a 1 -b 115200 -p none /dev/ttymx2
```


2. Start `modbus_client_simulator` serial connection, slave address is 1, function code is Write Single Register, register start address is 1, baud rate is 115200, verify type is none, device is `/dev/ttymxc2`, write date is 1. Function: Change the baud rate of `/dev/ttymxc2` from 9600 to 115200.

```
# modbus_client_simulator --debug -m rtu -a 1 -t 0x06 -r 1 -b 1 -p none /dev/ttymxc2
```

Note: For i.MX 6ULL, do the following:

J1704: pin7 - GND

J1703: pin1 - TX pin2 - RX

```
==> setenv fdtfile imx6ull-14x14-evk-lpuart.dtb
==> run bootcmd
```

Device name: `/dev/ttymxc1`

For i.MX 93, do the following:

J1001: pin25 - GND pin27 - TX pin28 - RX

```
==> setenv fdtfile imx93-11x11-evk-lpuart.dtb
==> run bootcmd
```

Device name: `/dev/ttyLP4`

6.8.3.3 Commands for all features

The commands for all features of TCP are as follows:

1. Read the status of the LED light:

```
modbus_client_simulator --debug -m tcp -t 0x01 -r 0 -p 1502 127.0.0.1
```

2. Change the status of the LED light:

```
# modbus_client_simulator --debug -m tcp -t 0x05 -r 0 -p 1502 127.0.0.1
{0xFF00|0x0000}
```

3. Read the temperature of CPU

```
# modbus_client_simulator --debug -m tcp -t 0x04 -r 0 -p 1502 127.0.0.1
```

The commands for all features of RTU are as follows:

1. Read the status of the LED light:

```
modbus_client_simulator --debug -m rtu -a 1 -t 0x01 -r 0 -b 3 -p none /dev/ttymxc2
```

2. Change the status of the LED light:

```
# modbus_client_simulator --debug -m rtu -a 1 -t 0x05 -r 0 -b 3 -p none /dev/ttymxc2 {0xFF00|0x0000}
```

3. Read the temperature of CPU:

```
# modbus_client_simulator --debug -m rtu -a 1 -t 0x04 -r 0 -b 3 -p none /dev/ttymxc2
```

4. Change `modbus_device_simulator` slave address from 1 to 6.

```
# modbus_client_simulator --debug -m rtu -a 1 -t 0x06 -r 0 -b 3 -p none /dev/ttymxc2 6
```

5. Change `modbus_device_simulator` baud rate from 115200 to 9600.

```
# modbus_client_simulator --debug -m rtu -a 6 -t 0x06 -r 1 -b 3 -p none /dev/ttymxc2 1
```

6.8.4 Testing Modbus-Simulator

Use two i.MX 8MP boards to test the functions of the modbus-simulator.




Figure 185. Testing the functions of the modbus-simulator

Note: Learn about the startup parameters before starting.

6.8.4.1 Testing TCP functions

Read the status of the LED light on the board1.

1. Enter board1 and start up `modbus_device_simulator`.

```
# modbus_device_simulator --debug -m tcp -p 1502 0.0.0.0
```

2. Enter board2 and start up `modbus_client_simulator`.

```
# modbus_client_simulator --debug -m tcp -t 0x01 -r 0 -p 1502 10.193.21.104
```

```

root@imx8mp-lpddr4-evk:~# modbus_device_simulator --debug -m tcp -p 1502 0.0.0.0
Debug:
    Coils: 0x0001
    Digital inputs: 0x0001
    Holding registers: 0x0002
    Input registers: 0x0001
    Ip address: 0.0.0.0
    Port: 1502
    New connection from 10.193.21.105:14071 on socket 4
Waiting for an indication...
<00><01><00><00><00><06><01><01><00><00><00><01>
[00][01][00][00][00][04][01][01][01][01]

```

Figure 186. Debug log on i.mx8mp-lpddr4-evk board

6.8.4.2 Testing RTU functions

Preparation: Connect the serial ports on the two boards, as shown in [Figure 187](#).




Figure 187. Connection of two i.MX 8MP boards

The pin connection information of the two boards should be as follows:

- GND <---> GND
- TXD <---> RXD
- RXD <---> TXD

Read the temperature of CPU on the board1

1. Enter board1 and start up `modbus_device_simulator`.

```
# modbus_device_simulator --debug -m rtu -a 1 -b 115200 -p none /dev/ttymx2
```

2. Enter board2 and start up `modbus_client_simulator`.

```
# modbus_client_simulator --debug -m rtu -a 1 -t 0x04 -r 0 -b 3 -p none /dev/
ttymx2
```

Note: The default number of stop bits is 1.

6.9 UART 9-bit Multidrop mode (RS-485) support

6.9.1 Overview

The UART provides a 9-bit mode to facilitate multidrop (RS-485) network communication. When 9-bit RS-485 mode is enabled, UART transmitter can transmit the ninth bit (9th bit) set by TXB8. The UART receiver can differentiate between data frames (9th bit = 0) and address frames (9th bit = 1).

Two examples are provided to demo UART RS485 9-bit multidrop support:

- `9bit-uart-interrupt-transfer` for interrupt mode
- `9bit-uart-polling` for polling mode.

These two demos are supported on the i.MX 8M Mini LPDDR4 EVK board.

6.9.2 Building and running the demo

6.9.2.1 Building the demo

Refer to *RTEDGEYOCTOUG* to set up Yocto environment and build the `nxp-image-real-time-edge`. All demos are in the `/examples` directory of the rootfs.

The bellow command compiles the demo separately:

```
bitbake 9bit-uart-interrupt-transfer
```

The demo is located in the below directory:

```
tmp/deploy/images/imx8mm-lpddr4-evk/examples/
```

6.9.2.2 Hardware setup

In order to test this feature by using a single i.MX 8M Mini LPDDR4 EVK board, use external loopback of UART3 for testing. Refer to the figure below to connect PIN8 (UART3_TXD) and PIN10 (UART3_RXD) by using a flying wire.




Figure 188. i.MX 8M Mini LPDDR4 EVK UART Connection

6.9.2.3 Preparing the demo

1. Connect 12 V power supply, switch SW101 to power on the board.
2. Connect a USB cable between the host PC and the J901 USB port on the target board. Two UART connections appear on the PC, one is used for Linux console, and another is used for FreeRTOS console. Open two serial terminal on host PC with the following settings for each UART connection:
 - 115200 baud rate
 - 8 data bits
 - No parity
 - One stop bit
 - No flow control

6.9.2.4 Running the demo

Power up the board and start up the M-core firmware under U-Boot by using the commands listed in the below paragraph. Use `9bit_uart_interrupt_transfer` or `9bit-uart-polling` to replace the "DEMO_NAME" in the following commands:

If you choose to run the binary in DRAM:

```
=> ext4load mmc 1:2 0x80000000 /examples/mcux-sdk/DEMO_DIRECTORY/ddr_release/  
DEMO_NAME.bin  
=> dcache flush  
=> bootaux 0x80000000
```

If you choose to run the binary in TCM:

```
=> ext4load mmc 1:2 0x48000000 /examples/mcux-sdk/DEMO_DIRECTORY/release/  
DEMO_NAME.bin  
=> cp.b 0x48000000 0x7e0000 0x20000  
=> bootaux 0x7e0000
```

The demo configures the UART used for testing in 9-bit multidrop mode and sets its slave address to be "0xfe". Then, it sends two pieces of data. If the 9th bit of the data is zero, the bits are figured out to be "data". In general, you must send the "address" with 9-bit "1" to the physical line of UART before sending "data" with 9-bit "0". To do function verification, the demo does not send "address" before sending "data", so UART cannot receive this data when it is looped back to itself. The second piece of data is received when it is looped back, because the demo sends "address" with 0xfe first before sending this "data".

Below is a sample log displayed on the FreeRTOS console when the demo runs successfully:

```
UART will send first piece of data out without addressing itself:  
0x10 0x11 0x12 0x13 0x14 0x15 0x16 0x17  
  
UART will send second piece of data out with addressing itself:  
Address: 0xfe : 0x80 0x81 0x82 0x83 0x84 0x85 0x86 0x87  
  
RS-485 Slave Address has been detected.  
UART received data:  
  
Address: 0xfe : 0x80 0x81 0x82 0x83 0x84 0x85 0x86 0x87  
  
All data matches!
```

7 Acronyms and abbreviations

The [Table 100](#) lists the acronyms used in this document.

Table 100. Acronyms and abbreviations

Term	Description
AVB	Audio video bridging
AMP	Asymmetric multiprocessing
BC	Boundary clock
BLE	Bluetooth low energy
BMC	Best master clock
CA	Client application
CAN	Controller area network
CBS	Credit-based shaper
CDW	Concurrent Dual Wi-Fi
CMLDS	Common Mean Link Delay Service
DoS	Daniel-of-Service
DEI	Drop eligibility indication
DP	Display port
DSA	Distributed switch architecture
EtherCAT	Ethernet for control automation technology
EVK	Evaluation kit
ECU	Electronic control units
FDB	Forwarding database
FQTSS	Forwarding and queuing enhancements for time-sensitive streams
FMan	Frame manager
GPU	General processor unit
HSR	High-availability Seamless Redundancy
ICMP	Internet control message protocol
IEEE	Institute of electrical and electronics engineers
IETF	Internet engineering task force
IPC	Inter-processor communication
KM	Key management
LBT	Latency and bandwidth tester
MAC	Medium access control
MU	Message Unit
NFC	Near field communication
NCI	NFC controller interface
NMT	Network management

Table 100. Acronyms and abbreviations...continued

Term	Description
OC	Ordinary clock
OpenIL	Open industry Linux
OPC	Open platform communications
OP-TEE	Open portable trusted execution environment
OS	Operating system
OTA	Over-the-air
OTPMK	One-time programmable master key
PCP	Priority code point
PDO	Process data object
PHC	PTP hardware clock
PIT	Packet inter-arrival times
PLC	programmable logic controller
PTP	Precision time protocol
QSPI	Queued serial peripheral interface
RCW	Reset configuration word
REE	Rich execution environment
RPC	Remote procedure call
RPMSG	Remote processor messaging
RTEdge	Real-time edge
RTC	Real-time clock
RTT	Round-trip times
RX	Receiver
SABRE	Smart application blueprint for rapid engineering
SDO	Service data object
SOEM	Simple Open EtherCAT master
SPI	Serial periphery interface
SRP	Stream reservation protocol
SRTM	Simplified Real-time Messaging
SRK	Single root key
TA	Trusted application
TAS	Time-aware scheduler
TC	Traffic classification
TCP	Transmission control protocol
TEE	Trusted execution environment
TFTP	Trivial file transfer protocol
TSN	Time sensitive networking

Table 100. Acronyms and abbreviations...continued

Term	Description
TX	Transmitter
TZASC	Trust zone address space controller
UDP	User datagram protocol
VLAN	Virtual local area network

8 Revision history

[Table 101](#) summarizes the revisions to this document.

Table 101. Document revision history

Document ID	Release date	Topic cross-reference	Description
REALTIMEEDGEUG v.3.1	26 March 2025	-	Updated for Real Time Edge Software 3.1 release.
REALTIMEEDGEUG v.3.0	17 December 2024	-	Updated for Real Time Edge Software 3.0 release.
REALTIMEEDGEUG v.2.9	25 July 2024	-	<ul style="list-style-type: none"> Updated for Real Time Edge Software Rev 2.9 release. Added i.MX 93 14x14 EVK support.
REALTIMEEDGEUG v.2.8	29 March 2024	-	<ul style="list-style-type: none"> Updated for Real Time Edge Software Rev 2.8 release. 'What's New' section moved to <i>REALTIMEEDGEUG</i>.
REALTIMEEDGEUG v. 2.7	18 December 2023	Added Section 4	Updated for Real Time Edge Software 2.7 release.
REALTIMEEDGEUG v.2.6	28 July 2023	-	Updated for Real Time Edge Software 2.6 release.
		Section 6.1.5.4	Added the Section.
REALTIMEEDGEUG v.2.5	30 March 2023	-	Updated for Real Time Edge Software 2.5 release.
		Section 3.6	Added the Section "Heterogeneous Multicore VirtIO and networking sharing"
		Section 6.1.5 , Section 2.5.4.14	Added the Sections "Codesys EtherCAT Master", "Flextimer module".
		Section 3.5.6	Added the Section "Running the i.MX 93 demo".
		<ul style="list-style-type: none"> Section 3.4.3 Section 2.5.4.4 	Modified the Sections: <ul style="list-style-type: none"> RPMSG Cortex-A core and Cortex-M core GPIO file
REALTIMEEDGEUG v.2.4	16 December 2022	-	Added support for i.MX 8DXL and i.MX 93 EVK boards throughout the document.
		Configuring Flow Meter policy on stream	Added the Section.
		Section 5.1.3.4.1	Added the Section for preemptive verify.
		Section 2.5.4.15	Added the Math library support.
		Section 6.7	Added the Section.
		Section 6.1.3.4.5	Added the Section.
REALTIMEEDGEUG v.2.3	28 July 2022	What's new in Real-time Edge software v2.3	Updated the Section.
		Heterogeneous AMP Architecture	Added the Section.

Table 101. Document revision history...continued

Document ID	Release date	Topic cross-reference	Description
		Section 5.2	Updated the Section.
		-	Other updates throughout the document.
REALTIMEEDGEUG v.2.2	29 March 2022	What's new in Real-time Edge software v2.2	Added the Section. Updated for Real-time Edge Software Rev 2.2. Support for TLS protocol removed for NETCONF feature.
REALTIMEEDGEUG v.2.1	15 December 2021	What's new in Real-time Edge software v2.1	Added the Section. Updated for Real-time Edge Software Rev 2.1.
		Open, fixed, and closed issues	Updated the Section.
REALTIMEEDGEUG v.2.0	29 July 2021	-	<ul style="list-style-type: none"> 'Real-time Edge Software' introduced instead of 'Open Industrial Linux', with real-time feature support. Rearranged the document structure to include the Chapters: Real-time System, Real-time Networking, and Protocols.
		-	Added support for building Real-time Edge image using Yocto Project build environment. Details are provided in the <i>Real-time Edge Yocto Project User Guide</i> .
		What's new	Added the Section.
		-	Integrated Baremetal framework in the document.
OPENILUG v.1.11	26 April 2021	What's new	Added the Section that describes the new features of each release.
		-	Updated the Section 'Getting Open IL'.
		-	Deleted references to Edgescale, OP-TEE, OTA throughout the document and other minor updates.
OPENILUG v.1.10	22 December 2020	Section 5.2	Added the Chapter and related contents.
		-	Added the 'Camera' Section and related details.
		-	Added the Host setup for i.MX 8M Plus EVK board details.
OPENILUG v.1.9	15 September 2020	Section 5.3	Added the Section.
		GPU	Added the Chapter and related description
		Wayland and Weston	Added the Chapter and related description
		Section 6.1.3.4	Made it a part of the Chapter "EtherCAT"
OPENILUG v.1.8	29 May 2020	Section 2.4	Added the Section in Section 6 .
		-	Updated this Section Interface naming in Linux for LS1028ARDB.
		-	Updated the Section Host system requirements for Open IL.

Table 101. Document revision history...continued

Document ID	Release date	Topic cross-reference	Description
		-	Updated the Section Running Selinux demo.
OPENILUG v.1.7.1	20 February 2020	Section 6.4.3.7	Updated this Section.
OPENILUG v.1.7	17 January 2020	Section 6.1.3.4	Added the Chapter (nxp servo).
		Section 5.3	Added the Chapter.
		Getting Open IL	Updated the Section.
		Section 6.4	Other updates.
OPENILUG v.1.6	31 August 2019	Section 5.1.4	<ul style="list-style-type: none"> Information related to pcpmap command removed from the Section Section 5.1.4.1 and Section 5.1.4.2 Port names "eno/swp0" changed to "swp0" for few tsntool commands. Note added in Section Section "Stream identification" for usage of <code>nulltagged</code> and <code>streamhandle</code> parameters. Added the Section Section 5.1.4.2.8. Other minor updates
		-	Updated the table "Host system mandatory packages". Added <code>autogen</code> <code>autoconf</code> <code>libtool</code> and <code>pkg-config</code> packages
		Section 6.5.3	Added this Chapter
		-	Updated Section 6.4
		Section 6.4	<ul style="list-style-type: none"> Added the Section Section 6.4.3.1 and other updates.
OPENILUG v.1.5	01 May 2019	-	Added the Section to describe interface naming for U-Boot and Linux for LS1028 ARDB. <ul style="list-style-type: none"> Added the Chapter 'EdgeScale Client' Updated the OpenIL version and Git tag in the Section 'Getting Open IL'.
		Section 5.1.4	Updated this Section in the Chapter Section 5.1
		Section 6.5.2	Added the Chapter
OPENILUG v.1.4	01 February 2019	Section 1.3	<ul style="list-style-type: none"> Added support for LS1028ARDB (64-bit and Ubuntu). Updated various Sections accordingly. Updated the OpenIL version and Git tag in the Section 'Getting Open IL'. Added LS1028ARDB support. Added the Chapter QT.
		Section 5.1	Reorganized this Chapter and added separate Section for Section 5.1.4
		Section 6.5.1	Added the Chapter.

Table 101. Document revision history...continued

Document ID	Release date	Topic cross-reference	Description
		Section 6.2	Minor updates in this Chapter. Also added the Section, Section 6.2.3.1 and Section 6.2.3.3 .
OPENILUG v.1.3.1	15 October 2018	-	Updated the OpenIL version and Git tag in the Section 'Getting Open IL'
OPENILUG v.1.3	31 August 2018	EtherCAT and FlexCAN	Added the Chapters.
		i.MX6QSabreSD support.	Added the Section in Chapter 'NXP OpenIL platforms'. Updated other Sections for i.MX6Q Sabre support.
		Getting Open IL	Updated the Section.
		Selinux demo	Added the Section for enabling SELinux and updated Basic setup. Updates in other Sections.
OPENILUG v.1.2	31 May 2018	-	Updated the Section, "Hardware requirements" for RTnet.
		-	Updated the Section, "Software requirements" for RTnet.
OPENILUG v.1.1.1	18 April 2018	-	Added the Section, "RTnet".
		-	Added a note for LS1043A switch setting
OPENILUG v.1.1	30 March 2018	-	Added support for industrial IoT Baremetal framework in this Section.
		-	Added a note for steps to be performed before booting up the board.
		Section 1.5	Added the Section.
OPENILUG v.1.0	22 Decemeber 2017	Section 6.3	Added the Chapter.
		Section 5.1	Chapters for "1-board TSN demo" and "3-board TSN demo" replaced by a single Chapter, "TSN demo".
		Section 6	<ul style="list-style-type: none"> Updated the Section, 'Industrial Features'. -IEEE 1588 -'sja1105-ptp' support removed.
OPENILUG v.0.3	25 August 2017	-	Set up the OpenIL website http://www.openil.org/ .
		-	OTA - Xenomai Cobalt 64-bit and SJA1105 support added.
		Section 5.1	Qbv support added. <ul style="list-style-type: none"> SELinux support for LS1043 / LS1046 Ubuntu Userland added. OP-TEE support for LS1021ATSN platform added. 4G LTE module - 64-bit support for LS1043ARDB, LS1046ARDB, and LS1012ARDB added.

Table 101. Document revision history...continued

Document ID	Release date	Topic cross-reference	Description
		Section 5	Ubuntu Userland support for 64-bit LS1043ARDB and 64-bit LS1046ARDB added.
OPENILUG v.0.2	26 May 2017	-	Initial public release.

9 Note about the source code in the document

Example code shown in this document has the following copyright and BSD-3-Clause license:

Copyright 2017-2025 NXP Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials must be provided with the distribution.
3. Neither the name of the copyright holder nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

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.

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.

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.

Evaluation products — This evaluation product is intended solely for technically qualified professionals, specifically for use in research and development environments to facilitate evaluation purposes. It is not a finished product, nor is it intended to be a part of a finished product. Any software or software tools provided with an evaluation product are subject to the applicable license terms that accompany such software or software tools. This evaluation product is provided on an "as is" and "with all faults" basis for evaluation purposes only and is not to be used for product qualification or production. If you choose to use these evaluation products, you do so at your risk and hereby agree to release, defend and indemnify NXP (and all of its affiliates) for any claims or damages resulting from your use. NXP, its affiliates and their suppliers expressly disclaim all warranties, whether express, implied or statutory, including but not limited to the implied warranties of non-infringement, merchantability and fitness for a particular purpose. The entire risk as to the quality, or arising out of the use or performance, of this evaluation product remains with user.

In no event shall NXP, its affiliates or their suppliers be liable to user for any special, indirect, consequential, punitive or incidental damages (including without limitation damages for loss of business, business interruption, loss of use, loss of data or information, and the like) arising out of the use of or inability to use the evaluation product, whether or not based on tort (including negligence), strict liability, breach of contract, breach of warranty or any other theory, even if advised of the possibility of such damages.

Notwithstanding any damages that user might incur for any reason whatsoever (including without limitation, all damages referenced above and all direct or general damages), the entire liability of NXP, its affiliates and their suppliers and user's exclusive remedy for all of the foregoing shall be limited to actual damages incurred by user based on reasonable reliance up to the greater of the amount actually paid by user for the evaluation product or five dollars (US\$5.00). The foregoing limitations, exclusions and disclaimers shall apply to the maximum extent permitted by applicable law, even if any remedy fails of its essential purpose and shall not apply in case of willful misconduct.

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.

Amazon Web Services, AWS, the Powered by AWS logo, and FreeRTOS — are trademarks of Amazon.com, Inc. or its affiliates.

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.

I2C-bus — logo is a trademark of NXP B.V.

Oracle and Java — are registered trademarks of Oracle and/or its affiliates.

Contents

1	Introduction	2	2.6.3	Booting Native RTOS on Cortex-A core	65
1.1	Real-time Edge software	2	2.6.4	Debugging Native RTOS on Cortex-A Core	65
1.2	Feature support matrix	3	2.6.5	Native RTOS real-time latency benchmark	65
1.3	Supported NXP platforms	7	2.6.5.1	rt_latency introduction	65
1.3.1	Switch settings	8	2.6.5.2	Building rt_latency benchmark application	67
1.3.2	Flashing pre-built images	9	2.6.5.3	Running Native RTOS rt_latency application	67
1.4	Real-time Edge Software Yocto Project	10	2.6.5.4	Output of rt_latency	71
1.5	Related documentation	10	2.7	RTOS on Cortex-A core with Jailhouse	72
2	Real-time system	12	2.7.1	Jailhouse	72
2.1	Overview	12	2.7.1.1	Overview	72
2.2	Real-time-system support matrix	15	2.7.1.2	Running PREEMPT_RT Linux in Inmate	72
2.3	Building, deploying, and releasing unified software	16	2.7.1.3	Running Jailhouse examples In Inmate	75
2.3.1	Yocto layer for Cortex-A core	16	2.7.2	Harpoon (RTOS on Cortex-A)	78
2.3.2	Yocto layer for Cortex-M core	17	2.7.2.1	Overview	78
2.3.2.1	Introduction to meta-rtos-industrial	17	2.8	RTOS and Baremetal on Cortex-M core	78
2.3.2.2	Integration of meta-rtos-industrial	18	2.8.1	Booting Cortex-M Core RTOS Image	78
2.3.2.3	Building meta-rtos-industrial	19	2.8.1.1	Using U-Boot Commands to Manage Life Cycle of RTOS on Cortex-M Core	78
2.4	Preempt-RT Linux	21	2.8.1.2	Using RemoteProc under Linux to Manage Life Cycle of RTOS on Cortex-M Core	79
2.4.1	System Real-time Latency tests	21	3	Heterogeneous Multicore Framework	81
2.4.1.1	Running Cyclictst	21	3.1	Overview	81
2.4.2	Real-time application development	21	3.2	Building Heterogeneous Multicore RTOS Application	85
2.5	Baremetal on Cortex-A core	22	3.2.1	Build with Yocto	85
2.5.1	BareMetal framework	22	3.2.2	Build with Standalone Mode	86
2.5.2	Getting started	25	3.3	Flexible Real-time System	92
2.5.2.1	Hardware and software requirements	25	3.3.1	Heterogeneous Multicore RAM Console	92
2.5.2.2	Hardware setup	26	3.3.1.1	Using RAM Console in FreeRTOS	92
2.5.2.3	Building the software	29	3.3.1.2	Use RAM Console in Zephyr	94
2.5.2.4	Booting up the Linux with Baremetal	31	3.3.1.3	Dump RAM Console Log	94
2.5.3	Running examples	31	3.3.2	Heterogeneous Multicore hello_world	95
2.5.3.1	Preparing the console	31	3.3.2.1	Overview	95
2.5.3.2	Running the Baremetal binary	32	3.3.2.2	Technical Points	96
2.5.4	Development based on Baremetal framework	32	3.3.2.3	Running flexible multicore hello_world application	98
2.5.4.1	Developing the Baremetal application	32	3.3.3	LwIP Networking Stack	119
2.5.4.2	Main file app.c	32	3.3.3.1	Overview	120
2.5.4.3	Common header files	33	3.3.3.2	Running lwIP Application	120
2.5.4.4	GPIO example	33	3.4	RPMSG data communication	122
2.5.4.5	I2C example	34	3.4.1	Overview	122
2.5.4.6	IRQ example	35	3.4.2	RPMSG performance evaluation	122
2.5.4.7	QSPI example	35	3.4.2.1	Building RPMSG performance evaluation application	123
2.5.4.8	IFC example	36	3.4.2.2	Running RPMsg performance application	123
2.5.4.9	Ethernet example	37	3.4.3	RPMSG between Cortex-A Core and Cortex-M Core	124
2.5.4.10	USB example	38	3.4.3.1	RPMSG with enhanced 8MB Vring buffer	125
2.5.4.11	PCIe example	39	3.4.4	RPMSG between Cortex-A Core and Cortex-A Core	127
2.5.4.12	ENETC example	40	3.4.4.1	RPMSG between Cortex-A Linux and Cortex-A RTOS	127
2.5.4.13	SAI example	40	3.4.4.2	RPMSG between Cortex-A RTOS and Cortex-A RTOS	130
2.5.4.14	FlexTimer example	45	3.4.5	Complex RPMSG on MPU	137
2.5.4.15	Newlib's math library	46			
2.5.4.16	ICC module	47			
2.5.4.17	Single hardware interrupt routed to multiple cores	52			
2.5.4.18	Hardware resource allocation	56			
2.6	Native RTOS on Cortex-A core	64			
2.6.1	Overview	64			
2.6.2	Building Native RTOS on Cortex-A core	65			

3.4.5.1	Overview	137	4.1	Introduction	183
3.4.5.2	Building the Complex str_echo application	137	4.2	Software architecture	183
3.4.5.3	Running the Complex str_echo application	137	4.2.1	Using one of the i.MX RT1180 switch ports as DSA CPU port	184
3.5	RPMSG based resource sharing	145	4.2.2	Using i.MX RT1180 ENETC port as DSA CPU port	185
3.5.1	Overview	145	4.3	Running the Heterogeneous Multi-SoC Framework	185
3.5.2	Software architecture and design	145	4.3.1	Building and running on MPU + i.MX RT1180 EVK	185
3.5.3	Resource sharing based on SRTM	145	4.3.1.1	Hardware setup for i.MX 8M Plus EVK and i.MX RT1180 EVK	185
3.5.3.1	UART sharing design details	146	4.3.1.2	Hardware setup for i.MX 93 EVK and i.MX RT1180 EVK	186
3.5.4	Source code files and configuration	148	4.3.1.3	Building the DSA switch application for i.MX RT1180 EVK	187
3.5.5	Building and running the demo on i.MX 8M Mini LPDDR4 EVK	151	4.3.1.4	Burning the image to flash on i.MX RT1180 EVK	187
3.5.5.1	Hardware setup for i.MX 8M Mini EVK	151	4.3.1.5	Bringing up MPU and i.MX RT1180 EVK	194
3.5.5.2	Building the demo images	151	4.3.2	Runtime usage on MPU and i.MX RT1180 EVK	196
3.5.5.3	Running the i.MX 8M Mini EVK demo	151	4.3.2.1	i.MX RT1180 NETC switch interface under Linux	196
3.5.6	Building and running the demo on i.MX 93 EVK and i.MX 93 14x14 EVK	153	4.3.2.2	Maximum frame size configuration	197
3.5.6.1	Hardware setup for i.MX 93 EVK and i.MX 93 14x14 EVK	153	4.3.2.3	Single port mode for i.MX RT1180 NETC switch ports	197
3.5.6.2	Building the demo images	153	4.3.2.4	Bridge mode for i.MX RT1180 NETC switch ports	199
3.5.6.3	Running the demo on i.MX 93 EVK and i.MX 93 14x14 EVK	154	4.3.2.5	i.MX RT1180 NETC switch port statistics counters	200
3.5.7	Building and running the demo on i.MX 93 QSB	156	4.3.2.6	VLAN configuration	201
3.5.7.1	Hardware setup for i.MX 93 QSB	156	4.3.2.7	FDB configuration	202
3.5.7.2	Building the demo images	156	4.3.2.8	PTP Stack	202
3.5.7.3	Running the demo on i.MX 93 QSB	156	4.3.2.9	802.1CB usage	206
3.6	Heterogeneous Multicore VirtIO and networking sharing	158	4.3.2.10	802.1 Qci usage	208
3.6.1	Heterogeneous Multicore VirtIO	158	4.3.2.11	802.1 Qbv usage	209
3.6.2	Heterogeneous Multicore VirtIO performance evaluation	160	4.3.2.12	802.1 Qbu usage	212
3.6.2.1	Overview	160	5	Real-time networking	213
3.6.2.2	Building VirtIO performance evaluation backend application	160	5.1	Time Sensitive Networking (TSN) on NXP platforms	213
3.6.2.3	Running VirtIO performance testing	160	5.1.1	TSN hardware capability	213
3.6.3	Heterogeneous Multicore VirtIO network sharing	164	5.1.2	TSN configuration	213
3.6.3.1	Overview	164	5.1.2.1	Using Linux traffic control (tc) to configure TSN	214
3.6.3.2	Building VirtIO network sharing backend application	165	5.1.2.2	Using Tsntool to configure TSN	214
3.6.3.3	Running VirtIO network sharing	165	5.1.2.3	Using NETCONF to configure TSN remotely	223
3.7	Unified Life Cycle Management	170	5.1.2.4	TSN configuration demo	223
3.7.1	Overview	170	5.1.3	TSN on i.MX 8DXL / i.MX 8M Plus / i.MX 93 ..	236
3.7.2	Using U-Boot Commands to Manage Life Cycle of RTOS on Cortex-A Core	171	5.1.3.1	Test environment	236
3.7.3	Using RemoteProc under Linux to Manage Life Cycle of RTOS on Cortex-A Core	176	5.1.3.2	Clock synchronization	237
3.7.4	Using U-Boot Commands to Manage Life Cycle of RTOS on Cortex-M Core	179	5.1.3.3	Qbv	238
3.7.5	Using RemoteProc under Linux to Manage Life Cycle of RTOS on Cortex-M Core	180	5.1.3.4	Qbu	239
3.8	Industrial applications	181	5.1.3.5	Qav	243
3.8.1	EtherCAT MainDevice Stack SOEM	181	5.1.4	TSN on LS1028A	244
3.8.1.1	Overview	181	5.1.4.1	TSN configuration on ENETC	244
3.8.1.2	Running SOEM Application on Cortex-A core	182	5.1.4.2	TSN configuration on Felix switch	255
3.8.1.3	Running SOEM Application on Cortex-M core	182	5.2	GenAVB/TSN stack	276
4	Heterogeneous Multi-SoC Framework	183	5.2.1	Introduction	276

5.2.1.1	GenAVB/TSN Roles	276	6.1.4.2	SOEM for i.MX 8M Mini LPDDR4 EVK platform	366
5.2.1.2	GenAVB/TSN Protocols	277	6.1.4.3	SOEM for i.MX 93 EVK platform	374
5.2.1.3	User Library	278	6.1.5	CODESYS EtherCAT MainDevice	378
5.2.1.4	AVB Kernel Module	278	6.1.5.1	Overview	378
5.2.1.5	Supported configurations	278	6.1.5.2	Features	378
5.2.1.6	AVB endpoint example applications	278	6.1.5.3	Building the image	378
5.2.1.7	TSN endpoint example application	279	6.1.5.4	Running CODESYS on the boards	380
5.2.2	GenAVB/TSN stack start/stop	280	6.1.5.5	Setting up CODESYS	383
5.2.3	Use cases description	281	6.2	FlexCAN and CANopen	395
5.2.3.1	AVB endpoint	281	6.2.1	Introduction	395
5.2.3.2	AVB Bridge	281	6.2.1.1	CAN bus	395
5.2.3.3	AVB Hybrid	288	6.2.1.2	CANopen	397
5.2.3.4	gPTP Endpoint	288	6.2.2	Introducing the function of CAN example code	398
5.2.3.5	gPTP Bridge	288	6.2.3	Running a CAN application	400
5.2.3.6	gPTP multiple domains	289	6.2.3.1	Hardware preparation for LS1028ARDB	400
5.2.3.7	TSN endpoint sample application	293	6.2.3.2	Running the SocketCAN commands	400
5.2.4	Configuration files	306	6.2.3.3	Testing CAN bus	401
5.2.4.1	System	306	6.3	OPC UA	402
5.2.4.2	gPTP	308	6.3.1	OPC introduction	402
5.2.4.3	SRP	313	6.3.2	The node model	402
5.2.5	Log files	314	6.3.3	Node Namespaces	403
5.2.5.1	gPTP Endpoint	314	6.3.4	Node classes	404
5.2.5.2	gPTP Bridge	316	6.3.5	Node graph and references	404
5.2.5.3	SRP Bridge	316	6.3.6	Open62541	405
5.2.5.4	TSN Endpoint example application	317	6.3.7	OPC UA Pub/Sub over TSN	406
5.3	IEEE 1588/802.1AS	320	6.3.7.1	OPC UA Pub/Sub introduction	406
5.3.1	Introduction	320	6.3.7.2	OPC UA PubSub over TSN	407
5.3.2	IEEE 1588 device types	320	6.3.7.3	OPC UA PubSub components	407
5.3.3	IEEE 802.1AS time-aware systems	320	6.3.7.4	OPC UA PubSub sample application	408
5.3.4	Software stacks	321	6.3.7.5	OPC UA PubSub sample application over TSN	409
5.3.4.1	linuxptp stack	321	6.3.8	OPC UA client installation and usage	419
5.3.4.2	NXP GenAVB/TSN gPTP stack	321	6.3.8.1	UaExpert	419
5.3.5	Quick Start for IEEE 1588	322	6.3.8.2	FreeOpcUa	423
5.3.5.1	Ordinary clock verification	322	6.4	NETCONF/YANG	424
5.3.5.2	Boundary clock verification	323	6.4.1	Overview	424
5.3.5.3	Transparent clock verification	324	6.4.2	Netopeer2	425
5.3.6	Quick Start for IEEE 802.1AS	325	6.4.2.1	Overview	425
5.3.6.1	Time-aware end station verification	325	6.4.2.2	Installing Netopeer2 on Ubuntu22.04	425
5.3.6.2	Time-aware bridge verification	325	6.4.2.3	Sysrepo	426
5.3.7	Virtual clock use case	326	6.4.2.4	Netopeer2 server	426
5.3.8	Long term test	327	6.4.2.5	Netopeer2 client	426
5.3.8.1	Virtual clock use case	327	6.4.2.6	Workflow in application practice	426
5.3.9	Known issues and limitations	327	6.4.3	Configuration	427
5.4	Networking	327	6.4.3.1	Enabling NETCONF feature	427
5.4.1	Q-in-Q on LS1028A Felix switch	327	6.4.3.2	Netopeer2-server	427
5.4.2	VCAP on LS1028A Felix switch	329	6.4.3.3	Netopeer2-cli	427
6	Protocols	333	6.4.3.4	Sysrepocfg	431
6.1	EtherCAT MainDevice	333	6.4.3.5	Sysrepocli	431
6.1.1	Introduction	333	6.4.3.6	List of yang models	432
6.1.2	EtherCAT protocol	333	6.4.3.7	Operation examples	433
6.1.3	IGH EtherCAT architecture	334	6.4.3.8	Application scenarios on LS1028A	434
6.1.3.1	IGH EtherCAT Device Drivers	335	6.4.4	Troubleshooting	437
6.1.3.2	IGH EtherCAT setup	336	6.5	Wireless on LS1028A	438
6.1.3.3	IGH EtherCAT User Space setup	340	6.5.1	NFC	438
6.1.3.4	The 'real-time-edge-servo' stack	342	6.5.1.1	Introduction	438
6.1.4	SOEM EtherCAT MainDevice	356	6.5.1.2	PN7120 features	438
6.1.4.1	SOEM for i.MX 8M Plus LPDDR4 EVK platform	357			

6.5.1.3	Hardware preparation	438
6.5.1.4	Software preparation	438
6.5.1.5	Testing the NFC click board	439
6.5.2	Bluetooth Low Energy	440
6.5.2.1	Introduction	440
6.5.2.2	Bluetooth Low Energy	440
6.5.2.3	Features	440
6.5.2.4	Hardware preparation	440
6.5.2.5	Software preparation	441
6.5.2.6	Testing the BLE P click board	442
6.5.3	BEE	444
6.5.3.1	BEE/ZigBEE	444
6.5.3.2	Features	444
6.5.3.3	Hardware preparation	444
6.5.3.4	Software preparation	445
6.5.3.5	Testing the BEE click board	446
6.6	SAI on LS1028ARDB	447
6.7	Wi-Fi on i.MX 8DXL EVK	448
6.7.1	Wi-Fi card information	448
6.7.2	Hardware Setup	448
6.7.3	Software Enablement	448
6.8	MODBUS	455
6.8.1	Libmodbus introduction	455
6.8.2	Modbus-Simulator introduction	455
6.8.3	Modbus-Simulator usage	455
6.8.3.1	Parameter description	455
6.8.3.2	Examples of TCP and RTU	456
6.8.3.3	Commands for all features	457
6.8.4	Testing Modbus-Simulator	458
6.8.4.1	Testing TCP functions	458
6.8.4.2	Testing RTU functions	459
6.9	UART 9-bit Multidrop mode (RS-485) support	459
6.9.1	Overview	460
6.9.2	Building and running the demo	460
6.9.2.1	Building the demo	460
6.9.2.2	Hardware setup	460
6.9.2.3	Preparing the demo	461
6.9.2.4	Running the demo	462
7	Acronyms and abbreviations	463
8	Revision history	466
9	Note about the source code in the document	471
	Legal information	472

Please be aware that important notices concerning this document and the product(s) described herein, have been included in section 'Legal information'.
