AN12572

Backplane Driver Support

Supports: LSDK 19.06

Rev. 1 — 09/2019 Application Note

Contents

1 Overview

This document describes how to enable backplane support for Layerscape and QorIQ devices with embedded support for this type of connection.

Ethernet operation over electrical backplanes, also referred to as “Backplane Ethernet,” combines the IEEE 802.3 Media Access Control (MAC) and MAC Control sublayers with a family of Physical Layers defined to support operation over a modular chassis backplane. Usually, there is no external PHY involved and the connection is made at the SoC’s PCS (Physical Coding Sublayer) level. Based on the link quality, a signal equalization is required. In cases where the link is realized based on passive direct attach cables, the link may need to be established with only the default (recommended) parameters for equalization. The standard states that a start-up algorithm should be in place in order to get the link up.

1.1 BaseKR support

Layerscape and QorIQ devices comes with embedded support for backplane connections at different baud rates.

- 10G is present in custom boards with the following devices: T2080, LS1046A, LS1088A, LS2088A, and LX2160A.

The enablement of backplane support is done in two parts. One refers to support from the device tree while the other is contained in the Linux kernel driver.

In the device tree, the following value is valid `backplane-mode`:

- `10gbase-kr`

In the Linux kernel driver, the implementation is different depending on each of the above mentioned cases. However, the following changes are common for all:

- Advertise the link partner with the correct working mode.
- Put the lane in the correct BaseKR mode.
- Use the recommended (if it is the case) parameters for pre- and post-tap coefficients in the lane initialization phase. This affects the starting point of the algorithm.
- Optionally, update the constraint relation between tap coefficients if this is needed.

1.2 Physical layer signaling system

The backplane Ethernet extends the family of 10GBASE-R physical layer signaling system to include the BASE-KR. This specifies 10/40 Gb/s operation over two differential, controlled impedance pairs of traces (one pair for transmit and one pair for receive). This system employs the 10GBASE-R PCS, the serial PMA, and the BASE-KR PMD sublayers.

The BASE-KR PMD’s control function implements the BASE-KR start-up protocol. This protocol facilitates timing recovery and equalization while also providing a mechanism through which the receiver can tune the transmit equalizer to optimize performance over the backplane interconnect. The BASE-KR PHY may optionally include Forward Error Correction (FEC).

Details about the aforementioned layers can be found in Clause 49, 51, and 74 of the IEEE Std 802.3.

1.3 Auto-negotiation

Auto-negotiation allows the devices at both ends of a link segment to advertise abilities, acknowledge receipt, and discover the common modes of operation that both devices share. It also rejects the use of operational modes not shared by both devices. Auto-negotiation does not test link segment characteristics.

1.4 Link training

Link training occurs after auto-negotiation has determined the link to be a Base-KR, but before auto-negotiation is done. It continuously exchanges messages (training frames) between the local and the remote device as part of the start-up phase. Link training also tunes the analog parameters of the remote and local SerDes transmitter to improve the link quality. Both LP (link partner/remote device) and LD (local device) perform link training in parallel. Link training stops when both sides decide that the link is passable. Then the link is considered up.

1.5 Backplane linux releases

Linux kernel with backplane support can be obtained from the following code aurora repository:

https://source.codeaurora.org/external/qoriq/qoriq-components/linux-extras/.

Different backplane releases are provided on top of LSDK base releases. Use the appropiate tag according to desired LSDK release, kernel version, and backplane release. The tags for backplane releases are created by using the following rules:

```
BACKPLANE - <LSDK_release> - <Kernel_version> - <Backplane_release>
```

2 Enable Backplane Support

2.1 Setup

Hardware setup

- Two custom boards (SoC from supported device list), with the XFI retimers removed.
- Passive direct attach cable (l <= 1M)

Software setup

- Linux kernel with backplane support enabled
- Device tree for custom boards with backplane PHY devices

NOTE

The QDS custom boards are used for this implementation/demonstration. You may use any custom board that enables access to the Backplane Ethernet feature.

2.2 Enable backplane connection from MC

This step is required only for devices based on DPAA2 architecture.

Use `MAC_LINK_TYPE_BACKPLANE` for all ports that will be used for backplane connections. In order to do that in the MC data path configuration file, add an entry like below for all ports used:

```
board_info {
        ports {
                        …
```

```
            mac@1 {
                link_type = "MAC_LINK_TYPE_BACKPLANE";
            };
            …
            };
        };
    };
```

Deploy this configuration file on the target board as per Data Path Configuration chapter from DPAA2 User Manual.

NOTE

Omitting this step can lead to an unreliable backplane connection. Random link-down or link-up events can be experienced. This is due to a concurrent access to MDIO bus between MC core (MC firmware) and GPP core (Linux kernel).

2.3 Enable backplane support in Linux kernel

2.3.1 Enable backplane PHY driver

Enable backplane support from:

```
Device Drivers -> Network device support -> Support for backplane on Freescale XFI interface ->
MDIO_FSL_BACKPLANE
```

2.3.2 Add backplane PHY devices in device tree

2.3.2.1 SerDes device and internal MDIO buses

The SerDes device and all internal MDIO buses should be listed in the SoC's common device tree source file:

```
<linux_kernel_repo>/arch/arm64/boot/dts/freescale/fsl-<device>xa.dtsi
```

To see if the SerDes module is listed, examine a block like the one below:

```
serdes1: serdes@1ea0000  {
         reg = <0x0 0x1ea0000 0 0x00002000>;
         compatible = "fsl,serdes-10g";
         little-endian;
};
```

If the SerDes module is listed then it means that the serdes1 (label for SerDes node) is registered and can be used. The only client of this node is the kernel backplane PHY driver which uses the node's unit address as a base address. The base address is mapped in the SOC's memory space to further access specific MDIO registers used to control the backplane connection.

In the DTS, there must be a `serdes1` node like the one represented above. If the node is not present, then it must be added. The device base address is listed in SoC's CCSR memory map. One thing that should be considered is endianness of SerDes module, which can be different than that of the GPP. In this case, the module's registers must be accessed using the endianness. Currently, the Linux kernel backplane PHY driver does not support access dependent on target endianness. One way to do this will be to use the endianness attribute in the device tree. This way the driver can be used on different targets without changes.

Next look after internal MDIO buses listed in the device tree. See the block below as an example:

```
pcs_mdio1:  mdio@0x8c07000 {
        compatibile = "fsl, fman-memac-mdio";
        reg = <0x0 0x8c07000 0x0 0x1000>;
        device_type = "mdio";
```

```
        little-endian;
};
```

The block above shows that `pcs_mdio1` is listed in the device tree. The unit address of this node (`0x8c7000`) is the WRIOP internal physical port 1 base address as it is mapped in the SoC memory space. The address `0x8c0000` is the WRIOP port block base address as it is listed in SoC reference manual. The address `0x7000` is the physical port offset in the WRIOP internal memory map. All `pcs_mdio` ports have an offset of `0x4000` between them, so the next port will be located at `0xb000` and so on. The attribute `fsl, fman-memac-mdio` means that the FSL MDIO driver will be used to access this MDIO bus. It is required to use a dedicated MDIO bus driver to access internal MDIO buses, because it uses proprietary MDIO control registers block and offset. See the DPAA2 User Manual for details about MDIO registers block.

The kernel MDIO driver used is:

```
<linux_kernel_repo>/drivers/net/Ethernet/freescale/xgmac_mdio.c
```

Internal MDIO buses should be listed for all PCS ports that support backplane KR connection, in the device tree. This is because for every port used, the management registers are accessed through the MDIO bus. See DPAA2 architecture for details on how internal MDIO registers block is mapped for every physical port and MDIO registers subchapter of SerDes chapter from SoC reference manual.

If no internal MDIO bus is listed then add one internal MDIO bus for every PCS port target that will be used in a backplane connection.

2.3.2.2 Backplane PHY devices

PCS ports are specific to each board. Backplane PHY devices should be added in board-specific device trees:

```
<linux_kernel_repo>/arch/arm64/boot/dts/freescale/fsl-<device>-<qds, rdb>.dts.
```

A backplane PHY device is registered on an internal MDIO bus. The block below is an example:

```
&pcs_mdio1 {
    pcs_phys1: ethernet-phy@0  {
        backplane-mode = “10gbase-kr”;
        compatible = “ethernet-phy-ieee802.3-c45”;
        reg = <0x0>;
        fsl, lane-handle = <&serdes1>;
        fsl, lane-reg = <0x9c0 0x40>; /* lane H */
};
```

Pcs-phys1 is listed on the MDIO bus and should be discovered when this bus is probed. The kernel backplane PHY driver should also register a PHY driver using PHY hardware ID (read using MDIO bus).

The `backplane-mode` attribute informs the kernel backplane PHY driver about how to configure a specific SerDes lane. Currently, a SerDes lane can be configured as: 10gbase-kr.

The `fsl, lane-handle` attribute is used to identify which SerDes lane the PCS port belongs to. In this case “serdes1” is used.

The `fsl, lane-reg` attribute is used to identify the SerDes lane used to send and receive data. `0x9c0` is the lane H offset in the SerDes1 internal memory map. See each platform's SoC Reference Manual for details and to find the other lane’s offsets.

For LS2088A boards backplane PHY devices are added already for use with lane from H to E.

NOTE

For SerDes 1 lane are numbered in the reversed order compared to WRIO physical ports and MACs.

If the backplane PHY device is not registered on the internal MDIO buses for a specific board, then it can be added in the DTS.

2.3.2.3 Connect with Backplane PHY device handle

The kernel PHY driver is instantiated by the kernel MAC driver, there should be a specified connection between the MAC and a specific PHY in the device tree. In the following example, the backplane PHY from SerDes lane H is used:

```
&dpmac1 {
        phy-handle = <&pcs_phy1>;
};
```

2.4 Enable backplane support in U-Boot

This step is only required for T2080 devices.

Specify all KR ports by using the property `fsl_10gkr_copper` in environment variable `hwconfig`. The values assigned to this property identifies the port that is to be enabled in KR mode: `fm1_10g1`, `fm1_10g2`.

For example: `hwconfig=fsl_10gkr_copper:fm1_10g1,fm1_10g2`

2.5 SerDes setup

- Enable XFI protocol on SerDes lane by using correct RCW (Reset Configuration Word)
- Initialize the SerDes lane registers with recommended values for mode:
  - Ethernet 10GBASE-KR

    SerDes lane registers can be initialized:
    - directly from initial RCW loaded

    or these registers can be updated later:
    - from U-Boot by using command: `mw` - memory write (fill)

      ```
      => mw.l <tecr0_address> <tecr0_hex_value>
      ```

    - from Linux by using command: `devmem`

      ```
      # devmem <tecr0_address> <tecr0_hex_value>
      ```

      NOTE

      Setup `AMP_RED` (amplitude reduction) at recommended value for 10GBase-KR (according to SerDes module RM): `0b000000`.
- Check the link capabilities with AN - software.
- Train the link - software.

2.6 Board configuration

Hardware adjustments

XFI retimers soldered on boards must be removed and PCS output signals should be routed directly to SFP+ cages pins. This is a very important operation and should be carried out carefully.

On SerDes1 module, XFI/Base-KR (XFI/10GBase-KR) protocol will be activated on desired lanes.

Connection cables

Connect with passive direct attach cable.

Two custom boards will be connected back to back with a passive direct attach copper cable, with a maximum length of 1m (for example: SFP-H10GB-CU1M).

3 Use Cases

ping

In order to run a backplane *ping use case*, two boards must be connected back to back with a passive direct attach copper cable. Start MC with specified DPC file. Apply DPL using `fsl_mc apply dpl` command from U-boot and then boot Linux on both boards. After booting Linux, the interfaces must be configured properly for the two ports connected together using two IP addresses from the same IP class.

For example, use:

- On first board: `ifconfig ni0 1.1.1.1`
- On second board: `ifconfig ni0 1.1.1.2`

Once the interfaces are configured, traffic can be sent between the two the boards through the backplane link:

- On first board: `ping 1.1.1.2`
- On second board: `ping 1.1.1.1`

netperf

The backplane *netperf use case* is similar to the ping usecase described above, and it is used for backplane performance benchmark. The board configuration must be done identically as described above. The difference is how traffic is sent between the two boards.

For example using UDP streams:

- On first board: `netperf -H 1.1.1.2 -l 60 -t UDP_STREAM -N &`
- On second board: `netperf -H 1.1.1.1 -l 60 -t UDP_STREAM -N &`

4 BaseKR Statistics

BaseKR algorithm statistics are available for the backplane driver by using ethtool PHY statistics counters. PHY statistics counters are displayed by using the following command:

```
ethtool --phy-statistics <intf>
```

Example: `ethtool --phy-statistics fm1-mac9`

This is an example of PHY statistics output. List of counter meanings are mentioned below:

Counters	Value	Description
LP detected	1	Link Partner detected
PCS Link up	1	Link state at the time of running ethtool command
PCS Link lost detected count	2	Number of times the link was detected as lost
AN Link up	0	AN Link state at the time of running ethtool command
AN Link lost detected count	2	Number of times AN link was detected as lost
Autonegotiation complete	1	Autonegotiation was successfully completed

Table continues on the next page...

Table continued from the previous page...

Counters	Value	Description
Autonegotiation restarted count	2	Number of times the Autonegotiation was restarted
PCS reporting high BER	0	PCS detected a high Bit Error Rate
BER counter	0	Bit Error Rate (BER) counter
Initial RATIO_PREQ	3	Initial ratio of full swing transition bit to pre-cursor used by the algorithm at startup
Initial RATIO_PST1Q	10	Initial ratio of full swing transition bit to post-cursor used by the algorithm at startup
Initial ADPT_EQ	41	Initial value of transmitter adjustment value used by the algorithm at startup
Current RATIO_PREQ	3	Current value of pre-cursor ratio at the time of running ethtool command
Current RATIO_PST1Q	10	Current value of post-cursor ratio at the time of running ethtool command
Current ADPT_EQ	41	Current value of transmitter adjustment value at the time of running ethtool command
Tuned RATIO_PREQ	3	Final value of pre-cursor ratio tuned by the training algorithm
Tuned RATIO_PST1Q	10	Final value of post-cursor ratio tuned by the training algorithm
Tuned ADPT_EQ	41	Final value of transmitter adjustment value tuned by the training algorithm
Initial TECR0	270741511	Initial value of TECR0 register used by the algorithm at startup
Tuned TECR0	270741511	Final value of TECR0 register tuned by the training algorithm
LT complete	1	Link training was successfully completed
LT duration	145	Total duration for all steps of Link training (in msec)
Link training steps	13	Total number of Link training steps
Link training restarted	39	Number of times the Link training was restarted
Link training fail count	26	Number of times the Link training failed
Link training timeout count	26	Number of times the Link training resulted in timeout
Remote Tx tuning cycle count	0	Total number of Remote Tx tuning cycles for all training steps
Local Tx tuning cycle count	0	Total number of Local Tx tuning cycles for all training steps
Coefficient Updates to LP	0	Total number of Coefficient Updates requests sent to link partner

Table continues on the next page...

Table continued from the previous page...

Counters	Value	Description
Coefficient Updates from LP	0	Total number of Coefficient Updates requests received from link partner
C(+1) increment count	2	Number of post-cursor increments
C(0) increment count	0	Number of main-cursor increments
C(-1) increment count	1	Number of pre-cursor increments
C(+1) decrement count	0	Number of post-cursor decrements
C(0) decrement count	0	Number of main-cursor decrements
C(-1) decrement count	0	Number of pre-cursor decrements
LD Preset count	1	Number of local device preset counts
LD Init count	1	Number of local device initialization counts
LD receiver ready	1	Local Device receiver is ready
LP receiver ready	1	Link Partner receiver is ready
Rx EQ Median Gaink2	15	Rx EQ Median Gaink2 value from all snapshots
PRBS sequence bit errors	0	PRBS Sequence bit errors counter

NOTE

On DPAA2 devices 'PHY statistics' must be collected on MAC interface by using the command: `ethtool -phy-statistics <macX>` and therefore linux kernel must be built with the following config options:

```
CONFIG_FSL_DPAA2_MAC=y
CONFIG_FSL_DPAA2_MAC_NETDEVS=y
```

5 BaseKR Algorithm Trace

BaseKR Algorithm Trace is based on Linux kernel `ftrace`. In order to use it, enable the following in Kernel:

- FTRACE
- Kernel Function Tracer

To facilitate early boot debugging, use the boot option, `trace_event=[event-list]` in `bootargs` environment variable.

The following trace events are available specifically for BaseKR Algorithm Trace:

- xgkr_debug_log - logs debug and trace information about KR algorithm
- xgkr_coe_update - logs information about KR coefficients update
- xgkr_coe_status - logs information about KR coefficients status
- xgkr_bin_snapshots - logs information about collected Bin snapshots: Bin1, Bin2, Bin3, Bin Offset, BinM1, and BinLong
- xgkr_gain_snapshots - logs information about collected Gain snapshots: GainK2, GainK3, and OSESTAT

Example:

```
setenv bootargs "console=ttyAMA0,115200 root=/dev/ram0 ramdisk_size=0x2000000
trace_event=xgkr_debug_log,xgkr_coe_update,xgkr_coe_status,xgkr_bin_snapshots,xgkr_gain_snapshots"
```

The traces are logged in the file: `/sys/kernel/debug/tracing/trace`.

How To Reach Us

Home Page:

nxp.com

Web Support:

nxp.com/support

Information in this document is provided solely to enable system and software implementers to use NXP products. There are no express or implied copyright licenses granted hereunder to design or fabricate any integrated circuits based on the information in this document. NXP reserves the right to make changes without further notice to any products herein.

NXP makes no warranty, representation, or guarantee regarding the suitability of its products for any particular purpose, nor does NXP assume any liability arising out of the application or use of any product or circuit, and specifically disclaims any and all liability, including without limitation consequential or incidental damages. “Typical” parameters that may be provided in NXP data sheets and/or specifications can and do vary in different applications, and actual performance may vary over time. All operating parameters, including “typicals,” must be validated for each customer application by customer's technical experts. NXP does not convey any license under its patent rights nor the rights of others. NXP sells products pursuant to standard terms and conditions of sale, which can be found at the following address: nxp.com/SalesTermsandConditions.

While NXP has implemented advanced security features, all products may be subject to unidentified vulnerabilities. Customers are responsible for the design and operation of their applications and products to reduce the effect of these vulnerabilities on customer’s applications and products, and NXP accepts no liability for any vulnerability that is discovered. Customers should implement appropriate design and operating safeguards to minimize the risks associated with their applications and products.

NXP, the NXP logo, Freescale, the Freescale logo, CodeWarrior, Layerscape, PowerQUICC, QorIQ, CoreNet, and QUICC Engine are trademarks of NXP B.V. All other product or service names are the property of their respective owners. Arm, Cortex, and TrustZone are registered trademarks of Arm Limited (or its subsidiaries) 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.

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

Date of release: 09/2019
Document identifier: AN12572

arm