

ADIN6310 Hardware and TSN Switch Evaluation User Guide

FEATURES

- ▶ 6 port TSN switch with RGMII or SGMII interface
 - ▶ 6 RGMII ports to 10 Mbps/100 Mbps/1000 Mbps ADIN1300 PHYs
 - ▶ RJ45 with integrated magnetics
 - ▶ 4 SGMII ports connected to on board SFP cages
- ▶ Host interface hardware strapping with jumpers, choice of
 - ▶ S/D/Q SPI interface
 - ▶ Ethernet port through RJ45 (Port 0)
- ▶ FMC (LPC) connector
 - ▶ Host port access through S/D/Q SPI interface or Port 0
- ▶ PHY Strapping through surface-mount configuration resistors
 - ▶ Default state is software power down from Port 1 to Port 5
 - ▶ Switch firmware manages PHY operation over MDIO
- ▶ Operates from a single, external 9 V to 17 V supply
- ▶ LED indicators on GPIO pins
- ▶ IEEE 802.1AS Time Synchronization
- ▶ Scheduled traffic (IEEE 802.1Qbv)
- ▶ Frame preemption (IEEE 802.1Qbu)
- ▶ Frame replication and elimination for reliability (IEEE 802.1CB)
- ▶ Per stream filtering and policing (IEEE 802.1Qci)
- ▶ VLAN table control (remapping, reprioritization)
- ▶ IGMP snooping
- ▶ GPIO/Timer control

EVALUATION KIT CONTENTS

- ▶ EVAL-ADIN6310EBZ evaluation board
- ▶ 9 V or 12 V, 18 W wall adapter with international adapters
- ▶ 1 Ethernet cable

EQUIPMENT NEEDED

- ▶ EVAL-ADIN6310EBZ evaluation kit
- ▶ Ethernet cables
- ▶ PC running Windows® 10 or Windows® 11

DOCUMENTS NEEDED

- ▶ [ADIN6310 data sheet](#)

SOFTWARE NEEDED

- ▶ TSN application suite (switch configuration GUI and web server)
- ▶ Npcap packet capture

GENERAL DESCRIPTION

The EVAL-ADIN6310EBZ is a flexible platform, which enables an efficient evaluation of the ADIN6310 industrial Ethernet Switch with time sensitive networking (TSN) capability. This user guide describes the hardware kit and software evaluation package (**TSN Switch Evaluation** application). It discusses how to use the kit to interface to one or more Switches to configure the Switch, TSN, or redundancy features to meet the requirements of an industrial network.

The **TSN Switch Evaluation** application enables the initial evaluation of the Switch and its functionality, which further enables the users to familiarize themselves with the Switch capability in advance of migrating to the driver library. Simply connect a PC through Ethernet port to Port 0 on the evaluation board and run the application. The **TSN Switch Evaluation** application can identify and allow configuration of a chain of up to 10 ADIN6310 devices. The application launches a PC-based web server and NETCONF server for each Switch device it finds. A user can interact with the web server to configure the Switch functionality or load YANG configurations from a NETCONF client. Once configuration completes, the user applications can communicate with other devices over the TSN network.

Figure 1 shows an overview of the evaluation board. Full specifications on the ADIN6310 are available in the ADIN6310 data sheet available from Analog Devices, Inc., and must be consulted with this user guide and hardware reference manual when using the EVAL-ADIN6310EBZ evaluation board.

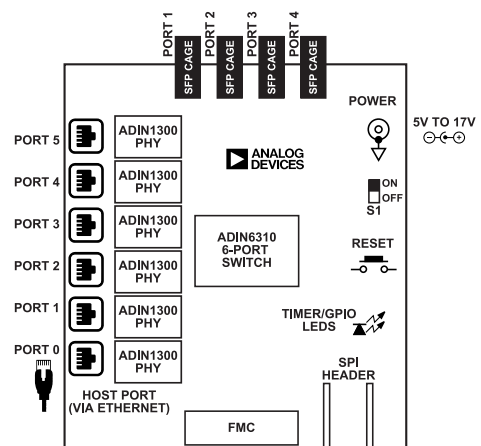


Figure 1. Hardware Overview

TABLE OF CONTENTS

Features.....	1	VLAN Reprioritization.....	40
Evaluation Kit Contents.....	1	Time Synchronization IEEE 802.1AS.....	42
Equipment Needed.....	1	Candidate Page.....	42
Documents Needed.....	1	PTP Configuration.....	42
Software Needed.....	1	Port Configuration.....	43
General Description.....	1	Common Services.....	45
Evaluation Board Hardware.....	6	PTP Instances.....	46
Power Supplies.....	6	Hardware Clock.....	48
Power Sequencing.....	6	External Port Config.....	48
Evaluation Board Use Cases.....	6	Status Page.....	49
Jumper And Switch Options.....	6	Time Sync Messaging.....	51
GPIO and TIMER Headers.....	7	Running Page.....	52
Clock Options.....	7	Startup Page.....	52
On-Board LEDs.....	7	Timer Pins, 1PPS Signal.....	52
Strapping And Configuration.....	7	Time Synchronization, IEEE 1588.....	53
MDIO Interface.....	8	Candidate Page.....	53
FMC Connector.....	8	Status Page.....	56
Software Installation.....	10	Frame Preemption.....	57
Installing the TSN Switch Evaluation		Candidate Page.....	57
Application Software.....	10	Status Page.....	58
Npcap Installation.....	11	Frame Preemption Example.....	59
TSN Switch Evaluation Contents Overview.....	12	Scheduled Traffic.....	60
Initial Evaluation Board Setup.....	17	Assigning Queues.....	60
Software Execution.....	18	Scheduled Traffic – Set Queue Max. SDU.....	63
TSN Switch Evaluation Web Page Overview.....	20	Scheduled Traffic – Schedule.....	65
Candidate/Running/Startup Pages.....	21	Schedule Enabled.....	65
Setup Page.....	22	Guard Bands.....	65
Save and Load Candidate Datastore.....	22	Cycle Time.....	66
Datastore Management.....	22	Base Time.....	66
Advanced.....	22	Cycle Time Extension.....	66
Performing a Reset.....	23	Gate States When Disabled.....	67
Port Statistics.....	24	Cut-Through Allowed.....	67
Port Configuration.....	25	Gate Control List, Time Intervals.....	67
Candidate Page.....	25	Hold En.....	67
Status Page.....	26	Candidate Page.....	68
MDIO Control.....	27	Running Page.....	69
GPIO and Timer Configuration.....	28	Startup Page.....	70
TSN Output Timer.....	28	Schedule on the Timer Pins.....	71
1PPS Periodic Output.....	28	LLDP Configuration.....	73
Periodic Output.....	28	Link Layer Discovery Protocol (LLDP).....	73
Capture Input.....	28	LLDP Candidate View.....	73
Other Modes.....	28	LLDP Status.....	75
Switching Table.....	30	LLDP Example	76
Candidate View.....	30	LLDP Example (Fast Tx).....	77
Status View – Dynamic Entries.....	33	Parallel Redundancy Protocol (PRP).....	78
Stream Table.....	34	Enabling PRP Example.....	79
VLAN Control.....	35	PRP Configuration Web Page Views.....	79
VLAN Table.....	35	PRP Status Page.....	83
VLAN Remapping.....	38	PRP – Supervision Frames.....	84

TABLE OF CONTENTS

PRP – Capture of PRP Tagged Traffic.....	87	Sequence Recovery.....	116
VLAN Table Operation in PRP Mode.....	88	Individual Recovery.....	116
Switching Table in PRP Mode.....	88	FRER Stream Table.....	116
High Availability Seamless Redundancy (HSR)..	89	FRER Configuration–Candidate View.....	117
HSR Operating Modes.....	89	FRER Status.....	119
Enabling HSR Example.....	89	Talker–Listener Configuration Example.....	122
HSR Candidate View.....	90	Internet Group Management Protocol (IGMP)	
HSR Status View.....	92	Snooping.....	124
HSR Running View.....	94	Router Timeout.....	124
HSR Startup View.....	94	Group Member Timeout.....	125
HSR – Supervision Frames.....	95	IGMP Versions.....	126
HSR – Capture of HSR Tagged Traffic.....	96	IGMP Snooping Example.....	126
Media Redundancy Protocol (MRP).....	97	NETCONF/YANG.....	130
MRP Stack on the Switch.....	97	Sysrepo Datastore.....	130
Recovery Profiles.....	97	YANG Models.....	130
Configuring MRP.....	97	Custom Leaf Nodes.....	130
Candidate Page.....	97	Startup Configuration.....	130
MRP Scenarios: MRM and MRC.....	99	Web Server Use and NETCONF.....	131
MRP Scenarios: MRA.....	101	YANG Model Examples.....	131
Per-Stream Filtering And Policing, Qci.....	103	Firmware Update.....	133
PSFP Candidate Page.....	104	Automatic Firmware Update.....	133
PSFP Status Page.....	107	Paired Firmware and Web Server.....	133
Multiple Spanning Tree Protocol (MSTP).....	108	Firmware Downgrade.....	133
MSTP Candidate Page.....	108	Troubleshooting.....	135
MSTP Status Page.....	111	GUI Does Not Find ADIN6310 Devices.....	135
Frame Replication and Elimination for		GUI Table Remains Blank.....	135
Reliability (FRER), 802.1CB.....	113	Web Page Fails to Load.....	135
Redundancy Tag.....	114	Firmware Did Not Update.....	136
Stream Identification.....	115	GUI Inconsistent at Finding Devices.....	136

REVISION HISTORY**7/2025—Rev. A to Rev. B**

Changes to Equipment Needed Section.....	1
Changes to ses-configuration File Section.....	13
Added PTP Specific Configuration Section.....	17
Changes to Figure 24.....	19
Changes to VLAN Table Configuration Section.....	35
Changes to Trunk/Access Configuration Section.....	35
Changes to Interval Times Section.....	45
Added Figure 58; Renumbered Sequentially	45
Changes to Parallel Redundancy Protocol (PRP) Section.....	78
Changes to Figure 112.....	79
Changes to PRP Candidate View Section and Figure 113.....	79
Changes to Figure 114.....	81
Changes to Figure 115.....	82
Changes to PRP Status Page Section and Figure 116.....	83
Added Figure 117.....	84
Changes to Proxy Node Table Section.....	84

TABLE OF CONTENTS

Changes to PRP – Supervision Frames Section and Figure 118 to Figure 121.....	84
Deleted GPIO/Timer Configuration Tab When Using PRP Mode Section.....	87
Changes to VLAN Table Operation in PRP Mode Section.....	88
Changes to Figure 162.....	113
Changes to Sequence Recovery Section.....	116
Changes to Individual Recovery Section.....	116
Changes to Stream Table Configuration Section and Figure 165.....	116
Changes to FRER Configuration–Candidate View Section.....	117
Changes to FRER Status Section and Figure 167.....	119
Added Table 20, Table 21, and Table 22; Renumbered Sequentially.....	119
Changes to Talker–Listener Configuration Example Section and Figure 169.....	122
Moved Figure 169, Figure 170, and Figure 171.....	122
Deleted Talker System Section, Listener System Section, Figure 167, Figure 168, and Figure 169; Renumbered Sequentially.....	123
2/2025—Rev. 0 to Rev. A	
Changes to ses-configuration File Section.....	13
Changes to Figure 17 to Figure 20.....	15
Changes to MRP Specific Configuration Section.....	16
Changed ADIN6310 and 10BASET1L Hardware Section to EVAL-ADIN6310T1L Hardware Section.....	17
Changes to EVAL-ADIN6310T1L Hardware Section.....	17
Added EVAL-ADIN3310 Hardware Section.....	17
Change to Initial Evaluation Board Setup Section.....	17
Changes to Software Execution Section.....	18
Changes to Figure 24.....	19
Added Figure 25; Renumbered Sequentially.....	19
Changes to TSN Switch Evaluation Web Page Overview Section.....	20
Changes to Figure 26.....	20
Changes to Port Statistics Section and Figure 29.....	24
Changes to Candidate Page Section.....	25
Changes to Candidate View Section.....	30
Changes to Dynamic Table Section.....	30
Added Stream Table Section and Figure 41.....	34
Changed Time Synchronization Section to Time Synchronization IEEE 802.1AS.....	42
Changes to Time Synchronization IEEE 802.1AS.....	42
Changes to PTP Configuration Section and Figure 52.....	42
Change to Table 11.....	45
Changes to Hardware Clock Section.....	48
Added Time Synchronization, IEEE 1588 Section and Figure 70.....	53
Added Candidate Page Section, Figure 71 to Figure 73, and Table 13 and Table 14; Renumbered Sequentially.....	53
Added Status Page Section and Figure 74.....	56
Changes to Frame Preemption Section.....	57
Changes to Scheduled Traffic Section.....	60
Changes to MRP Stack on the Switch Section.....	97
Changes to Configuring MRP Section.....	97
Changes to Figure 145.....	103
Changes to Stream Filter Section and Figure 146.....	104
Changes to Stream Gate Section and Figure 148.....	105

TABLE OF CONTENTS

Changes to PSFP Status Page Section.....	107
Deleted Figure 144.....	107
Changes to Figure 151.....	107
Added Multiple Spanning Tree Protocol (MSTP) Section and Figure 152.....	108
Added MSTP Candidate Page Section, Figure 153, and Figure 154.....	108
Added MSTP Status Page Section and Figure 155.....	111
Changes to Frame Replication and Elimination for Reliability (FRER), 802.1CB Section.....	113
Added Figure 157.....	113
Changed Stream Table Section to FRER Stream Table Section.....	116
Changes to FRER Stream Table Section.....	116
Changes to Stream Table Configuration Section.....	116
Changes to Figure 160 Caption.....	117
Changes to FRER Configuration—Candidate View Section.....	117
Changes to Talker-Listener Configuration Example Section.....	122
Changes to Talker System Section and Figure 164.....	122
Replaced Figure 165.....	123
Changes to GUI Does Not Find ADIN6310 Devices Section.....	135
Added Figure 185.....	135

10/2024—Revision 0: Initial Version

EVALUATION BOARD HARDWARE

POWER SUPPLIES

The EVAL-ADIN6310EBZ operates from a single, external, 5 V to 17 V supply rail. A 9 V or 12 V wall adapter is supplied as part of the kit.

Apply the wall adapter to P2 connector or alternatively 5 V to 17 V to the P1 plug. Switch BRD_ON_OFF to the ON position. The LED DS4 lights up to indicate a successful power up of the main power rails.

The EVAL-ADIN6310EBZ power requirements are generated from the input power rail by an on-board [LTM4668A](#) μ Module regulator, which provides the four rails required for operation of the [ADIN6310](#) Switch, the six [ADIN1300](#) Ethernet PHYs and other support circuitry. The default nominal voltages are listed in [Table 1](#).

By default, the VDDIO_A and VDDIO_B share the same voltage rail and default to 1.8 V with the installed components and jumper settings.

Table 1. Default Device Power Supply Configuration

LTM4668A Output	Nominal Voltage	ADIN6310 Switch	ADIN1300 PHY
V _{OUT1}	3.3 V	VDD3P3	AVDD3P3
V _{OUT2}	1.8 V	VDDIO_A/B	VDDIO
V _{OUT3}	1.1 V	VDDCORE	N/A ¹
V _{OUT4}	0.9 V	N/A ¹	VDD0P9

¹ N/A means not applicable.

The VDDIO_A rail provides a separate voltage domain for the Switch interface pins that can connect to a Host interface. This includes SPI interface, TIMER, GPIO, and Port 0 MAC interface pins. The motivation for partitioning the VDDIO_A/B voltage rails is to ensure flexible Host interface I/O voltage while helping to reduce overall power consumption for the Switch ports and PHY devices. For normal operation of the evaluation hardware, the default voltage rail should be sufficient. If user is connecting own Host interface over SPI or FMC connector, flexibility to change the VDDIO_A rail may be beneficial.

If a different VDDIO_A voltage is required, user can adjust by changing the placement of configuration jumpers. The VDDIO_A rail can be changed from 1.8 V default to either 2.5 V or 3.3 V. To change the VDDIO_A rail to 2.5 V, the LDO, U3 must be used. The jumpers to reconfigure this are P3, P4, P5, and P33.

For more details, see [Table 2](#) and the evaluation board schematics.

Table 2. VDDIO_A Configuration

VDDIO_A	Jumper Setting
1.8 V	P3 (1-2), P4 (1-2), P33 (OPEN)
2.5 V	P3 (OPEN), P4 (2-3), P33 (1-2)
3.3 V	P3 (1-2), P4 (2)-P5(1), P33 (OPEN)

[Table 3](#) shows an overview of the EVAL-ADIN6310EBZ current for various operating modes.

Table 3. Board Quiescent Current (P2 = 9 V)

Board Status	Typical Quiescent Current
On Power-Up (S1 on)	104 mA initially
In Hardware Power-Down (RESET_N Held Low)	72 mA
1000BASE-T, 2 RGMII + HOST Port	250 mA
1000BASE-T, 5 RGMII + HOST Port	360 mA

POWER SEQUENCING

The ADIN6310 device does not have any power supply sequencing requirements, however the preferred power up sequence is to bring up VDDCORE last and removed first on power down. There are no power sequence requirements for the ADIN1300 devices. The evaluation board is configured to bring up the power rails in the following order VDD3P3 and VDD0P9 -> VDDIO_A/B -> VDDCORE.

EVALUATION BOARD USE CASES

The EVAL-ADIN6310EBZ can be used in two general modes. The default and expected use case utilize Port 0 as the Host interface port through the RJ45 connector. Port 0 is connected to a PC running the TSN evaluation software package for network configuration and control. Port 0 can still be used for data traffic, but it is not a part of the time aware network as it is connected to the PC. In this use case, the other five RGMII ports and four SGMII ports on the EVAL-ADIN6310EBZ can be used to evaluate IEEE802.3 and TSN features of the ADIN6310, establish links with other link partners and evaluate the performance of the chip.

Alternatively, the user can connect their own Host directly to the EVAL-ADIN6310EBZ. If Host interface is SPI, then option to connect directly through the SPI header or the FMC LPC connector (FPGA mezzanine card low pin count). The FMC connector can be plugged into an FPGA development board. When the Switch hardware is used with an FPGA board, the media independent interfaces (MIIs) for Port 0, SPI interface, GPIO, and TIMER signals can be connected to the FPGA. In this use case, a MAC-MAC type Host interface can be used on Port 0 or the SPI (quad, dual, or single) interface can be used for control and configuration with the FGPA as the Host processor for evaluation of the ADIN6310 in a full system. With the SPI interface as the HOST interface, the system can have six TSN capable ports.

JUMPER AND SWITCH OPTIONS

Several jumpers on the EVAL-ADIN6310EBZ must be set for the required operating setup before using the EVAL-ADIN6310EBZ for evaluation. The default settings and functions of these jumper options are described in [Table 4](#).

Table 4. Default Jumper, Switch Options and Descriptions

Link	Position	Function
BRD_ON_OFF	OFF	Power ON/OFF Switch
S1	3	Reset options
P3	Inserted	

EVALUATION BOARD HARDWARE

Table 4. Default Jumper, Switch Options and Descriptions (Continued)

Link	Position	Function
P4	1-2 Inserted	VDDIO_A = VDDIO_B = 1.8 V; runs off the Switching regulator
P33	Open	Enable for VDDIO_A LDO
TIMER2	Open	Host strapping (RGMII No Tx Rx Delay)
SPI_SS, TIMER0, TIMER1, TIMER3	1-2 Inserted	Host strapping (RGMII No Tx Rx Delay 1000 Mbps)
P28	1-2 Inserted	Power to TIMER/GPIO LEDs
P41	1-2 Inserted	Connect VCCIO supply of FTDI to VDDIO_A
P11, P13, P17, P18	1-2 inserted	PortX link from PHY
P36	Open	Connect power to U26

GPIO AND TIMER HEADERS

The EVAL-ADIN6310EBZ provides a header (P10) for observation of all Timer and GPIO signals. In addition to the header, there are LEDs on these pins. When using the **TSN Switch Evaluation** application, TIMER2 is configured for a 1 pulse per second (1PPS) signal by default and the LED connected to TIMER2 pin can be observed to blink at a 1 second rate when the board is powered and has been successfully configured using the **TSN Switch Evaluation** application.

If the Switch Host strapping is changed to SPI interface (default is Ethernet Host - RGMII), the TIMER0 pin functionality changes to be an Interrupt signal to the Host and TIMER0 is no longer available for timer or TSN functionality.

CLOCK OPTIONS

A crystal oscillator, Y8, is used to provide the [ADIN6310](#) a clock signal. It is a 25 MHz crystal connected across the XTAL_I pin and XTAL_O pin of the ADIN6310 on the board. The clock for the [ADIN1300](#) Ethernet PHYs can be provided from a buffered 25 MHz clock from the ADIN6310 or alternatively from a dedicated 25 MHz crystal local to each PHY (default). If the buffered clock option is selected, once the ADIN6310 has successfully powered up it generates a 25 MHz clock on CLK_OUT_1 pin. This clock is routed to a clock buffer chip, SI5330F-B00214-GMR (U31), which provides a buffered version of 25 MHz clock to each of the six ADIN1300 transceivers on the board.

ON-BOARD LEDs

The EVAL-ADIN6310EBZ has one LED, DS4, that lights up to indicate a successful power up of the circuit. There are eight LEDs, that are controlled by GPIO (0-3) and the Timer (0-3) signals when link P28 is inserted.

For the ports that support SGMII interface, there are LEDs (DS1, DS2, DS3, DS5) close to the SFP modules. When an SFP module is inserted and the link is up, the LOS signal from the SFP module is used to indicate optical activity/link status.

STRAPPING AND CONFIGURATION

ADIN6310 Host Port Strapping

The ADIN6310 Switch supports stack Processor/Host control over SPI or any of the six Ethernet ports. There is no stack processor/microcontroller used on this board, instead use a Windows PC as the Host with the **TSN Switch Evaluation** package.

When using this hardware, the user can connect a Host in a few different ways, firstly, via Port 0 Ethernet Port, alternatively, connected to RMII/RGMII directly over the FMC connector or otherwise via SPI through the dedicated headers (P39, P40). The Host Hardware strapping jumpers must be set according to the Host interface required.

The default Host port strapping configuration for this hardware is using Ethernet interface with Port 0 as the Host interface. The Switch port is configured for RGMII with no TXC or RXC delays and port speed of 1000 Mbps. In a typical application, with an MII interface to the Host, the Switch MAC port is directly connected to the Host MAC interface without a PHY in the path. As a result, when the Switch is configured for MAC interface Host, the Switch does not expect a PHY and does not perform any PHY configurations for that port. The EVAL-ADIN6310EBZ hardware does include a PHY on Port 0 (default RGMII Host interface), but the **TSN Switch Evaluation** application does not configure this PHY directly. As a result, the link brought up by the PHY on the Host port must match the Switch port speed set by strapping jumpers, default 1000 Mbps. The PHY is hardware strapped to auto-negotiate all speeds, if it brings up a lower speed link, there is a link mismatch between the Switch port and the PHY, which blocks the communication between the Host and Switch.

The Host port and Host port interface selection are configured via jumpers labeled TIMER_0/_1/_2/_3 and SPI_SS.

The Timer and SPI pins have internal pull-up/-down resistors, as shown in [Table 5](#), the strapping jumpers provide user with ability to reconfigure the strapping to select alternative Host port types. For more details on all options available, refer to the Host Strapping section in the data sheet.

Table 5. Host Port Selection Jumpers

Host Port	SPI_SS	TIMER3	TIMER2	TIMER1	TIMER0
Internal Pull up (PU)/Pull down (PD)	PU	PD	PD	PU	PU
SPI (Single)	OPEN	OPEN	OPEN	OPEN	OPEN
SPI (dual)	OPEN	INSERT	OPEN	OPEN	OPEN
SPI (quad) (low drive strength)	INSERT	OPEN	INSERT	OPEN	OPEN
SPI (quad) (high drive strength)	INSERT	INSERT	INSERT	OPEN	OPEN

EVALUATION BOARD HARDWARE

Table 5. Host Port Selection Jumpers (Continued)

Host Port	SPL_SS	TIMER3	TIMER2	TIMER1	TIMER0
RGMII 1000M (default H/W config)	INSERT	INSERT	OPEN	INSERT	INSERT

ADIN1300 Strapping

There are six **ADIN1300** devices on this evaluation board. The PHY on Port 0 is hardware strapped for auto-negotiation for all speeds (10 Mbps/100 Mbps/1000 Mbps), which allow it to bring a link up with a remote partner without any configuration from the Switch/Host. By default, the Switch Host strapping is configured for Port 0 as Host interface, the PHY needs to be able to bring up a link to enable communication path between the Host and the Switch so the **TSN Switch Evaluation** package can configure the Switch.

The other five PHYs (on Port 1 to Port 5) are configured for the same speeds (10 Mbps/100 Mbps/1000 Mbps), but power up in software power down mode with the Switch bringing them out of software power down and configuring them over the MDIO interface.

The PHY strapping upon power up is shown in [Table 6](#).

Table 6. ADIN1300 PHY Port Configuration

Function	PHY Port 0	PHY Port (1-5)
MAC Interface	RGMII With Tx& Rx DLL Enabled	RGMII With Tx& Rx DLL Enabled
MDI Mode	AutoMDI, Pref MDI	AutoMDI, Pref MDI
Speed	10/100 HD/FD, 1000 FD Target	10/100 HD/FD, 1000 FD Target, SftPd

ADIN1300 Link Status Polarity

The ADIN1300 LINK_ST output pin is active high by default, whereas the P0_LINK input of the **ADIN6310** is active low by default, therefore the EVAL-ADIN6310EBZ hardware includes an inverter in the path between the Port 0 PHY LINK_ST and the P0_LINK of the Switch. The other five ports do not include this inverter, instead the PHY link polarity is changed to default low during the initial configuration.

As a result of this hardware difference on the Host Port 0, when the board first powers up, prior to configuration, the right LED of Port 1 to Port 5 lights. Once the configuration has been loaded from the **TSN Switch Evaluation** application, the PHY link signal as seen at the LEDs matches for all ports. The RJ45 right LEDs light to indicate link up, the left LEDs are on for link up and blink for traffic activity.

ADIN1300 Link Status Voltage Domain

The ADIN1300 LINK_ST is primarily intended to drive the Switch Px_LINK input signal, therefore, resides on the VDDIO_A/B voltage domain (default voltage rail is 1.8 V). If using the LINK_ST pin to drive an LED to indicate link active, a level shifter must be used to

provide voltage and drive capability for the LED function. The LED anode is connected to 3.3 V through a 470 Ω resistors.

ADIN1300 PHY Addressing

The ADIN1300 PHY addresses are configured by sampling their RXD pins after power on, when they come out of reset. The ADIN6310 Switch has internal pull-up/-down resistors on its RXD pins to support assignment of unique PHY addresses to each PHY per port. As a result, external PHY address strapping resistors are not necessary, unless different PHY addressing is required. The default PHY addresses assigned to the ADIN1300 devices is shown in [Table 7](#).

Table 7. Default PHY Addressing (set by ADIN6310)

Port Number	PHY Address
0	0
1	1
2	2
3	4
4	8
5	9

MDIO INTERFACE

The MDIO bus of the ADIN6310 connects to the MDIO bus of each of the six PHYs on the evaluation board. Configuration of the PHYs is done by the Switch firmware via this MDIO bus. The **TSN Switch Evaluation** application supports read and write access of the PHYs on all ports.

FMC CONNECTOR

This evaluation board is fitted with a low pin count FPGA mezzanine card (LPC FMC) connector on the back of the board. This allows it to interface directly with a compatible FPGA board. All port 0 signals, SPI, TIMER, and GPIO signals are brought directly to the connector. This allows users to directly interface with the ADIN6310 with any one of the three Host interface options SPI, RGMII, and RMII. To use the FMC connector to interface with an FPGA or a processor board, make the changes to the resistor set, as shown in [Table 8](#).

Table 8. Resistor configuration for FMC Use

Signal	Remove	Install
	RGMII/RMII	
P0_TXC	R239	R227
P0_TXCTL	R240	R228
P0_TXD0	R238	R226
P0_TXD1	R237	R225
P0_TXD2	R236	R224
P0_TXD3	R235	R217
P0_RXC	R242	R219
P0_RXCTL	R241	R218
P0_RXD0	R243	R220

EVALUATION BOARD HARDWARE

Table 8. Resistor configuration for FMC Use (Continued)

Signal	Remove	Install
RGMIIRMI		
P0_RXD1	R244	R221
P0_RXD2	R245	R222
P0_RXD3	R246	R223
SPI		
SPI_SS	R485	R372
SPI_SCLK	R484	R329
SPI_SIO0	R493	R358
SPI_SIO1	R492	R365
SPI_SIO2	R499	R373
SPI_SIO3	R501	R379
TIMER0	R494	R378

SOFTWARE INSTALLATION

INSTALLING THE TSN SWITCH EVALUATION APPLICATION SOFTWARE

The evaluation package runs on Windows 10. To use the **TSN Switch Evaluation** software, first run the installer package to install the GUI and PC based web server. The installation steps are listed in the following section. The default location for the TSN Switch software install is **C:\Analog\ADINx310EVKSW-Relx.x.x** folder.

When the **TSN Switch Evaluation** software installation is complete, install Npcap if not already present on the machine. Download from Packet capture library for windows on the Npcap website. Npcap is recommended over WinPcap.

TSN Switch Evaluation Software Installation

To install the **TSN Switch Evaluation** software package, do the following steps:

1. Launch the installer file to begin the **TSN Switch Evaluation** software installation.
2. If a window appears asking for permission to allow the program to make changes to the PC, click **Yes**.
3. The installation process starts, see [Figure 2](#).

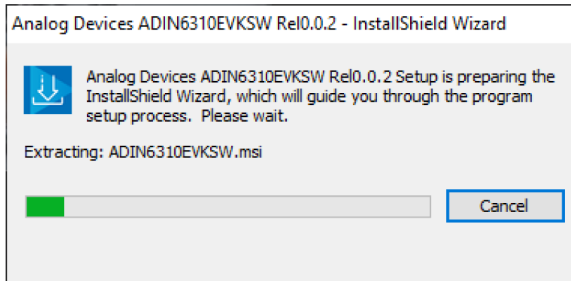


Figure 2. Installation Begins

4. The welcome window appears (see [Figure 3](#)), with prompts that user must separately install Npcap, click **Next**.

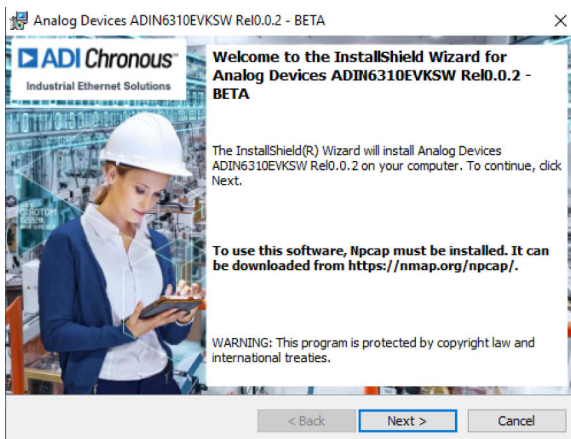


Figure 3. Welcome Message

5. A license agreement appears. Read the agreement and click **I accept the terms in the license agreement** to allow the installation to proceed, as shown in [Figure 4](#), click **Next**.

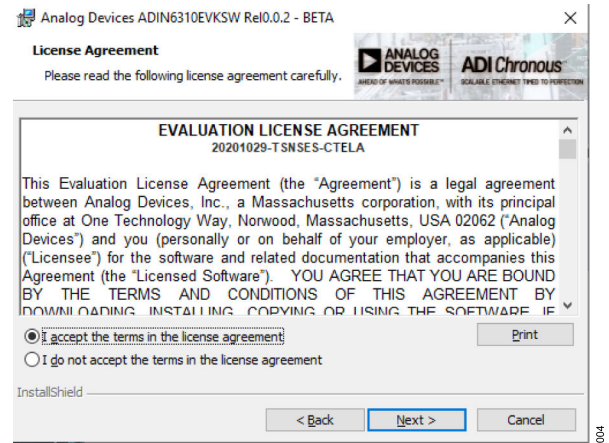


Figure 4. License Agreement

6. Select the location to install the **TSN Switch Evaluation** software and click **Next** (see [Figure 5](#)).




Figure 5. Select Destination Folder

7. At the next step, click **Install** (see [Figure 6](#)).

SOFTWARE INSTALLATION

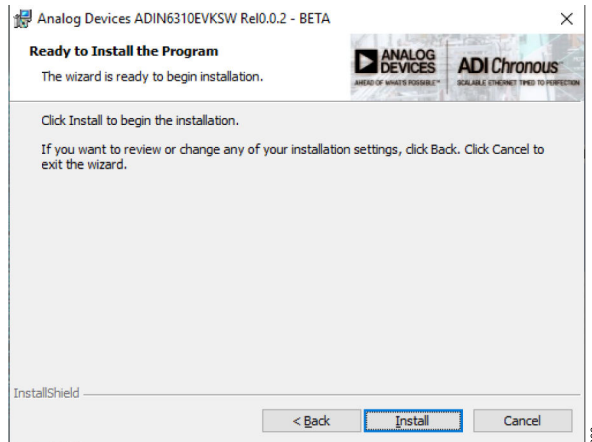


Figure 6. Installation Begins




Figure 8. Remove or Repair Installation

8. A window appears, which shows the progress of the installation. When installation is complete, click **Finish**. (see Figure 7).

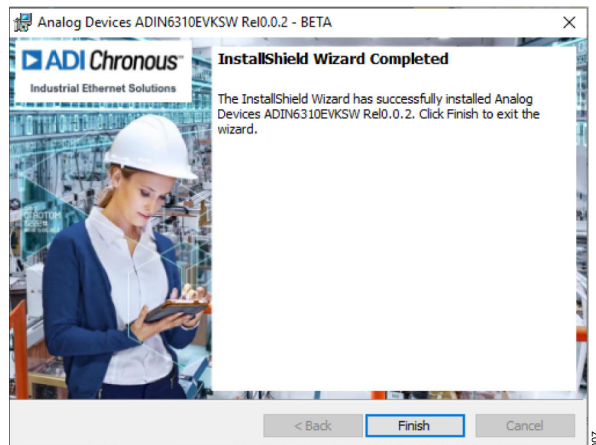


Figure 7. Installation Completes

2. Follow the steps until complete, click **Finish** (see Figure 9).

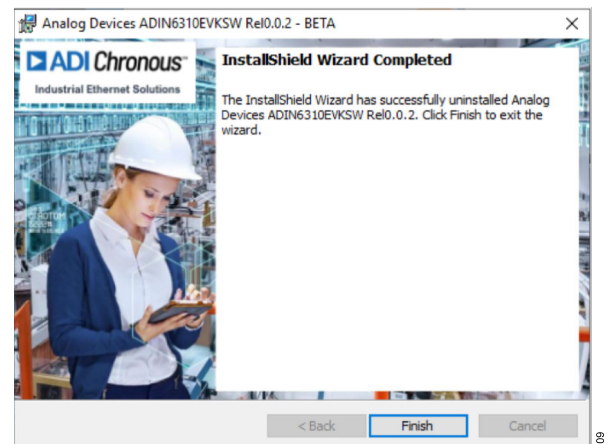


Figure 9. Remove or Repair Completes

Repair/Remove Installation

The installer also supports removing or repairing the installation. Launch the installer to remove or repair and do the following steps:

1. Select **Repair** or **Remove** and click **Next** (see Figure 8).

NPCAP INSTALLATION

Install Npcap if not already present on the machine. Npcap is recommended over WinPcap. Download from Packet capture library for Windows on the Npcap website.

When installing NPCAP, ensure that the **Install Npcap in WinPcap API-compatible Mode** check box is selected, as shown in Figure 10.

SOFTWARE INSTALLATION




Figure 10. Npcap Installation Option (WinPcap API-Compatible Mode Selected)

TSN SWITCH EVALUATION CONTENTS OVERVIEW

The software consists of GUI used to identify the Switch or chain of Switches and launch a PC based web server for each ADIN6310 device connected. The following section shows the different portions of the software. The default location for the TSN Switch Evaluation software install is C:\Analog\ADINx310EVKSW-Relx.x.x folder (see Figure 11).




Figure 11. Main Folder

This main folder contains the following sub-folders:

- ▶ The **bin** folder contains the firmware. New versions of the **TSN Switch Evaluation** package take care of automatically updating the latest firmware when initially run.
- ▶ The **doc** folder contains release note, schematics, and layout for the evaluation board in PDF format.
- ▶ The **exe** folder contains the executable (GUI), configuration files, and the web server file system (see Figure 12).
- ▶ The **lic** contains the license files (ELA license).




Figure 12. Contents of the exe Sub-Folder

ADIN6310-tsn-evaluation-util

The **TSN Switch Evaluation** application GUI is named as **ADIN6310-tsn-evaluation-util.exe**. This application executes on a Windows PC platform and is used to query ADIN6310 boards that are present on a network. When an ADIN6310 board is found, the GUI configures the device primary MAC address, and allow the user to launch the **TSN Switch Evaluation** web page.

Process Application (windows-tsn-io-app)

The process tool runs automatically in the background for each instance of SES device found and does not need to be launched by the user.

Modules Folder

The **modules** folder contains yang models and start-up configuration.

FileSystemFolders

The **FileSystemFolders** folder (see Figure 13) contains the PC-based web server pages for each instance of the Switch that can be supported by the GUI (up to 10 max). Each device has its own file system, which is emulated on the PC by having a unique folder to act as the file system root.




Figure 13. Contents of the FileSystemFolders

SOFTWARE INSTALLATION

When the application is first run, the process needs to create a repository inside the file system instance, this can take some time, on order of 30 seconds to complete. Creating the repository is done first, prior to communicating with the Switch. Once the repository is successfully created, only then the process starts to communicate with the Switch and load the default start up configuration.

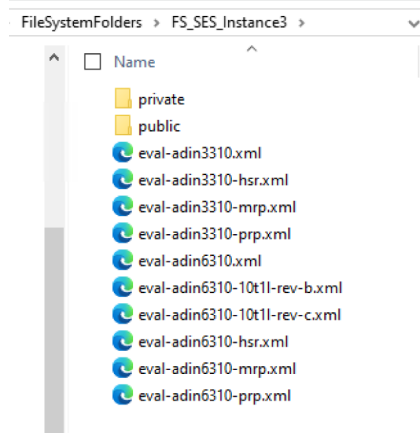


Figure 14. Contents of an FS_SES_Instance_0 Folder Prior to First Run of Application

Once the application runs successfully, additional folders can be observed in the File System Folders, specifically the eventLog, log, and repository folders (see Figure 15).




Figure 15. Contents of an FS_SES_Instance_0 Folder After Running Application and Successfully Configuring Device

ses-configuration File

The ses-configuration.txt file shown in Figure 16 contains configuration parameters for the PC based web server, such as IP address, Port, NETCONF server port, location of file system, and hardware configuration XML file:

- ▶ **IP and Port address:** Specifies the IP and Port address used by the process application instance web pages. For the ADIN6310 evaluation kit the IP address must be set to use the local Host,

otherwise known as the loop back address, which is fixed to 127.0.0.1. Given the IP must remain the same for all process instances a port number must be used to identify which process instance the web pages belong to. This allows multiple instances of the process application to execute while controlling each board independently.

- ▶ **FsName:** Name of file system folder for each device.
- ▶ **NetconfPortSsh:** Port on which NETCONF server is listening (SSH), different port for each SES device.
- ▶ **ImageType:** Pass **Production**. Production applies for all released material or sample material with -U3 branding on the package.

There are 10 instances included in the folder, one instance for each possible Switches in the network (up to 10 maximum supported by GUI). The StartupFileName points to the board specific configuration for the device, and for this example is using the EVAL-ADIN6310EBZ evaluation board. The software also supports operation with the EVAL-ADIN3310 and EVAL-ADIN6310T1L versions of hardware.

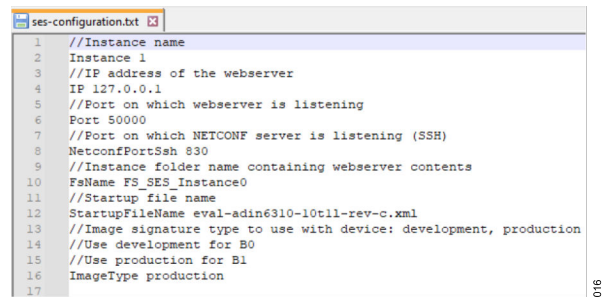


Figure 16. Contents of ses-configuration

The device/hardware specific configuration is contained in XML files within in each FS_SES_Instance folder inside the File System Folders.

Pass the matching xml file name to the ses-configuration.txt instance for the version of hardware being used and the required mode of operation.

Example XML files are provided for various configurations and parameters can be modified within the XML files. The Switch configuration is volatile and power cycling the Switch requires reconfiguration:

- ▶ **Default configuration** is for TSN functionality (for example, file names eval-adin6310 and eval-adin6310-10t11.xml all support TSN capability).
- ▶ **Redundancy** configuration examples are provided for HSR, PRP, and MRP.

Syntax and case are important when modifying parameters in the XML file. Errors or passing incorrect parameters is not supported and affect the operation of application.

SOFTWARE INSTALLATION

Per-port configuration parameters, including MII mode and PHY related specifics, see [Figure 18](#).

- ▶ **MII/Port MAC Interface Selection:** All ports support RMII/RGMII, additionally Port 1 to Port 4 support the following MAC interfaces, however hardware must be configured to match the required MAC interface:
 - ▶ SGMII
 - ▶ 1000base-SX/LX
 - ▶ 1000Base-KX
 - ▶ 100BASE-FX
- ▶ **If-type:** The default configuration for the EVAL-ADIN6310EBZ is RGMII interface to the [ADIN1300](#) PHYs. The hardware does not support RMII interface to the PHYs. Port 0 (Ethernet Host) is always configured in **unmanaged** mode and the Switch does not configure that PHY directly. Hardware must be capable of the MII configured by software, for example, EVAL-ADIN6310 and EVAL-ADIN3310 hardware can support MII modes: **rgmii**, **sgmii**, **sgmii-1000base-sxlx**, or **sgmii-100base-fx**. EVAL-ADIN6310T1LEBZ supports RGMII interface for all PHYs and can optionally support **sgmii**, **sgmii-1000base-kx**, or **sgmii-100base-fx** options on Port 2 and Port 3.
- ▶ **Phy-rx-delay-supported/phy-tx-delay-supported:** RxDelay/TxDelay: RXC and TXC delays configuration for the Port.
- ▶ **Phy-type:** Per port identification of what PHY is connected. Choice of [ADIN1100](#), [ADIN1200](#), [ADIN1300](#), or Unmanaged for ports that either have no PHY or have a different PHY. Hardware must match/support. By default **Unmanaged** is passed to Port 0.
- ▶ **clock-selection:** For use with RMII mode only. A setting of 0 enables a 50 MHz clock to be output onto the Port TXC pin for use by the PHY. Only use RMII mode where hardware is configured appropriately, EVAL-ADIN6310EBZ evaluation board supports RGMII mode by default for all ports.
- ▶ **PHY Address:** PHY address as configured by internal/external strapping. EVAL-ADIN6310 evaluation board uses ADIN6310 internal strapping to provide unique PHY address to each PHY, see [ADIN1300 PHY Addressing](#) section.
- ▶ **Link-polarity:** ADIN6310 expects Port `_LINK` pin to be driven low for link up, high for link down. The default polarity of the ADIN1200/ADIN1300/ADIN1100 PHYs is for the `LINK_ST` pin to be active high with link up, however the polarity can be inverted via MDIO write if needed as part of the port initialization. In the **ses-configuration.txt** file, for this parameter, a setting of active-low indicates the default is active low (no inversion needed), while passing active-high instructs the ADIN6310 to

perform a MDIO write to invert polarity of LINK signal in PHY as part of the initialization routine. For the EVAL-ADIN6310EBZ evaluation hardware, there are six ADIN1300 PHYs, the PHY on Port 0 includes an inverter in the path between the `LINK_ST` and the Switch `P0_LINK` pin, therefore the inversion is already done for that port. For the remaining PHYs on Port 1 to Port 5, there is no inverter in the path, instead the ADIN6310 configuration needs to invert the polarity of the PHY `LINK_ST` pin by writing over MDIO to configure the PHY.

- ▶ **Phy-pull-up-control:** Options of: internal, external, do-not-disable. Allows configuration of whether the PHY address strapping uses the internal pull resistors from the Switch RXD lines or uses external strapping resistors for PHY addressing. With EVAL-ADIN6310EBZ, the internal or do-not-disable options must be used. Do not use the **external** option as there are no external PHY address strapping resistors and this results in all PHYs defaulting to Address 0.
 - ▶ **Internal:** Internal pulls are enabled. Default setting for the EVAL-ADIN6310EBZ evaluation board, the Switch sets unique PHY addresses for each PHY. No external strapping resistors are required for PHY addressing as a result. The strapping resistors are enabled until the PHY is brought out of reset and then disabled.
 - ▶ **External:** Internal pulls are disabled. Use with EVAL-ADIN6310T1LEBZ evaluation board, external resistors are used to configure PHY addresses, therefore the internal pulls are disabled.
 - ▶ **Do not disable:** Internal pulls are enabled and left enabled even after the PHYs are configured.
- ▶ **Speed:** Choice of 1, 0.1, 0.01 (Gbps).
- ▶ **Duplex:** Choice of full or half. The default is full.

Device: Configuration specifics such as device MAC Address, what redundancy capability is enabled and configuration parameters for the PTP stack are passed next. See [Figure 19](#).

- ▶ **MAC address:** Specifies the mac address the ADIN6310 hardware uses. The MAC address is unique to each device found and is also used by the process application to establish point-to-point communication with each device.
- ▶ **MSTP:** MSTP is enabled by default in the EVAL-ADIN6310EBZ.xml example.
- ▶ **PTP Instance configuration:** Pass the various parameters for the PTP configuration per port. The XML file includes the default startup values.

SOFTWARE INSTALLATION

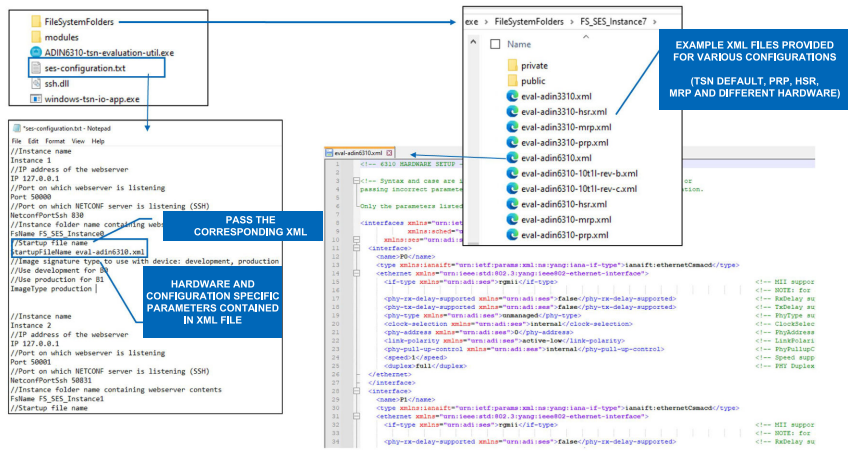


Figure 17. Ses-configuration.txt File Overview and Location of Hardware Configuration XML



Figure 18. EVAL-ADIN6310EBZ: Per-Port Specific Configuration

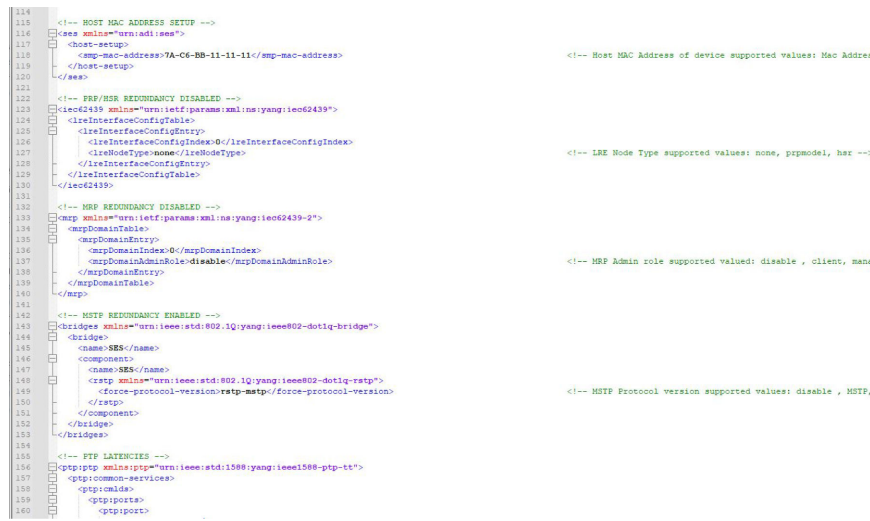


Figure 19. Device MAC Address, Redundancy Configuration, MSTP, PTP/PHY Latency

SOFTWARE INSTALLATION

```

121 <!-- HSR REDUNDANCY MODE -->
122 <iec62439 xmlns="urn:ietf:params:xml:ns:yang:iec62439">
123 <interfaceConfigTable>
124 <interfaceConfigEntry>
125 <interfaceConfigIndex>0</interfaceConfigIndex>
126 <lreNodeType>none</lreNodeType>
127 <lreSwitchingEndNode>hsrnode</lreSwitchingEndNode>
128 <lreDupListResideMaxTime>625</lreDupListResideMaxTime>
129 <lreRedundancyDevice xmlns="urn:adi:ses">damp</lreRedundancyDevice>
130 </interfaceConfigEntry>
131 </interfaceConfigTable>
132 <port-assignment xmlns="urn:adi:ses">
133 <lrePortA xmlns="urn:adi:ses">P2</lrePortA>
134 <lrePortB xmlns="urn:adi:ses">P1</lrePortB>
135 <lrePortC xmlns="urn:adi:ses">P2</lrePortC>
136 <lrePortX xmlns="urn:adi:ses">P1</lrePortX>
137 <redboxInterlinkPortC1 xmlns="urn:adi:ses">none</redboxInterlinkPortC1>
138 <redboxInterlinkPortC2 xmlns="urn:adi:ses">none</redboxInterlinkPortC2>
139 <redboxInterlinkPortC3 xmlns="urn:adi:ses">none</redboxInterlinkPortC3>
140 <redboxInterlinkPortC4 xmlns="urn:adi:ses">none</redboxInterlinkPortC4>
141 </port-assignment>
142 </iec62439>

```

Figure 20. EVAL-ADIN6310EBZ-HSR: HSR Specific Configuration

HSR Specific Configuration

Figure 20 shows an example of `eval-adin6310-hsr.xml` file where HSR is enabled. All TSN functionality is disabled when PRP/HSR is enabled, therefore the **TSN Switch Evaluation** web server only exposes HSR functionality and any TSN related functionality is hidden. The configuration specific parameters for HSR functionality are:

- ▶ **lreNodeType**: LRE node type supported: **none** (redundancy disabled), **prpnode1** for PRP operation or **hsr** to configure the device for HSR mode. Pass the relevant parameter to this field.
- ▶ **lreSwitchingEndNode**: Defines the type of functionality, use **hsrnode** for a DANH or **hsrredboxsan** for HSR redbox.
- ▶ **lreDupListResideMaxTime**: Duplicate list reside max time in second fraction units.
- ▶ **lreMacAddress**: MAC address of the LRE device, this must be the MAC address of the Host interface.
- ▶ **lrePortX**: Pass which ports are A, B ports.
- ▶ **lreDanPortC**: Pass which port is used as Port C. If using SPI Host interface, pass none to this parameter.
- ▶ **RedboxInterlinkPortCx**: For redbox configurations, identify which ports are interlink ports.

```

<!-- LRE Node Type supported values: none, prpnode1, hsr -->
<!-- PRP/HSR LRE Switching End Node Supported values: prpnode, hsrnode, hsrredboxsan -->
<!-- LRE Duplicate List Reside Max Time supported values: 0 to 2554 (corresponding to 16 Pp)
<!-- LRE MAC Address of device supported values: Mac Address format -->

<!-- LRE Port A supported values: none, P0, P1, P2, P3, P4, P5 -->
<!-- LRE Port B supported values: none, P0, P1, P2, P3, P4, P5 -->
<!-- LRE Port C supported values: none, P0, P1, P2, P3, P4, P5 -->
<!-- LRE Redbox Interlink Port C1 supported values: none, P0, P1, P2, P3, P4, P5 -->
<!-- LRE Redbox Interlink Port C2 supported values: none, P0, P1, P2, P3, P4, P5 -->
<!-- LRE Redbox Interlink Port C3 supported values: none, P0, P1, P2, P3, P4, P5 -->
<!-- LRE Redbox Interlink Port C4 supported values: none, P0, P1, P2, P3, P4, P5 -->

```

- ▶ **lreNodeType**: LRE node type supported: **none** (redundancy disabled), **prpnode1** for PRP operation.
- ▶ **lreSwitchingEndNode**: Defines the type of functionality, use **prpnode**.
- ▶ **lreDupListResideMaxTime**: Duplicate list reside max time in second fraction units.
- ▶ **lreMacAddress**: MAC address of the LRE device, this must be the MAC address of the Host interface.
- ▶ **lrePortX**: Pass which ports are A, B ports.
- ▶ **lreDanPortC**: Pass which port is used as Port C. If using SPI Host interface, pass none to this parameter.
- ▶ **RedboxInterlinkPortCx**: For redbox configurations, identify which ports are interlink ports.

MRP Specific Configuration

MRP can be enabled up front or alternatively use the default `eval-adin6310.xml` configuration and enable the function through the MRP web server page. Other redundancy features such as PRP or HSR must be disabled if MRP is enabled.

Figure 22 shows an example of `eval-adin6310-mrp.xml` file where MRP is enabled. TSN functionality is supported with MRP, so the full web server configuration is exposed.

PRP Specific Configuration

Figure 21 shows an example of `eval-adin6310-prp.xml` file where PRP is enabled. All TSN functionality is disabled when PRP is enabled, therefore the PC-based web server only exposes the PRP related functionality and all TSN related functionality is hidden.

```

<!-- PRP REDUNDANCY MODE -->
<iec62439 xmlns="urn:ietf:params:xml:ns:yang:iec62439">
<interfaceConfigTable>
<interfaceConfigEntry>
<interfaceConfigIndex>0</interfaceConfigIndex>
<lreNodeType>prpnode1</lreNodeType>
<lreSwitchingEndNode>prpnode</lreSwitchingEndNode>
<lreDupListResideMaxTime>625</lreDupListResideMaxTime>
<lreRedundancyDevice xmlns="urn:adi:ses">damp</lreRedundancyDevice>
</interfaceConfigEntry>
</interfaceConfigTable>
<port-assignment xmlns="urn:adi:ses">
<lrePortA xmlns="urn:adi:ses">P1</lrePortA>
<lrePortB xmlns="urn:adi:ses">P2</lrePortB>
<lrePortC xmlns="urn:adi:ses">P0</lrePortC>
<redboxInterlinkPortC1 xmlns="urn:adi:ses">none</redboxInterlinkPortC1>
<redboxInterlinkPortC2 xmlns="urn:adi:ses">none</redboxInterlinkPortC2>
<redboxInterlinkPortC3 xmlns="urn:adi:ses">none</redboxInterlinkPortC3>
<redboxInterlinkPortC4 xmlns="urn:adi:ses">none</redboxInterlinkPortC4>
</port-assignment>
</iec62439>

```

Figure 21. PRP Configuration

- ▶ **Domain ID**: Unique domain ID for the MRP ring.
- ▶ **MRP OUI**: MRP OUI, defaults to 0x080006 (Siemens OUI).
- ▶ **Domain Name**: Domain name for the ring.
- ▶ **MRP Role**: Choice of client (default), manager or auto-manager.
- ▶ **Ring Ports 1, 2**: Default Port 1 and Port 2, choice of any port.
- ▶ **Domain VLANID**: Defaults to untagged/4095.
- ▶ **React on Link Change**: For faster recovery, use react on link change enabled for which the manager does not wait for test frames to timeout, instead, reacts on the link change frames.
- ▶ **Recovery rate**: Recovery profile choice of 30 ms, 200 ms, and 500 ms.
- ▶ **MRP Port Tx Priority**: Default Queue 7 is highest priority. PTP traffic also egresses in Queue 7. If using lowest recovery profile, change default PTP queue from 7 to a lower priority in the **Time Synchronization IEEE 802.1AS** page.

The configuration specific parameters for PRP functionality are:

SOFTWARE EXECUTION

Start the application by double-clicking the **ADIN6310-tsn-evaluation-util.exe** executable. The GUI application window appears, as shown in [Figure 24](#).

1. The GUI automatically detects the available network adapters. Select the adapter that is connected to the ADIN6310 board Host (Port 0) by double-clicking the description line for that adapter. Once the adapter is selected, the device configuration information pulled from the **ses-configuration.txt** and XML files load and populate the lower window.
2. The GUI supports configuration of a chain of up to 10 devices from one network adapter, alternatively, the user can control multiple switches from individual network adapters. Select the chosen adapter in the upper window and then assign it to the switch **Ethernet Adapter** field by clicking in the GUI in that area. The GUI also supports selection of different startup files through the dropdown in the **Startup File Name** column. Click **Find and Configure SES Devices** button to start searching for connected Switch boards.
3. The GUI searches for and configures the MAC address for any ADIN6310 device it finds. Each Switch powers up with the same default MAC address (7a:c6:bb:ff:fe:00). The first thing the GUI application does during configuration is to assign a primary MAC address (based on XML configuration). If observing the traffic from Host to Switch using Wireshark, initially messages are sent from the PC to the default multicast address (79:c6:bb:ff:fe:00) and responses come from default MAC address 7a:c6:bb:ff:fe:00 until the primary MAC address gets assigned. An LED turns green for each board found. Clicking on the LED for each connected ADIN6310 device launches a browser for each board as shown in [Figure 24](#). Once the web server is launched, the LED color changes to orange. Keep the PC application open, it needs to stay running while interacting with the web server. The GUI application continues to search for more ADIN6310 devices, so if all connected devices have been identified, stop the application searching by clicking the **Find and configure** button again. The find LED then stops flashing.
4. If boards are power cycled or reset button is pressed, the device reverts to the default MAC address and if the GUI application is searching, it sees them as new devices (additional LED lights go green). To avoid this, close the older processes associated with those instances of boards on the keyboard, use **Ctrl** and **Close All Running Processes**, as shown in point 4 of [Figure 24](#).

Note that the first time the application launches the web page, a user may receive a security warning regarding Windows firewall settings. Ensure that the firewall settings are configured to allow communications to pass through the firewall.

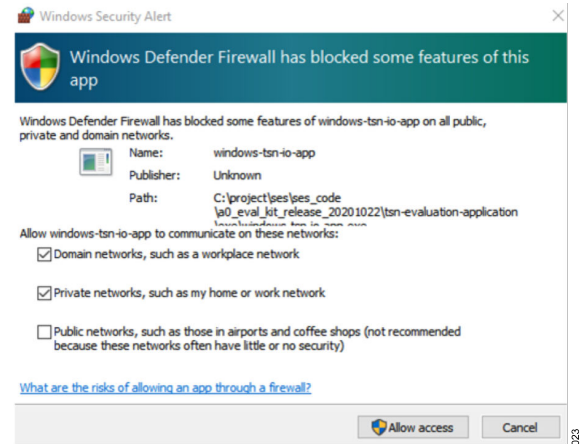


Figure 23. Firewall Security Pop-Up

SOFTWARE EXECUTION

The application keeps searching for devices. In the event it does not discover more switches, there may be a pop-up window indicating that it found not other devices, like shown in Figure 25. Click

OK and proceed to launch the web server for the devices the application discovered.



Figure 24. Starting the GUI Application




Figure 25. GUI Pop-Up When the Application Does Not Find a Switch

TSN SWITCH EVALUATION WEB PAGE OVERVIEW

The **TSN Switch Evaluation** software package contains a set of web pages to configure the Switch for use in a TSN network or with redundancy features (see [Figure 26](#)).

A separate instance of the web server is used for each evaluation board connected to the PC and identified by the GUI.

The **TSN Switch Evaluation – Home** page provides access to the following web pages:

- ▶ **Setup:** Which allows the user to perform global actions, such as loading, storing, and managing the overall TSN database.
- ▶ **Port Statistics:** Provides an overview of each port transmit and receive information and any errors observed.
- ▶ **Port Configuration:** Provides user ability to control the port configuration and change port speed, interface type (hardware must support). It is not supported to change MAC interface modes during run-time. MAC interface must be configured during initial configuration. User can also communicate directly with the Ethernet PHYs via an MDIO read/write from this page.
- ▶ **GPIO and Timer Configuration:** Configure the functionality of the GPIO and Timer pins.
- ▶ **IGMP Snooping Configuration:** Provides user ability to enable and configure timeouts for IGMP snooping in the Switch.
- ▶ **Switching Table:** Provides user ability to install static entries in the lookup table, install extended table entries and flush the dynamic table. Status view provides insight into the learned dynamic entries. Per stream filtering and policing filters can be linked with static and extended entries in the Switching table.
- ▶ **Stream Table:** Provides user ability to install entries into the Stream table.
- ▶ **VLAN Table:** Provides user ability to configure the port behavior for virtual LAN (VLAN) IDs. Choice of standard VLAN configuration or configuring ports as Trunk or Access ports.

- ▶ **VLAN Remapping:** Provides the ability to remap VLAN IDs for each port.
- ▶ **VLAN Re prioritization:** Gives user ability to configure remapping of VLAN priority on a port basis.
- ▶ **Time Synchronization:** Provides ability to configure and observe status of time synchronization (IEEE802.1AS).
- ▶ **Frame Preemption:** Provides ability to configure frame preemption on each port and observe preemption statistics.
- ▶ **Scheduled Traffic – Assign Queue:** Provides user the ability to configure the mapping of VLAN priorities to the available queues for each port.
- ▶ **Scheduled Traffic – Set Queue Max. SDU:** Provides ability to configure the maximum SDU transmission size for each port and each queue.
- ▶ **Scheduled Traffic – Schedule:** Provides ability to set up schedules per port and also configure a schedule for the hardware Timer pins.
- ▶ **LLDP Configuration:** Provides LLDP configuration.
- ▶ **PSFP Configuration:** Provides ability to configure Per-Stream filtering and policing, Qci.
- ▶ **MSTP Configuration:** Provides ability to configure multiple spanning tree protocol.
- ▶ **FRER Configuration:** Frame replication and elimination for Reliability, 802.1CB, configuration page and stream entry for FRER streams.
- ▶ **Firmware Update:** Provides ability to update/check version of device firmware.

Click any of these links to go to the required page. Once in a page, use the menu on the left to navigate to any of the other pages at any time. Ensure the GUI is kept running while navigating the web pages.

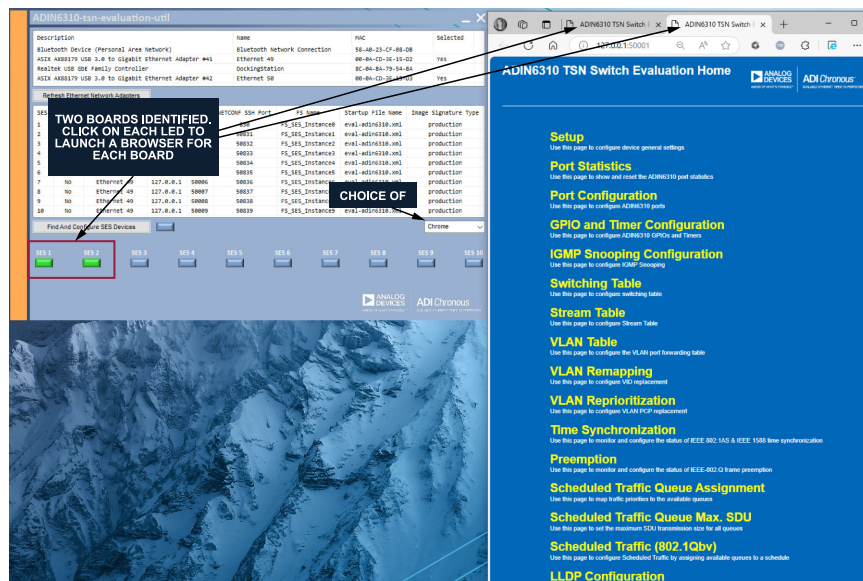


Figure 26. TSN Switch Evaluation – Home Page

TSN SWITCH EVALUATION WEB PAGE OVERVIEW

CANDIDATE/RUNNING/STARTUP PAGES

All configuration pages have **Candidate/Running/Startup** views and are linked to the sysrepo repository. To tune the way a function performs, users can change several parameters in the **Candidate** pages. Once the user has a new set of values for the candidate configuration, click **Save** followed by **Commit** to send the candidate configuration entries to the **Running** configuration. Click **Discard** to revert the candidate configuration back to current running configuration. The **Startup** page shows the current startup configuration. This may be the default startup configuration or user may have saved a previous configuration to **Startup**.

SETUP PAGE

This page is used to perform global operations on the **Candidate**, **Running**, and Startup configurations. [Figure 27](#) shows these three configurations and which commands act on each configuration from the **Setup** page. Click the following command labels to perform the following actions:

SAVE AND LOAD CANDIDATE DATASTORE

- ▶ **Save Candidate as:** Save Candidate in JSON or XML format. The file gets saved to **Downloads** folder.
- ▶ **Load Candidate from file:** Select JSON or XML file to load.

DATASTORE MANAGEMENT

- ▶ **Save current Running as Startup:** To store the running configuration to the startup configuration.
- ▶ **Commit All:** To push saved configuration to the device.
- ▶ **Discard All:** To discard configuration and revert to startup.

ADVANCED

- ▶ **Save Status as JSON:** The operational file gets saved in JSON format to **Downloads** folder.

- ▶ **Restore default values:** Revert to default.
- ▶ **Hardware Reset:** Provides ability to do a reset of the ADIN6310 over the Ethernet Port. This also resets all the **ADIN1300** PHYs (except for the Host Port PHY on Port 0). When this reset is used, this requires that any previous application processes running on the PC (running the web server) need to be closed. To close the process instances, press the keyboard **Ctrl** key, click **Close all running processes**. Release the **Ctrl** key, click **Find and Configure SES Devices** to resume operation, as shown in [Performing a Reset](#) section.

Port 0 to Port 5 Status: The LEDs on the left of page visually show which ports have established a link, these LEDs do not update automatically and require a refresh of the page.




Figure 27. TSN Switch Evaluation – Setup Page

SETUP PAGE

PERFORMING A RESET

After performing a reset, either using the RESET push button on the Evaluation kit or alternatively through the **Hardware Reset** button in the **Setup** page, the Switch reverts to its power on reset configuration and the device MAC address reverts to default, therefore if the GUI is searching for devices, it likely finds it as a new device, not one of the previously found devices. Either reset the GUI or do the following steps (see [Figure 28](#)):

1. To reestablish communication with the device, return to the GUI. Using the keyboard **Ctrl** button, click **Close All Running Processes**. All LEDs should turn off on the GUI.
2. Click **Find And Configure SES Devices** to identify and connected boards (shown with green LEDs) the devices again.

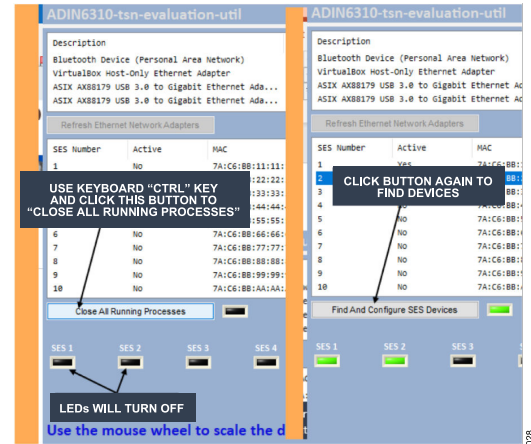


Figure 28. After a Reset – Close Processes and Find Again

PORT STATISTICS

The **Port Statistics** page provides visibility into the various transmit and receive statistics reported for each port. As shown in [Figure 29](#), click **Port Statistics** in the menu item on the **Home** page or in the menu on the left of the page to access the **Port Statistics** page. The upper table on this page shows traffic counts for packets that have been transmitted and received on each port and provides insight into any errors observed during reception or transmission. The Clear buttons enable clearing of individual ports or all port statistics.

The second table shows the extended statistics providing insight into the per port, per queue drop counters, and information about the buffer utilization and the interrupts (IRQ) available. Note that once an IRQ statistic is set, it remains set until actively cleared (clear button on the web page). The components of the second table are as follows:

- ▶ **Queue 0 Drop Count to Queue 7 Drop Count:** These drop counters show the number of dropped packets for transmit on Queue 0 through Queue 7.

- ▶ **Highest Number of Buffers Used:** Number of the highest buffer reached since reset. The value ranges from 0 to 219.
- ▶ **Current Number of Buffers in Use:** Indicates how many buffers are currently being used. The value ranges from 0 to 219.
- ▶ **Buffer Limit IRQ Status:** Set to 1 if the buffer limit of 219 is reached.
- ▶ **Buffer Parity Error IRQ Status:** Set to 1 if a buffer parity error occurred.
- ▶ **Buffer Allocate IRQ Status:** Set to 1 if a buffer cannot be allocated
- ▶ **Buffer Return Error IRQ Status:** Set to 1 if a buffer cannot be returned.

The **Port Statistics** page updates automatically on a refresh rate of 5 seconds. To update on demand, refresh the page in the browser.

There is a **Download as CSV** option on the bottom right of each table on this page, where the current snapshot of statistics can be saved to an excel file.

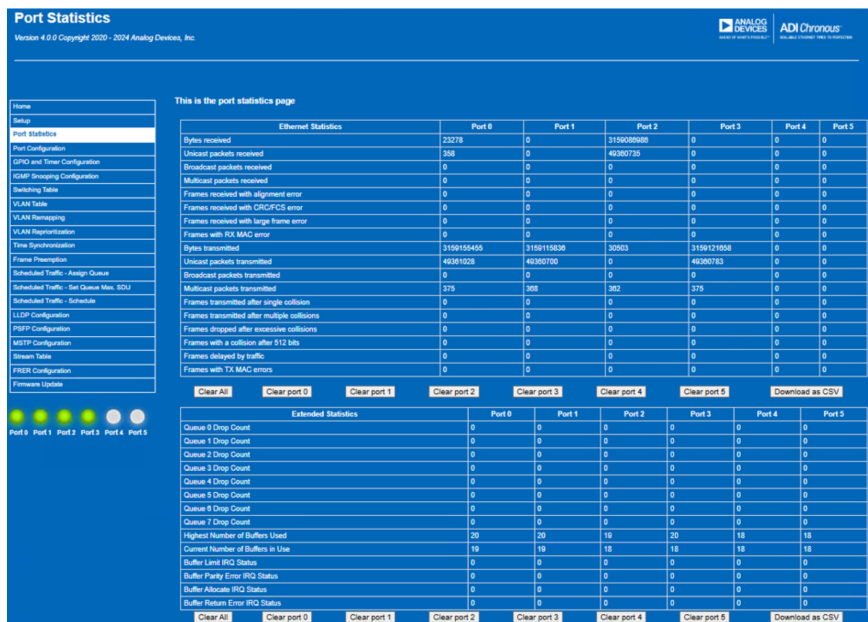


Figure 29. Port Statistics Page

PORT CONFIGURATION

CANDIDATE PAGE

The **Port Configuration** page provides the user ability to configure port specific parameters in addition to seeing the current status of the ports.

As shown in [Figure 26](#) or [Figure 27](#), in the menu item on the **Home** page or in the menu on the left of the page, click **Port Configuration**. Similar to other pages, there are **Status**, **Candidate**, **Running**, and **Startup** views for this page.

The **Candidate** page provides user ability to configure some parameters for the port operation. Note that the XML file in the package is the primary opportunity for port configuration, but some additional run-time configuration is possible within the **Candidate** page.

Each port can be configured independently and saved, or alternatively, there is a **Save** button at top of page. Click the **Commit** button to push any changes to the device.

As shown in [Figure 30](#), the configuration provided here as follows:

- ▶ **Enable Port:** This check box allows user to enable or disable ports. By default, all ports are enabled.
- ▶ **MAC Address:** The default MAC addresses shown corresponds to the MAC addresses assigned to each port based on the

primary MAC address set by the XML configuration file. Changes to this field are supported within the web page, enter the required MAC address and click **Save** button.

- ▶ **PHY Type:** This shows what is provided in the XML configuration file.
- ▶ **PHY Auto-Negotiation:** This check box is enabled by default and if disabled, indicates that the PHY is in Forced Speed mode, therefore only speeds 10 Mbps/100 Mbps are available.
- ▶ **Speed:** For Auto-Negotiation enabled options of 10 Mbps/100 Mbps/1000 Mbps, 10 Mbps/100 Mbps or 10 Mbps. When Auto-Negotiation disabled, options of 10 Mbps or 100 Mbps only.
- ▶ **PHY Duplex:** Full duplex by default. PHY duplex can be configured for speeds of 10 Mbps or 100 Mbps.
- ▶ **PHY Crossover Config:** Enables user to decide the cable crossover configuration of the **ADIN1300** PHY on each port. Defaults to Auto MDIX. The user can select the following options:
 - ▶ Auto
 - ▶ MDI
 - ▶ MDIX
- ▶ **RGMII Strength:** Configuration of the drive strength of the RGMII from the Switch side.

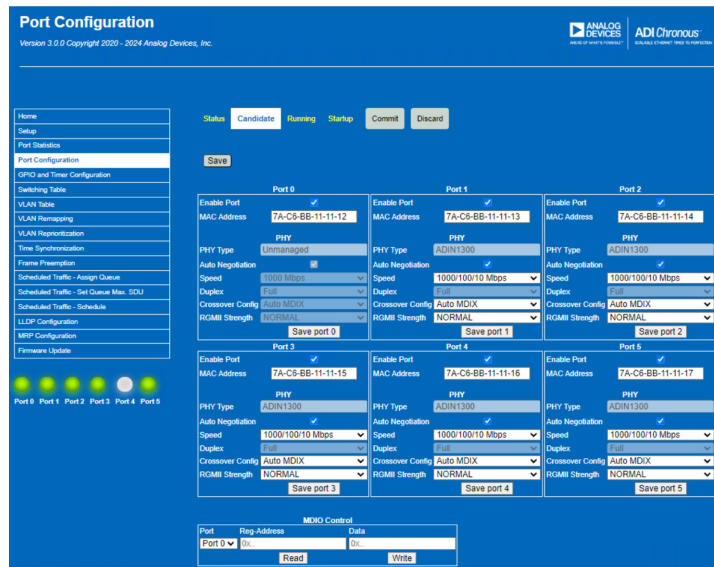


Figure 30. Port Configuration Page Overview – Candidate View

PORT CONFIGURATION

STATUS PAGE

The **Status** page provides user snapshot of the current port configuration status.

As shown in [Figure 31](#), the configuration provided here as follows:

- ▶ **Interface Type:** Shows the MAC interface as configured by the XML file. When using EVAL-ADIN6310EBZ, this hardware supports RGMII on all ports and SGMII interfaces on Port 1 to Port 4. SGMII modes need to be configured during initial configuration by editing the XML configuration. This hardware does not have any PHYs connected via RMII, therefore no RMII connectivity is possible.

- ▶ **MAC Address:** Shows the assigned MAC address to the port.
- ▶ **PHY Type:** Shows what PHY is connected.
- ▶ **Crossover:** Shows the actual crossover configuration.
- ▶ **Link:** Shows whether the link is up or down.
- ▶ **Speed (Mbps):** Shows the speed of the established link.
- ▶ **PHY Delay:** Shows the PHY Tx delays (of ADIN1300 PHY), which depends on the speed of the link established.
- ▶ **RGMII Strength:** Shows the configured drive strength of RGMII from the Switch side.

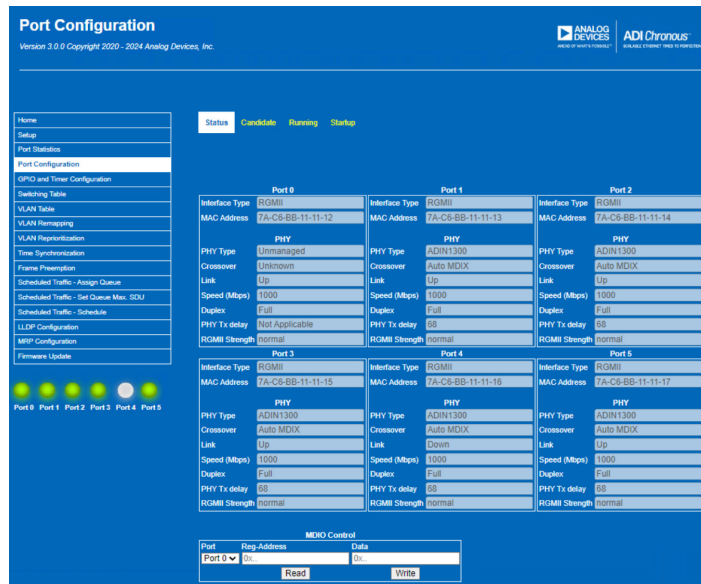


Figure 31. Port Configuration Page Overview – Status View

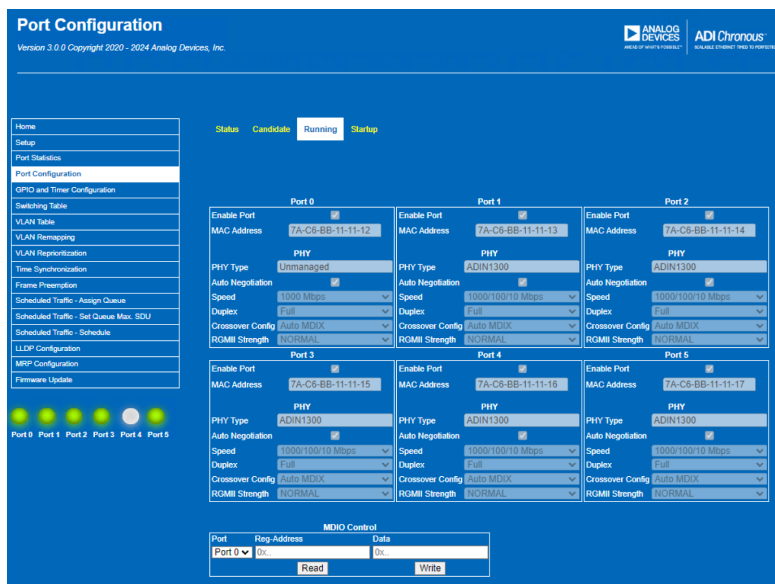


Figure 32. Port Configuration Page Overview – Running View

PORT CONFIGURATION

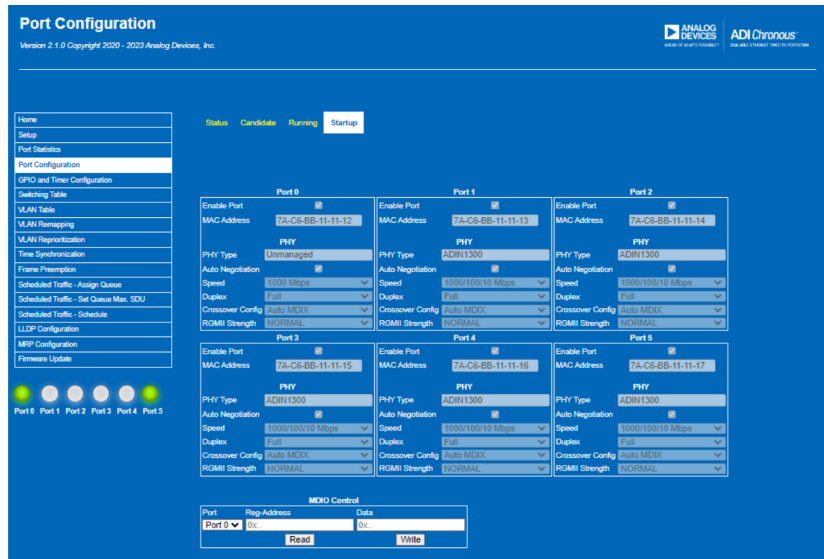


Figure 33. Port Configuration Page Overview – Startup View

MDIO CONTROL

The MDIO Control field is shown at the bottom of the Port Configuration page and provides user ability to interrogate any of the six ADIN1300 PHYs on the evaluation board.

Clause 22 read/writes are supported to the standard IEEE802.3 registers and vendor specific registers up to 0x1F. As shown in Figure 34, to read a register, in the Port field, select the port, in the Reg-Address field, enter the register address, and then click the Read button. The Switch communicates over MDIO bus to the appropriate PHY and the data field appears with the register information returned.

Similarly to write a PHY register, in the Port field, select the port, in the Reg-Address field, enter the register address, and then click the Write button to load.

Access to Clause 45 or Extended registers is supported. Register address input format is 0xHEX.




Figure 35. MDIO Control – Access of Extended Register Space




Figure 34. MDIO Control – Communication with the PHYs on the Evaluation Board

GPIO AND TIMER CONFIGURATION

There are four GPIO pins and four Timer pins. This page provides user ability to control the function of these hardware pins. There are **Status**, **Candidate**, **Running**, and **Startup** pages for this functionality.

All pins are enabled by default. The GPIOs are enabled as Outputs. Timer0 is enabled as a GPIO by default, Timer1 is enabled for TSN timer function, Timer2 is enabled as a 1 pulse per second (1PPS) timer signal, and Timer3 is configured to be a Capture Input.

The available configurations and default configuration for these pins is shown in [Table 9](#).

When changing GPIO or Timer operation, each change must be saved individually, otherwise, the user loses the change.

When SPI mode is selected as Host interface, Timer0 automatically configures as an Interrupt for the SPI interface to the Host and does not available to configure as a Timer/GPIO pin.

TSN OUTPUT TIMER

This is the default operation for Timer1. When TSN Output Timer function is selected in this page, then a user needs to navigate to the [Scheduled Traffic – Schedule](#) page. The TSN Output Timer functionality allows the user to control the Timer pins with specific cycle times and is configured through the [Scheduled Traffic – Schedule](#) page.

1PPS PERIODIC OUTPUT

Timer2 and Timer3 can support a 1 pulse per second (1PPS) output. As shown in [Figure 36](#), in the **Mode** drop-down box, select the **1PPS_PERIODIC_OUT** option. The low/high pulse-width fields fix at 500 ms.

Table 9. GPIO and Timer Pin Functionality

Hardware Pin	Available Mode
GPIO0	GPIO
GPIO1	GPIO
GPIO2	GPIO
GPIO3	GPIO
GPIO4/TIMER0	GPIO, TSN Output Timer (Default), Interrupt (SPI INT)
GPIO5/TIMER1	GPIO, TSN Output Timer (Default)
GPIO6/TIMER2	GPIO, TSN Output Timer, Periodic Output, 1PPS Output (Default)
GPIO7/TIMER3	GPIO, TSN Output Timer, Periodic Output, 1PPS Output, Capture In (Default)

PERIODIC OUTPUT

Timer2 and Timer3 also support a user-configurable periodic output. As shown in [Figure 36](#), in the **Mode** drop-down box, select the **PERIODIC_OUT** option and enter the required high/low pulse-width for required pulse. The minimum value of high/low pulse-width is 16 ns and the time period must not exceed 1 second.

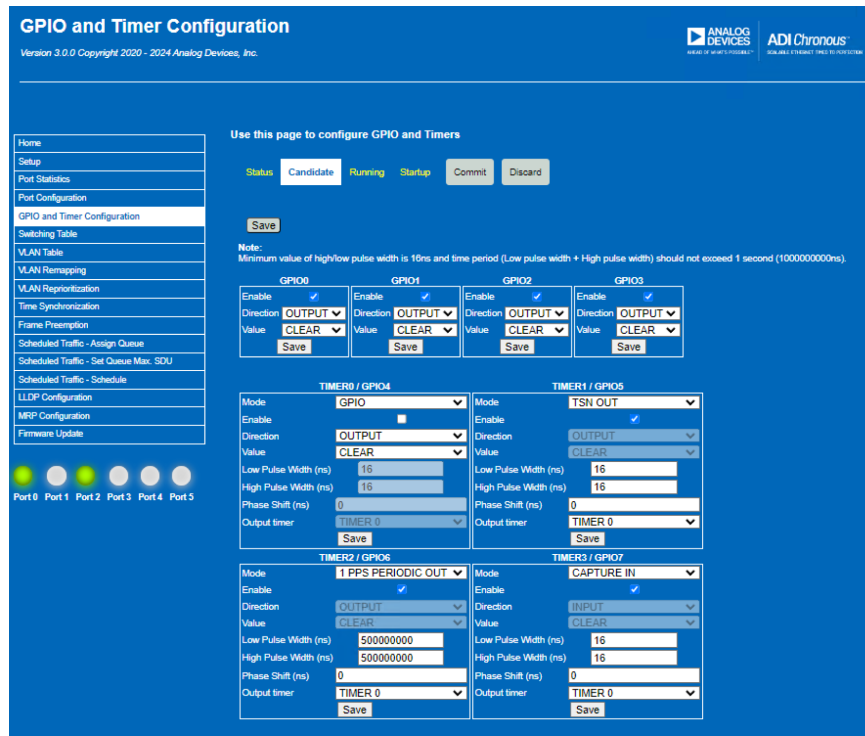
CAPTURE INPUT

Timer2 and Timer3 can also support configuration as a capture Input. By default, Timer3 is a capture input. A possible usage for the capture input is to trigger the Switch to capture a hardware timestamp in response to a transition on the Timer3 and send that timestamp information to the Host. Note the web server does not support this configuration, the driver APIs need to be used to enable this and send the message to the Host.

OTHER MODES

Any greyed out options are not available yet and intended for future releases.

GPIO AND TIMER CONFIGURATION



GPIO and Timer Configuration
Version 3.0.0 Copyright 2020 - 2024 Analog Devices, Inc.

ANALOG DEVICES
ADI Chronous

Use this page to configure GPIO and Timers

Status: **Candidate** Running Startup Commit Discard

Save

Note: Minimum value of high/low pulse width is 16ns and time period (Low pulse width + High pulse width) should not exceed 1 second (1000000000ns).

GPIO0	GPIO1	GPIO2	GPIO3
Enable: <input checked="" type="checkbox"/>	Enable: <input checked="" type="checkbox"/>	Enable: <input checked="" type="checkbox"/>	Enable: <input checked="" type="checkbox"/>
Direction: OUTPUT	Direction: OUTPUT	Direction: OUTPUT	Direction: OUTPUT
Value: CLEAR	Value: CLEAR	Value: CLEAR	Value: CLEAR
Save	Save	Save	Save

TIMER0 / GPIO4	TIMER1 / GPIO5
Mode: GPIO	Mode: TSN OUT
Enable: <input type="checkbox"/>	Enable: <input checked="" type="checkbox"/>
Direction: OUTPUT	Direction: OUTPUT
Value: CLEAR	Value: CLEAR
Low Pulse Width (ns): 16	Low Pulse Width (ns): 16
High Pulse Width (ns): 16	High Pulse Width (ns): 16
Phase Shift (ns): 0	Phase Shift (ns): 0
Output timer: TIMER 0	Output timer: TIMER 0
Save	Save

TIMER2 / GPIO6	TIMER3 / GPIO7
Mode: 1 PPS PERIODIC OUT	Mode: CAPTURE IN
Enable: <input checked="" type="checkbox"/>	Enable: <input checked="" type="checkbox"/>
Direction: OUTPUT	Direction: INPUT
Value: CLEAR	Value: CLEAR
Low Pulse Width (ns): 500000000	Low Pulse Width (ns): 16
High Pulse Width (ns): 500000000	High Pulse Width (ns): 16
Phase Shift (ns): 0	Phase Shift (ns): 0
Output timer: TIMER 0	Output timer: TIMER 0
Save	Save

Port 0 Port 1 Port 2 Port 3 Port 4 Port 5

Figure 36. GPIO and Timer – Candidate Page

SWITCHING TABLE

CANDIDATE VIEW

The **Switching Table** page provides user ability to install entries into the switch forwarding table. Static entries can be installed in addition to extended entries. The learned dynamic entries can also be viewed through the **Status** page.

Dynamic Table

Entries in the **Dynamic Table** are entries learned by the Switch based on traffic crossing the Switch. The Switch learns based on Source MAC address and if the VLAN configuration is enabled for learning, the Switch automatically installs an entry in the table with an age value based on when the entry is updated. The table ages out frames if they are no longer seen within the configured aging period. The default configuration is for learn and forwarding on untagged traffic. VLAN tagged traffic is not learned or forwarded unless user configures the VLAN table accordingly, see [VLAN Table](#). The **Switching Table** page provides the user ability to configure the aging period of the **Dynamic Table** entries, simply enter the aging period in ms in the field and click the **Save** button to adjust the aging (range of 1000 ms to 10000000 ms). The default setting for aging is 300 seconds. New age out times will be applied to subsequent learned entries, existing entries will retain the aging time from when they were first learned.

As shown in [Figure 37](#), a user can flush the Dynamic table on-demand, by clicking the **Flush Dynamic Table** button.

Source Port Lookup Modes

The default behavior on all ports is to perform a destination MAC and VLAN lookup.

User can configure the lookup behavior on a port basis to instruct the Switch to perform other lookup options. Checking the bit 0 field for a port enables a Source lookup on all traffic to ingress that port. Setting bit 1 enables extended lookup on all frames for that port and setting bit 2 enables a Destination MAC address lookup (802.1D). Combinations of lookups are supported.

Static Table Entries

The **Static Table** allows user to install/remove entries in the lookup table. When the Switch is configured for TSN mode, the startup configuration installs an entry in the table for LLDP multicast addresses. This static entry can be seen as the first row of the table. Do not interfere or overwrite this entry.

To install a new entry, first add a row, then fill in the **Destination MAC Address**, **VLAN identifier**, and **Egress Ports**. For untagged traffic use **4095** as a **VLAN identifier** to indicate no VLAN identified associated with entry. For tagged traffic, ensure to also configure the VLAN table to support the VLAN IDs of interest for specific

ports. The format of the **Destination MAC Address** must be entered as xx-xx-xx-xx-xx and the **Egress Port** must be entered in hex.

[Figure 37](#) shows examples of adding various entries with different VLAN tags destined to egress on specific ports.

The **Static Table** also gives user the ability to add or remove VLAN tags from traffic. To insert a tag, add the table entry with the **Add Tag Option** and define the **VLAN ID** and **Priority** to add. To remove the tag as the frame egresses, select the **Remove tag**. The standards indicate a minimum sized frame for a VLAN tagged frame is 68 bytes (64 bytes + 4 byte VLAN tag). If user is ingressing frames of 64-bytes including VLAN tag and configuring the Switch to remove the VLAN tag directly or using VLAN access port, the Switch deliberately corrupts the frame on egress.

By default, only untagged or VID 0 frames crosses the Switch, the VLAN table must be configured to forward other VIDs.

Extended Table Entries

Similarly, this page allows the user to install extended table entries and define how they are handled. A VLAN tag can be inserted or removed. Note that configuring the extended table to install a VLAN tag in traffic that has an existing VLAN tag results in two VLAN tags. This operation is a misconfiguration by user. Two VLAN tags are visible in the frame, upper layers need to handle accordingly.

The extended table input fields in the web server currently only support basic lookups up to 14-bytes. Installing lookups for EtherTypes such as IPv4, IPv6, and PTP are not yet supported and rejected by the web server. These type of entries are supported using the Driver APIs directly, for more details, refer to the [ADIN6310 Hardware Reference Manual](#).

Cut Through Enable

When installing a **Static Table** entry, user can install with cut through enabled/disabled by selecting the **Cut Through Enable** check box.

Stream Filter

When installing a **Static Table** entry or **Extended Table** entries, a **Stream Filter** can be associated with this entry. **Stream Filter** is a part of PSFP functionality. To use this feature, select the **Stream Filter Enable** check box and pass the ID of the **Stream Filter** to apply, then go to the PSFP web page to configure the **Stream Filter**, **Stream Gate**, or **Flow Meter** as required. Stream filters can only be applied to static entries that are configured for Store and forward mode, ensure that the **Cut Through Enable** check box is not selected when using PSFP.

SWITCHING TABLE

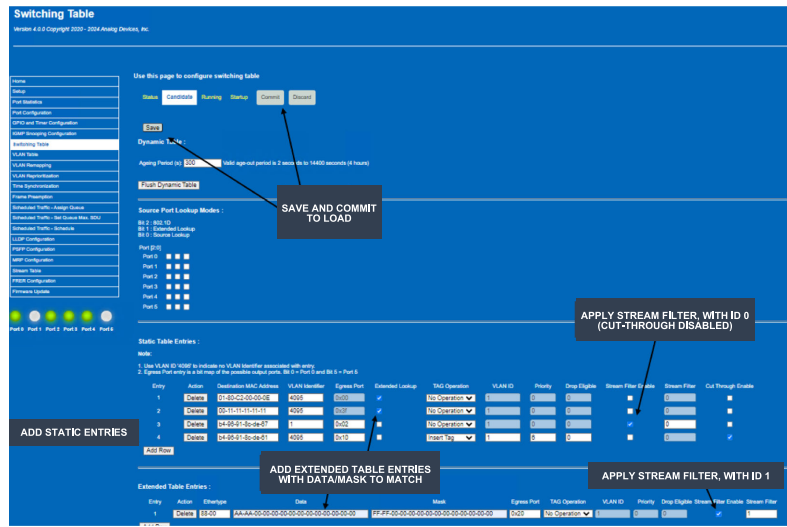


Figure 37. Switching Table – Candidate View – Adding Static Entries and Extended Table Entries

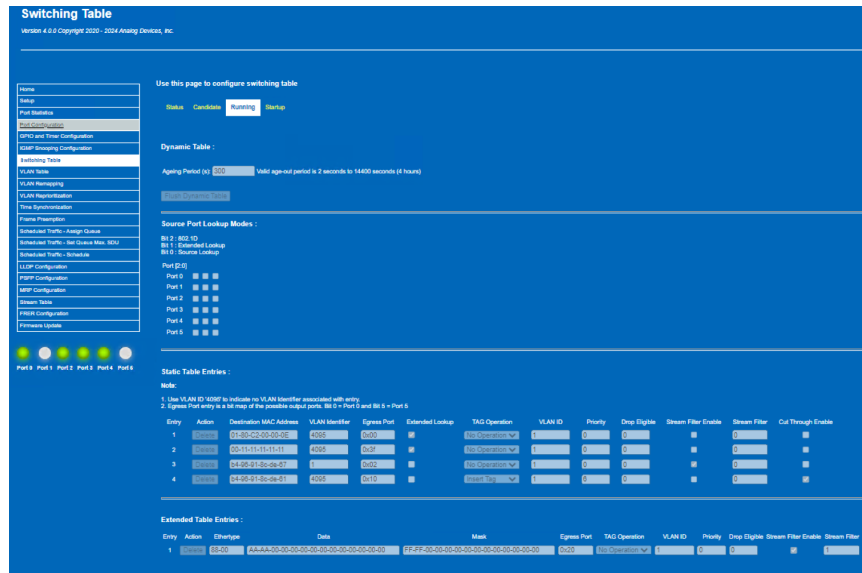


Figure 38. Switching Table – Running View with Added Static Entries

SWITCHING TABLE

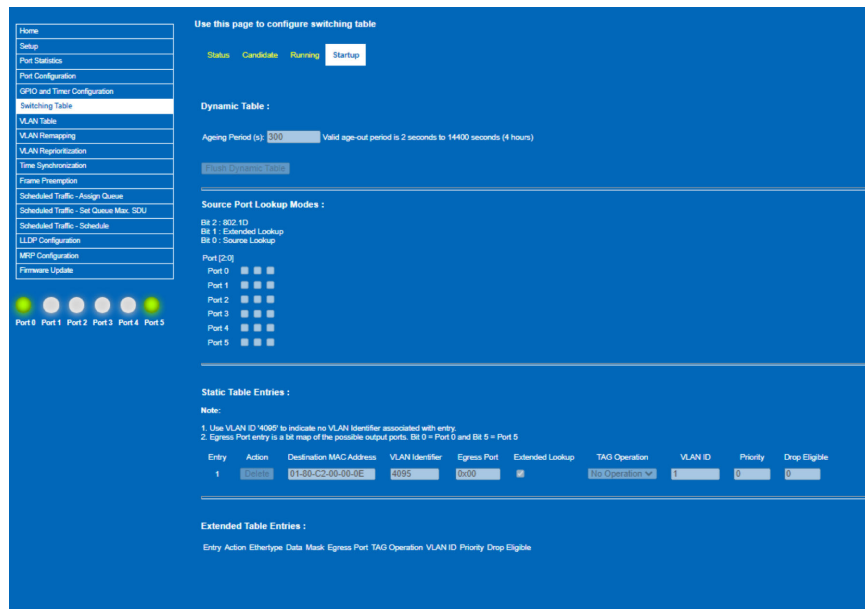


Figure 39. Switching Table – Startup View

SWITCHING TABLE

STATUS VIEW – DYNAMIC ENTRIES

The **Status** view allows user to readback the MAC addresses learned by the Switch. **Figure 40** shows the untagged traffic entries learned as the user is ingress Port 0.



Figure 40. Switching Table – Status View with Learned Entries

STREAM TABLE

The **Stream Table** page allows the user to place entries into the stream table. The switch stream table supports up to 16,000 entries and is partitioned into 16 × 1000 blocks. The web server can support display of up to 60 entries.

A stream block base destination address must first be created, saved, and committed before adding individual stream table entries

for that block. When deleting entries, the process must be done in reverse: delete all stream table entries of a block address first, followed by deleting that stream block base destination address. They are located in a separate area of the forwarding table to the static and dynamic table entries.

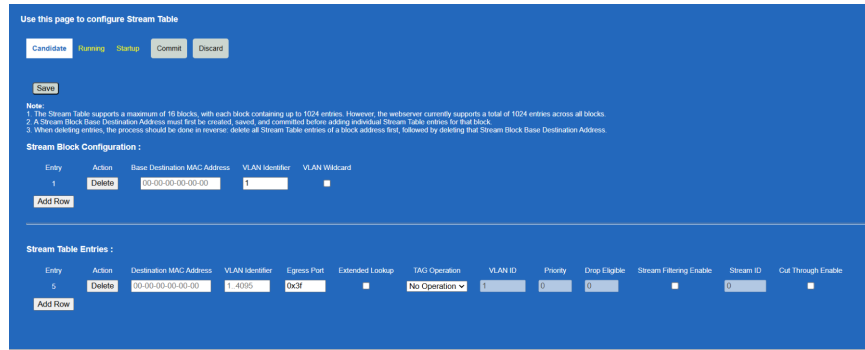


Figure 41. Stream Table Candidate View

The **Stream Table** page contains only **Candidate**, **Running**, and **Startup** views. The **Status** page does not provide any information and is blank.

VLAN CONTROL

As shown in [Figure 42](#), the **VLAN Table** has **Candidate**, **Running**, and **Startup** pages. There is no Status page for VLAN function. To see how the VLANs are configured based on changes in the **Candidate** view, see the **Running** page.

VLAN TABLE

The **VLAN Table** page provides user ability to configure the port learning and forwarding operational mode for each VLAN IDs (1 through 4094).

The default behavior is **No Learn and No Forward** for all VLANs with exception of VLAN ID 0/untagged traffic.

There are two modes of operation within the **VLAN Table**, **Candidate** page: **Trunk/Access Port configuration** or **VLAN Table Configuration**. The default is **VLAN Table**, but this can be changed using the **VLAN Table/Mode Table Switch** check box.

VLAN Table Configuration

To configure each individual port behavior, enter the **VLAN ID** and select the appropriate behavior for each port. Ensure the **Save** button is clicked following the port selection for each **VLAN ID**. Click the **Commit** button once all port selections are complete. The configuration is loaded and the web page moves automatically to show the **Running** view. To read the configuration for a specific **VLAN ID**, enter the ID of interest in the **Running** page.

The configuration mode choices for each port are: **Learn and Forward**, **Learn and No Forward**, **No Learn and Forward**, or **No Learn and No Forward**.

Trunk/Access Configuration

The Switch ports can be configured as Trunk or Access ports. Trunk ports can support multiple VLAN IDs or ranges of VLAN IDs, whereas access ports support only 1 VLAN ID. When a port is configured as a Trunk port, it transmits and receives traffic with VLAN tags in the configured range of VLAN IDs excluding untagged

traffic. When a port is configured as an Access port, only VLAN traffic matching the configured VLAN ID is transmitted and the tag is removed as the frame egresses the port. Untagged traffic ingressing the Access port has the configured VLAN tag inserted.

The Switch handles the insertion and removal of VLAN tags where required when traffic is crossing between ports. When removing a VLAN tag on an access port, the switch expects a minimum sized frame for a VLAN tagged frame to be 68 bytes (64 bytes + 4 byte VLAN tag). If user is ingressing frames of 64-bytes including VLAN tag and configuring the Switch to remove the VLAN tag directly or using VLAN access port, the Switch sees this frame as a runt frame and deliberately corrupts the frame on egress. This is documented in the data sheet as a silicon anomaly.

To use this feature, first enable the check box **VLAN Table/Mode Table Switch**.

Then configure the ports as Trunk or Access ports with the VLAN IDs or ranges of interest.

In the example shown in [Figure 43](#), Port 0 is configured as a **Trunk** port for VLAN IDs 1 to 5, but VLAN ID 2 is not disabled.

Port 1 to Port 4 are configured as **Access** ports for individual VLANs and Port 5 is another **Trunk** port subscribing to VLAN IDs in range of 1 to 5 including VLAN ID 2. VLAN Priority can also be configured for the access port, so any traffic ingressing the access port have a VLAN tag inserted with the VID and the priority configured for the access port.

There is an upper limit on the number of different VLAN IDs that can be active with a max of 62 different VLAN IDs, therefore, when configuring Trunk ports, user must avoid enabling the full range of VLAN IDs. In practice, for managed switch use cases only a small number of VLAN IDs are used.

The **Running** view in [Figure 44](#) shows the configured VIDs across ports. Note that VID 2 is not shown for Port 0, but is for Port 5, which matches what is configured.

VLAN CONTROL

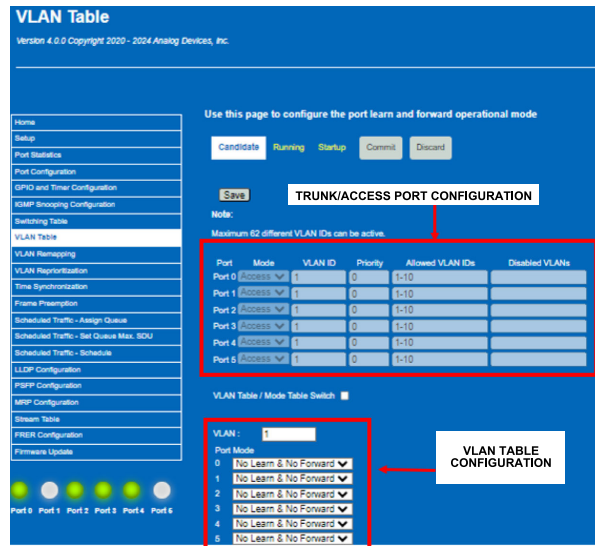


Figure 42. VLAN Table for Port Configuration

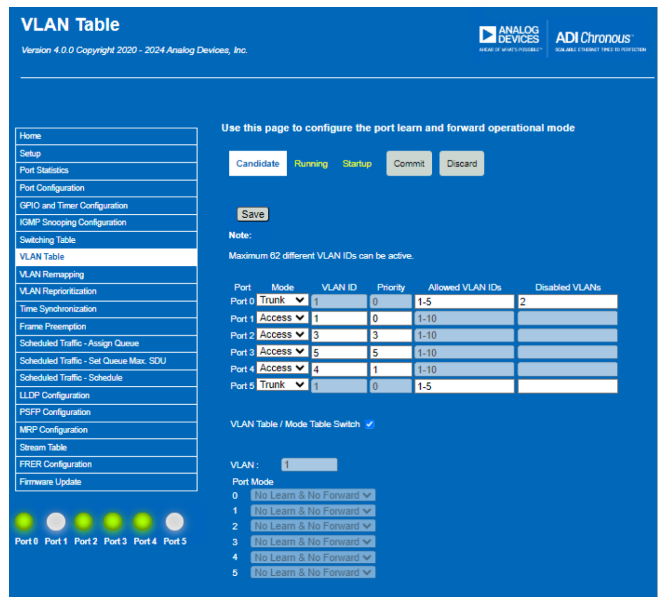


Figure 43. Using Trunk/Access Port Configuration – Candidate View

VLAN CONTROL

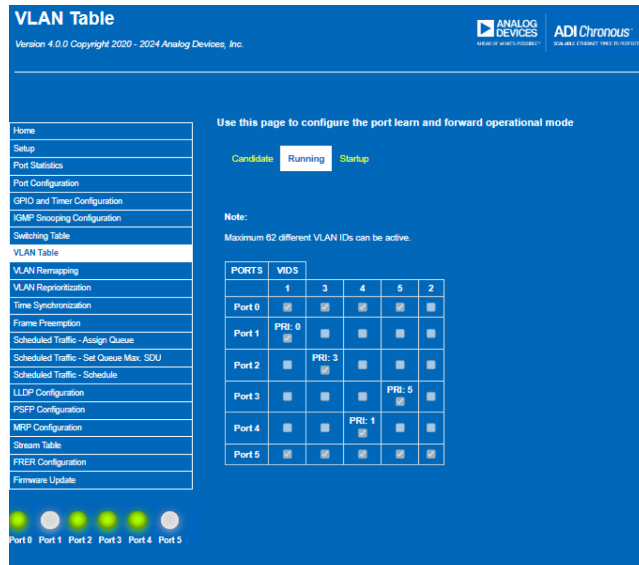


Figure 44. Using Trunk/Access Port Configuration – Running View

VLAN CONTROL

VLAN REMAPPING

As shown in Figure 45, the **VLAN Remapping** page provides user ability on a per port basis to add entries to remap incoming VLAN IDs to a different VLAN ID. Remapping is achieved by replacing the source VID in an incoming VLAN tagged frame with a destination VID. Per port, a table with 16 entries (slots) is used to configure the remapping. To add an entry, select the port of interest, select the **Enable remapping** check box, add the **Source VLAN ID** and

the **Target VLAN ID**, click the **Save** button followed by the **Commit** button. These remap entries are then saved and loaded to the device. Traffic ingressing a port with a corresponding VLAN ID can be observed to egress on the defined port with the remapped/target ID. To remove an entry, choose **Delete** button (see Figure 46). To view status of other ports, select the other port and any existing entries are displayed in the table.

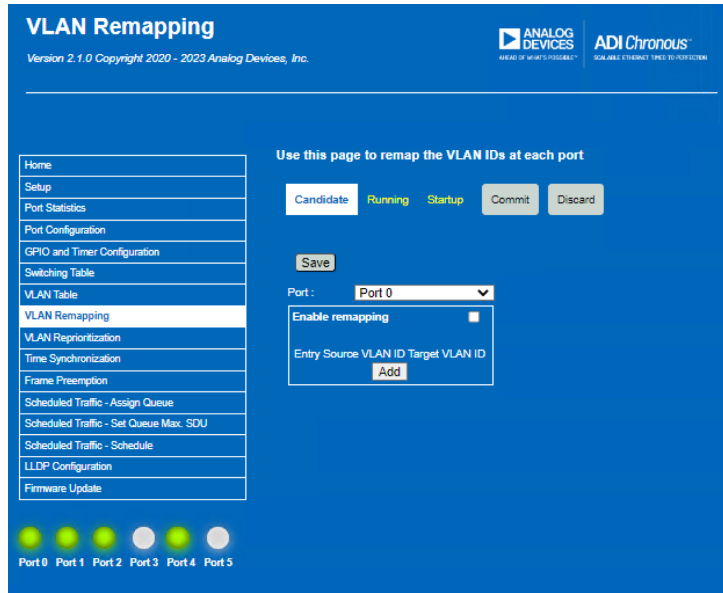


Figure 45. VLAN Remapping Page – Candidate View

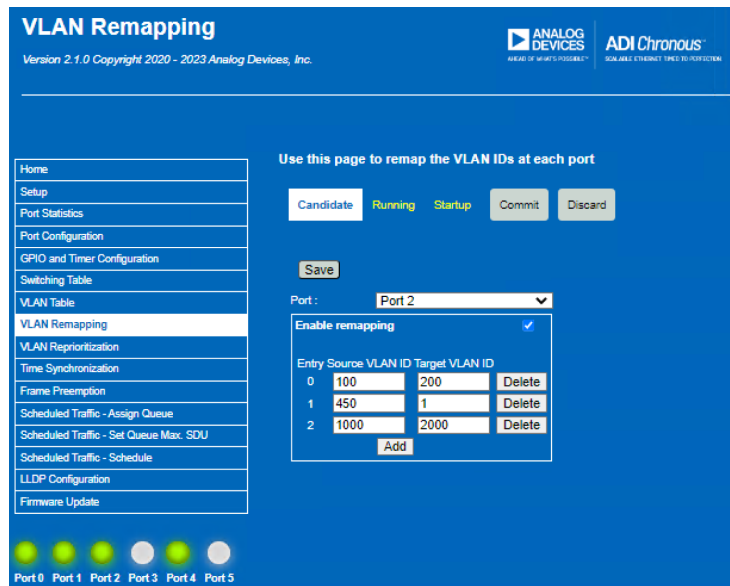


Figure 46. VLAN Remapping Page – Candidate View – Adding Entries for Port 2

VLAN CONTROL

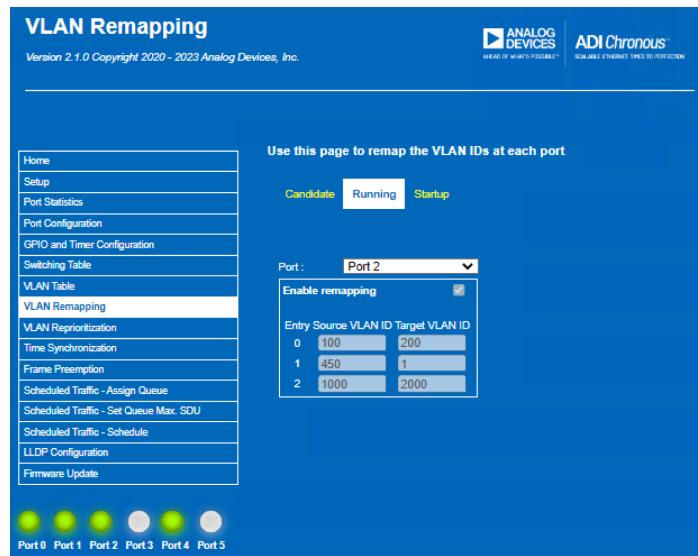


Figure 47. VLAN Remapping Page – Running View – Displays Entries for Port 2

VLAN CONTROL

VLAN REPRIORITIZATION

As shown in Figure 48, the VLAN Reprioritization page gives user ability to remap the priority of the VLAN traffic on a port basis. There is a **Candidate**, **Running**, and **Startup** view for these pages. There is no status page for VLAN Reprioritization.

All configuration happens in the **Candidate** page. To select a different priority for a VLAN ID, select the port of interest, enable prioritization on that port by enabling the **Enable Reprioritization** check box, then select the appropriate remapping IDs, use the

individual **Save** buttons or the main page **Save** button to save the changes to the web server and click the **Commit** button to load the changes to the device. When the **Commit** button is clicked, the changes are loaded and the web server automatically changes to show the **Running** view, where user can confirm programmed changes are applied.

The **Discard** button allows user to revert changes in the **Candidate** field, by copying the running configuration back to the **Candidate**.

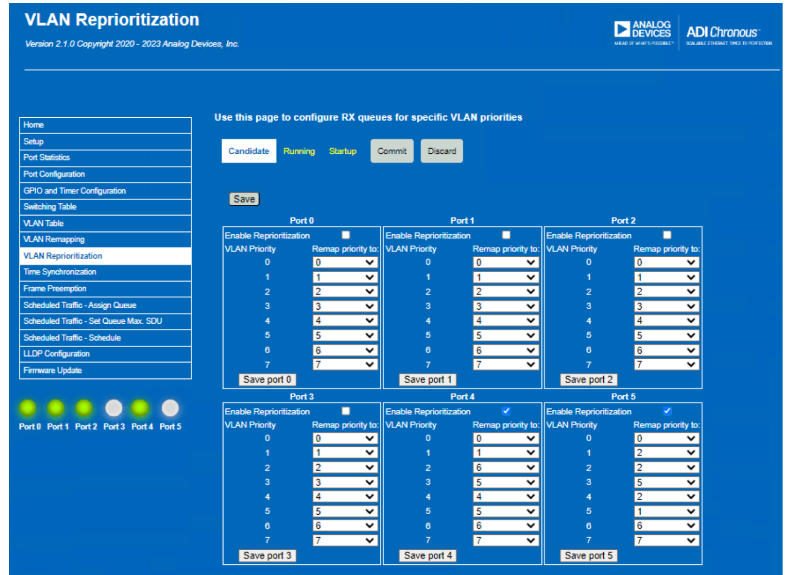


Figure 48. VLAN Reprioritization Page – Candidate View

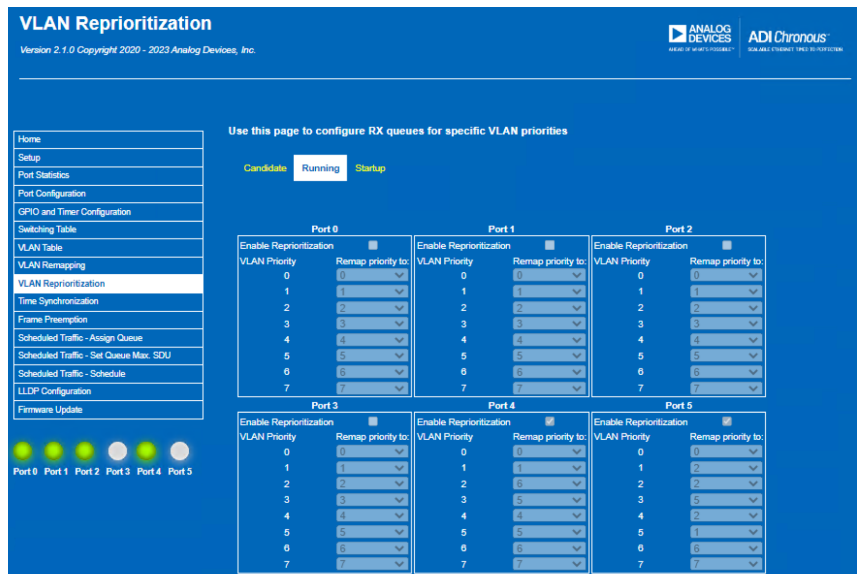


Figure 49. VLAN Reprioritization Page – Running View

VLAN CONTROL

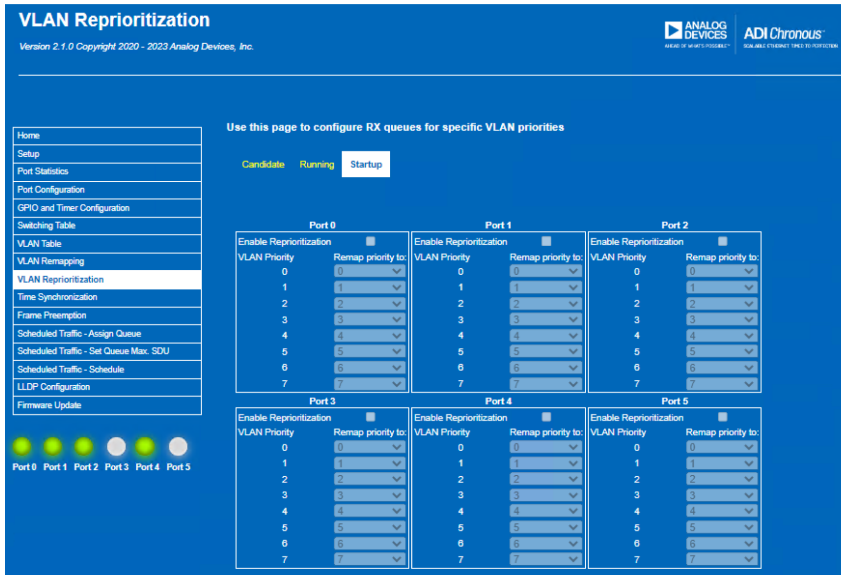


Figure 50. VLAN Reprioritization Page – Startup View

TIME SYNCHRONIZATION IEEE 802.1AS

Time synchronization is the process of ensuring that the clocks of different systems or devices have the same sense of time. The switch supports configuration of three different profiles of time synchronization. IEEE 802.1AS-2020 is the time-sensitive networking (TSN) standard for time synchronization within industrial Ethernet networks. It provides a precise and scalable method for synchronizing clocks across devices in an Ethernet-based network, allowing them to work in a coordinated manner for time-sensitive applications. IEEE 1588-2019 is the precision time protocol (PTP) standard for precision clock synchronization protocol for networked measurement and control systems. This profile is used in industrial, instrumentation, and power grid applications. The third profile is IEEE C37.238-2017 which is used for power system applications. This profile is not yet supported and will follow in subsequent software updates.

As shown in [Figure 26](#) or [Figure 27](#), click **Time Synchronization** in the menu item on the **Home** page or in the menu on the left of the page to access the **Time Synchronization** pages.

By default, when the web server runs, the PTP stack running on the switch is automatically enabled with one instance of IEEE 802.1AS, Domain 0, and PTP is enabled for all ports using the instance-specific peer-to-peer delay mechanism. The instance-specific peer-to-peer delay mechanism supports backward compatibility with IEEE 802.1AS 2011. To configure this profile for 2020 version of IEEE 802.1AS, enable CMLS delay mechanism. Additional instances can be enabled through the web page and the device supports configuration for up to four separate instances.

CANDIDATE PAGE

The **Candidate** page provides user ability to modify the operation of the PTP instance or add additional instances. Any changes must be saved and then committed. When changes are successfully committed, the **Running** page appears and shows the updated configuration. If an update is unsuccessful, a pop-up appears that shows the user update failed, and the **Running** page displays the last successful configuration.




Figure 51. Time Synchronization Candidate Page – Update Unsuccessful

PTP CONFIGURATION

The PTP stack supports up to 4 PTP instances. By default, one instance of IEEE 802.1AS, Domain 0, is enabled for all 6 ports.

As shown in [Figure 52](#), the page shows the mapping of PTP ports to Link Port Numbers. In the IEEE802.1AS standard, the Port assignment starts with Port 1. The web server Time Synchronization Port numbering aligns with this, but elsewhere in the web server, the port number starts at Port 0.

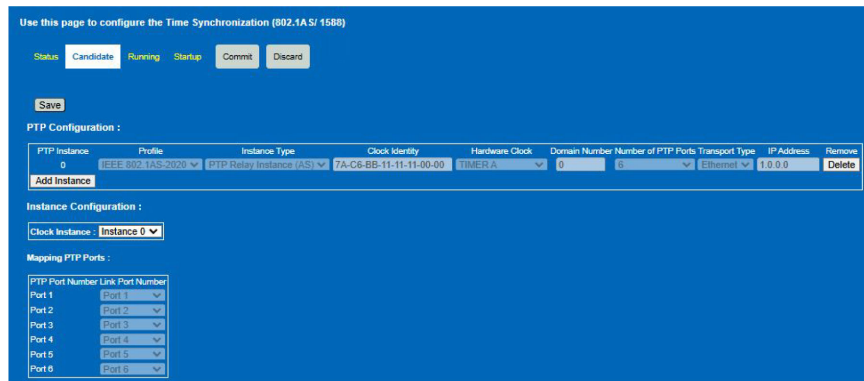


Figure 52. Time Synchronization Candidate Page – PTP Configuration

Traffic Priority: By default, PTP messages go into the highest priority transmit queue which is queue 7. The queue that PTP messages use for each port can be changed via the corresponding **Transmit Priority** fields, see [Figure 53](#). Using MRP with the fastest Recovery profiles may motivate the user to make changes to the priority of PTP messaging.

TIME SYNCHRONIZATION IEEE 802.1AS

Priorities

The default behavior for the enabled instance has Priority1 and Priority2 values set to 248. Priority values are among the parameters used as part of the best timeTransmitter clock algorithm (BTCA).

Lower values in the **Priority1** or **Priority2** fields, increase the chance that device becomes the Grandmaster. A service in a TSN network should not try to claim Grandmaster functionality unless it is by design. A typical Grandmaster is a node with a time normal receiver, a global positioning system (GPS) receiver, or an atomic clock. In industrial automation, an infrastructure Switch or a controller can cover Grandmaster functions. The priority value range is 0 to 255.




Figure 53. Time Synchronization Candidate Page – Traffic Priority and Default Dataset

PORT CONFIGURATION

The **Port Configuration**, shown in Figure 55, provides ability to configure various parameters associated with the PTP instance per port. The web server provides ability to change each individual port individually with the **PTP Port Number** drop-down menu, as shown in Figure 54. When changing port configuration, remember the port numbering for PTP is offset by one.

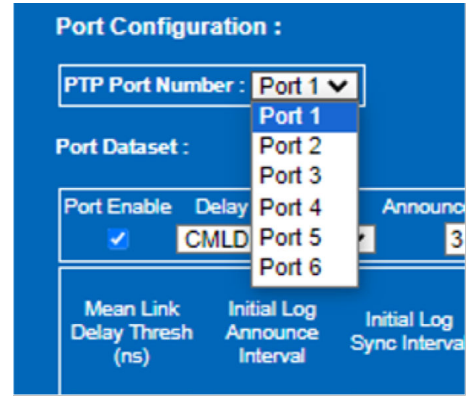


Figure 54. Time Synchronization Candidate Page – Port Configuration per Port Selection

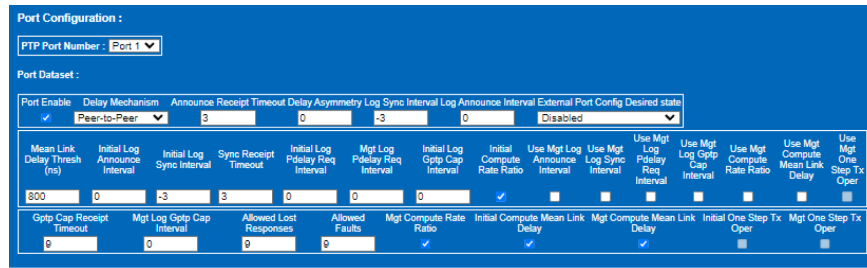


Figure 55. Time Synchronization Candidate Page – Port Configuration

TIME SYNCHRONIZATION IEEE 802.1AS

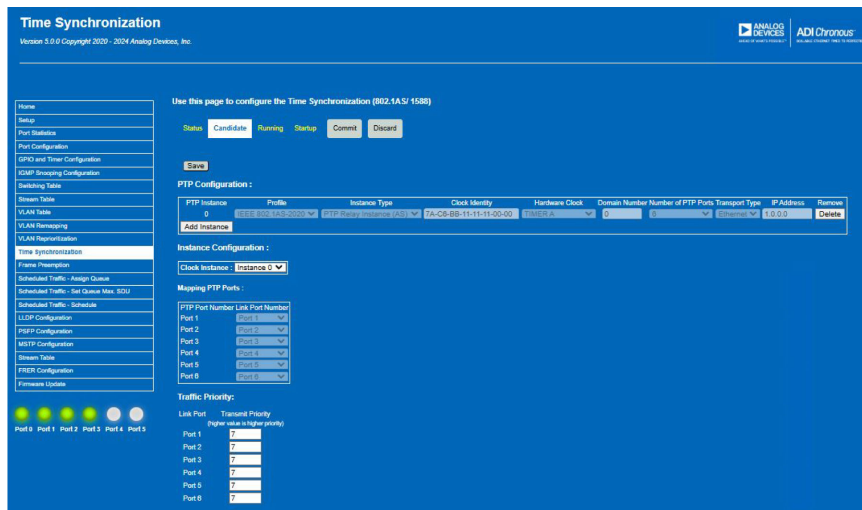


Figure 56. Time Synchronization – Candidate Page (Top)

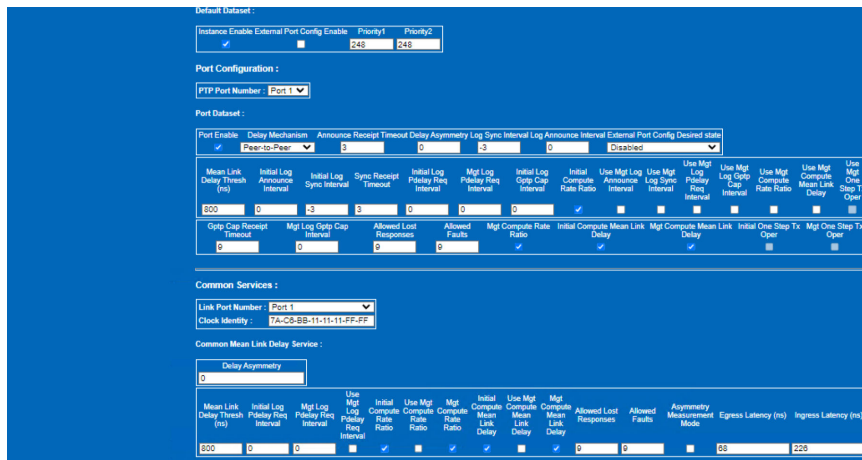


Figure 57. Time Synchronization – Candidate Page (Bottom)

TIME SYNCHRONIZATION IEEE 802.1AS

Delay Mechanism

By default, the PTP instance is enabled on all ports with a peer-to-peer delay mechanism, which supports backward compatibility to IEEE802.1AS 2011.

If only one time domain is enabled, user has choice of instance-specific peer-to-peer delay mechanism or common mean link delay service (CMLDS). CMLDS provides the mean propagation delay and neighbor rate ratio to all active domains.

For any PTP instance with a domain number that is not zero, CMLDS is enabled. Attempting to enable the instance-specific peer-to-peer delay mechanism on any one instance with a domain number that is not zero is not accepted. The update is rejected and the **Running** page shows the previous successful update.

To synchronize with devices running IEEE802.1AS 2020 and to add additional PTP instances, change the delay mechanism for the relevant ports to CMLDS.

Interval Times

This section of web page gives user ability to adjust the interval messaging for Sync, Announce and Peer Delay request messages.

Changing interval settings can result in improved tuned application behavior of the synchronized clocks. For example, lowering the sync interval (to a smaller value) can improve the precision of the synchronization.

Table 10. Port Delay Message Options

Parameter	Description	Value Range	Default
Log Pdelay_Req Interval	The interval of peer delay requests sent from the timeReceiver to the timeTransmitter.	+5 to -5	0 (1 sec)
Log Sync Interval	The interval of sync messages sent out by the timeTransmitter.	+5 to -5	-3 (125 ms)
Log Announce Interval	The interval in which the timeTransmitter announces its leadership.	+5 to -5	0 (1 sec)

The interval time is given in log2 values, as the 802.1AS standard suggests:

$$t_{INTERVAL} = 2^{\log INTERVAL} \tag{1}$$

Table 11. Interval Time Setting

LogInterval	t _{INTERVAL}
-5	31.25 ms
-4	62.5 ms
-3	125 ms
-2	250 ms
-1	500 ms
0	1 sec

Table 11. Interval Time Setting (Continued)

LogInterval	t _{INTERVAL}
+1	2 sec
+2	4 sec
+3	8 sec
+4	16 sec
+5	32 sec

Within the range specified, users can make changes to any of these values.

To change the interval times, user must modify the **Log Interval** value in addition to the **Use Mgt Interval** for the relevant parameter. In **Figure 58**, the sync interval is changed from default of -3 (every 125 ms) to 0 (every second) by adjusting the **Log Sync Interval** field and also checking the **Use Mgt Log Sync Interval** field. Similarly, the **Log Announce Interval** is adjusted from default of 0 to -3 and the **Mgt** field is checked to enable this change.



Figure 58. Interval Time Adjustment

Mean Link Threshold

The mean link threshold defaults to 800 ns. If using an Ethernet tap inline with the time aware link or attempting to synchronize over a 10BASE-T1L link, larger threshold values are required. In event the link delay is in excess of the programmed threshold, devices are not able to synchronize.

COMMON SERVICES

Figure 59 shows the configuration for **Common Mean Link Delay Service**. This section only applies when using CMLDS as the **Delay Mechanism** (selected from the drop-down in the **Port Dataset** view).

Note that the clock identity for common services differs from the PTP instance Clock identity.

Egress/Ingress Latency

The Egress/Ingress Latency values shown are specific to the Ethernet physical layer device (PHY). These are hardware dependent parameters based on the PHYs used. The default values shown in the web page are based on the **ADIN1300** PHYs that are connected over RGMII interface to the Switch on the EVAL-ADIN6310EBZ hardware.

The RGMII latencies based on the ADIN1300 configuration are listed in **Table 12**. For Link Port 2 to Port 6, the ADIN1300 Rx/Ingress latency defaults to the lowest latency mode when the cable length is estimated to be <100 m (CDIAG_CBL_LEN_EST (0xBA25)) and the PHY MSE (mean squared error) is <14 on

TIME SYNCHRONIZATION IEEE 802.1AS

all four dimensions (MSE_A (0x8402), MSE_B (0x8403), MSE_C (0x8404), and MSE_D (0x8405)).

When Link Port 1 (Physical port 0) is used as the Host interface, the PHY is treated as unmanaged, therefore the ingress/egress latency for that port always shows the higher Ingress latency of 226 ns.

Table 12. ADIN1300 PHY RGMII Actual Rx/Tx Delay/Latency

Speed	Tx/Egress	Rx/Ingress	Comment
1000 Mbps	68 ns	178 ns	Low Latency mode when cable length <100 m or MSE values <14.

Table 12. ADIN1300 PHY RGMII Actual Rx/Tx Delay/Latency (Continued)

Speed	Tx/Egress	Rx/Ingress	Comment
1000 Mbps	68 ns	226 ns	Standard latency mode when cable length >100 m or MSE values >14.
100 Mbps	92 ns	250 ns	Standard latency mode.
10 Mbps	124 ns	250 ns	Standard latency mode.




Figure 59. Time Synchronization Candidate Page – Common Services

PTP INSTANCES

To add an additional instance, click the **Add Instance** button, as shown in Figure 60. Another row appears on the web page with different **Clock Identity**, **Hardware Clock**, and **Domain number**.




Figure 60. Time Synchronization Candidate Page – Adding a Second Instance

Select the number of ports with which this instance should be used. By default, only one is selected. Configure any specific other parameters associated with this instance. When more than one instance is configured, the delay mechanism CMLDS is automatically used. To remove an instance, in the **Remove** area, click the **Delete** button.

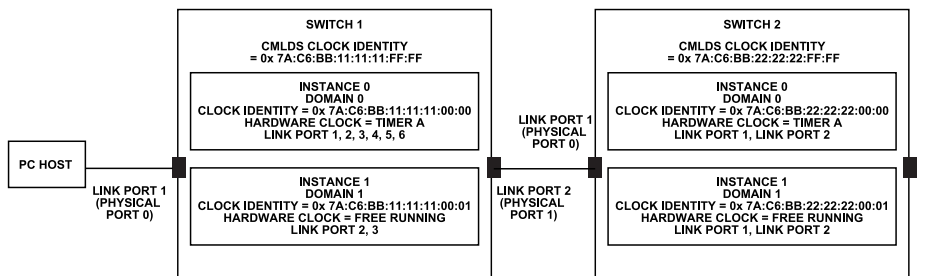


Figure 61. Time Synchronization with Two Instances

TIME SYNCHRONIZATION IEEE 802.1AS

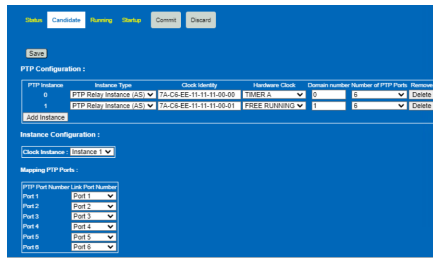


Figure 62. Time Synchronization Candidate Page – Adding a Second Instance for All 6-Ports

TIME SYNCHRONIZATION IEEE 802.1AS

HARDWARE CLOCK

The Switch supports two hardware clocks (TIMER A and TIMER B) and a free running clock. Currently only **TIMER A** and **FREE RUNNING** options are available. By default, the first instance is configured with **TIMER A**. When a second instance is added, it runs from the **FREE RUNNING** clock automatically. The switch hardware functions such as scheduled traffic, 1PPS, and PSFP run off of Timer A. Either timers can be used to synchronize the switch in a network. However, to observe synchronization through the 1PPS signal on the timer pins or have scheduled traffic and stream gate functions synchronized, then Timer A must be used as a hardware clock for one instance. Only one clock instance can use Timer A. If Timer A is used for the first instance, each subsequent instance will default to use the free running clock.

EXTERNAL PORT CONFIG

The External Port configuration enable is used where user does not want to use BTCA to decide who is Grandmaster in the network. Instead, user configures each device and port accordingly. The **External Port Config Enable** is used in conjunction with the **External Port Config Desired state** drop-down.



Figure 63. Time Synchronization Candidate Page – External Port Config Enable

Figure 64 shows an example where the first Switch is configured to be Grandmaster by configuring all its ports with timeTransmitter as the desired state. The following applies to Switch 2 and Switch 3. The **External Port Config Desired state** for the ports connected to Switch 1 is configured as **timeReceiver** and all other ports as **timeTransmitter**.

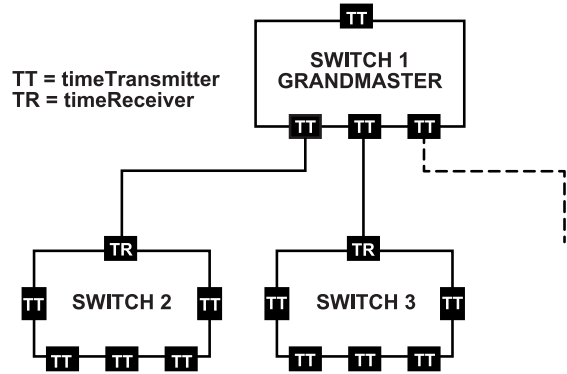


Figure 64. Time Synchronization Example for External Port Config

TIME SYNCHRONIZATION IEEE 802.1AS

STATUS PAGE

As shown in Figure 65, the Status page provides insight into the current status of PTP instances and indicates whether the device is synchronized.

The status information is available per configured instance, showing information such as who is Grandmaster.

The remaining parameters displayed are those defined by IEE 802.1AS and provide information regarding the operation of the time synchronization.

The following parameters are displayed for each port in the Port Configuration > Port Dataset:

- ▶ **Port State:** which is either **timeTransmitter**, **timeReceiver**, or **Disabled**.
- ▶ **Mean Link Delay (ns):** Measured the link delay across the cable.
- ▶ **AS Capable:** Which is either **Enable** or **Disable**.

The Status page also shows the detailed Port Statistics for the PTP instance, such as counts for PTP messaging, see Figure 67 and Figure 68 for CMLDS dataset (only shows valid information if CMLDS is active).

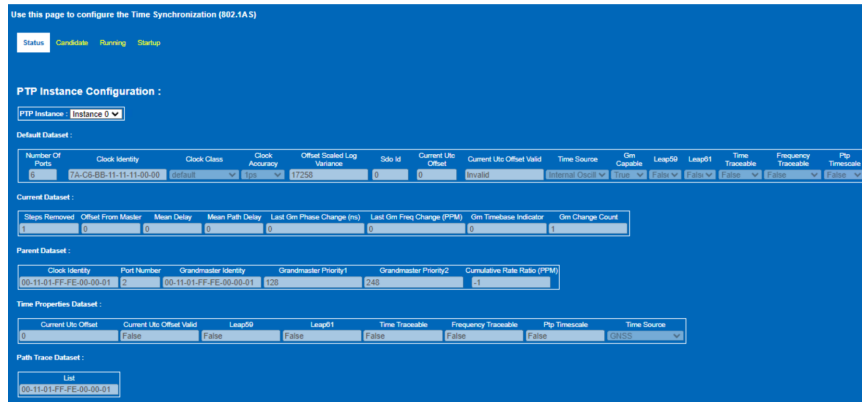


Figure 65. Time Synchronization Status Page – PTP Instance Configuration

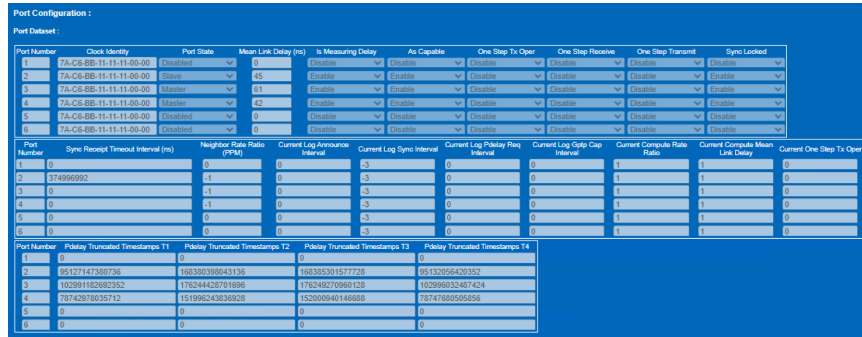


Figure 66. Time Synchronization Status Page – Port Configuration Dataset



Figure 67. Time Synchronization Status Page – Port Statistics Dataset

TIME SYNCHRONIZATION IEEE 802.1AS

Common Services Port Dataset:

Enable Link Port Number:

1
2
3
4
5
6

Common Mean Link Delay Service:

Default Dataset:

Clock Identity: 7A-CS-BB-11-11-11-FF-FF
Number Link Ports: 6

Link Port Dataset:

Port Number	Clock Identity	Domain Number	Mean Link Delay (ns)	Scaloid Neighbor Rate Ratio	Enable Link Port Enabled	Is Measuring Delay	As Capable Across Domains	Use Mgt Log Policy Req Interval	Use Mgt Compute Rate	Use Mgt Compute Mean Link Delay	Asymmetry Measurement Mode
1	7A-CS-BB-11-11-11-FF-FF	0	0	0	1	1	1	1	1	1	1
2	7A-CS-BB-11-11-11-FF-FF	0	0	0	1	1	1	1	1	1	1
3	7A-CS-BB-11-11-11-FF-FF	0	0	0	1	1	1	1	1	1	1
4	7A-CS-BB-11-11-11-FF-FF	0	0	0	1	1	1	1	1	1	1
5	7A-CS-BB-11-11-11-FF-FF	0	0	0	1	1	1	1	1	1	1
6	7A-CS-BB-11-11-11-FF-FF	0	0	0	1	1	1	1	1	1	1

Port Number	Mean Link Delay Threshold (ns)	Initial Log Policy Req Interval	Current Log Policy Req Interval	Mgt Log Policy Req Interval	Initial Compute Rate Ratio	Current Compute Rate Ratio	Mgt Compute Rate Ratio	Initial Compute Mean Link Delay	Current Compute Mean Link Delay	Mgt Compute Mean Link Delay	Allowed Lost Responses	Allowed Faults	Egress Latency (ns)	Ingress Latency (ns)
1	500	0	0	0	1	1	1	1	1	1	0	0	68	226
2	500	0	0	0	1	1	1	1	1	1	0	0	68	178
3	500	0	0	0	1	1	1	1	1	1	0	0	68	178
4	500	0	0	0	1	1	1	1	1	1	0	0	68	178
5	500	0	0	0	1	1	1	1	1	1	0	0	68	226
6	500	0	0	0	1	1	1	1	1	1	0	0	68	226

Port Number Policy Truncated Timestamps T1 Policy Truncated Timestamps T2 Policy Truncated Timestamps T3 Policy Truncated Timestamps T4

1	0	0	0	0
2	0	0	0	0
3	0	0	0	0
4	0	0	0	0
5	0	0	0	0
6	0	0	0	0

Port Statistics Dataset:

Port Number	Rx Policy Req Count	Rx Policy Resp Count	Rx Policy Resp Follow Up Count	Rx Packet Discard Count	Policy Allowed Lost Exceeded Count	Tx Policy Req Count	Tx Policy Resp Count	Tx Policy Resp Follow Up Count
1	0	0	0	0	0	0	0	0
2	0	0	0	0	0	0	0	0
3	0	0	0	0	0	0	0	0
4	0	0	0	0	0	0	0	0
5	0	0	0	0	0	0	0	0

Figure 68. Time Synchronization Status Page – Common Mean Link Delay Service Dataset and Statistics (Returns Information if CMLDS is Enabled)

TIME SYNCHRONIZATION IEEE 802.1AS

TIME SYNC MESSAGING

Using an Ethernet Tap between two Time Aware devices with just one instances enabled, user can view the gPTP messaging.

Figure 69 shows an example of the messaging between two devices with the default Time Sync parameters. The messaging intervals can be modified through the **Candidate** web page.

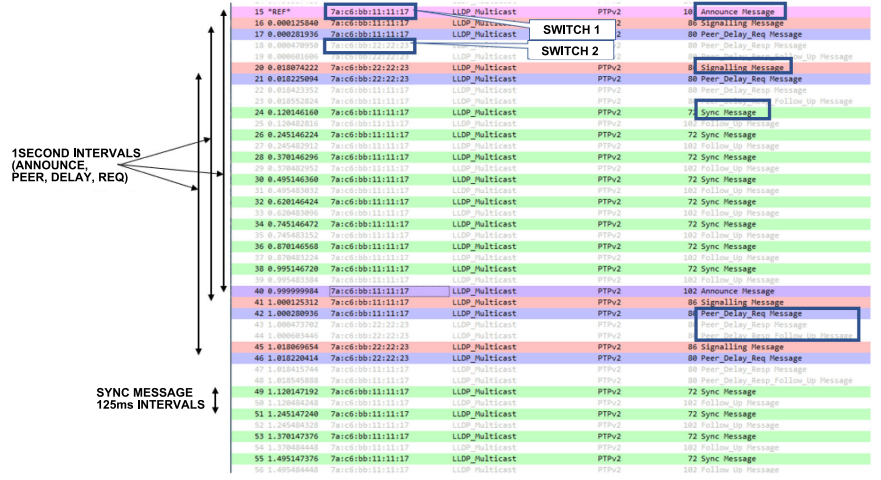


Figure 69. Time Synchronization View of Messaging through Wireshark Using an Ethernet Tap between Two Devices

TIME SYNCHRONIZATION IEEE 802.1AS

RUNNING PAGE

The **Running** page shows the running configuration on the device.

The fields on this page cannot be edited. Return to the **Candidate** configuration to change configuration.

STARTUP PAGE

The **Startup** page shows the startup configuration. These parameters are displayed to verify the values of the **Startup** configuration only.

TIMER PINS, 1PPS SIGNAL

The TIMER2 pin is used to provide a 1PPS (one pulse per second) signal. Probing the TIMER2 pin with a logic analyzer shows the 1PPS Time Synchronization pulse, as shown in [Figure 70](#). It is also visible on the evaluation board via the blinking of LED TIMER2.




Figure 70. 1PPS Signal on TIMER2 Pin for Two Switch Boards

TIME SYNCHRONIZATION, IEEE 1588

IEEE 1588-2019 is the precision time protocol (PTP) standard for precision clock synchronization protocol for networked measurement and control systems. This profile is used in industrial, instrumentation, and power grid applications. Figure 71 illustrates a typical network setup using IEEE 1588 PTP for synchronizing clocks across networked devices, a common use case in complex network environments. The switch supports configuration as a boundary clock, ordinary clock, or an end-to-end transparent clock. The propagation time is measured using delay request-response mechanism. Support for measuring the propagation time using peer-to-peer delay mechanism, and support for peer-to-peer transparent clock will be available in future software updates.

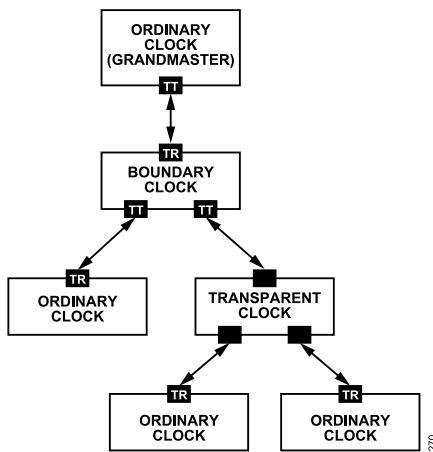


Figure 71. Overview of IEEE 1588 Devices Showing Grandmaster, Boundary, Ordinary, and Transparent Clocks

CANDIDATE PAGE

The **Candidate** page provides user ability to modify the operation of the PTP instance or add additional instances. Any changes must be saved and then committed. When changes are successfully committed, the **Running** page appears and shows the updated configuration. By default, the switch is configured with one instance of the IEEE 802.1AS time synchronization profile enabled.

The following section focuses specifically on one instance of the IEEE 1588 profile. To configure the IEEE 1588 profile, first, delete the existing profile and then add an IEEE 1588 profile. The switch supports up to four time sync instances and can support a mix of IEEE 802.1AS and IEEE 1588 profiles. The following selections are provided in the web server as shown in Figure 72:

- ▶ **Profile:**
 - ▶ Using the dropdown menu, choose the IEEE 1588 profile. This results in a dropdown selection for instance type of **Ordinary**, **Boundary**, or **End-to-end Transparent Clock**.
- ▶ **Instance Type:** The dropdown menu provides the following selections:
 - ▶ **Ordinary Clock**—A clock that functions as either a timeTransmitter or a timeReceiver, but not both. It can serve as a time

source if elected as a timeTransmitter and synchronizes to a timeTransmitter clock when operating as a timeReceiver. Usually an end device.

- ▶ **Boundary Clock**—Runs PTP on two or more interfaces. The clock can act as a timeReceiver and a timeTransmitter within the synchronization process. It synchronizes its local clock to an upstream grandmaster or a boundary clock while providing synchronization to downstream devices or clocks
- ▶ **Transparent Clock**—A clock that does not alter the time information in synchronization messages but measures and accounts for the residence time (delay) of these messages as they pass through it. This allows end devices or boundary clocks to compensate for network induced delays accurately. The device configured as a transparent clock is not part of the synchronized network. Currently, only end-to-end transparent clock is supported. Peer-to-peer transparent clock will be added in future software updates.
- ▶ **Clock Identity:** Clock identity for the instance. By default, the clock identity uses the MAC address assigned to the device and two additional bytes. For the first instance, the additional bytes is 00 00. If additional instances are added, the clock identity for new instances increments by one.
- ▶ **Hardware Clock:** Choice of free running or Timer A. The switch hardware functions such as scheduled traffic, 1PPS, and PSFP run off of Timer A. Either timers can be used to synchronize the switch in a network. However, to observe synchronization through the 1PPS signal on the timer pins or have scheduled traffic and stream gate functions synchronized, Timer A must be used as a hardware clock for one instance. Only one instance can use Timer A. If Timer A is used for the first instance, each subsequent instance defaults to use the free running clock. When the evaluation package is run, the IEEE 802.1AS profile defaults to use Timer A. If the user does not require IEEE 802.1AS profile, the user can delete this profile and add an IEEE 1588 profile.
- ▶ **Domain Number:** Sets the domain number for the clocks.
- ▶ **Number of PTP Ports:** For ordinary clocks, only one port can be selected for each instance. For boundary or transparent clocks, up to six ports can be selected.
- ▶ **Transport Type:** PTP messages can be sent over Ethernet or IPv4.
- ▶ **IP Address:** If transport type is configured to IPv4, PTP messages transmit with this IP address.
- ▶ **Instance Configuration:** Clock instance dropdown allows user to navigate between different instances when multiple instances have been created. Ensure to choose the correct instance when modifying other parameters.
- ▶ **Mapping PTP Ports:** Allows user to map PTP ports to the link ports.

TIME SYNCHRONIZATION, IEEE 1588

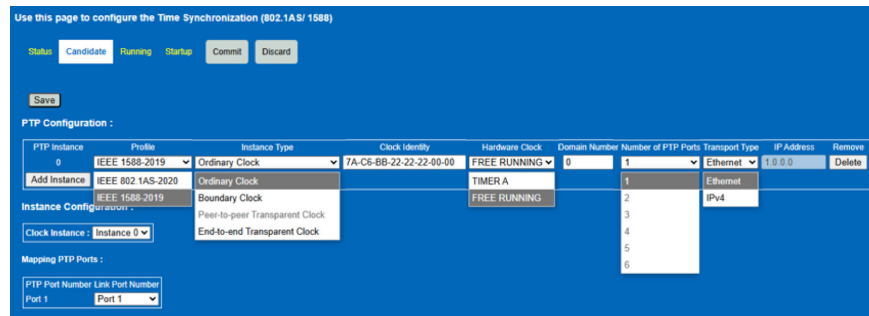


Figure 72. Time Synchronization Candidate Page—IEEE 1588 Profile Selection

Figure 73 shows the following additional configuration parameters for the IEEE 1588 instance:

- ▶ **Traffic Priority:** Provides the user ability to change with transmit queue the PTP messages egress, by default, they go into the highest priority transmit queue which is Queue 7. To modify egress queue, enter the queue number in the field for the corresponding port, choice of 0 to 7, lower number is lower priority.
- ▶ **Default Dataset:** Allows the user to enable the instance and configure **Priority 1** and **Priority 2** values. The default priorities are set to 248. Priority values are among the parameters used as part of the best time Transmitter clock algorithm (BTCA) to determine who is the grandmaster in the network. Note that **External Port Config Enable** is not yet supported.
- ▶ **Port Configuration:** In this section, the **Port Dataset** provides configuration of per PTP port parameters. Ensure to select the correct port from the PTP port number dropdown menu before making changes. Changes are made per port, not grouped.
- ▶ **Delay Mechanism:** Currently, End-to-End is supported.
- ▶ **Announce Receipt Timeout:** Specifies the number of announce intervals that have to pass without the receipt of an announce message before the PTP receiving node considers the time synchronization is lost. Default of 3.
- ▶ **External Port Config Desired State:** Not yet supported by software. Port configuration, if external, port config is enabled. Used for manual configuration of PTP ports (BTCA is disabled). Choices are **Disabled**, **TimeTransmitter**, **Passive**, or **TimeReceiver**.
- ▶ **Egress/Ingress Latency:** The values shown are specific to the Ethernet PHY used and reflect the latency from the [ADIN1300](#) PHY used on the EVAL-ADIN6310EBZ/EVAL-ADIN3310EBZ evaluation boards. By default, the latency values used in the web server reflect the ADIN1300 specified latency from the data sheet for RGMII interface. During configuration of the managed ADIN1300 PHYs (link ports 2 to 6, physical ports 1 to 5), the [ADIN6310](#) modifies a configuration in the PHY, which results in a reduction of ingress/receiver latency from 226 ns to 178 ns, if the cable and signal quality conditions described in [Table 12](#) are met. As a result, the parameters in the web page overwrites the ingress latency values used by the PTP stack, resulting in the PTP stack using ingress latency of 226 ns instead of the actual

178 ns. This impacts the synchronization between devices and can be observed on 1PPS signal as an increased variation/offset in the accuracy of the 1PPS signals. To account for this error until addressed in software future updates, the user must modify the ingress latency value in the **Candidate** page (per port for managed PHYs), save, and commit. When the host interface is configured for Ethernet interface over Port 0, the PHY connected to that switch port is not directly managed, therefore the reduced latency does not apply for that port. The ingress latency for that port must remain at 226 ns.

- ▶ **Message Intervals:** The PTP message intervals for sync, announce, and delay request messages are detailed in [Table 13](#).

TIME SYNCHRONIZATION, IEEE 1588

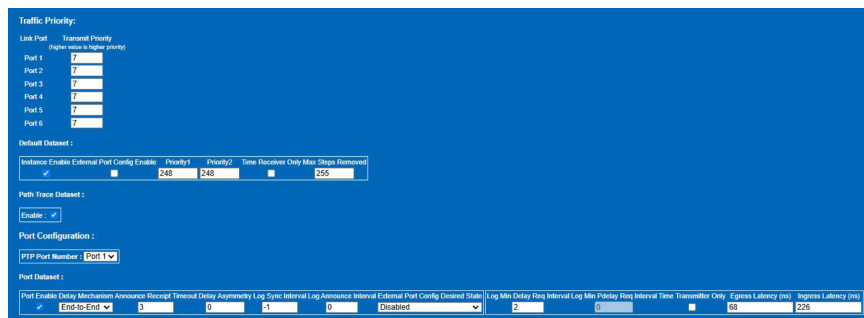


Figure 73. Time Synchronization Candidate Page—IEEE 1588 Dataset Configuration

Table 13. IEEE 1588 Port Delay Message Options

Parameter	Description	Value Range	Default
Log Delay_Req Interval	The interval of delay requests sent from the timeReceiver to the timeTransmitter.	0 to +5	2 (4 sec)
Log Sync Interval	The interval of sync messages sent out by the timeTransmitter.	-1 to +1	-1 (500 ms)
Log Announce Interval	The interval in which the timeTransmitter announces its leadership.	0 to +4	0 (1 sec)

The interval time is given in log2 values, as the standard suggests:

$$t_{INTERVAL} = 2^{\log INTERVAL} \tag{2}$$

Table 14. Interval Time Setting

LogINTERVAL	t _{INTERVAL}
-1	500 ms
0	1 sec
+1	2 sec
+2	4 sec
+3	8 sec
+4	16 sec
+5	32 sec

The candidate page for IEEE 1588 transparent clock has fewer configuration parameters as shown in Figure 74.

TIME SYNCHRONIZATION, IEEE 1588

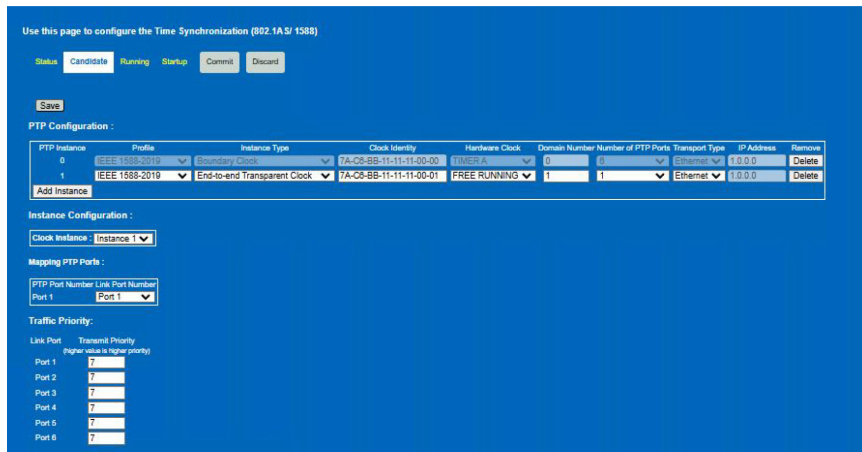


Figure 74. Candidate Page for Transparent Clock

STATUS PAGE

The **Status** page shows the current state of the PTP instances and whether the ports are synchronized. The status information is available per configured instance, showing information such as who is grandmaster in the network, the port states, and mean path delay information.



Figure 75. Time Synchronization Status Page—IEEE 1588 Showing Status for Boundary Clock

The user can also observe the 1PPS signal from each device configured as an ordinary or boundary clock in the network to assess how well they are synchronized, similar to that shown in Figure 70.

Note that a switch configured as a transparent clock is not part of a synchronized network. It does not generate PTP messages or respond to messages. Instead, a transparent clock adds its

residence time to the correction field of the synchronization messages that pass through it. This adjustment ensures that the delay introduced by the transparent clock itself is accounted for in the timing data for the other devices to allow them maintain accuracy of the synchronization process across the network. As a result, the 1PPS signal of a transparent clock will not be synchronized to the other network 1PPS signals.

FRAME PREEMPTION

Frame preemption is a mechanism that allows higher priority traffic to interrupt or preempt lower priority traffic during transmission. This feature is crucial for ensuring that time-sensitive data can be delivered with minimal delay, even in congested networks. Frame preemption is defined in IEEE 802.1Qbu/IEEE 802.3Qbr.

Frame preemption is designed to be used as part of the TSN suite of features but not with redundancy features such as PRP or HSR. Frame preemption relies on LLDP being enabled first as LLDP messages are used to advertise preemption capability to the link partner through the additional Ethernet capabilities TLVs of the LLDP frames per port. When the web server is enabled, it automatically enables LLDP function.

Click **Frame Preemption** in the menu item on the **TSN Switch Evaluation – Home** page (see [Figure 26](#)) or in the menu on the left of the page to access the **Frame Preemption** page, as shown in [Figure 77](#). Similar to **Time Synchronization**, the **Frame Preemption** page has **Status**, **Candidate**, **Running**, and **Startup** views.

CANDIDATE PAGE

To configure the way **Frame Preemption** operates, users can configure each port through the **Candidate** page, see [Figure 76](#).

The following control parameters are provided for each port:

- ▶ **Preemption Support:** Check box to enable or disable the function, default is disabled.
- ▶ **Ignore Peer Preemption Status:** Check box to allow port ignore the peer preemption capabilities. This bypasses the checks for peer preemption. This must be used in conjunction with the **Disable Verify Message Transmit**.

- ▶ **Minimum Non-Final Fragment Size (bytes):** Provides control of the fragment size, drop-down with choice of **64**, **28**, **192**, or **256** bytes.
- ▶ **Disable Verify Message Transmit:** By default, this check box is cleared, which is the expected operation. Preemption requires that a port sending a verification frame must get a response to allow frame preemption be enabled. This check box provides ability to disable the verify message if required to force preemption on.
- ▶ **Verify Message Period (ms):** Sets the verify frame transmit retry timer with a range of 1 ms to 128 ms, default 10 ms.
- ▶ **Express Queues:** Defaults to all queued marked as express. Select the required check boxes to enable preemption on that queue. Queues map directly to VLAN priorities.

Once the user has a new candidate configuration, click **Save** button followed by **Commit** button to send the **Candidate** configuration entries to the **Running** configuration. Click **Discard** button to revert the **Candidate** configuration back to current **Running** configuration.

When committing the **Candidate** configuration to the **Running** configuration, the current **Running** configuration saves to a running backup configuration before the **Candidate** configuration saves to the **Running** configuration. The purpose of this save to the running backup configuration is to allow the user to undo the **Commit** action in the event that the committed **Candidate** configuration results in a catastrophic effect on the TSN operation.

The example configuration shown in [Figure 76](#) has Port 5 with preemption enabled and all queues except queue 5 are configured as preemptable.

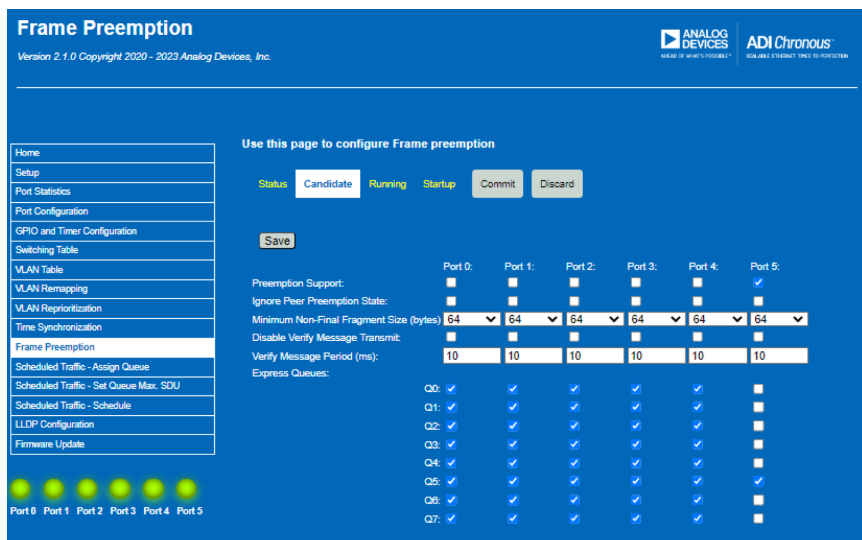


Figure 76. Frame Preemption – Candidate Page View

FRAME PREEMPTION

STATUS PAGE

The Preemption **Status** page is shown in [Figure 77](#). The following status information is provided per port:

- ▶ **Preemption Active:** Reports status check box. Selected indicates active, cleared indicates inactive.
- ▶ **Verify Status:** Shows state (**Initial, Active**).
- ▶ **Peer Supported:** Shows whether the peer is capable of Preemption.
- ▶ **Peer Enabled:** Shows whether the peer has preemption enabled
- ▶ **Peer Active:** Shows whether the peer has preemption active.
- ▶ **Hold Advance (nsec):** Shows the maximum number of nanoseconds that can elapse between issuing a Hold to the MAC and the MAC ceasing to transmit any preemptable frame that is in the process of transmission or any preemptable frames that are queued for transmission, including any MAC specific delay before transmission of an express frame can start once preemptable frame transmission has ceased.

- ▶ **Release Advance (nsec):** Shows the maximum number of nanoseconds that can elapse between issuing a Release to the MAC and the MAC being ready to resume transmission of preemptable frames, in the absence of there being any express frames available for transmission.
- ▶ **Preemption Statistics:** Provides overview of the various statistics associated with Transmit and Receive processing:
 - ▶ **Frame Assembly Error Count**
 - ▶ **Frame SMD Error Count**
 - ▶ **Frame Assembly OK Count**
 - ▶ **Fragment Count Rx**
 - ▶ **Fragment Count Tx**
- ▶ **Hold Count:** Associated with use of Hold_EN with Scheduled Traffic, returns a count of the number of times the HOLD enable transitions from FALSE to TRUE.

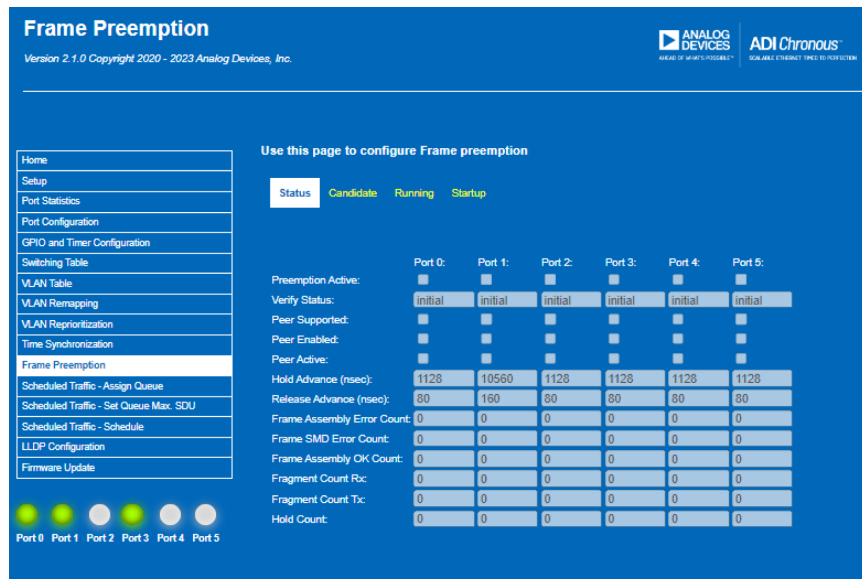


Figure 77. Frame Preemption Status Page

FRAME PREEMPTION

FRAME PREEMPTION EXAMPLE

The following example works through enabling preemption in a configuration with two EVAL-ADIN6310EBZ evaluation boards. Port 3 of Switch 1 is connected to Port 0 of Switch 2.

- ▶ The Preemption settings for each device are configured through the individual web pages.
- ▶ In Switch 1, Preemption is enabled on Port 3, express and preemptable queues configured accordingly Queue 5 is assigned as the only express queue and all other queues are cleared and therefore preemption can be applied to these queues. Once the

changes are made, then click the **Save** button followed by the **Commit** button to load the settings.

- ▶ In Switch 2, enable Preemption support on Port 0. A transmitting port only sends frames with an SMD-S/C (frames to which preemption has been applied) only after it has been established that the link partner supports preemption and the transmitting port has been instructed to enable preemption on the Tx Queues for this link. LLDP frames are used to exchange capabilities.
- ▶ Enabling Preemption in Switch 2 allows preemption to become active as shown in [Figure 79](#).

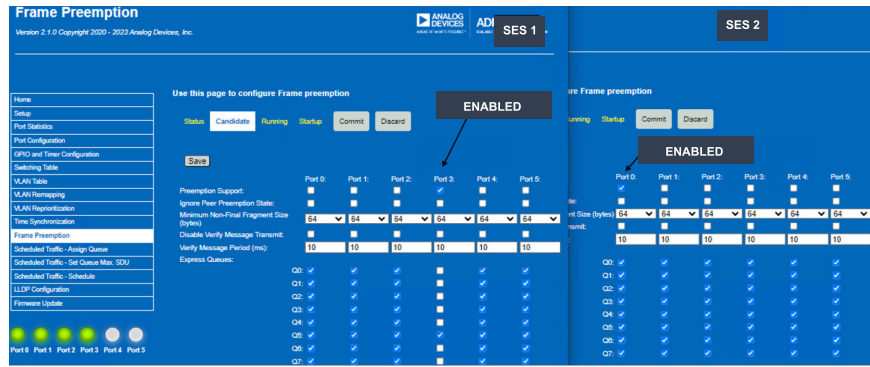


Figure 78. Candidate Page View to Enable and Configure Preemption

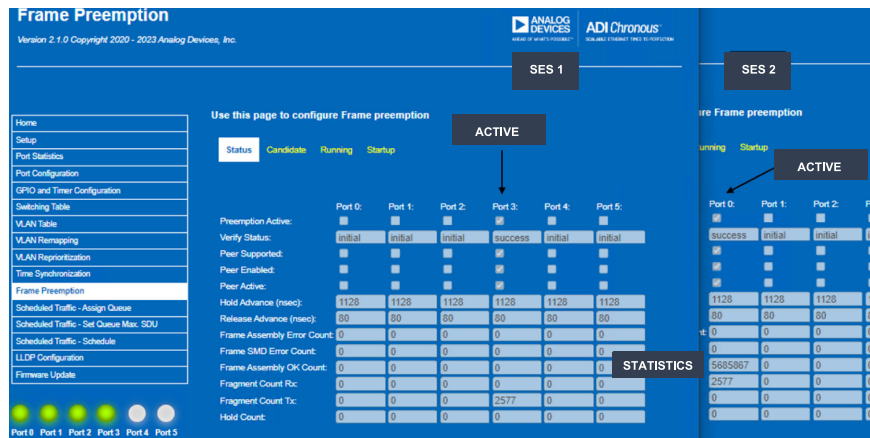


Figure 79. Status View with Preemption Enabled and visibility into Statistics

SCHEDULED TRAFFIC

Scheduled traffic is defined by **IEEE 802.1Qbv** and is a method for managing and controlling time-sensitive data traffic in Ethernet networks, specifically for real-time communication. It belongs to the suite of TSN standards and allows for precise scheduling of traffic to ensure critical traffic can be transmitted with minimal delay.

Configuration of Scheduled Traffic is done on three main pages: **Scheduled Traffic – Assign Queue**, **Scheduled Traffic – Set Queue Max SDU**, and **Scheduled Traffic – Schedule** (see [Figure 80](#)).

Click **Scheduled Traffic – Assign Queue** menu item on the Home page or in the menu on the left of the page to start configuration, this opens the **Candidate** view, as shown in [Figure 81](#).




Figure 80. Scheduled Traffic Pages

ASSIGNING QUEUES

Traffic is scheduled on a TSN network using VLAN Priority. By assigning **VLAN Priority** to an Ethernet message, that message can be assigned to a queue in the Switch hardware. There are 8 queues in this hardware and any of the 8 VLAN Priorities can be assigned to any of the queues.

By default, PTP and LLDP traffic has been designated to go into **Q7**, therefore **Q7** must be enabled for at least 10 μ s to provide bandwidth for time synchronization messages.

Q0 is designated **best effort**, untagged traffic is forced to this queue.

The **Candidate** page is used to map **VLAN Priority** to a queue. Click the white dot under a queue corresponding to the required **VLAN Priority**. The default mapping is, for example, **Q0** to **VLAN Priority 0**, **Q1** to **VLAN Priority 1**.

In [Figure 81](#), the configuration for **Port 1** has been remapped to the following, by clicking **Q0** for **VLAN Priority 0** and 1, **Q1** for **VLAN Priority 2** and 3, **Q2** for **VLAN Priority 4** and 5, **Q3** for **VLAN Priority 6** and 7. Click the **Save** button followed by **Commit** button to send assignments to the **Running** configuration. Click the **Discard** button to revert to current **Running** configuration.

Click the **Running** button to display the current **Running** configuration. This is shown in [Figure 82](#) for the Queue Assignment Running page.

Click the **Startup** button to display the configuration of the **Startup** configuration. This is shown in [Figure 83](#) for the Queue Assignment Startup page.

SCHEDULED TRAFFIC

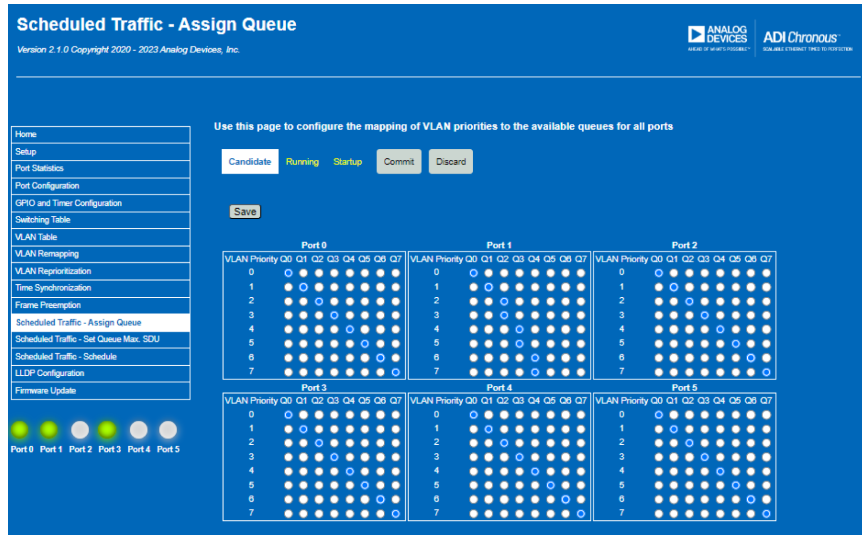


Figure 81. Scheduled Traffic – Queue Assignment Candidate Page

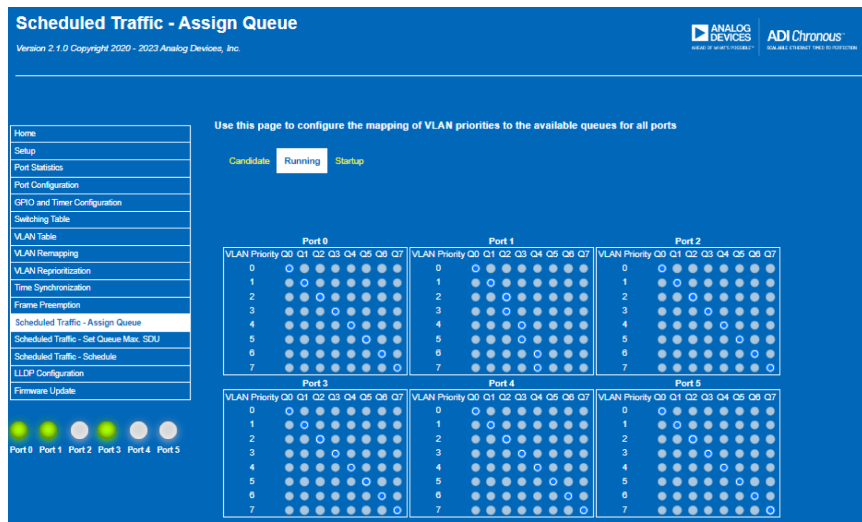


Figure 82. Scheduled Traffic – Queue Assignment Running Page

SCHEDULED TRAFFIC

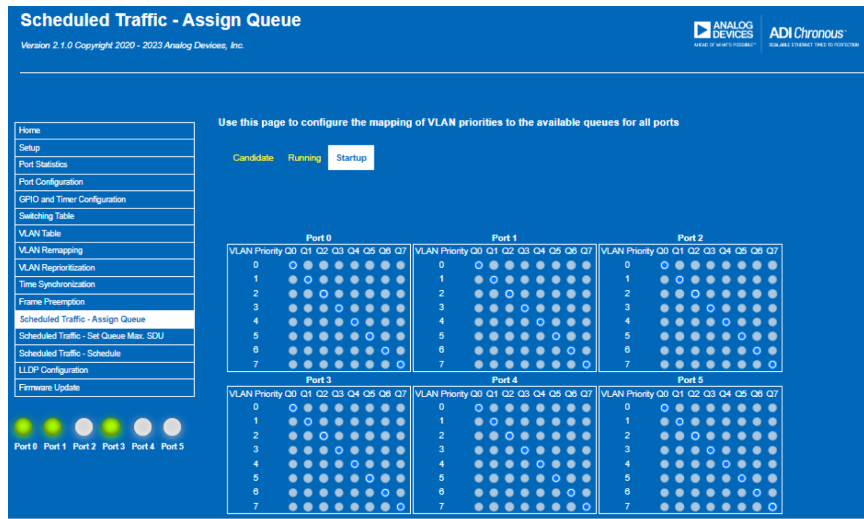


Figure 83. Scheduled Traffic – Queue Assignment Startup Page

SCHEDULED TRAFFIC

SCHEDULED TRAFFIC – SET QUEUE MAX. SDU

This page provides ability to adjust the service data unit (SDU) size of the frames allowed to egress per queue per port. The web server startup default setting is 1536 bytes, while the hardware defaults to 10,000 bytes. Adjusting the SDU size allows the user to fine tune the timing of the scheduled traffic. These values only need to be adjusted if the user knows precisely how they want to configure the timing. QueueMaxSDU does not include MAC addresses or FCS (QueueMaxSDU = Frame Size – 16 bytes).

The Queue Max. SDU **Candidate** page (see Figure 84) has a **Max. SDU [bytes]** field that can be defined for each of the 8 queues per port. Use this page to edit to the SDU byte sizes, by changing the

values. Click the **Save** button and then click the **Commit** button to load the new values. To return to the previously used parameters, click the **Discard** button.

When using Scheduled traffic with guard bands enabled, the guard band calculation uses the Max. SDU value to determine the duration of guard band to implement.

Note that Queue Max. SDU limits only apply to traffic forwarding in Store and Forward mode. When the Switch is cutting frames through, the frame has already started to egress before the frame size is known.

Similar to the other pages, there are **Candidate**, **Running**, and **Startup** views.

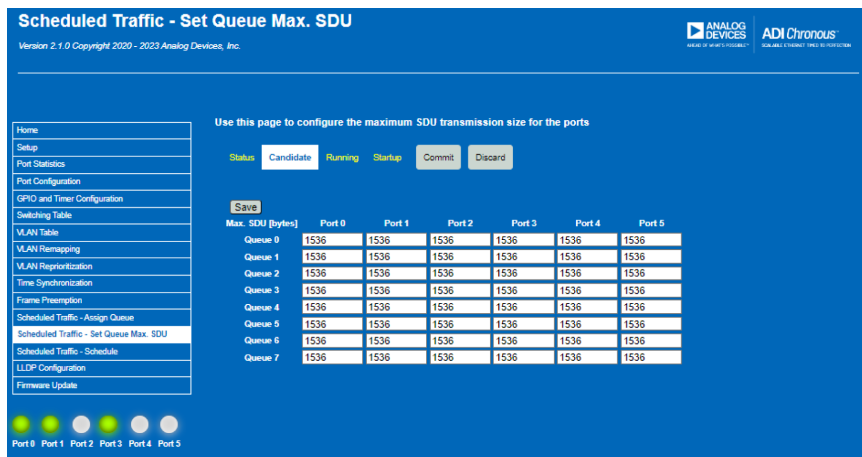


Figure 84. Scheduled Traffic – Queue Max. SDU Candidate Page

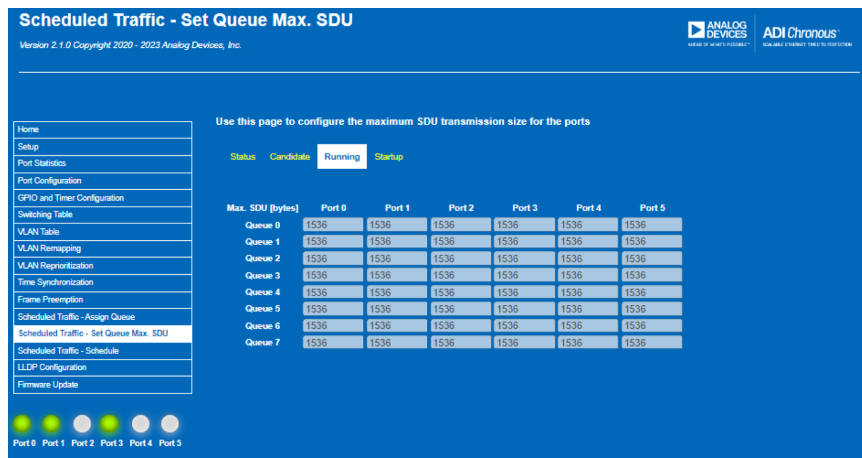


Figure 85. Scheduled Traffic – Queue Max. SDU Running Page

SCHEDULED TRAFFIC

Scheduled Traffic - Set Queue Max. SDU

Version 2.1.0 Copyright 2020 - 2023 Analog Devices, Inc.

Use this page to configure the maximum SDU transmission size for the ports

Status
Candidate
Running
Startup

Max. SDU (bytes)	Port 0	Port 1	Port 2	Port 3	Port 4	Port 5
Queue 0	1536	1536	1536	1536	1536	1536
Queue 1	1536	1536	1536	1536	1536	1536
Queue 2	1536	1536	1536	1536	1536	1536
Queue 3	1536	1536	1536	1536	1536	1536
Queue 4	1536	1536	1536	1536	1536	1536
Queue 5	1536	1536	1536	1536	1536	1536
Queue 6	1536	1536	1536	1536	1536	1536
Queue 7	1536	1536	1536	1536	1536	1536

- Firmware Update
- LLDP Configuration
- Scheduled Traffic - Schedule
- Scheduled Traffic - Set Queue Max. SDU
- Scheduled Traffic - Assign Queue
- Frame Preemption
- Time Synchronization
- VLAN Reprioritization
- VLAN Remapping
- VLAN Table
- Switching Table
- GPIO and Timer Configuration
- Port Configuration
- Port Statistics
- Setup
- Home

Port 0

Port 1

Port 2

Port 3

Port 4

Port 5

Figure 86. Scheduled Traffic – Queue Max. SDU Startup Page

SCHEDULED TRAFFIC – SCHEDULE

As shown in [Figure 26](#) or [Figure 27](#), click the **Scheduled Traffic – Schedule** menu item on the Home page or in the menu on the left of the page to access the **Scheduled Traffic – Schedule** pages. The first page that is navigated to is the **Scheduled Traffic Candidate** page. The **Scheduled Traffic Candidate** page provides a means to set the gate open events for each of the queues to support 802.1Qbv Scheduled Traffic.

Schedules can be configured on a per port basis. [Figure 87](#) shows the controls for **Port 0** only.

SCHEDULE ENABLED

To enable a Schedule, select the **Schedule Enabled** check box. Clear to disable scheduled traffic on this port. Note that any schedule must be saved and committed to load it to the device.

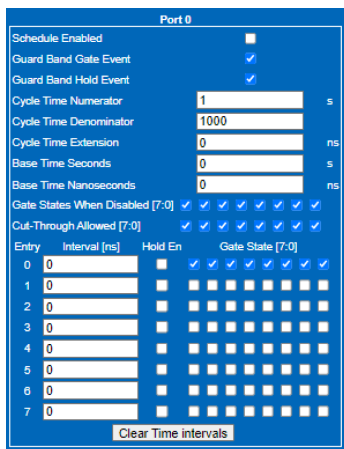


Figure 87. Scheduled Traffic Page (Showing Port 0 Only)

GUARD BANDS

As shown in [Figure 87](#), two check boxes (**Guard Band Gate Event** and **Guard Band Hold Event**) are associated with Guard band capability. The Switch supports automatic insertion of guard bands when these check boxes are enabled.

Guard bands are used with scheduled traffic to protect transmission of the schedule gate open times. Consider the scenario shown in [Figure 88](#) where no guard bands are used. An Ethernet port that has started transmission of a frame must complete transmitting that frame before another transmission can start. Consider a scenario where a new frame transmission starts just before the end of the first cycle, with a frame size too large to complete before the second cycle is due to begin, this results in a delayed start of the second cycle. The impact of this is that potentially lower priority traffic can be infringing on the start of time critical time slice, meaning real-time frames delays, which impact the application requirements.

Scheduled traffic can use guard bands in front of every time slice that carries time critical traffic. During the guard band duration, no new Ethernet transmissions can be started, only ongoing transmissions can complete. The duration of the guard band is sized for as long as it takes the maximum frame size to safely transmit.

When the **Guard Band Gate Event** check box is enabled, the Switch automatically inserts a guard band between the step that has the gate open for a traffic class and the step that has the gate closed. The length of the guard band is the product of the QueueMaxSDU value of the queue associated with the gate and the current link speed. The guard band time value is subtracted from the gate close time. This ensures that the start of the time slots do not get delayed.

As the different queues can have different QueueMaxSDU values, the guard bands for the different queues are calculated accordingly, as shown in [Figure 90](#).

Different QueueMaxSDU values do consume entries in the internal Gate control list. In the event the automatic guard band insertion fails, the driver package reports a return error. Exotic schedules with many different time slots and different QueueMaxSDU values can result in failure of guard bands to be inserted, but the GUI prompts in this event and user can review their schedule and revise accordingly.

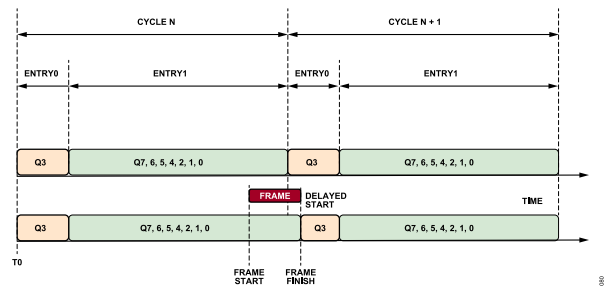


Figure 88. Scheduled Traffic Affect of No Guard Band




Figure 89. Scheduled Traffic with Guard Band




Figure 90. Scheduled Traffic with Different Max. SDUs per Queue

When using the evaluation package and PC based web server, the QueueMaxSDU values are set to 1536 bytes per queue per port. With a 1 Gbps link speed, this corresponds to 12.29 μs guard band.

The Switch hardware defaults to QueueMaxSDU setting of 10,000 bytes, therefore, when interfacing directly to the driver from own

SCHEDULED TRAFFIC – SCHEDULE

stack processor, configure the QueueMaxSDU values as required to avoid having excessive guard bands.

The second Guard band check box, **Guard Band Hold Event** is only relevant where Scheduled traffic and Frame preemption co-exist and the **Hold En** check box in the gate control list is enabled. When **Guard Band Hold Event** is enabled, a guard band of hold advance length is inserted between the step that has the Hold_En signal asserted and the previous step and the release advance length is inserted between the steps where the Hold_En signal transitions from asserted to deasserted. The value of hold advance can be seen in the **Frame Preemption Status** page and vary depending on the speed of the link established.

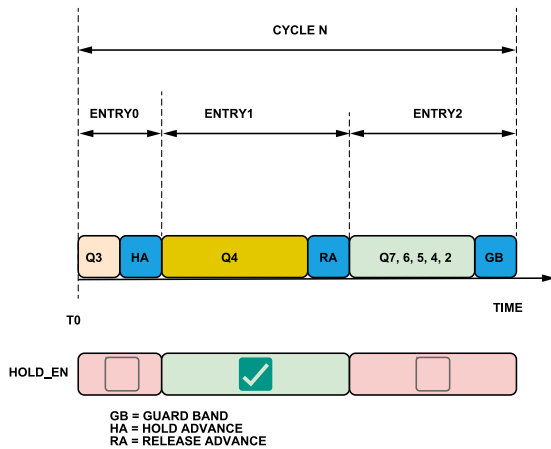


Figure 91. Guard Band Hold Event – Hold Advance and Release Advance

Table 15. Hold Advance, Release Advance (1 Gbps)

Fragment Size	Hold Advance (ns)	Release Advance
64	1128	80
128	1640	80
192	2152	80
256	2664	80

Table 16. Hold Advance, Release Advance (100 Mbps)

Fragment Size	Hold Advance (ns)	Release Advance
64	10560	160
128	15680	160
192	20800	160
256	25920	160

In all cases, for guard bands to be successfully inserted, the Port link must be up, as the speed of the established link is part of the calculation to determine the guard band duration required.

Note that the guard bands do rely on the Switch operating in Store and Forward mode. When the Switch cuts through traffic, frames egress an empty port before the frame size is known, therefore a frame larger than the Max. SDU setting can egress and the start of a time slot can be delayed in this scenario.

When a schedule has been saved and committed, confirm that the schedule is accepted by viewing the **Running** page. If the **Running** page does not show what is loaded, then there is an issue with the loaded schedule.

CYCLE TIME

The next controls available in Figure 87 are related to the **Cycle Time** for scheduled traffic. The first check box is the **Cycle Time Numerator**, the second is the **Cycle Time Denominator** expressed in seconds. The ratio of the cycle time numerator and denominator must be an integer multiple of 1 ns. Values that do not result in integer multiples of 1 ns are not loaded to the device, with the schedule being rejected. If an invalid cycle time is entered, when a user clicks the **Save** button, followed by the **Commit** button, if the schedule is not accepted, the **Running** page is not updated. The default values in the web page of 1/1000 results in a cycle time of 1 ms.

BASE TIME

The **Base Time** value is the absolute time at which a new schedule is required to take effect. A new schedule takes effect at the programmed base time. If the base time is in the past, then the Switch takes the base time for the new schedule and projects forward based on the new schedule cycle times to get past the current time and apply the new schedule at the next new cycle boundary.

CYCLE TIME EXTENSION

The **Cycle Time Extension** value defines the maximum amount of time by which the old cycle for the port is permitted to be extended when Switching to a new schedule. When changing from an old schedule to a new schedule, without cycle time extension, the new cycle can result in a partial or runt cycle of the old schedule directly before the transition to the new cycle, as shown in Figure 92.

Using the **Cycle Time Extension** ensures a more seamless transition between schedules. Now instead of a partial old schedule, the last valid cycle is extended with the old schedule gate states being

SCHEDULED TRAFFIC – SCHEDULE

retained until the new schedule is implemented at the programmed base time, thereby, bridging the Switchover between the two schedules, as shown in Figure 93.




Figure 92. Schedule Switchover with no Cycle Time Extension Setting




Figure 93. Schedule Switchover with Cycle Time Extension Setting

GATE STATES WHEN DISABLED

These check boxes only apply to the gates after a schedule has been disabled.

CUT-THROUGH ALLOWED

The Switch operates in cut-through mode by default. To configure queues per port to Store and Forward mode, clear the corresponding check box, click **Save** button followed by **Commit** button. This functionality does not require a valid schedule to be running. Note that guard bands do rely on the Switch operating in Store and Forward mode. When the Switch cuts-through traffic, frames egress an empty port before the frame size is known, therefore, a frame larger than the Max. SDU setting can egress and the start of a time slot can be delayed in this scenario.

GATE CONTROL LIST, TIME INTERVALS

The Switch supports a gate control list of 32 entries per port through the driver. By default, the web page displays 8 entries, but it is possible to extend this to 32 entries per port. When the last time interval has a value entered into the **Interval [ns]** field, the web page automatically increases the number of entries displayed. The last entry always needs to be 0 ns.

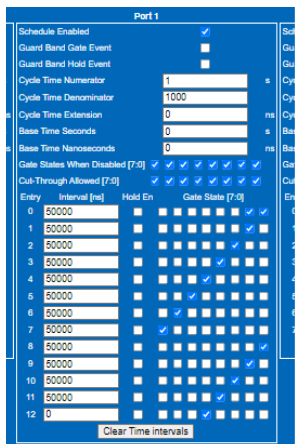


Figure 94. Adding Additional GCL Entries

Next to the time interval entries are **Gate State** check boxes for each of the 8 queues. The gate state corresponds to Queue number, with lowest priority on the right. When a box is selected, the gate for that queue is open from the completion of the last time slot for the duration specified in the entry field. If a check box is not selected the gate, for that queue is closed for that duration. Up to 8 entries (Entry 0 to 7) can be entered for the queues by default, the web server allows additional fields to be added. These entries make up a queue's **Gate Control List**. The entries are relative, meaning they are additive from the previous entry. Entry 0 is from Time = 0, so entering a value of 100000 ns means the gate control value for the first entry is from 0 μ s to 100 μ s. Entering 100000 ns in Entry 1 means that gate control value starts at 100000 ns with 100000 ns duration, so finish at 200000 ns or 200 μ s.

For queues that are checked for Entry 0, their gate opens at the start of the cycle. For queues that are checked for Entry 1, their gate opens at 100 μ s. For any entry where a queue is not checked, those queues have their gates closed at that entry duration. For example, if Q0, Q1, Q2, and Q3 are all checked and 100000 ns is entered at Entry 0, all 4 queues open at 0 μ s. The Entry 1 Gate States become active at 100 μ s. Then, if Q0 and Q1 are checked, Q2 and Q3 are unchecked, and 100000 ns is entered at Entry 1, Q0 and Q1 continue to have their gates open for another 100 μ s and Q2 and Q3 have their gates closed. And at Entry 2 closes the gates for Q0 and Q1 if their queues are not checked. This is shown in Figure 95.

Note that gPTP and LLDP messages use Queue7 by default, therefore Gate 7 must always be open for some duration of the cycle.

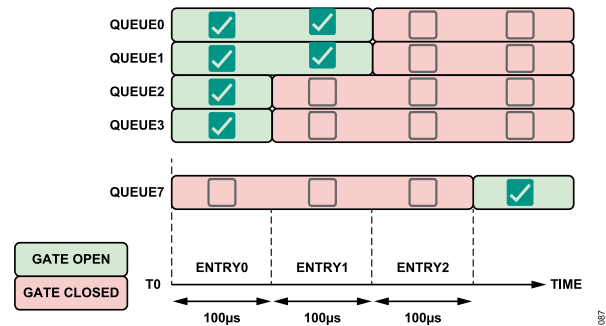


Figure 95. Simplified Schedule

Schedules with time entry intervals in excess of the programmed cycle time are accepted, but the duration and gate states in excess of the cycle time are ignored.

HOLD EN

There is an additional check box shown adjacent to the **Gate State**. This check box provides the ability to enable **Hold EN** for each entry. This feature can be enabled when Scheduled traffic and frame preemption are used in combination. When **Hold EN** is enabled for a time slot, no preemptable traffic is allowed to start egressing the port in that window.

SCHEDULED TRAFFIC – SCHEDULE

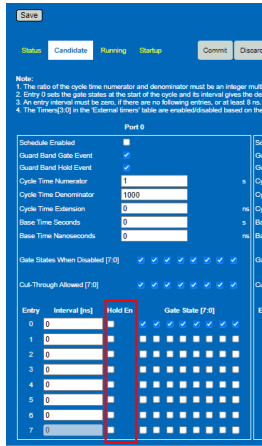


Figure 96. Hold EN Control

possible to configure a schedule for the hardware Timer pins. Once the user has a new set of values for the **Candidate** configuration, click the **Save** button, followed by **Commit** button to send the **Candidate** configuration entries to the **Running** configuration. Click **Discard** button to revert the **Candidate** configuration back to current **Running** configuration.

When committing the **Candidate** configuration to the **Running** configuration, the current **Running** configuration saves to a running backup configuration before the **Candidate** configuration saves to the **Running** configuration. The purpose of this save to the running backup configuration is to allow the user to undo the **Commit** action in the event that the committed **Candidate** configuration results in a catastrophic effect on the TSN operation.

CANDIDATE PAGE

The default page is the **Candidate** tab, see Figure 97, where user can configure the schedule for each port individually. It is also

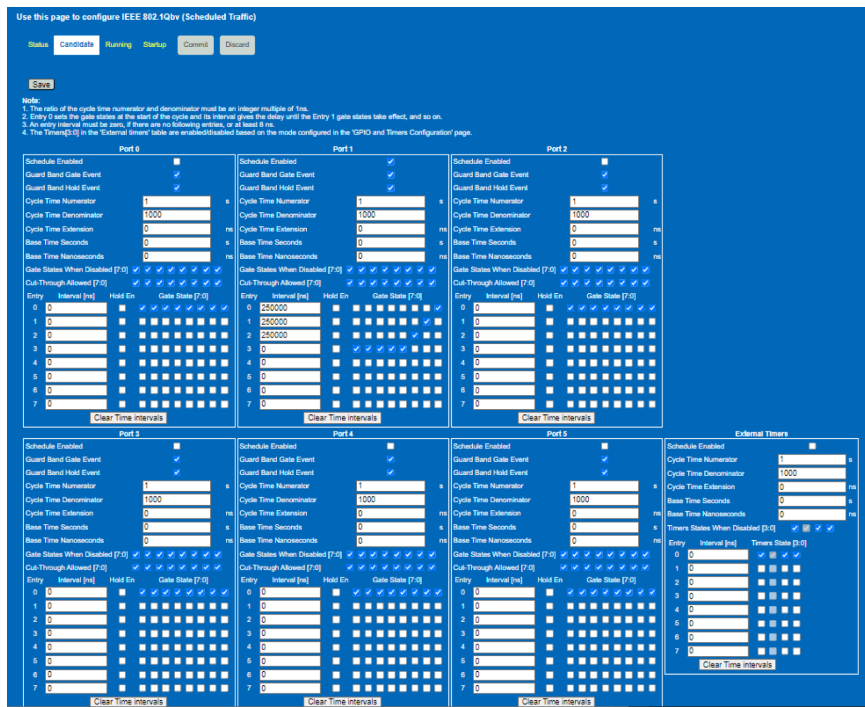


Figure 97. Scheduled Traffic – Candidate Page

SCHEDULED TRAFFIC – SCHEDULE

RUNNING PAGE

Click **Running** to display the **Running** configuration, as shown in [Figure 98](#). The fields on this page cannot be edited. Return to the **Candidate** configuration to change configuration.

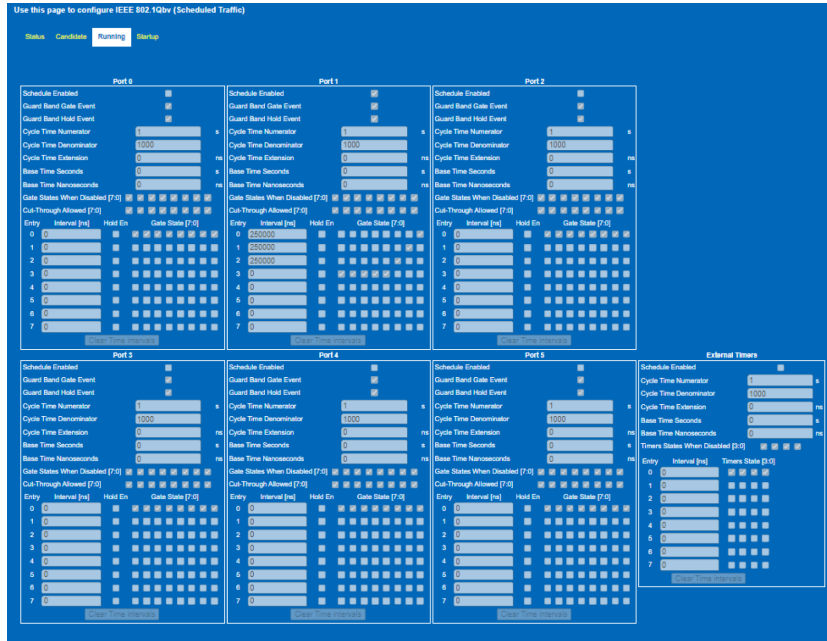


Figure 98. Scheduled Traffic – Running Page

SCHEDULED TRAFFIC – SCHEDULE

STARTUP PAGE

The **Startup** page displays the **Startup** configuration, see [Figure 99](#). These parameters are displayed to verify the values of the **Startup** configuration only.

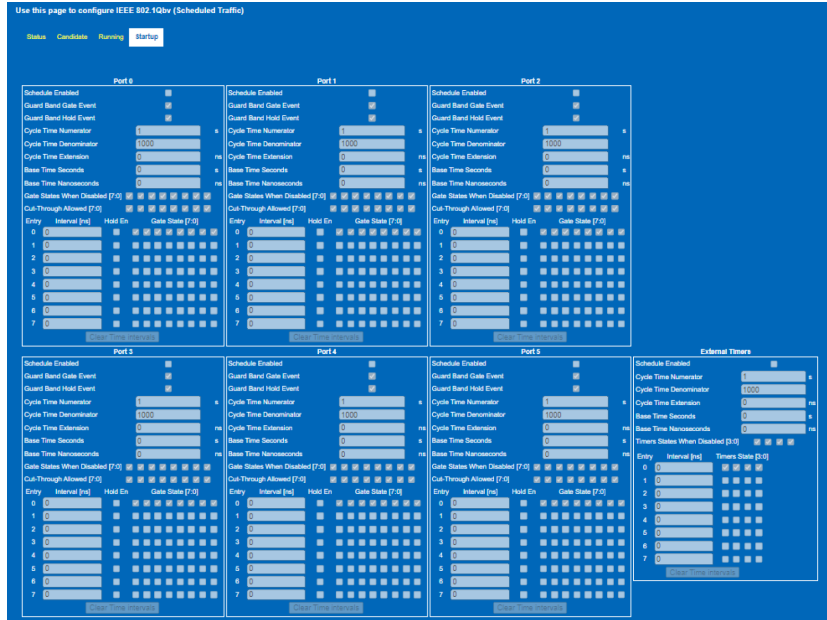


Figure 99. Scheduled Traffic – Startup Page

SCHEDULED TRAFFIC – SCHEDULE

SCHEDULE ON THE TIMER PINS

The Switch has four hardware Timer pins. A schedule can be created on all four pins. The functionality of the timer pins can be configured in the [GPIO and Timer Configuration](#) page. By default, Timer0 and Timer1 are configured to allow a schedule be created, while Timer2 is configured to provide a 1PPS signal and is shown greyed out in this page, Timer3 defaults as a Capture input. To apply a TSN schedule on Timer2 or Timer3, first change the configuration in the GPIO and Timer Configuration page. In the examples below, two different schedules have been applied to the two devices for the Timer0, Timer1, and Timer3 pins.

Figure 100 and Figure 101 show the two different scheduled for two sets of ADIN6310 Timers pins.

For Switch 1 timers, the Cycle time is 1 ms and there are four time slots. Each Timer is enabled for a window of 200 μs, starting with Timer0, followed by Timer1, then Timer3. The remaining time of the 1 ms cycle time, all timers are in the Off state. Figure 102 and Table 17 show visually how the schedule looks in terms of time.

For the second devices, Switch 2, the Cycle time is still 1 ms and there are eight time slots with a binary pattern enabled for a slot duration of 100 μs, using 700 μs of the cycle time. The remaining time of the 1 ms cycle time, all timers are in the Off state. Table 18 shows visually how the schedule looks in terms of time and Figure 103 shows a capture of the scheduled activity on the Timer pins for both devices using a logic analyzer.

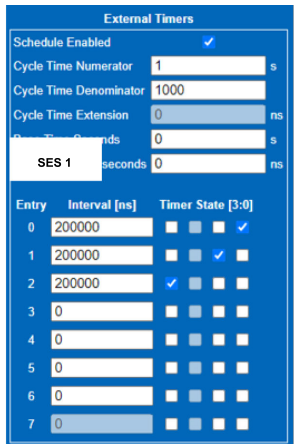


Figure 100. Configured Schedule for Switch 1 Timers

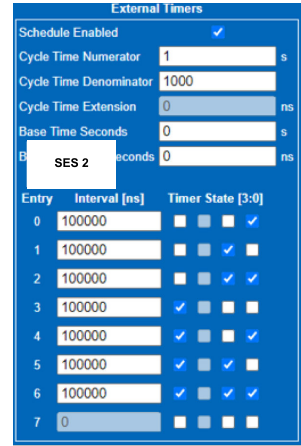


Figure 101. Configured Schedule for Switch 2 Timers

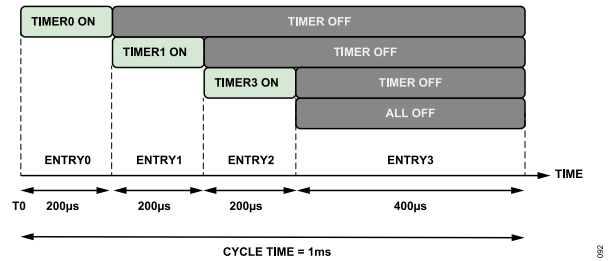


Figure 102. Schedule for Switch 1 Timer Pins

Table 17. Switch 1 Schedule (1 ms Cycle Time)

Entry	Start time			Timer3	Timer1	Timer0
	(μs)	End Time (μs)	(μs)			
0	0	200	0	0	1	
1	200	400	0	1	0	
2	400	600	1	0	0	
3	600	Remainder	0	0	0	

Table 18. Switch 2 Schedule (1 ms Cycle Time)

Entry	Start time			Timer3	Timer1	Timer0
	(μs)	End Time (μs)	(μs)			
0	0	100	0	0	1	
1	100	200	0	1	0	
2	200	300	0	1	1	
3	300	400	1	0	0	
4	400	500	1	0	1	
5	500	600	1	1	0	
6	600	700	1	1	1	
7	700	Remainder	0	0	0	

SCHEDULED TRAFFIC – SCHEDULE

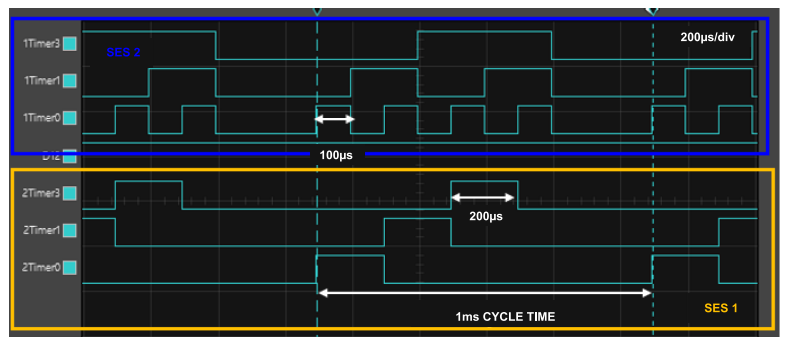


Figure 103. Logical Analyzer View of a Schedule on the Hardware Timer Pins

LLDP CONFIGURATION

LINK LAYER DISCOVERY PROTOCOL (LLDP)

LLDP is a protocol devices use to advertise information about their capabilities between peers. The configuration exposed in the web server is limited to basic configuration and visibility into statistics. Additional capability is exposed in the TSN driver library APIs, for more details, refer to the [ADIN6310 Hardware Reference Manual](#).

LLDP CANDIDATE VIEW

The LLDP stack runs on the Packet Assist engine of the Switch and is enabled during initial configuration of the device from the GUI application when using TSN functionality or HSR functionality (not currently enabled with PRP operation). The default configuration is shown in the **Candidate** page in [Figure 104](#). Configure the required changes, click the **Save** button followed by **Commit** button to load changes to the Switch. The admin configuration included in the web server are as follows:

- ▶ **Admin status:** Choice of **Tx and Rx**, **Tx only**, **Rx only**, or **disabled**.
- ▶ **Message Fast Tx:** Time intervals (in ticks) between transmissions during fast transmission periods. Default is 1, range of 1 to 3600. Fast transmission periods are initiated when a new neighbor is detected and results in LLDP packets to be transmitted on a shorter time interval than the normal message Tx interval.
- ▶ **Message Tx Hold Multiplier:** Used as a multiplier of msgTxInterval to determine the value of txTTL (Time to Live), $txTTL =$

$((\text{Message Tx Interval} \times \text{Message Tx Hold}) + 1)$. Default is 4, intended range is 1 to 100, but web page currently limits field to 2 to 10, this needs to be addressed in future release.

- ▶ **Message Tx Interval:** Time interval in ticks between transmission during normal transmission periods. Default is 30, range of 1 to 3600.
- ▶ **Reinit Delay:** Amount of delay from when **Admin Status** becomes **disabled** until reinitialization is attempted. Default value is 2 seconds.
- ▶ **Tx Credit Max:** TxCredit is the number of consecutive LLDPDUs that can be transmitted at any time. The parameter is the maximum value of txCredit. Default is 5, range of 1 to 10.
- ▶ **Tx Fast Init:** Used as the initial value for txFast. Default is 4, range of 1 to 8.
- ▶ **Number of peer supported:** Per port number of peers supported.
- ▶ **Enable end of LLDPDU TLV:** Enable or disable end of LLDPDU TLV in Tx LLDP frames, which marks the end of the LLDPDU frame.
- ▶ **Override MAC address:** The default MAC address for the LLDP stack is derived from the Port MAC address. Overriding the MAC address changes the source MAC address, PortID and/or ChassisID in the LLDP frames egressing from the given port.

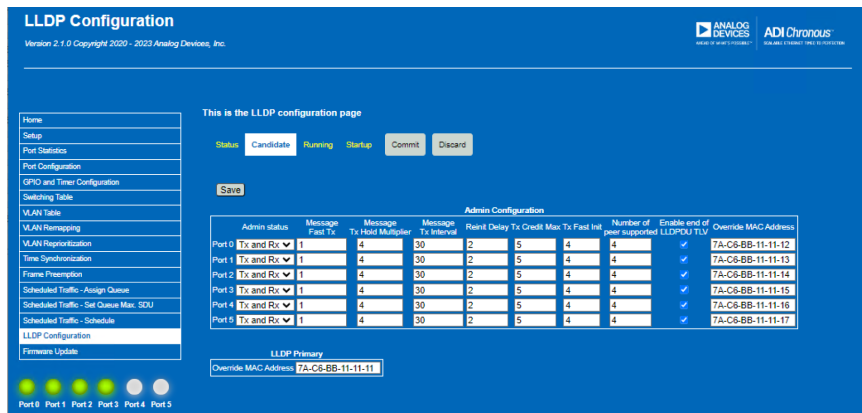


Figure 104. LLDP Candidate Page

LLDP CONFIGURATION

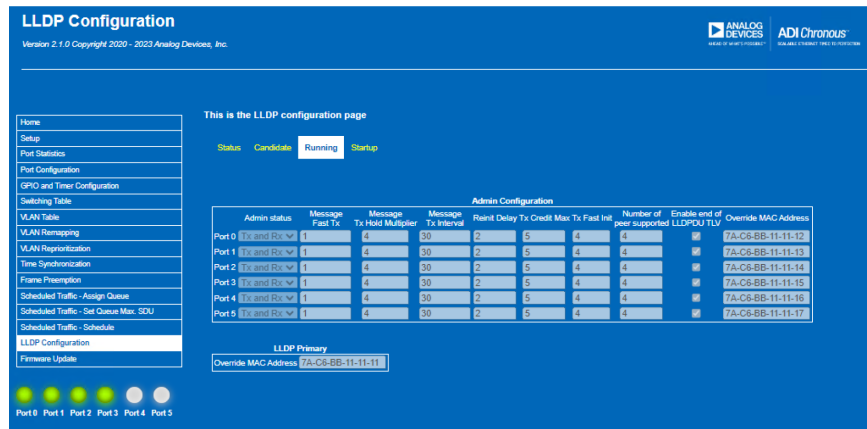


Figure 105. LLDP Running Page

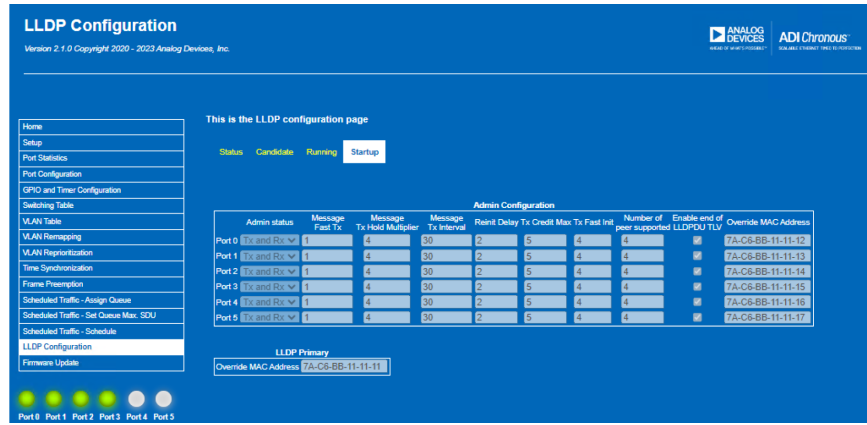


Figure 106. LLDP Startup Page

LLDP CONFIGURATION

LLDP STATUS

The **Status** view shows an overview of the Remote, Local, and Port based statistics for the LLDP feature. This includes a capture of the

LLDP frames transmitted and received and additional information related to error scenarios, ageouts, inserts, and deletes.

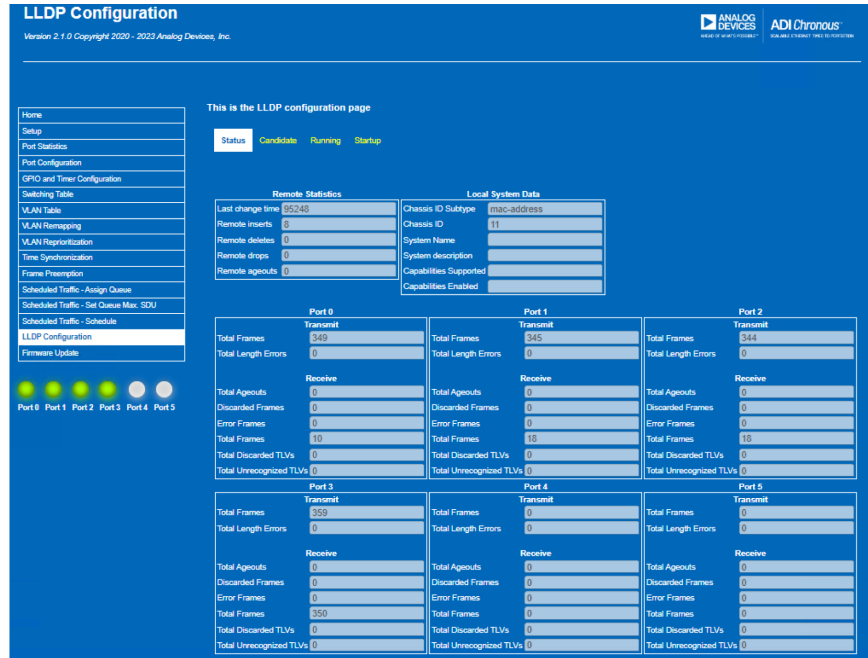


Figure 107. LLDP Status Page

LLDP CONFIGURATION

LLDP EXAMPLE

Figure 108 shows a Wireshark capture of the LLDP messages exchanged between two ADIN6310 devices (Switch 1 - Port 3 to Switch 2 - Port 0). The messages are targeted at the LLDP multicast address 01:80:c2:00:00:0e, and originate with a source MAC of the Switch Port MAC address. The LLDP protocol message contents can be observed in the capture, with information describ-

ing the Chassis Subtype, Port Subtype, Time to live ((message Tx Hold x message Tx Interval) + 1 = (4 x 30) + 1 = 121), and additional Ethernet capabilities. The Switch uses LLDP to exchange capability for Frame Preemption with its peer.

The LLDP messages can be observed every 30 seconds (Message Tx Interval).

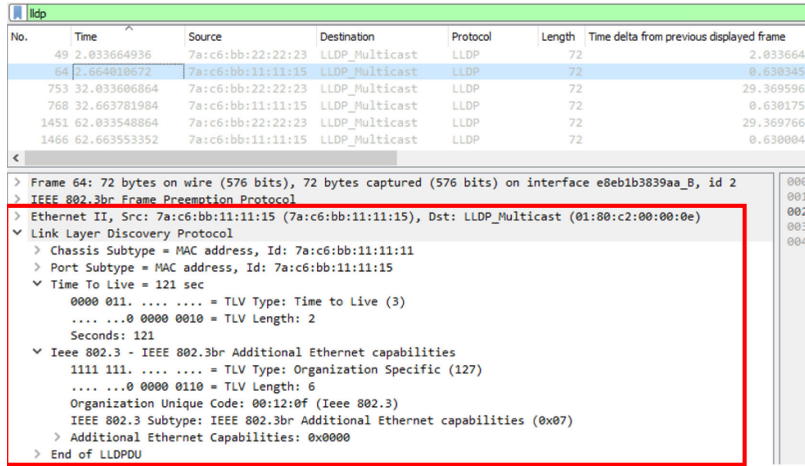


Figure 108. Wireshark Capture of LLDP Frames between Two ADIN6310 Devices (Default Configuration)

LLDP CONFIGURATION

LLDP EXAMPLE (FAST TX)

Fast transmission periods are initiated when a new neighbor is detected and results in LLDP packets to be transmitted on a shorter time interval than the normal message Tx interval. The default setting for Message Fast Tx is 1 second. As shown in [Figure 109](#), LLDP is disabled on Port 3 after time 12 seconds and then

reenabled after approx 73 seconds. At that time, both SES 1 and SES 2 start transmitting fast Tx messages at a 1 second interval before returning to the normal Tx interval of 30 seconds. They each send 5 LLDP messages as the default value of the Maximum Tx Credit parameter is set to 5.

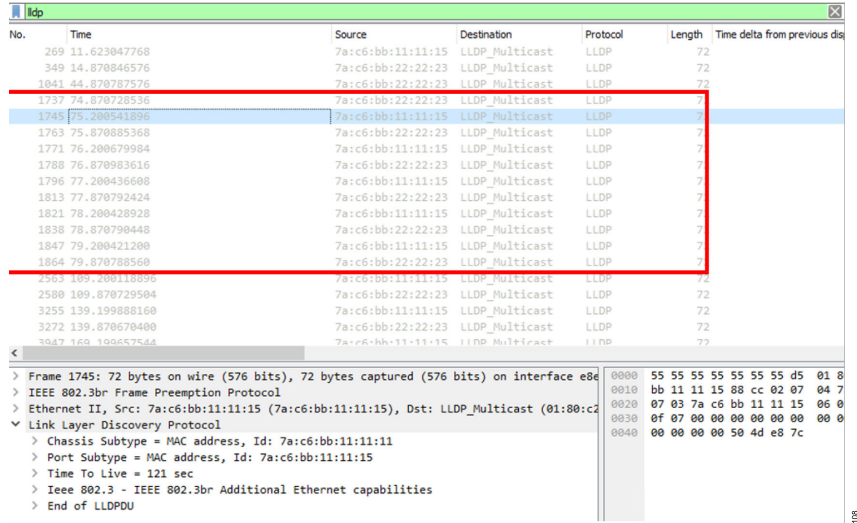


Figure 109. Wireshark Capture of LLDP Frames between Two ADIN6310 Devices when LLDP is Reenabled on Switch 1

PARALLEL REDUNDANCY PROTOCOL (PRP)

The Switch hardware supports PRP per IEC62439-3 (2021 ed4) standard. The capability exposed in the Switch is ability to support one instance of a doubly attached node obeying PRP (DANP) or Redundancy Box (Redbox) function configured over either an SPI or Ethernet connected Host (web server only supports configuration over Ethernet Host). The Host configures the Switch for the PRP function, defining which ports are PRP network ports, sets the link redundancy entity (LRE) MAC address and enables the PRP function.

The Switch hardware takes care of duplicating the outgoing traffic onto LAN A/B and inserting the RCT tag to the end of the frame. On reception of PRP traffic, the Switch consumes the first frame, removes the tag, and discards duplicates. The Switch generates supervision frames, which are sent out on the LAN A/B ports and maintain a nodes table of other PRP DANP, Redbox, and SAN entities in the network. PRP supervisory frames are generated periodically with or without VLAN tag every 2 seconds. The device maintains a nodes table, recording the last time a frame is received from a node. Node entries are removed from the table if no frames are received for over 60 seconds. The node table can support 1024 entries max. The Switch supports operation of one instance of PRP on the 6-port ADIN6310 Switch, multiple instances of PRP running on the 6-port device is not supported.

The Switch PRP functionality can be configured to support:

- ▶ PRP operation as a DANP or PRP Redbox.
- ▶ PRP configured over Ethernet or SPI Host.

Enabling LLDP and VLAN functionality with PRP is supported. Enabling PTP functionality with PRP is not currently supported, future software updates include this capability. Using TSN features such as Scheduled Traffic, Frame Preemption, Per Stream Filtering, and Policing or Frame Replication and Elimination for Reliability with PRP is not supported.

Figure 110 shows a simple configuration of the Switch with Ethernet Host (Port C) configured as a DANP and connected over a PRP network (LAN A/B) to another PRP capable Switch. The duplicate network, LAN A/B, provides the redundant path ensuring seamless redundancy. Three Ethernet ports are used in a PRP DANP device, Port A, Port B are network facing ports, while Port C is connected to the Host/End node over Ethernet interface and is used for control plane configuration of the Switch and PRP data plane traffic to that node. PRP Port C can also be connected over SPI interface.




Figure 110. Example of Switch Configuration as PRP-DANP (Host Connected over Ethernet)

PARALLEL REDUNDANCY PROTOCOL (PRP)

ENABLING PRP EXAMPLE

To enable PRP, pass an XML file that includes PRP configuration to the ses-configuration.txt file as PRP needs to be enabled as part of the initial configuration. The XML file must include all relevant PRP configuration, how PRP is configured, which ports are PRP Port A, Port B, and Port C for DANP configuration and interlink ports if configuring the device for PRP Redbox. The LRE MAC address must be the Host MAC address (or if connected over Ethernet to a PC Network adapter, it must be the MAC of the NIC). In the configuration shown in Figure 111, Port 0, Port 1, and Port 2 are the LRE Port C, Port A, Port B, and Port 3 to Port 5 are configured as Interlink ports. For details on the XML configuration for PRP, see the PRP Specific Configuration section.




Figure 111. PRP Configuration as a Redbox

After the configuration file has been edited for the relevant PRP configuration, launch the GUI, click the **Find and Configure** to search for a connected Switch. Once the Switch has been configured and the GUI LED turns green, the web server can be opened. When using PRP functionality, the web server shows the features supported with PRP, which is a reduced feature vs. when operating in TSN mode, see Figure 112.

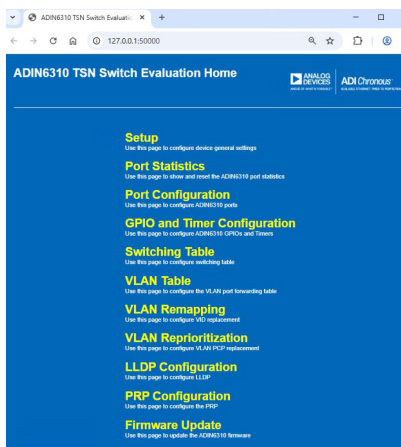


Figure 112. Web Page when PRP Function is Enabled

PRP CONFIGURATION WEB PAGE VIEWS

PRP Candidate View

The default PRP configuration is based on the XML configuration as discussed in the ses-configuration File section. There are addi-

tionally some run-time configurable parameters for PRP, as shown in the **Candidate** view, see Figure 113. To change the PRP configuration during operation, make the required changes, click the **Save** button followed by **Commit** button to load them to the device. The configuration included are as follows:

- ▶ **Redundancy Device:** Shows the type of redundancy device configured as per the XML file. Choice of **PRP DANP** or **PRP Redbox**.
- ▶ **Duplicate mode:** The Switch receiving hardware can detect duplicates based on information in the RCT tag in the frame. When Duplicate Discard mode is enabled, it only forwards the first frame of a pair to the upper layers. Duplicate Accept is typically only used for testing purposes and allows the Switch to forward both duplicate frames to upper layers. The default mode is for Duplicate Discard. In the event a frame is received with the wrong LANID (ID 0xA on Port B or ID 0xB on port A), the Switch performs a Duplicate Discard and strip off the PRP RCT trailer, this applies to DANP and PRP Redbox use cases.
- ▶ **Port-A Admin State:** Shows if the port is active or not, choice of **On** or **Off**, default is **On**.
- ▶ **LRE MAC Address:** Shows the LRE MAC address as configured through the XML file.
- ▶ **Max Reside Time:** Sets the maximum time an entry may reside in the duplicate list. The default is 10 ms (15 μs × 625). The range of possible values is 15 μs to 400 ms (corresponding to 0 to 26214).
- ▶ **Evaluate Supervision:** By default, the Switch evaluates supervision frames in the network and add nodes to its nodes table. This can be disabled by clearing the check box.
- ▶ **Transparent Reception:** By default, the Switch removes the RCT tag from the frame before it passes it to upper layers. Select **Pass** to leave the PRP RCT tag in the frame.
- ▶ **Supervision VLAN ID (0-4095):** By default the supervision frames are sent untagged (VLAN 4095). To send supervision frames with a VLAN tag, enter a valid VLAN ID and PCP (priority code point) in each field. The supervision frames can be sent with different VLAN tag on each LAN.

Reside Time: Consideration needs to be given to the programmed duplicate list reside time (IreDupListResideMaxTime). The switch maintains a list in the forwarding table (max 2000 entries) for the duplicate discard algorithm. It stores the sequence number and source MAC for the duration of the programmed IreDupListResideMaxTime. The value programmed indicates the maximum time an entry can reside in the LRE duplicate list, expressed as a number of seconds divided by 65536. The standard defines a default time of 400 ms (corresponding to 26214), the programmable range is 1 to 26214 (15 μs to 400 ms). The age out time is based on a calculation using the programmed reside time. The minimum age out time is 1 ms (IreDupListResideMaxTime = 33 to 98).

```
SECOND_FRACTION_NS = 15259;
ONE_HALF_MS = 500000;
```

PARALLEL REDUNDANCY PROTOCOL (PRP)

```
ONE_MS = 1000000;
ageOutMs = (uint32_t)((ResideMaxTime * SEC
OND_FRACTION_NS) + ONE_HALF_MS) / ONE_MS);
```

When operating at Gigabit speed, with minimum frame size (64 bytes + preamble + IFG = 90 bytes or 720 ns per frame), the table

fills in $2000 \times 720 \text{ ns} = 1.44 \text{ ms}$. For medium sized frames, for example 326 bytes (including Ethernet and PRP overhead), the table fills in $326 \times 8 \text{ ns} = 2608 \text{ ns}$ per frame and 2000 frames in 5.2 ms.

Larger frame sizes of 1532 fills the table in approximately 24 ms.

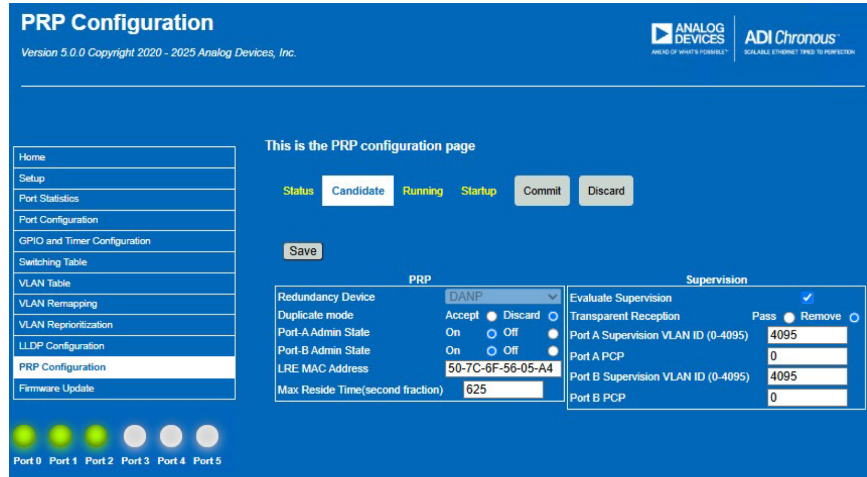


Figure 113. PRP Candidate Page

PARALLEL REDUNDANCY PROTOCOL (PRP)

PRP Running View

PRP **Running** page shows the configuration loaded to the device. No changes can be made in the **Running** page. See [Figure 114](#).

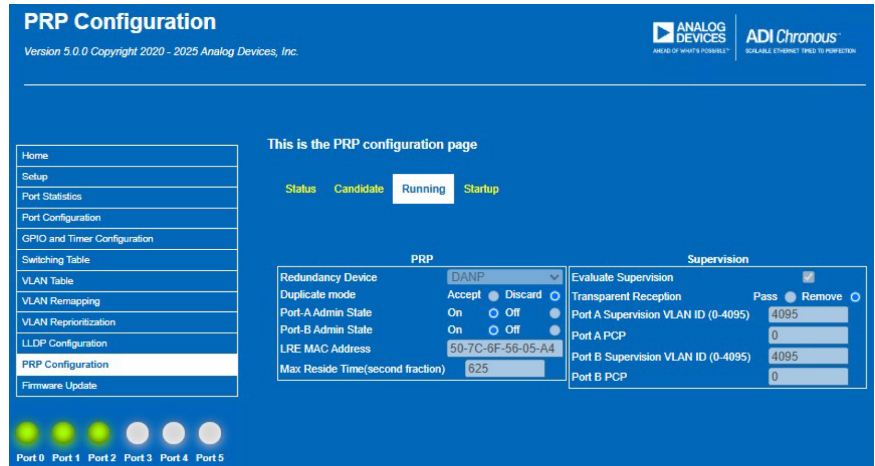


Figure 114. PRP Running Page

PARALLEL REDUNDANCY PROTOCOL (PRP)

PRP Startup View

PRP **Startup** page shows the **Startup** configuration. See [Figure 115](#).

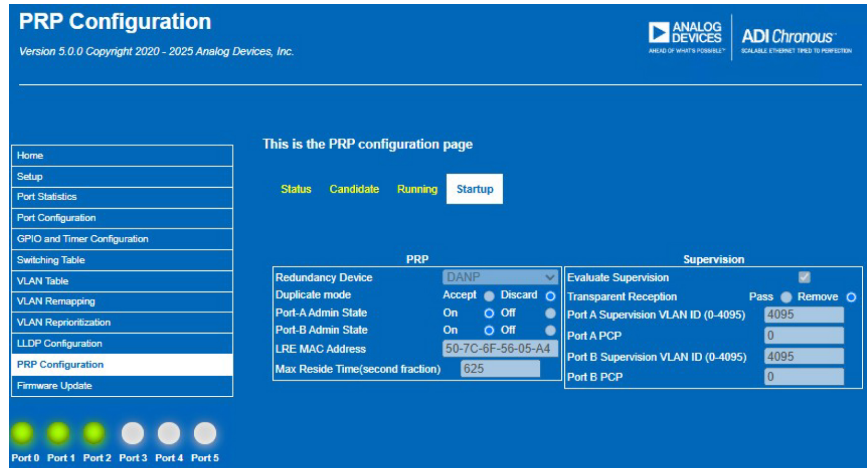


Figure 115. PRP Startup Page

PARALLEL REDUNDANCY PROTOCOL (PRP)

PRP STATUS PAGE

The **Status** page shows the **LRE Statistics**, the **Node Table Statistics**, and the **Proxy Node Table**. See [Figure 116](#).

The **LRE Statistics** section shows the traffic statistics and error counters associated with each PRP LAN that have been observed by the Switch in addition to showing how many nodes are in the network. The configuration included are as follows:

- ▶ **Rx Count:** Shows the number of frames received by Port A or Port B that have PRP RCT trailers added.
- ▶ **Tx Count:** Shows the number of frames transmitted by Port A or Port B that have PRP RCT trailers added.
- ▶ **Error Count:** Shows the number of frames with errors received on the LRE Port A or Port B.
- ▶ **Wrong LAN error count:** Shows the number of frames with the wrong LAN identifier received on LRE Port A or Port B.
- ▶ **Duplicate Count:** Shows the number of entries in the duplicate detection mechanism on Port A or Port B for which one single duplicate is received.

- ▶ **Multi Count:** Shows the number of entries in the duplicate detection mechanism on Port A or Port B for which more than one duplicate is received.
- ▶ **Unique Count:** Shows the number of entries in the duplicate detection mechanism on Port A or Port B for which no duplicate is received.
- ▶ **Node count:** Returns the number of nodes detected in the system.
- ▶ **Proxy Node count:** Returns the number of proxy nodes connected to a Redbox device. Not used for DANP devices.

[Figure 116](#) reports the node and proxy node entries for the hardware configuration shown in [Figure 117](#), where both devices are configured as Redboxes with at least one Interlink port. A switch's Proxy node table contains entries on behalf of the SAN devices connected to it. The node table for each switch is built up of MAC addresses from the supervision frames it receives from LAN A/B network.

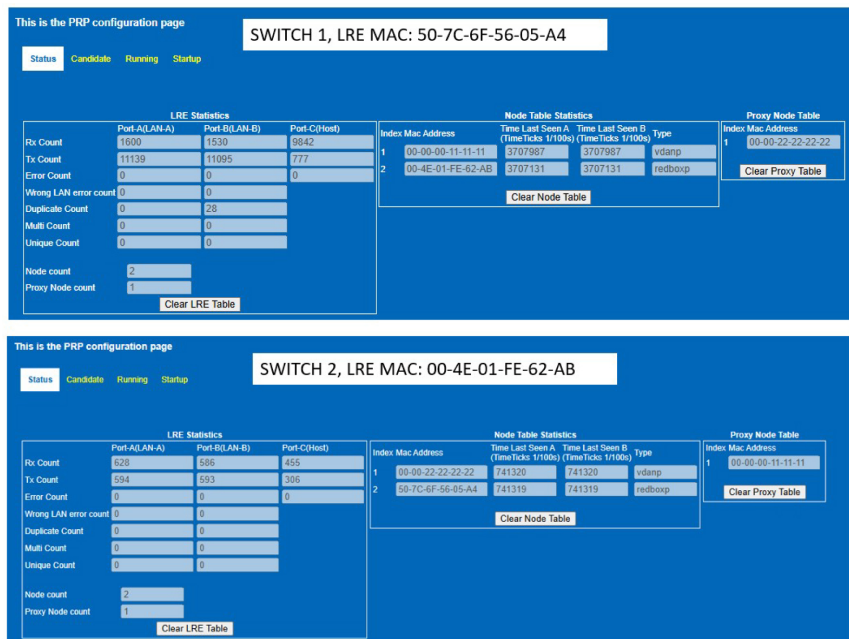


Figure 116. PRP Status Page

PARALLEL REDUNDANCY PROTOCOL (PRP)




Figure 117. PRP Configuration Example with Two Redboxes and Proxy Nodes

Node Table Statistics

The **Node Table Statistics** shows the MAC addresses of other PRP devices in the network. The nodes table is built up of entries based on Supervision frames received by the Switch from other PRP devices. The nodes table can hold 1024 entries. The Switch also reports the time (in TimeTicks 1/100 seconds) a node is last seen on either Port A or Port B and what type of device it is. The Host can access this information to understand what is happening and whether there are any issues in the network. The nodes table refreshes every 60 seconds, therefore, node entries remove from the table if traffic from that address is no longer seen.

Proxy Node Table

The **Proxy Node Table** captures information when the Switch is configured as a PRP Redbox. A Switch **Proxy Node Table** shows the LRE MAC addresses of the detected SAN devices connected to its Interlink ports. The **Proxy Node Table** can support up to 8 entries. The entries are learned based on the traffic coming into the interlink port. The table refreshes every 60 seconds, therefore, node entries are removed from the table if traffic from that address is no longer seen.

PRP – SUPERVISION FRAMES

The Wireshark capture shown in [Figure 118](#) is a supervision frame generated by the Switch configured as a DANP and transmitted

on Port B (LAN B). By default, supervisory frames are transmitted by each device at a 2 second interval and without VLAN tags. The Switch sends a supervisory frame for its LRE MAC address and on behalf of any proxy nodes connected to its interlink ports if configured as a Redbox. The supervision frame contents include information about the type of PRP device as shown in the Redbox supervision frame in [Figure 119](#). Supervisory frames are sent out to the PRP network, therefore, only visible on Port A and Port B. The PRP RCT tag has a suffix 0x88fb. This tag and the other contents such as sequence number and LSDU size can be seen at the end of the frame.

[Figure 120](#) shows a supervision frame which has a VLAN tag with ID 10 and priority 5 added. The supervision frame is sent out on the network even if the port is not configured for this VLAN. However, to receive by another switch, those switch ports must belong to the VLAN ID of the supervision frame. Otherwise, this VLAN tagged supervision frame is not received.

The Redbox device also sends supervision frames on behalf of any Proxy nodes connected to it. [Figure 121](#) shows a supervision frame from the Redbox on behalf of a device with MAC address 00:00:22:22:22:22 connected to one of its ports. This MAC address will also be added to the Proxy Node table.

PARALLEL REDUNDANCY PROTOCOL (PRP)

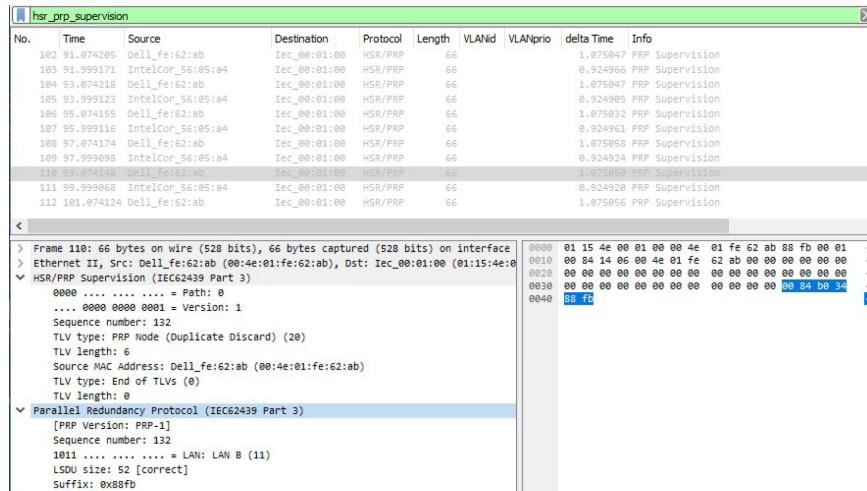


Figure 118. Wireshark Capture of DANP Supervision Frames in LAN B

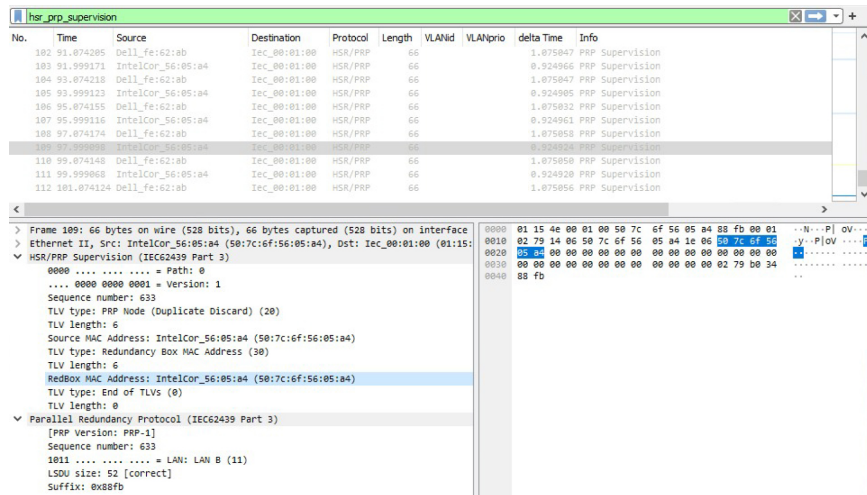


Figure 119. Wireshark Capture of Redbox Supervision Frame in LAN B

PARALLEL REDUNDANCY PROTOCOL (PRP)

PRP – CAPTURE OF PRP TAGGED TRAFFIC

The Wireshark capture shown in [Figure 122](#) is traffic sent into Port C of the Switch and observed on Port B of the PRP network. The

PRP RCT tag can be seen at the end of the frame, with LAN information, SDU size, and Sequence number.

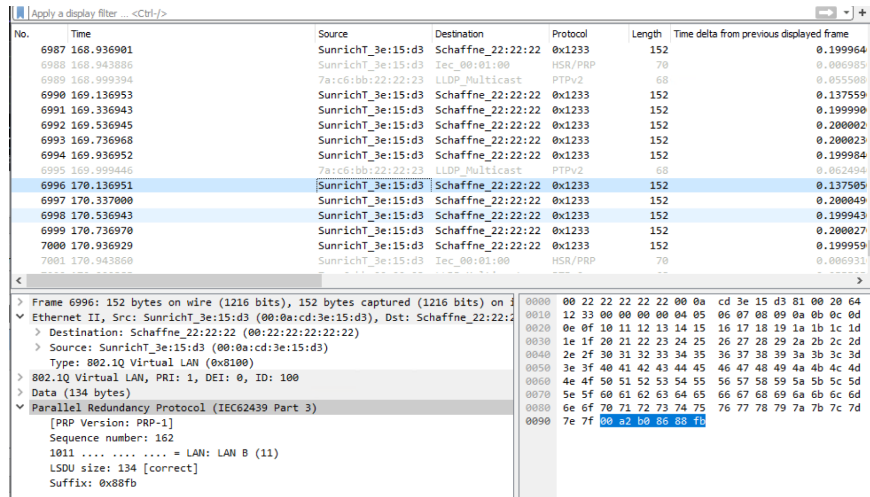


Figure 122. Wireshark Capture of PRP Tagged Frames in LAN B

PARALLEL REDUNDANCY PROTOCOL (PRP)

VLAN TABLE OPERATION IN PRP MODE

VLAN configuration is supported in PRP mode. The web server VLAN pages can be used to configure VLAN functionality as described in [VLAN Table](#) section. In PRP mode, the default VLAN behavior is for forwarding on VLAN ID 0x0 and untagged (0xFFFF).

SWITCHING TABLE IN PRP MODE

Dynamic Table, Learning Operation

Normal learning is disabled when in PRP mode.

Static Table Entries

Entries can be placed into the static table in the usual way and used to direct traffic from the DANP/End Node Host to a SAN on one of the LANs. For the ADIN6310, static entries can be used to route traffic from the Host to ports not involved in PRP network or from the other ports to SANs on the network, as shown in [Figure 123](#) and [Figure 124](#). This traffic is not duplicated, and do not have PRP RCT tags added and only egress on the port(s) defined by the installed table entry.

By default, broadcast entries do not forward in the PRP device, therefore, user needs to install a broadcast entry in the Switching table to support broadcast frames crossing from Port C to Port A/Port B. This is required to ping across a PRP device.

Extended Table Entries

In PRP mode, extended table is available and entries can be installed similar to the static table.




Figure 123. Host Routing to SAN on One of the LANs

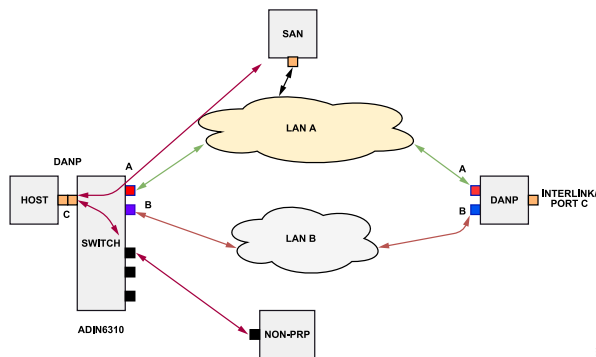


Figure 124. Host Routing to SAN Device on One of the LANs or to Another Port

HIGH AVAILABILITY SEAMLESS REDUNDANCY (HSR)

HSR is a ring protocol that provides seamless fail-over in event of a single failure in the network. The Switch supports being configured as a doubly attached node obeying HSR (DANH) or HSR Redundancy box (Redbox). Following initial device configuration, the Host configures the Switch for the required HSR function. In the case of a DANH, Port A/Port B/Port C are defined. In the case of a RedBox, Port A/Port B/Port C are defined as well as any other Interlink ports used to bridge singly attached node (SAN) devices to the HSR ring. The Host sets the link redundancy entity (LRE) MAC address (same as Host MAC address) and enables the HSR function.

Once configured for the HSR mode, the Switch hardware takes care of HSR functionality, duplicating the outgoing traffic onto each of its ring ports with the HSR tag inserted into the frame. On receipt of HSR frames from the ring, the receiving device consumes the first frame, removes the tag on reception and discards duplicates. The Switch generates supervisory frames and maintains a nodes table that lists other HSR entities in the network based on the supervision frames it received from the ring ports. The HSR supervisory frames are generated periodically with or without VLAN tag every 2 seconds. The hardware records the last time a frame is received from a node, refreshing the nodes table. Each device in the HSR ring maintains its own nodes table. Node entries are removed from the table based on the NodeForgetTime default of 1 minute. The node table is currently capable of supporting up to 1024 entries. The nodes table records entries for DANH, RedBox, and VDAN devices connected to the ring, based on the supervision frames circulating the ring.

When operating as a RedBox, the Switch maintains a **Proxy Node Table** in addition to the Nodes table. The **Proxy Node Table** is a list of the detected SANs that are connected to the RedBox and the last time they are seen. The **Proxy Node Table** learns the SAN/VDAN MAC based on ingressing traffic on ports configured as Interlink ports. Like the nodes table, the **Proxy Node Table** keeps its table refreshed based on incoming frames and ages out entries after 60 seconds. The maximum size of the **Proxy Node Table** for HSR Redbox is 8. The Switch supports operation of one instance of HSR on the 6-port ADIN6310 Switch, multiple instances of HSR running on the 6-port device is not supported.

The Switch HSR functionality can be configured to support:

- ▶ HSR as DANH
- ▶ HSR as RedBoxSAN
- ▶ HSR with LLDP and VLAN Table

Enabling PTP functionality with HSR is not currently supported, future software updates include this capability. Using TSN features such as Scheduled Traffic, Frame Preemption, Per Stream Filtering,

and Policing or Frame Replication and Elimination for Reliability with HSR is not supported.

HSR OPERATING MODES

The Switch supports the various HSR Modes. Mode H is the default operating mode.

- ▶ **MIB_PRP_HSR_modeH**: Default mode, the DANH inserts the HSR tag on behalf of its Host and forwards the ring traffic, except for frames sent by the node itself, duplicate frames, and frames for which the node is the unique destination.
- ▶ **MIB_PRP_HSR_modeN**: No forwarding, node behaves as Mode H with the exception that it does not forward ring traffic from Port to Port.
- ▶ **MIB_PRP_HSR_modeT**: Transparent forwarding, removes the HSR tag before forwarding the frame to the other Port and sends a frame from the Host to both Ports, untagged, and without discarding duplicates.
- ▶ **MIB_PRP_HSR_modeU**: Unicast forwarding, the node behaves as in Mode H, except that it also forwards traffic for which it is the unique destination.
- ▶ **MIB_PRP_HSR_modeM**: Mixed forwarding, the DANH inserts the HSR tag depending on local criteria when injecting frames into the ring.
- ▶ **MIB_PRP_HSR_modeX**: No sending on counter-duplicate, node behaves as in Mode H, except that a Port does not send a frame that is a duplicate of a frame that it received completely and correctly from the opposite direction.

ENABLING HSR EXAMPLE

HSR needs to be enabled as part of the initial configuration. To enable HSR, pass an XML file that includes HSR configuration to the ses-configuration.txt file. The XML file must include all relevant HSR configuration. This entails the HSR mode (DANH/RedBox), the identification of which ports are Port A/Port B/Port C and the identification of any interlink ports connected if a RedBox. The LRE MAC address must be the Host MAC address. In case of connection over Ethernet to a PC Network adapter, the LRE MAC address must be the MAC of the network interface controller (NIC). In this example, the LRE MAC is shown as 00:11:11:11:11 and matches the Host MAC, as shown in the RedBox in [Figure 125](#). For more details, see the [HSR Specific Configuration](#) section.

After the configuration file has been edited for the relevant HSR configuration, launch the GUI and configure the Switch by clicking **Find and Configure**. Once the device has been configured and the GUI LED turns green, the web server can be opened and shows a reduced feature set, as shown in [Figure 126](#).

HIGH AVAILABILITY SEAMLESS REDUNDANCY (HSR)

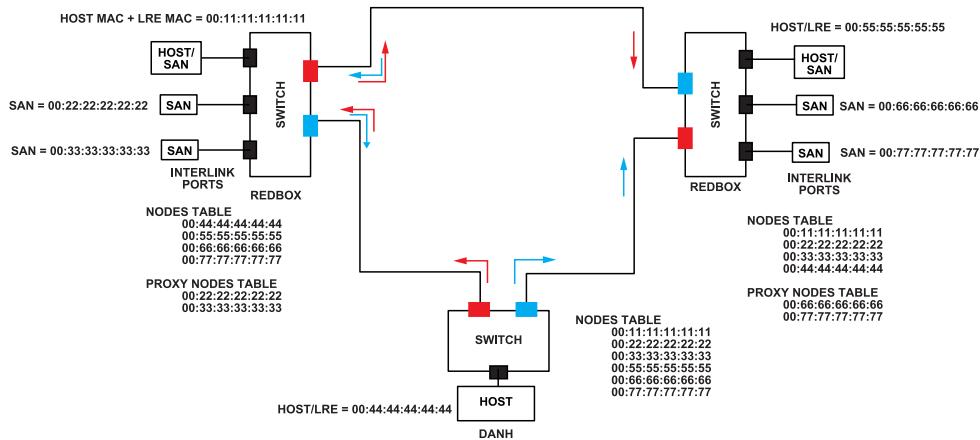


Figure 125. HSR Configuration with Two HSR Redbox Devices and One DANH




Figure 126. Web Page when HSR Function is Enabled

HSR CANDIDATE VIEW

The default HSR configuration originates from the initial configuration passed from the XML file. In addition, there are some run-time configurable parameters, as shown in the **Candidate** view, see Figure 127. Make the required changes, click the **Save** button followed by **Commit** button to load them to the device. In the **Mode Transition** field, **Mode H** is the default mode.

HSR Configuration

The **HSR Configuration** included as follows:

- ▶ **Cut-Through**: By default, the Switch is configured for Store and Forward mode. Note that when cut-through operation is enabled for HSR, the device applies it to all Ports. However, Interlink Ports must always operate Store and Forward as the HSR device needs to understand frame size to correctly calculate the LSDU size. Cut-through operation must not be applied to Interlink Ports and result in frames sent from Interlink Ports to

Ring Ports with the wrong LSDU size. For now, the Switch must only be used in Store and Forward mode when HSR is enabled. This is addressed in future software updates.

- ▶ **Port-A Admin State**: Shows if the Port is active or not, choice of **On** or **Off**, default is **On**.
- ▶ **LRE MAC Address**: Shows the MAC address as set in the XML configuration file.
- ▶ **Max Reside Time**: Sets the maximum duration for which an entry may reside in the duplicate list. The default is 625 expressed in second fraction, which corresponds to 15µs × 625. The range of possible values is 0 to 26214 (corresponding to 15 µs to 400 ms).
- ▶ **Evaluate Supervision**: By default, the Switch evaluates supervision frames in the network and add nodes to its nodes table. This can be disabled by clearing the check box.
- ▶ **Supervision VLAN ID (0-4095)**: By default, the supervision frames are sent untagged (VLAN 4095). To send supervision frames with a VLAN tag, enter a valid VLAN ID in this field.

HIGH AVAILABILITY SEAMLESS REDUNDANCY (HSR)

► **Supervision Address:** The supervision frames have a multicast destination MAC address 01:15:4E:00:01:xx. The last byte is

programmable. The default is 0x00, but can be configured to use any value between 0x00 and 0xFF.

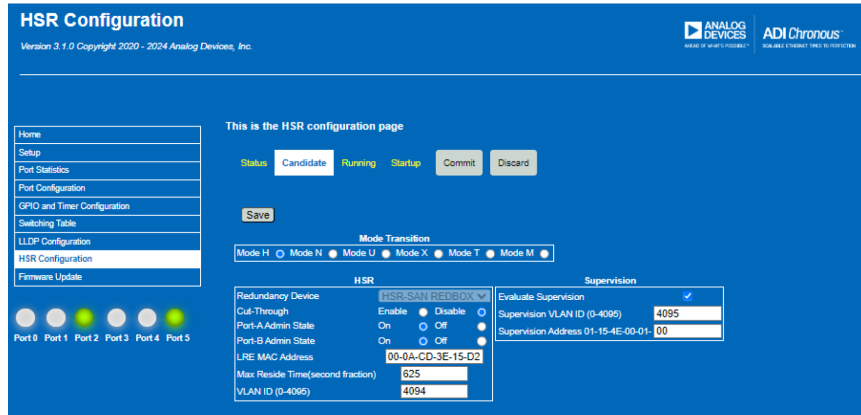


Figure 127. HSR Candidate View

HIGH AVAILABILITY SEAMLESS REDUNDANCY (HSR)

HSR STATUS VIEW

The **Status** page shows the **LRE Statistics** and the **Node Table Statistics**. See [Figure 128](#).

The **Status** page shows the **LRE Statistics**, which shows the traffic statistics and error counters associated with each HSR port. The **Node Table Statistics** provides a list of the other HSR entities in the network, in addition to how many nodes in the network. When the device is configured as a Redbox, the **Proxy Node Table** provides a list of SANs connected to the device Interlink Ports.

The **LRE Statistics** fields are as follows:

- ▶ **Rx Count:** Shows the number of frames received by Port A or Port B that have a HSR tag.
- ▶ **Tx Count:** Shows the number of frames transmitted by Port A or Port B that have a HSR tag inserted.

- ▶ **Error Count:** Shows the number of frames with errors received on the LRE Port A or Port B.
- ▶ **Duplicate Count:** Shows the number of entries in the duplicate detection mechanism on Port A or Port B for which one single duplicate is received.
- ▶ **Multi Count:** Shows the number of entries in the duplicate detection mechanism on Port A or Port B for which more than one duplicate is received.
- ▶ **Unique Count:** Shows the number of entries in the duplicate detection mechanism on Port A or Port B for which no duplicate is received.
- ▶ **Node Count:** Returns the number of nodes detected in the system.

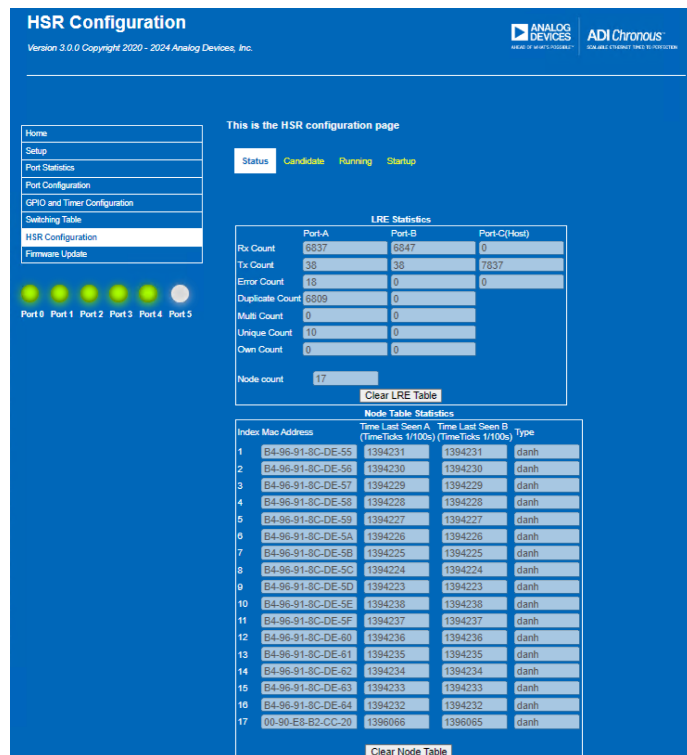


Figure 128. HSR Status Page

HIGH AVAILABILITY SEAMLESS REDUNDANCY (HSR)

Nodes Table Statistics

The **Nodes Table Statistics** shows the MAC addresses of other HSR devices in the network. The nodes table is built up of entries based on Supervision frames received by the Switch from other HSR devices. The nodes table can hold maximum of 1024 entries. By analyzing the incoming the supervision frames, the Switch reports the time (in TimeTicks 1/100 seconds) at which other devices were last seen on either Port A or Port B and reports the type of each device. The Host can access this information to gain insight into the workings of the network/network functionality and to identify any issues. The nodes table continuously refreshes its content. Node entries are automatically removed from the table if traffic from

that address has not been seen by the Switch within a duration of 60 seconds.

Proxy Node Table

As shown in [Figure 129](#), the **Proxy Node Table** shows the MAC addresses of the detected SAN devices connected to the Interlink Ports when the device is configured as a Redbox. The **Proxy Node Table** can support up to 8 entries. The entries are learned based on the traffic sent into the Interlink Port. The table refreshes each entry every 60 seconds, with entries getting removed from the table if traffic from that address is not seen again within this duration.

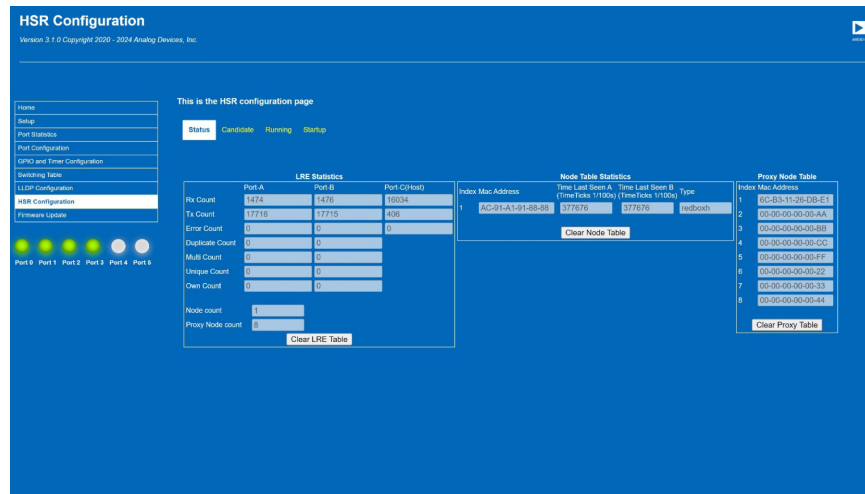


Figure 129. HSR Status Page for HSR Redbox with Proxy Node Table Entries

HIGH AVAILABILITY SEAMLESS REDUNDANCY (HSR)

HSR RUNNING VIEW

HSR **Running** page shows the configuration loaded to the device. See [Figure 130](#).

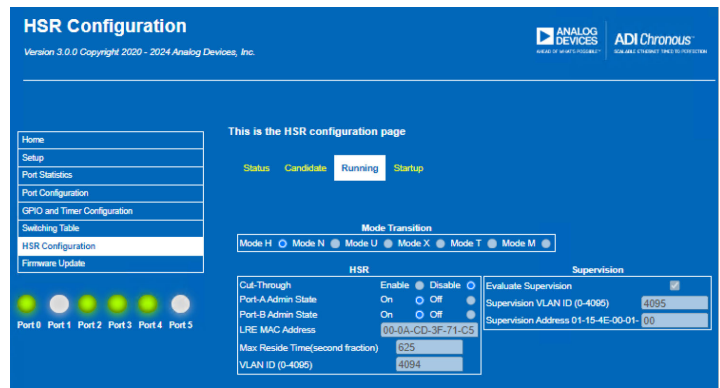


Figure 130. HSR Running Page

HSR STARTUP VIEW

HSR **Startup** page shows the **Startup** configuration. See [Figure 131](#).

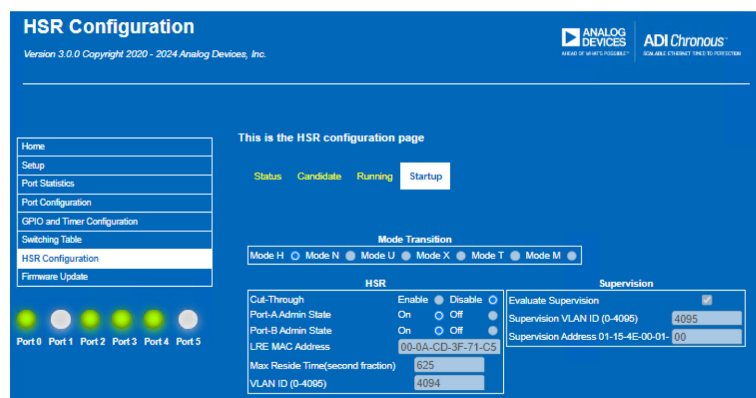


Figure 131. HSR Startup Page

HIGH AVAILABILITY SEAMLESS REDUNDANCY (HSR)

HSR – SUPERVISION FRAMES

The Wireshark capture shown in [Figure 132](#) is a supervision frame sent out from the Switch on Port B. By default, supervision frames are sent out on both ring ports at a 2 second interval and without

VLAN tags. The Switch sends a supervision frame for LRE MAC address for the DANH. Supervision frames are sent out to the HSR ring, so only visible on Port A and Port B. The HSR tag can be seen in the frame with the suffix 0x892f.

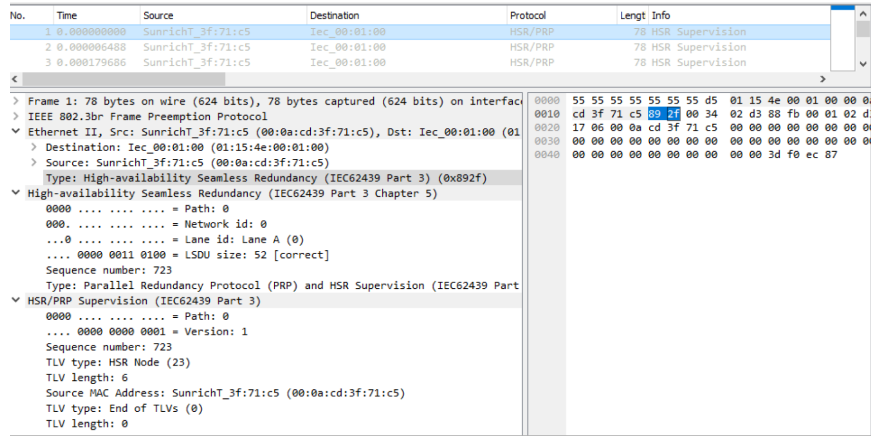


Figure 132. HSR Supervision Frame

HIGH AVAILABILITY SEAMLESS REDUNDANCY (HSR)

HSR – CAPTURE OF HSR TAGGED TRAFFIC

The Wireshark capture shown in [Figure 133](#) is traffic sent into Port C of the DANH and observed on Port A. The HSR tag can be seen

in the frame with type 0x892f in addition to the LAN information, SDU size and sequence number.

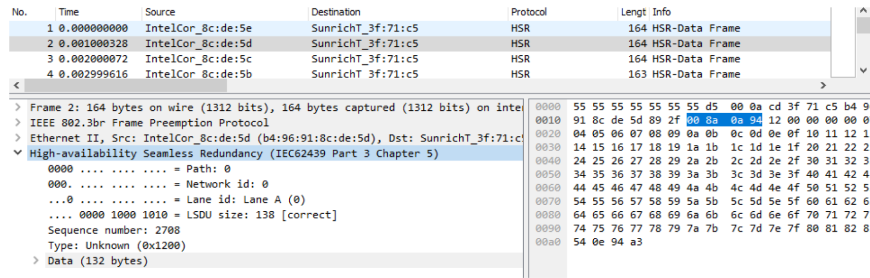


Figure 133. Wireshark Capture of HSR Tagged Frames in the HSR Ring Seen at Port A

MEDIA REDUNDANCY PROTOCOL (MRP)

MRP is a redundancy protocol used to avoid single points of failure in industrial communications networks. The MRP protocol is based on a ring topology and in accordance with IEC 62439-2 2021 standard. For full details on MRP protocol review the detailed standard. The following descriptions provide an overview of MRP to help describe the functionality provided by the Switch and does not intend to be a full overview of MRP function.

MRP STACK ON THE SWITCH

MRP can be configured on startup of the Switch. The MRP stack is running on the Packet Assist Engine, thereby, offloading MRP overhead from the Host. The Switch supports operation as a media redundancy client (MRC), a media redundancy manager (MRM), or a media redundancy automanager (MRA).

The Switch does not support interconnected rings.

One instance of MRA/MRC/MRM is supported on a 6-port device. All TSN functionality is supported with MRP. Other redundancy protocols are not supported with MRP.

RECOVERY PROFILES

When the Switch is configured for MRP operation, it supports recovery profiles of 500 ms, 200 ms, or 30 ms. In practice, all MRP devices in the ring may be configured with the same recovery profile, but it is also possible to have different recovery profiles. For example, the MRM can be configured with a 30 ms recovery profile and the clients with 200 ms profile.

CONFIGURING MRP

MRP must be configured up front using parameters in the xml file (see [MRP Specific Configuration](#) section).

By default (XML), the Switch is configured as an MRC with a recovery profile of 500 ms and with Port 1 and Port 2 used as ring ports.

CANDIDATE PAGE

The **Candidate** page (see [Figure 134](#)) provides user ability to configure how the MRP capability of the device should operate. The configuration included are as follows:

- ▶ **MRP Role:** Choice of **Client** (default), **Manager**, or **Auto-Manager**.
- ▶ **Domain ID:** Unique Domain ID for the MRP ring.
- ▶ **Domain Name:** Domain name for the ring.
- ▶ **OUI:** MRP OUI, defaults to **0x080006** (Siemens OUI).
- ▶ **Domain VLANID:** Defaults to untagged/4095.
- ▶ **Recovery Profile:** Choice of **30 ms**, **200 ms**, or **500 ms** (default).
- ▶ **Ring Port 1/Ring Port 2:** Default **Port 1** and **Port 2**, choice of any pair of ports.
- ▶ **Ring 1 Priority/Ring 2 Priority:** Default Queue **7** is highest priority. PTP traffic also egresses in Queue 7. If using lowest recovery profile, change default PTP queue from 7 to a lower priority in the **Time Synchronization** page.
- ▶ **React on Link Change:** For faster recovery, enable **React on Link Change**. This allows the manager to react on the link change frames instead of waiting for test frames to timeout.

The screenshot shows the 'MRP Candidate' configuration page. At the top, there are tabs for 'Status', 'Candidate', 'Running', and 'Startup', with 'Candidate' selected. Below the tabs are 'Commit' and 'Discard' buttons. A 'Save' button is located below the tabs. The main configuration area is titled 'Configuration' and contains the following fields:

- MRP Role: Client (dropdown menu)
- Domain ID (hex): FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF (text input)
- Domain Name: MRP (text input)
- OUI: 080006 (text input)
- Domain VLANID (0-4095): 4095 (text input)
- Recovery Profile: 500ms (dropdown menu)
- Ring Port 1: Port 1 (dropdown menu)
- Ring Port 2: Port 2 (dropdown menu)
- Ring 1 Priority: 7 (text input)
- Ring 2 Priority: 7 (text input)
- Manager Priority: 0 (text input)
- React on Link Change: (checkbox)

Figure 134. MRP Candidate Page – Default Configuration

The MRP **Status** page (see [Figure 135](#)) provides insight into the MRP configuration, the state of the ring, the forwarding/blocked position of the ports (for MRM), and statistics related to the MRP operation.

MEDIA REDUNDANCY PROTOCOL (MRP)

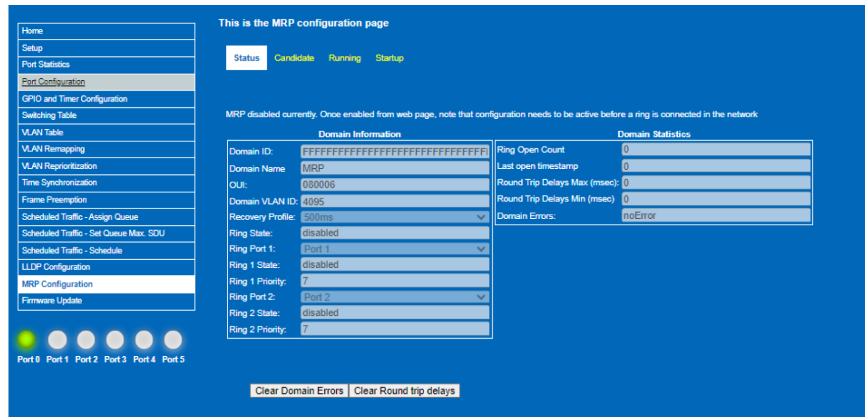


Figure 135. MRP Status Page – Default Configuration – MRP Disabled

MEDIA REDUNDANCY PROTOCOL (MRP)

MRP SCENARIOS: MRM AND MRC

The following example shows a simple MRP configuration with two Switch devices in a ring topology. Switch 1 is configured as MRM and Switch 2 configured as MRC through the web server.




Figure 136. MRP Configuration with Two Switches, One MRM, One MRC

After configuration, the ring cables are connected and from the **Status** page shown in Figure 139, Switch 1 configured as MRM reports the ring is closed, it is forwarding on Ring Port 1 with Ring Port 2 blocked. Figure 140 shows that Switch 2 as MRC is forwarding on both ports. Per the IEC standard, the MRC does not report the ring state, instead shows it as undefined.

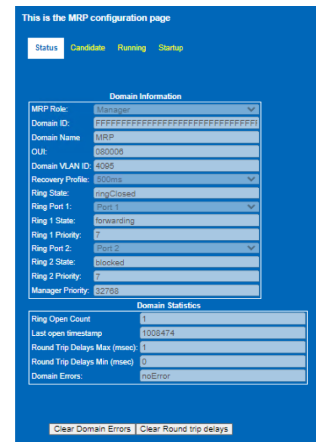


Figure 139. MRP Status Page: Switch 1, MRM (Manager) – Ring Closed

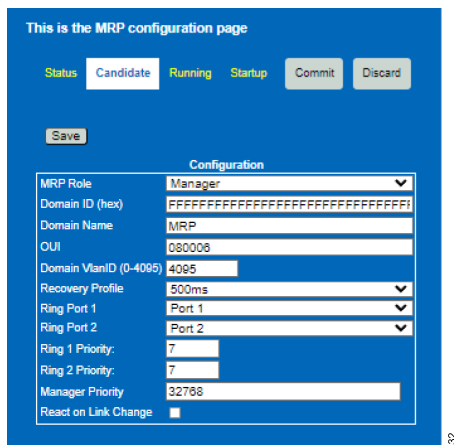


Figure 137. MRP Candidate Page: Switch 1, MRP Enabled as MRM (Manager)

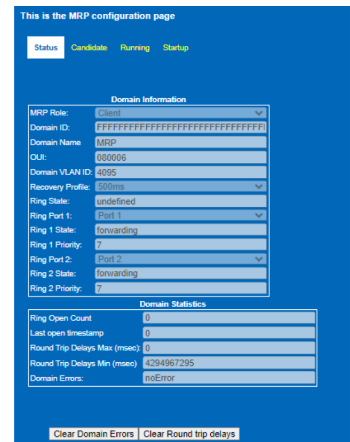


Figure 140. MRP Status Page: Switch 2 MRC (Client) – Ring Closed

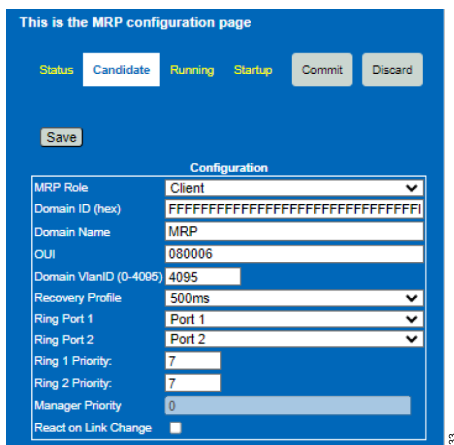


Figure 138. MRP Candidate Page: Switch 2, MRP Enabled as MRC (Client)

Figure 141 shows the MRP test frames sent out each MRM ring port. When the recovery profile is 500 ms, these frames are sent every 50 ms on both ports.

In event of a fault in the ring, the MRM device no longer sees the test frames on the other ring port and opens the blocked port to forward traffic around the ring. MRP devices send other MRP frames with information related to what is happening in the ring such as MRP_LinkChange and MRP_TopologyChange, see Figure 142 and Figure 143.

MEDIA REDUNDANCY PROTOCOL (MRP)

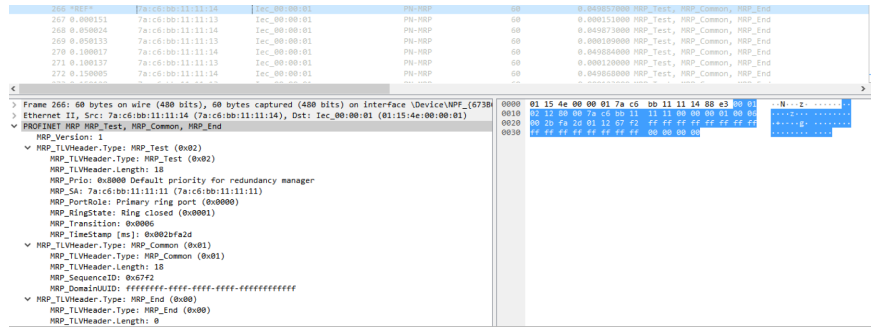


Figure 141. MRP Test Frames from Both Ports of MRM Every 50 ms

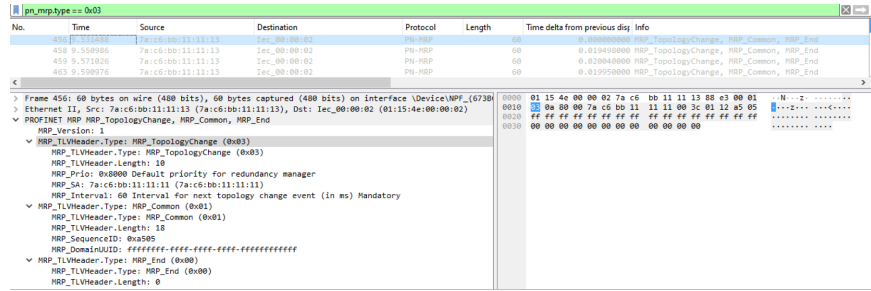


Figure 142. MRP Topology Change Frame – Link Opened Somewhere in the Ring Port 1 Side of MRM

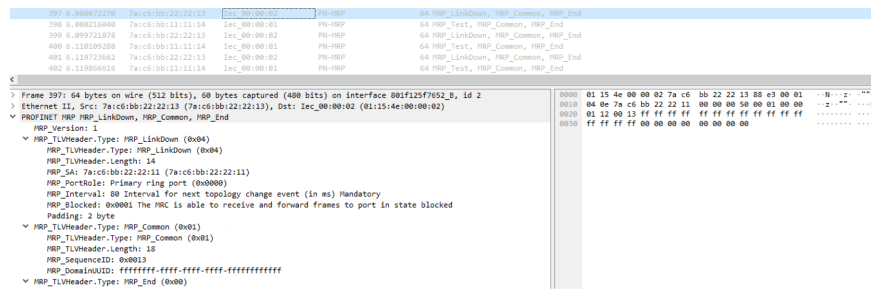


Figure 143. MRP Topology Change Frame – Link Down Reported by Switch 2 MRC Device

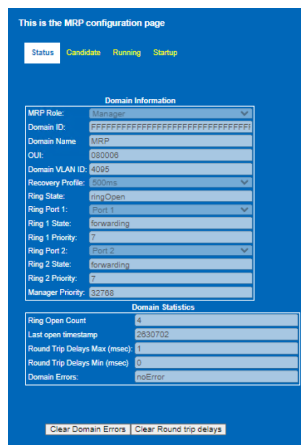


Figure 144. MRP Status Page: Switch 1 MRM (Manager) – Ring Port 1 Open

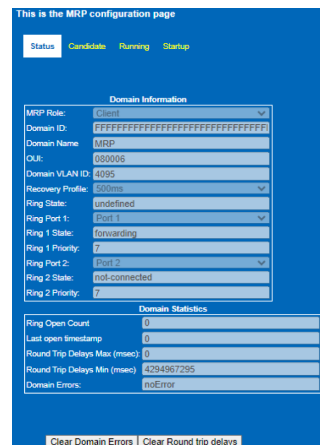


Figure 145. MRP Status Page: Switch 2 MRC (Client) – Ring Port 1 Open

MEDIA REDUNDANCY PROTOCOL (MRP)

MRP SCENARIOS: MRA

A ring can only have one MRM. In a configuration where there is more than one MRA, the MRA devices use the MRP Voting process to decide who is the rings MRM.

When configuring a ring with multiple MRAs, other devices must be configured with MRA or MRC roles. Per the IEC62439-2 standard, it is not supported to have MRA and MRM combinations in one ring.

The **Status** page shows additional information for the MRA configuration, for example, reporting the best Manager MAC. Voting is based on priority and MAC address. The MRP priority is 0xA000 (40960) for MRA devices. Since both devices have the same priority, voting is based on MAC address.




Figure 146. MRP Configuration with Two Switches, Both Configured as MRA

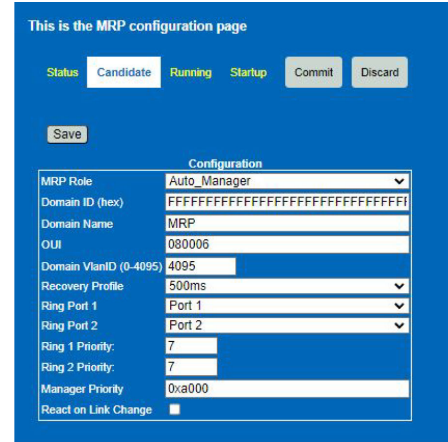


Figure 147. MRP Configuration with Two Switches, Both Configured as MRA Using Web Server

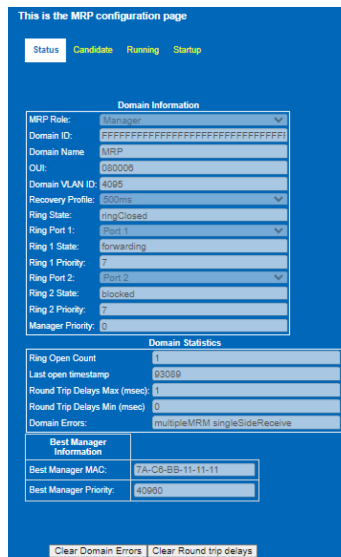


Figure 148. MRP Configuration with Two Switches, Switch 1 Becomes MRM

MEDIA REDUNDANCY PROTOCOL (MRP)

This is the MRP configuration page

Domain Information	
MRP Role:	Client
Domain ID:	FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
Domain Name:	MRP
OUI:	080006
Domain VLAN ID:	4095
Recovery Profile:	500ms
Ring State:	undefined
Ring Port 1:	Port 1
Ring 1 State:	forwarding
Ring 1 Priority:	7
Ring Port 2:	Port 2
Ring 2 State:	forwarding
Ring 2 Priority:	7

Domain Statistics	
Ring Open Count	0
Last open timestamp	0
Round Trip Delays Max (msec)	0
Round Trip Delays Min (msec)	4294967295
Domain Errors:	multipleMRM

Best Manager Information	
Best Manager MAC:	7A-C6-BB-11-11-11
Best Manager Priority:	40960

Figure 149. MRP Configuration with Two Switches, Switch 2 Becomes MRC

PER-STREAM FILTERING AND POLICING, QCI

Per-stream filtering and policing (PSFP), as defined by IEEE 802.1Qci standard provides filtering policing for a stream.

The purpose of this feature is to prevent traffic overload conditions from affecting the receiving node. This is done by filtering traffic on a per-stream basis.

PSFP applies to Store and Forward Switching as frame size must be known to apply a filter. A cut-through frame starts forwarding before the full frame size is known.

The filtering and policing capabilities apply on a stream basis to the receive path:

- ▶ Size based filters – 32 per port
- ▶ Time based filters – 16 gates per port
- ▶ Rate based filters – 8 per port

Any stream can be assigned to any combination of filters and the device can support up to 32 combinations of filters.

PSFP relies on Switching table entries (static entries and extended table entries) mapped to a stream filter. Stream table entries with PSFP are not yet supported.

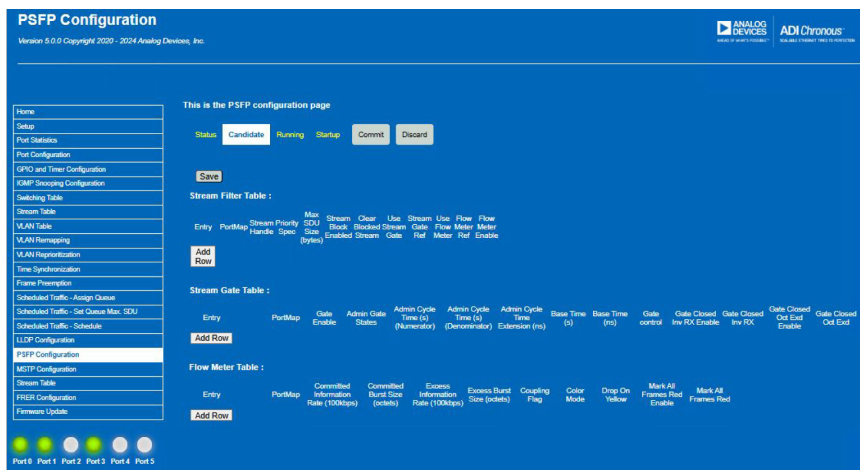


Figure 150. PSFP Candidate View

PER-STREAM FILTERING AND POLICING, QCI

PSFP CANDIDATE PAGE

Stream Filter

Filtering of frames is based on VLAN Priority (PCP), stream ID and frame size. Frames matching the filter in excess of the programmed frame size are discarded. The stream filter can be blocked or unblocked. The stream filter can optionally be associated with a stream gate and flow meter. Frames that fail the filter are discarded. Stream filter entries can also be later deleted. The configuration included are as follows:

- ▶ **PortMap:** Ingress port map to apply the filter to. Can be an individual port or group of ports.
- ▶ **Stream Handle:** Stream ID, map stream ID to Static table entry. Uniquely identifies the stream filter to be used.
- ▶ **Priority Spec:** VLAN Priority/PCP, 0 to 7, Wildcard indicates any priority.

- ▶ **Max SDU Size (bytes):** Frames in excess of MaxSDU value are discarded. MaxSDU definition includes the MAC addresses and FCS.
- ▶ **Stream Block Enabled:** Provides ability to completely block stream if the stream violates the programmed MaxSDU. If disabled (default), any frames that exceed MaxSDU for the filter are discarded, smaller frames are allowed to be received. If enabled, any frame that exceeds the MaxSDU for the filter are discarded and all subsequent matching frames are dropped unconditionally.
- ▶ **Clear Blocked Stream:** Allows user to clear a blocked stream (blocked/unblocked status is visible in the **Status** page).
- ▶ **Use Stream Gate and Stream Gate Ref:** To associate a stream gate, enable the check box and pass the stream gate ID.
- ▶ **Use Flow Meter, Flow Meter Ref, Flow Meter Enable:** To associate a flow meter with the stream filter, select the use check box, provide the flow meter ID and enable the flow meter.



Figure 151. Stream Filter Configuration




Figure 152. Stream Filter with Multiple Entries and Flow Meters Associated

PER-STREAM FILTERING AND POLICING, QCI

Stream Gate

The Stream Gate is either open or closed. The Stream Gate monitors the arrival time of frames on that stream and uses the port Timer control unit to control the gate, similar to that of Scheduled traffic on the Transmit side. If a stream arrives when the gate is open, accept the frame and perform the required lookups and handle as required. Alternatively, if the stream arrives when the gate is closed, discard the frame.

Internal Priority Vector (IPV) – The Stream Gate can change the IPV of a frame.

The Stream Gate can allow ports support a defined amount of traffic in a certain amount of time, Octets per unit time based on the port Timer control unit (TCU).

Note that the cycle time for the Stream Gate must match that of Scheduled Traffic for the port.

The stream gate configuration parameters are as follows:

- ▶ **Gate ID/Ref:** Entry value corresponds to the Gate reference. This uniquely identifies the Stream Gate instance to be used and linked to the stream filter.
- ▶ **PortMap:** Ingress port map to apply the filter to. Can be an individual port or group of ports.
- ▶ **Gate Enable:** Enable the stream gate function.
- ▶ **Admin Gate States:** If open, frames are permitted to pass through the gate. If closed, frames are not permitted to pass.
- ▶ **Admin Cycle Time (Numerator, Denominator):** Sets cycle time for the Stream Gate. Default 1 ms. Cycle time must be same as the Cycle time configured in Scheduled traffic for that port.
- ▶ **Admin Cycle Time Extension:** The Cycle Time Extension value defines the maximum amount of time by which the old cycle is

permitted to be extended when Switching to a new schedule for the Stream Gate.

- ▶ **Base Time (s, ns):** Absolute time at which a new Stream Gate schedule is required to take effect.
- ▶ **Gate Control:** Provides additional configuration of the Stream Gate schedule parameters.
- ▶ **Gate Closed Invalid RX Enable:** When enabled, if any frame is discarded because the gate is in the closed state, then the gate remains in the closed state and all subsequent frames are discarded. If disabled, has no effect.
- ▶ **Gate Closed Octet Exceeded Enable:** When enabled, if any frame is discarded because there is insufficient value of IntervalOctetsleft, then the gate closes and remains in the closed state and all subsequent frames are discarded. During an interval, each frame size is compared to IntervalOctetsleft, if the frame size is smaller than IntervalOctetsleft, the frame passes the gate and the frame size is subtracted from IntervalOctetsleft.

Gate Control fields provide four entries as follows:

- ▶ **Time Interval Value (ns):** Time interval.
- ▶ **Gate-State-Value: Open or closed.**
- ▶ **IPV Spec:** Use Ignore IPV to leave stream route based on admin IPV. Alternatively, if reprioritization is required, pass the IPV value with the traffic class the stream should egress.
- ▶ **Interval Octet Max:** Maximum number of octets allowed to pass during a time interval, if exceeded, subsequent frames are discarded.
- ▶ **Interval Octet Enable:** When set, indicates the Interval Octet value must be used.



Figure 153. Stream Filter and Stream Gate

PER-STREAM FILTERING AND POLICING, QCI

Flow Meter

Stream data per unit time allows a certain amount of traffic through the port. This feature uses a token bucket or a bandwidth profile where it compares a frame size to how many tokens in a one or two buckets (commit, excess).

The flow meter configuration parameters are as follows:

- ▶ **Flow Meter ID/Ref:** Entry value corresponds to the Flow meter reference. This uniquely identifies the flow meter instance to be used and linked to the stream filter.
- ▶ **PortMap:** Ingress port map to apply this filter to. Can be an individual port or group of ports.
- ▶ **Committed Information Rate:** CIR expressed in units of 100 kbps. Rate at which tokens are added to commit bucket.
- ▶ **Committed Burst Size:** CBS, expressed in octets, maximum capacity of the commit bucket.
- ▶ **Excess Information Rate:** EIR, expressed in units of 100 kbps. Rate at which tokens are added to excess bucket.
- ▶ **Excess Burst Size:** EBS, expressed in octets, maximum capacity of the excess bucket.
- ▶ **Coupling Flag:** Shows whether any overflow in the commit bucket tokens must be added to excess bucket. Select the check box to enable coupling the overflow to the excess bucket.

- ▶ **Color Mode:** Option of **Color Aware** or **Color Blind**. If **Color Aware** is selected, the VLAN tag drop eligible indicator (DEI) bit of incoming frames is used in metering decisions. Incoming frames, with DEI bit set to one, are dropped if drop on yellow is enabled. If **Color Blind**, then incoming frame color is ignored.
- ▶ **Drop on Yellow:** If enabled, discard yellow frames. If disabled, forward yellow frames if bandwidth allows. Yellow frames have their DEI bit set to 1 before transmitting. **Drop on Yellow** control is only available when color mode is **Color Aware**. Disabled by default.
- ▶ **Mark All Frames Red Enable:** If enabled, mark all incoming frames red, discard all frames. When disabled (default) flow meter works as normal.
- ▶ **Mark All Frames Red:** Enables the **Mark All frames Red**.

In the example shown in Figure 154, three separate flow meters are configured for use with ingress traffic on Port 2. The first entry at Flow meter reference 0, configures a CIR rate of 10, corresponding to 1 Mbps (10 × 100 kbps). The second entry configures a 100 kbps rate. The third entry configures 200 kbps from the commit bucket and uses the excess bucket with an additional 100 kbps. Thus if bandwidth allows, the flow meter potentially grants this stream up to 300 kbps bandwidth.

Entry	PortMap	Committed Information Rate (100kbps)	Committed Burst Size (octets)	Excess Information Rate (100kbps)	Excess Burst Size (octets)	Coupling Flag	Color Mode	Drop On Yellow	Mark All Frames Red Enable	Mark All Frames Red
0	0x07	10	1500	0	0	zero	color-aware	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
1	0x07	1	1500	0	0	zero	color-aware	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
2	0x07	2	1000	1	100	zero	color-aware	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Figure 154. Flow Meter with Different CIR Rates per Stream

Figure 155. Stream Filter + Flow Meter

PER-STREAM FILTERING AND POLICING, QCI

PSFP STATUS PAGE

The **Status** page (see [Figure 156](#)) shows statistics for the PSFP functionality. For the stream filter, there are a number of statistics captured and available to read for a number of parameters. PSFP statistics have an associated errata as shown in the **Silicon Anomaly** topic included in the [ADIN6310](#) data sheet.

The statistics available for each stream filter instance are as follows:

- ▶ A count of frames **matching** both the stream_handle and priority specifications.
- ▶ A count of frames that passed the maximum SDU size filter (see the **Passing SDU Count** column in [Figure 156](#)).
- ▶ A count of frames that did not pass the maximum SDU size filter (see the **Not Passing SDU Count** column in [Figure 156](#)).
- ▶ Whether the stream is blocked because of oversized frame (see the **Stream Blocked Due to Oversized Frame** column in [Figure 156](#)).

The statistics available for each stream gate instance are as follows:

- ▶ A count of frames that **passed** the Stream Gate.
- ▶ A count of frames that **did not pass** the Stream Gate.
- ▶ Statistics related to the Stream Gate instance.

The statistics available for each instance of the flow meter are as follows:

- ▶ A count of frames that were discarded as a result of the operation of the flow meter (see the PSFP Flow Meter Statistics section in [Figure 156](#)).

Statistics are available on a per port basis, select the port of interest to view the statistics. The statistics counts accumulate but refreshes of the page are required to get an update.

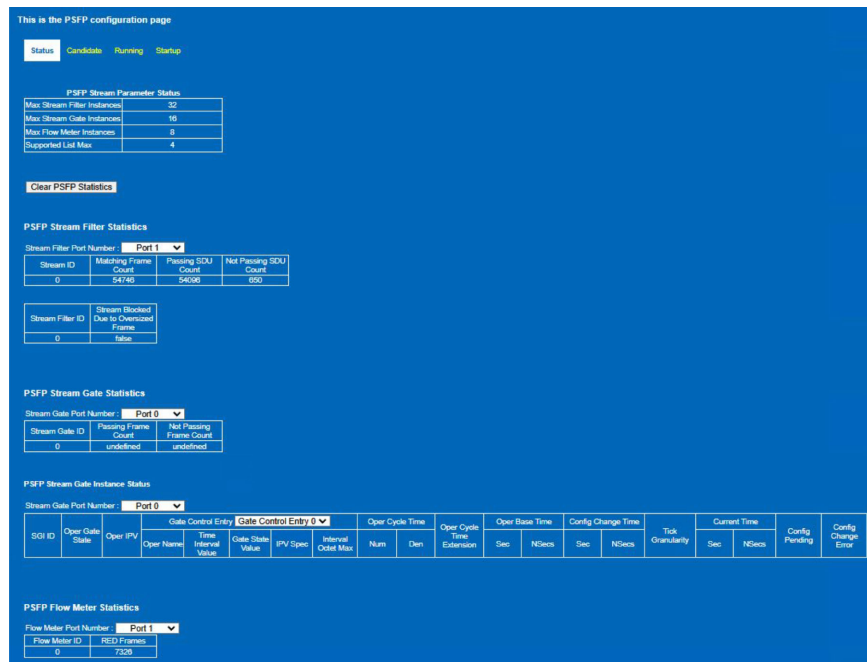


Figure 156. PSFP Status View

MULTIPLE SPANNING TREE PROTOCOL (MSTP)

Spanning tree protocols were developed to prevent broadcast storms by eliminating loops in network topologies. STP is the original spanning tree protocol, rapid spanning tree protocol (RSTP) was defined after significantly improving convergence speed, and multiple spanning tree protocol (MSTP) is the most recent version defined in IEEE 802.1Q-2022. STP and RSTP networks use a single spanning tree instance for the entire network, whereas MSTP supports grouping and mapping of VLANs into different spanning tree instances. This approach reduces the number of spanning tree calculations, improving resource utilization.

When operating in MSTP mode, the switch is interoperable with other switches running MSTP, STP, or RSTP. It supports up to four instances, three multiple spanning tree instances (MSTIs), and the common and internal spanning tree (CIST). Each MSTI operates independently within an MST region and elects its own root bridge to optimize traffic flow for the VLANs assigned to that instance. Any number of VLANs can be mapped to an instance. However, only 64 VLANs can be active. By default, all VLANs (0 to 4095) are mapped to the CIST. Mapping does not indicate that VLANs are active. It just serves as a lookup table when VLANs are activated. For example, if 1000 to 2000 VLAN IDs are mapped to MSTI 1, and VLAN ID 1500 gets activated later, it will be subjected to port states of MSTI 1.

Across the entire MST region, the CIST selects a single root bridge. The CIST spans the entire network, connecting MST regions and integrating with non-MST networks. The CIST is the default spanning tree instance for MSTP; all VLANs that are not explicitly configured into another MSTI are part of the CIST. When MSTP is enabled, the port role and port state are calculated for the CIST and separately for each other instance. The bridge with the best bridge ID is chosen as the CIST root. Additionally, MSTI regional roots are identified for each region. Bridge priority, path cost, and port priority are used by the spanning tree instances to determine the topology.

In the example shown in [Figure 157](#), without MSTP, any traffic would just circulate in the network. Switch 1 acts as the root bridge for the CIST spanning tree instance. Two additional MSTIs are created with different VLAN ID mappings. Switch 2 is configured to be the root for MSTI1 and sends traffic to Switch 1 and Switch 3. The path for this particular traffic between Switch 1 and Switch 3 is blocked to prevent a loop. Similarly, Switch 3 is configured as the root for MSTI2, traffic is allowed to flow from Switch 3 to Switch 1 and Switch 2, but not directly between Switch 2 and Switch 1. If an existing active path failed after the network converged, because of a link dropping or change in switches, MSTP activates an alternative path.

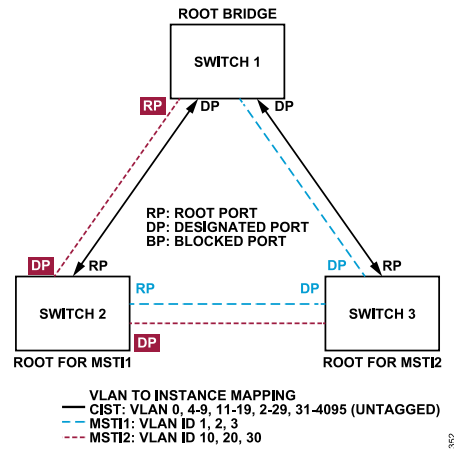


Figure 157. Simplified Example of MSTP with CIST and Two MSTIs

Running the MSTP stack on the switch is supported with other network protocols, such as LLDP, time synchronization, scheduled traffic, FRER, PSFP, and frame preemption. MSTP is not intended to operate with other redundancy protocols such as MRP, PRP, or HSR.

MSTP uses bridge protocol data units (BPDUs) to exchange spanning tree information across the network. Each port in MSTP assumes a specific role within an instance, ensuring proper network behavior and redundancy.

MSTP CANDIDATE PAGE

The **Candidate** page provides user ability to configure MSTP for desired operation.

The MSTP stack runs on the switch, with the packet assist engine managing the port states and roles based on the network. When the device is configured for TSN functionality (using one of the eval-adin6310.xml, eval-adin3310.xml, or eval-adin6310-10t11-rev-c.xml configuration files) MSTP is enabled by default for all ports. Only one instance is created (CIST) and all active VLANs mapped to the CIST. Any additional MSTIs must be configured through the web server **Candidate** page.

MSTP automatically manages all the VLANs that are activated after MSTP is started. By default, VLAN 0 and VLAN 4095 (untagged) become active, and the MSTP state machine applies the relevant state (forwarding or discarding) based on the port role and port state transition state machines. If another VLAN is made active later, this VLAN gets mapped to the CIST by default and the relevant forwarding/discarding state is applied. The user can also map a set of VLANs to an MSTI. VLANs that are part of an MSTI are not part of the CIST because a VLAN cannot be mapped to multiple spanning trees.

MULTIPLE SPANNING TREE PROTOCOL (MSTP)

Use this page to configure MSTP

Status: Candidate Running Startup

STP Bridge Configuration :

Force Protocol Version	MSTP
Bridge Hello Time (seconds)	2
Max Hops	20
Bridge Forward Delay (seconds)	15
Bridge Max Age (seconds)	20
Migrate Time (seconds)	3
Tx Hold Count	6
CIST Bridge ID Priority	32768

Exclude VLAN IDs from MST
(Warning: Excluded VLANs cannot be removed from the exclusion list or modified!)

STP Port Configuration

Port Number	MSTP Enable	Admin Edge Port	Auto Edge Port	Auto Isolate Port	Restricted Role	Restricted TCN	Point to Point
Port 0	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Auto
Port 1	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Forced True
Port 2	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Auto
Port 3	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Auto
Port 4	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Auto
Port 5	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Auto

STP Port Configuration

Port Number	Path Cost		Priority
	Auto for external	Auto for internal	
Port 0	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	128
Port 1	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	128
Port 2	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	128
Port 3	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	128
Port 4	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	128
Port 5	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	128

Figure 158. MSTP Candidate Page (Part 1)

The candidate page contains the following configuration parameters:

- ▶ **STP Bridge Configuration:** Configuration parameters for the spanning tree protocol and CIST.
 - ▶ **Force Protocol Version:** The protocol version to use. Choice of STP, RSTP, and MSTP. When STP or RSTP is chosen, only one instance is available for configuration. The default value is **MSTP**.
 - ▶ **Bridge Hello Time:** The time interval between BDPUs messages. The default value is **2** seconds
 - ▶ **Max Hops:** The maximum hops of BDPUs information. The default value is **20** hops.
 - ▶ **Bridge Forward Delay:** The duration spent in listening and learning state. The default value is **15** seconds.
 - ▶ **Bridge Max Age:** The maximum age of the configuration BDPUs maintained in the device. The default value is **20** seconds.
 - ▶ **Migrate Time:** Fixed to **3** seconds.
 - ▶ **Tx Hold Count:** The maximum number of BDPUs per second that the switch can send from an interface. The default value is **6**.
 - ▶ **CIST Bridge ID Priority:** Bridge priority. The value ranges from 0 to 61440 in increments of 4096. The default value is **32768**. The device with the highest priority (lowest number) is chosen as the root bridge of the CIST.
- ▶ **Exclude VLAN IDs from MST**
 - ▶ Defines VLANs that must be excluded from spanning tree. These VLANs may share links with spanning trees but not be subject to the spanning tree state. When excluding VLANs, user must ensure loops are handled correctly to protect against duplicate paths causing storms.
 - ▶ Note that excluded VLANs cannot be removed from the exclusion list.
- ▶ **STP Port Configuration:** Per port configuration parameters.
 - ▶ **MSTP Enable:** Enable or disable MSTP on the port. The default value is enabled
 - ▶ **Admin Edge Port:** Ports that connect to end points must be administratively configured as edge ports. These type of ports have no other bridges attached, therefore, no possibility of loops and do not need to run spanning tree protocols. The default value is disabled.
 - ▶ **Auto Edge Port:** This feature automatically detects whether a port should be an edge port. If a port does not receive any BDPUs for a duration of MigrateTime (3 seconds), the port classifies itself as an edge port and transitions directly to the forwarding state. Enabled by default.
 - ▶ **Auto Isolate Port:** If true, it causes a designated port to transition to discarding if both **Admin Edge Port** and **Auto Edge Port** are false and the other bridge presumed to be attached to the same point-to-point LAN does not transmit periodic BDPUs. The default value is disabled.

MULTIPLE SPANNING TREE PROTOCOL (MSTP)

- ▶ **Restricted Role:** Enabling restricted role on a port causes it not to be selected as root port for the CIST or any MSTI. Restricted role ports are selected as alternative ports after the root port has been selected. The default value is disabled.
- ▶ **Restricted TCN:** If enabled, it causes the port not to propagate received topology change notifications and topology changes to other ports. The default value is disabled.
- ▶ **Point to Point:** Option of **Auto**, **Forced True**, or **Forced False**. If set to **Auto**, the port automatically determines whether the link is a point-to-point connection based on the duplex mode of the link. Full duplex is treated as a point-to-point link and half duplex is treated as a shared medium link. If set to **Forced True**, the port is treated as connected to a point-to-point link. If **Forced False**, the port is treated as having a shared media connection.
- ▶ **STP Port Configuration**
 - ▶ **Path Cost:** The subparameters are **Auto** and **Value**. MSTP calculates the path cost for each link based on the established link speed. The user can modify the default path cost type and configure a cost administratively. The port with the lowest path cost to the root bridge is selected as the root port. If auto is true, the path cost is selected based on link speed. A 1 Gbps link speed equals a path cost of 20,000, while a 100 Mbps link speed results in a cost of 200,000 and 10 Mbps is 2,000,000. The default value is true.
 - ▶ **Priority:** The priority of the port value ranges from 0 to 240 in increments of 16. The default value is 128.
- ▶ **MSTP Configuration Identification:** In a spanning tree network, devices in the same MST region must have the same VLAN to instance mapping, configuration revision number, and name. When using the web server with cascaded switches, the **Configuration Name** for each switch in the chain is, by default, based

- on the MAC address assigned from sysrepo during initialization. This results in each device being in a region of its own. Even so, MSTP configuration continues to protect against loops and allows both tagged and untagged traffic in the active path. To assign devices to the same region, the user must ensure the **Configuration Name** is identical for all devices intended for that region.
 - ▶ **Configuration Name:** The name of a region.
 - ▶ **Configuration Revision:** The revision level in the MST configuration identifier, which ranges from 0 to 65535. The default value is 0.
- ▶ **MSTI Configuration:** Three MST instances can be created for a total of four instances including the default CIST.
 - ▶ **Bridge Priority:**
 - ▶ **VLAN IDs Mapped:** VLAN IDs mapped to the MST instance. Any one VLAN ID can only be mapped to one instance.
- ▶ **MST Instance Port Configuration** (per instance drop-down)
 - ▶ **Path Cost:** For a port, the path cost setting can be different for different MSTIs it belongs to.
 - ▶ **Auto:** Automatically sets the port path cost based on the speed of established link. A 1 Gbps link speed equals a path cost of 20,000, while a 100 Mbps link speed results in a cost of 200,000 and 10 Mbps is 2,000,000. The default value is true.
 - ▶ **External:** External path cost is the path cost between regions. If the path cost has to be modified inside the same region, external must be set to false. The default value is false.
 - ▶ **Value:** Based on link speed if Auto is true.
 - ▶ **Priority:** The priority of the port value ranges from 0 to 240 in increments of 16. The default value is 128.

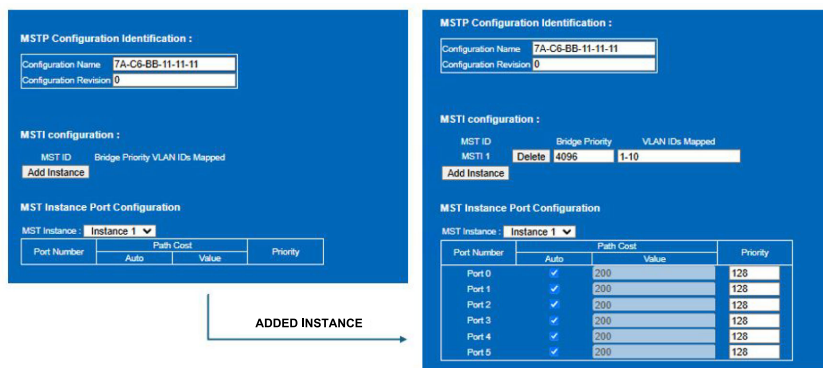


Figure 159. MSTP Candidate Showing MST Instance

Similar to the other web pages, all changes in the **Candidate** page must be saved and then committed. The **Running** page shows the configuration loaded to the device and the **Startup** page shows the startup configuration for the function.

MULTIPLE SPANNING TREE PROTOCOL (MSTP)

MSTP STATUS PAGE

MSTP Configuration
Version: 5.0.0 Copyright 2020 - 2024 Analog Devices, Inc.

Use this page to configure MSTP

Status: **Candidate** | Running | Startup

General STP status:

Forward Delay (seconds)	15
Hello time (seconds)	2
Max Age (seconds)	20
Max Hops	10
Migrate Time (seconds)	3

Status Info:

CIST Bridge ID Priority	32768
CIST Bridge MAC Address	7A-C8-BB-11-11-11
Time Since Topology Change (seconds)	0
Topology Change Count	0
Root Path Cost	0
CIST Root Port Number	0
CIST Root Priority	32768
CIST Root MAC Address	7A-C8-BB-11-11-11
CIST Regional root priority	32768
CIST Regional root MAC Address	7A-C8-BB-11-11-11
VLAN to CIST mapping	0-4095

Port Configuration

Port Number	Port State	Port Role	Port ID	External Path Cost	Designated Bridge ID Priority	Designated Port ID	Open Edge Port	Isolate Port
0	forwarding	designated-port	32768	20000	128	7A-C8-BB-11-11-11	32768	true
1	forwarding	designated-port	32770	20000	128	7A-C8-BB-11-11-11	32770	true
2	forwarding	designated-port	32771	20000	128	7A-C8-BB-11-11-11	32771	true
3	disabled	disabled-port	32772	200000000	128	7A-C8-BB-11-11-11	0	false
4	disabled	disabled-port	32773	200000000	128	7A-C8-BB-11-11-11	0	false
5	disabled	disabled-port	32774	200000000	128	7A-C8-BB-11-11-11	0	false

Figure 160. MSTP Status Page

The MSTP **Status** page shows the status information and port information for the switch based on what is happening in the network. Information is provided for each configured MSTI with a dropdown selection, allowing the user to view the status information for each different instance.

General STP Status: These parameters provide information on the forwarding delay, max age and hops, and migrate time as configured in the **Candidate** page.

Status Info: Provides insight into the CIST bridge information and identifies the CIST root bridge.

- ▶ **Time Since Topology Change:** Time in seconds since the topology changed.
- ▶ **Topology Change Count:** How many times the topology has changed.
- ▶ **VLAN to CIST Mapping:** Returns the VLAN IDs mapped to the CIST.

Port Configuration

- ▶ **Port State:** The switch supports the various port states, including disabled, blocking, learning, and forwarding but does not support the listening port state (used in STP).
 - ▶ **Disabled:** If the port link is down.
 - ▶ **Discarding:** The port sends and receives BPDUs but does not learn MAC addresses and does not forward user traffic.
 - ▶ **Learning:** The port learns MAC addresses and does not forward user traffic. It sends and receives BPDUs.

- ▶ **Forwarding:** The port learns MAC addresses and forwards user traffic. It sends and receives BPDUs.
- ▶ **Port Role:** MSTP assigns roles to ports, and the switch supports all MSTP defined port roles: disabled, root, leader, designated, alternate, and backup.
 - ▶ **Root Port:** The port that has been determined to be the least cost path back to the root bridge. Forwards data for a non-root bridge to a root bridge.
 - ▶ **Leader Port:** The leader port is a port on the shortest path from the local MST region to the common root bridge. It is not always located on the regional root, instead it is a root port on the CIST and a leader port on the other MSTIs.
 - ▶ **Designated Port:** The port that is in a forwarding state, forwards data away from root port to the downstream device.
 - ▶ **Alternate Port:** Acts as the backup port for a root port or leader port. In the event of the root port or leader port becomes blocked, the alternate port takes over.
 - ▶ **Backup Port:** Normally seen if a hub is used in the network, that is, a bridge receives its own BPDUs on different ports. If the designated port is invalid, the backup port becomes the new designated port. When two ports of the same spanning tree device are connected, a loop is formed; therefore, the device blocks one of the ports. The blocked port acts as the backup.
- ▶ **Port ID:** ID assigned to port.
- ▶ **External Path Cost:** Configured path cost.

MULTIPLE SPANNING TREE PROTOCOL (MSTP)

- ▶ **Designated Bridge ID:** A 64-bit bridge identifier consists of the first four MSB of the priority, then a system ID extension in the next 12 bits and the bridge address in the remaining 48 bits.
- ▶ **Designated Port ID:** A 16-bit port identifier consists of the first four MSB of the priority and a 12-bit bridge port number.
- ▶ **Oper Edge Port:** Indicates if port is configured as an edge port.
- ▶ **Isolate Port:** Indicates if port isolate is active.

If MSTIs are configured, two additional tables are displayed on the **Status** page, showing MSTI status info and MSTI per port status info as discussed below:

MSTI Status Info (per port with dropdown to select different MST instances)

- ▶ **Bridge Priority & MAC Address:** Returns the configured priority and MAC address as configured in the **Candidate** page.
- ▶ **TC Time:** Time since topology change.
- ▶ **TC Count:** Topology change count.
- ▶ **Internal Root Path Cost:** Root path cost.
- ▶ **Root Port ID:** Port ID of root port.

- ▶ **Regional Root Priority:** Priority of regional root.
- ▶ **Regional Root ID MAC Address:** MAC address of regional root.

MSTI Per Port Status Info (per port with dropdown)

- ▶ **Port State:** Disabled, blocking, learning, and forwarding.
- ▶ **Port Role:** Disabled, root, leader, designated, alternate, and backup.
- ▶ **Internal Path Cost:** Returns the path cost, a 1 Gbps link speed equals a path cost of 20,000, while a 100 Mbps link speed results in a cost of 200,000 and 10 Mbps is 2,000,000.
- ▶ **Designated Bridge ID:** A 64-bit bridge identifier consists of the first four MSB of the priority, then a system ID extension in the next 12 bits and the bridge address in the remaining 48 bits.
- ▶ **Designated Port ID:** A 16-bit port identifier consists of the first four MSB of the priority and a 12-bit bridge port number,
- ▶ **Disputed Port Status:** It will be set if a BPDU is received with a worse message priority with learning flag set and the port role indicates that it is designated (see 13.21 partial and disputed connectivity in IEEE 802.1Q 2022 standard)

FRAME REPLICATION AND ELIMINATION FOR RELIABILITY (FRER), 802.1CB

FRER aims to improve network reliability by reducing packet loss due to equipment failures as defined by the IEEE 802.1CB standard.

The Switch supports the replication of frames in the talker (source) and elimination of frames in the listener (destination). The talker sends replicated streams along two or more redundant paths with a redundancy tag added to the frame. The point of having these redundant paths is to minimize packet loss due to link failure, device failure, or stream congestion. The listener is then responsible for eliminating the duplicate packets. The network ensures that no matter which of the paths the stream takes, it arrives where and when it is supposed to. The Switch includes the sequence generation and recovery algorithms for FRER support.

Note that the FRER needs loop protection to prevent non-FRER/unknown traffic from creating a storm when connecting multiple ports of an FRER system. By default, MSTP is enabled and will protect the network against loops. The default behavior of the Switch is to flood unknown traffic on all ports when the lookup returns a miss. The miss behavior is configurable for Unicast and Multicast/Broadcast traffic. When using FRER function, multiple ports can be connected together between Switches, thereby, cre-

ating potential for a loop. MSTP prevents non-FRER traffic from circulating. By default, all VLANs are part of the CIST and therefore part of the spanning tree. To prevent MSTP blocking VLANs of interest to FRER, ensure to configure the VLAN exclusion list in the MSTP page. Consideration must be given to other traffic in these VLANs that do not have FRER applied, as neither MSTP or FRER will be active on that traffic; therefore, potential for a loop does exist for this type of traffic.

Figure 162 shows the **Candidate** page, which provides user ability to configure the FRER functionality and configure streams to apply FRER to within the same web page.




Figure 161. Frame Replication and Elimination Overview

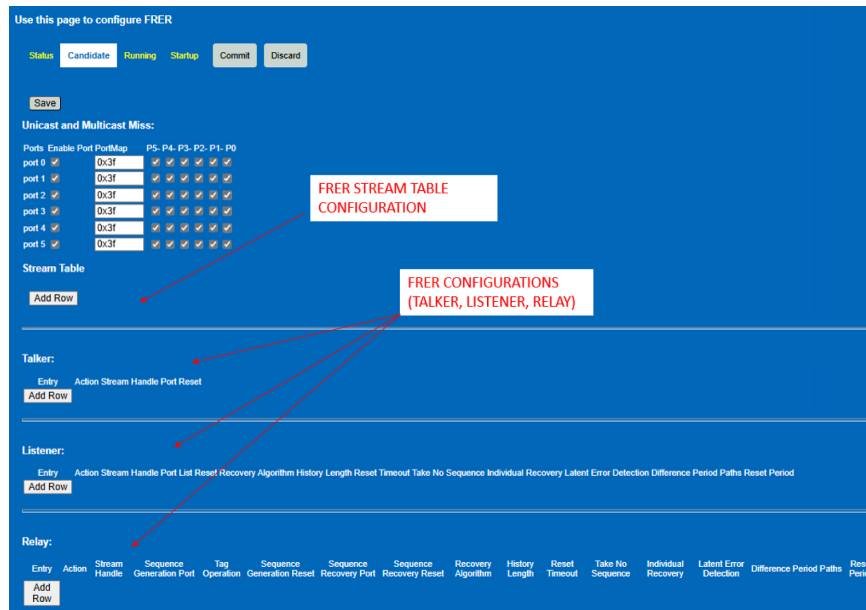


Figure 162. FRER Candidate Page Including Stream Table Configuration

FRAME REPLICATION AND ELIMINATION FOR RELIABILITY (FRER), 802.1CB

REDUNDANCY TAG

The Redundancy Tag (R-Tag) includes an Ether type (F1-C1) and a two-byte sequence number, which gets incremented for each frame

received by the generator. The sequence number is used in the listener recovery function to identify and eliminate duplicates of a frame. Figure 164 shows an R-Tag inserted into a frame.



Figure 163. FRER R-Tag




Figure 164. FRER R-Tag in Frame Observed through Wireshark

FRAME REPLICATION AND ELIMINATION FOR RELIABILITY (FRER), 802.1CB**STREAM IDENTIFICATION**

The stream identification function is used to determine the port routing for the streams. There are two types of stream identification: active or passive. Passive identification only examines the packets of the stream, while active identification modifies data parameters

of the packet to be transmitted. The Switch can be configured to handle all the different stream identifications described in IEEE 802.1CB. [Table 19](#) shows the different types of stream identification, the parameters they examine in a packet, and where applicable, the parameters they overwrite.

Table 19. Stream Identification Types

Stream Identification	Active/Passive	Examines	Overwrites
Null Stream	Passive	DA, VLAN ID	None
Source MAC and VLAN Stream	Passive	SA, VLAN ID	None
Active Destination MAC and VLAN Stream	Active	DA, VLAN ID	DA, VLAN ID, PCP
IP Stream	Passive	DA, VLAN ID, IP Source, Destination, DSCP, IP next protocol, source port, destination port	None
Mask-and-match Stream	Passive	DA, SA, MAC SDU	None

FRAME REPLICATION AND ELIMINATION FOR RELIABILITY (FRER), 802.1CB

SEQUENCE RECOVERY

There are two sequence recovery algorithms defined by IEEE802.1CB, Match algorithm or Vector algorithm. The Switch supports both recovery algorithms.

When using the Match algorithm, the Switch accepts the first packet received as valid. The sequence number in subsequent packets are evaluated based on their match status with the sequence number of the last accepted packet. If the sequence number matches the last accepted one, it is considered a duplicate and is discarded. If the sequence number does not match, the frame is accepted and forwarded, subsequent frames are compared against this new sequence number. Each accepted sequence number resets a timer. Expiration of the timer resets the algorithm and the next sequence number is accepted. The Match algorithm is ideal for intermittent streams where no more than one packet is in flight on any given path compared to other paths. The Match algorithm keeps count of the numbers of packets passed, discarded, and out-of-order. It also tracks the number of times the reset timer has expired.

The Vector algorithm provides a more robust duplicate elimination. Upon packet arrival, the Switch checks whether the sequence number falls within the range of the sequence number of the previously accepted packet. The acceptable range is defined as plus or minus the history length parameter. The default history length is 2 with a maximum length of 16. Any packets outside of the programmed range are discarded and duplicated packets within the history length are also discarded. Each time a packet is accepted, the timer restarts. When the timer expires, the algorithm resets, which allows acceptance of any sequence number in the next arriving packet. Increasing the history length of the Vector algorithm makes it suitable for scenarios of bulk streams, where there can be more than one packet in flight on any given path.

The Switch allows tracking of the various packet type for the Vector algorithm such as Passed Packets, Discarded Packets, Out-of-Order Packets, Rogue Packets, and Lost Packets.

INDIVIDUAL RECOVERY

The individual recovery addresses specific errors, such as a stuck transmitter, which repeatedly sends the same packet. When a transmitter gets stuck, it may send duplicate packets with the same sequence number. The duplicates can disrupt the reliability of the network. The individual recovery identifies repeating sequence

number within a single member stream and removes them early on. This allows early detection of errors. The individual recovery can be applied to each port of the Switch. Individual recovery is applied at the ingress port and performs the recovery function on each of the two member streams prior to them being presented to the sequence or compound recovery function that is applied to the egress port.

FRER STREAM TABLE

Using **FRER** requires configuration of the stream table section within the **FRER** page. The streams configured here are only used with **FRER** function and installed entries are only active when associated with an **FRER** configuration.

Stream Table Configuration

[Figure 165](#) shows an overview of the Stream table configuration and streams must be configured prior to configuring **FRER** functionality. This table is used to configure streams used by the **FRER** function:

- ▶ **Handle:** Number to differentiate between different streams.
- ▶ **Handle Alias:** Alias to nickname the Stream and easily map the handle to the **FRER** configuration page.
- ▶ **Port Map:** Determines where stream egress.
- ▶ **Identification Type:** Drop-down selection for the different stream identification type.

When configuring a Stream entry, the second row of the stream identification changes based on the identification type. [Figure 165](#) shows the different fields for different identification types based on [Table 19](#). The parameters to be examined need to be configured here. For example, in [Figure 165](#), the first entry is a Null Stream, where the VLAN ID and the destination MAC of interest are configured.

When installing stream entries, where **Stream Table** is false, the entry is installed in the regular static table entry space. Setting stream table to true, installs the entry into the **Stream Table** space. The **Stream Table** can support 16,000 entries, (16 blocks of 1024 entries). Enabling the stream table groups streams together in a single block based on the base MAC address, which is defined by the first 38 bits of the MAC address. The **Stream Table** is ideal for groups of devices in the same MAC address range.

FRAME REPLICATION AND ELIMINATION FOR RELIABILITY (FRER), 802.1CB

Entry	Action	Handle	Handle Alias	Port Map	Identification Type
1	Delete	100	nul_stream	PS: P4, P5, P0, P1, P6	Null Stream Identification
VLAN Identifier: useStreamtable Destination MAC Address: 00-00-00-11-11-11					
2	Delete	101	smac_stream	PS: P4, P5, P0, P1, P6	Source MAC and VLAN Stream
VLAN Identifier: 100 Source MAC Address: 00-00-00-22-22-22					
3	Delete	102	dmac_stream	PS: P4, P5, P0, P1, P6	Active Destination MAC and V
Down VLAN Identifier: 100 Up VLAN Identifier: 200 Up Priority: 1 Down Destination MAC Address: 00-00-00-33-33-33 Up Destination MAC Address: 00-00-44-44-44-44 useStreamtable: False					
4	Delete	103	IP_stream	PS: P4, P5, P0, P1, P6	IP Stream Identification
VLAN Identifier: 100 Destination MAC Address: 00-11-11-11-11-11 IP Source Address: 192.168.1.3 IP Destination Address: 192.168.1.15 Drop: 0 Next Protocol: RFC 768 Source Port: 0 Destination Port: 0 IPv4: checked					
5	Delete	104	mask_match_stream	PS: P4, P5, P0, P1, P6	Mask and Match Identification
Mask Length: 4 Mask: FF-FF-00-00 Match: FF-FF-00-00 Destination MAC Mask Address: FF-FF-FF-FF-FF-FF Destination MAC Match Address: 00-66-66-66-66-66 Source MAC Mask Address: FF-FF-FF-FF-FF-FF Source MAC Match Address: 00-77-77-77-77-77					

Figure 165. Overview of the Stream Table Section of the FRER Candidate Page

FRER CONFIGURATION—CANDIDATE VIEW

The Switch can be configured as a **Talker** for some streams, a **Listener** for others, or a **Relay** system (see Figure 166). The FRER configuration can be configured in the **FRER Configuration Candidate** page. The configuration parameters for the FRER function are broken into three distinct sections, as follows:

Talker:

- ▶ **Stream Handle:** Drop-down option of existing stream entries. This is mapped to the **Handle Alias** set in the **Stream Table** entry. **Stream Table** changes must be saved and committed to appear as an option in the **Stream Handle**.
- ▶ **Port:** Stream to apply FRER to, stream ingressing this port is replicated and R-Tag inserted.
- ▶ **Reset:** Set to enable (**True**) or disable (**False**) sequence generation reset.

Listener:

- ▶ **Stream Handle:** Drop-down option of existing stream entries. This is mapped to the **Handle Alias** set in the **Stream Table** entry. **Stream Table** changes must be committed to appear as an option in the **Stream Handle**.
- ▶ **Port List:** Ports where R-Tagged stream ingress. This defines which port to apply individual recovery.
- ▶ **Reset:** Set to enable (**True**) or disable (**False**) sequence recovery reset.
- ▶ **Recovery Algorithm:** Set to **Match** or **Vector** algorithm.
- ▶ **History Length:** Applies to **Vector** Algorithm, defines the **History Length** range.
- ▶ **Reset Timeout:** Set duration of reset before timeout.
- ▶ **Take No Sequence:** Determines whether frames without a sequence number are to be accepted (**True**) or not (**False**).
- ▶ **Individual Recovery:** Determines whether to enable (**True**) or disable (**False**) individual recovery.

- ▶ **Latent Error Detection:** Currently set to **False**, this is not yet available.

Relay:

- ▶ **Stream Handle:** This is mapped to the **Handle Alias** set in the **Stream Table**. **Stream Table** changes must be committed to appear as an option in the **Stream Handle**.
- ▶ **Sequence Generation Port:** Stream entering this port is replicated and R-Tag added. Set to NULL (-) if **Tag Operation** is set to **No Operation**.
- ▶ **Tag Operation:** Set whether to **Insert R-Tag**, **Remove R-Tag**, or **No Operation**.
- ▶ **Sequence Generation Reset:** Set to enable (**True**) or disable (**False**) sequence generation reset.
- ▶ **Sequence Recovery Port:** Port where R-Tagged stream ingress. This defines which port to apply recovery.
- ▶ **Sequence Recovery Reset:** Set to enable (**True**) or disable (**False**) sequence recovery reset.
- ▶ **Recovery Algorithm:** Set to **Match** or **Vector** algorithm.
- ▶ **History Length:** Applies to **Vector** Algorithm, defines the **History Length** range.
- ▶ **Reset Timeout:** Set duration of reset before timeout.
- ▶ **Take No Sequence:** Determines whether frames without a sequence number are to be accepted (**True**) or not (**False**).
- ▶ **Individual Recovery:** Determines whether to enable (**True**) or disable (**False**) individual recovery.
- ▶ **Latent Error Detection:** Currently set to **False**, this is not yet available.

FRAME REPLICATION AND ELIMINATION FOR RELIABILITY (FRER), 802.1CB

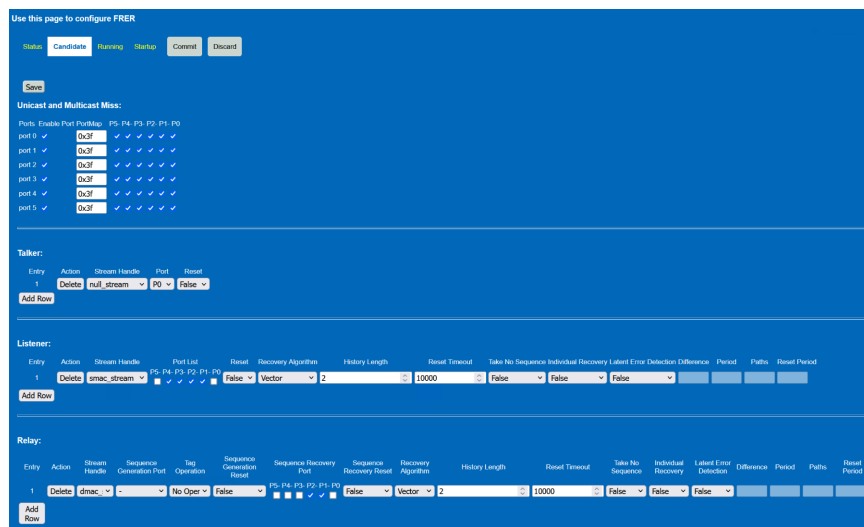


Figure 166. FRER Configuration Page

Unicast and Multicast Miss:

The default behavior of the Switch is to flood unknown traffic on all ports. The miss behavior is configurable for Unicast and Multicast/Broadcast traffic. When using **FRER** function, multiple ports may be connected together between Switches, thereby, creating potential for a loop. If MSTP is enabled, it can take care of loops for non-FRER traffic. Alternatively, to support evaluation of **FRER** functionality without MSTP, configuration of the default miss behavior can be used to prevent a loop. The following examples include configuration for the Miss return:

- ▶ **Enable Port:** Check box to enable Miss Return on specific port.
- ▶ **PortMap:** Set bit to 1 to route miss values to certain ports.
- ▶ **Port Check Box:** Check box to route miss values to this port.

FRAME REPLICATION AND ELIMINATION FOR RELIABILITY (FRER), 802.1CB

FRER STATUS

FRER statistics are shown in the **Status** view of the **FRER Configuration** page (see [Figure 167](#)). The **Status** page currently clears the count on read and does not accumulate, therefore, refreshing the page or navigating to another page and back to **Status** page clears any previous statistics captured. Statistics counts are provided for sequence recovery algorithm configurations for listener and relay functions. The **Status** page shows the following parameters on a per-stream per port basis:

- ▶ **Passed-pkts:** Count of packets successfully accepted.
- ▶ **Discarded-pkts:** Count of packets discarded.
- ▶ **Out-of-order-pkts:** Count for received packet with a sequence number older than a previous packet and not previously accepted.
- ▶ **Rogue-pkts:** Count of packets with sequence number beyond the history window length.
- ▶ **Lost-pkts:** Count of packets lost, the number of sequence numbers that were not received within the vector window.
- ▶ **Tagless-pkts:** Count of packets received without an FRER tag
- ▶ **Resets:** Count of the number of times the sequence recovery reset function is called.
- ▶ **Encoded-errored-pkts:** Count of packets that are unable to be decoded successfully.

The statistics provided depends on what recovery algorithm is active. Both vector and match algorithms track the number of passed,

discarded, out-of-order, and tagless packets in addition to capturing the number of times the reset timer has expired. Additionally, the vector algorithm captures rogue and lost packets.

To better understand the statistics counters, you may examine a few scenarios where multiple frames are processed by each algorithm and explore how each algorithm handles the incoming frames and updates its respective counters. These examples are intended to illustrate how the counters are updated in practice.

[Table 20](#) shows an example using the Match Algorithm. All frames listed ingress in the order shown and all match the stream identification parameters (for example, destination address and VLAN). The first arriving frame does not have an R-Tag and is forwarded, with both the tagless and encoder error counters incrementing. The next frame has an R-Tag of 0xF1C1 and a sequence number of 4. It is accepted, and the Match Algorithm stores this sequence number. The following frame has the same sequence number (4). Because it matches the stored sequence, it is deemed a duplicate and is discarded. As a result, the discard counter increments. The next frame has a sequence number of 8. Although it does not match the previous one and is therefore accepted, it is not in sequence. This causes the out-of-order counter to increment. The following frame has a sequence number of 9. It is accepted as the sequence progresses correctly. The final frame in the table has a sequence number of 7. Although it is not a duplicate, it is out of order, so it is accepted and the out-of-order counter increments again.

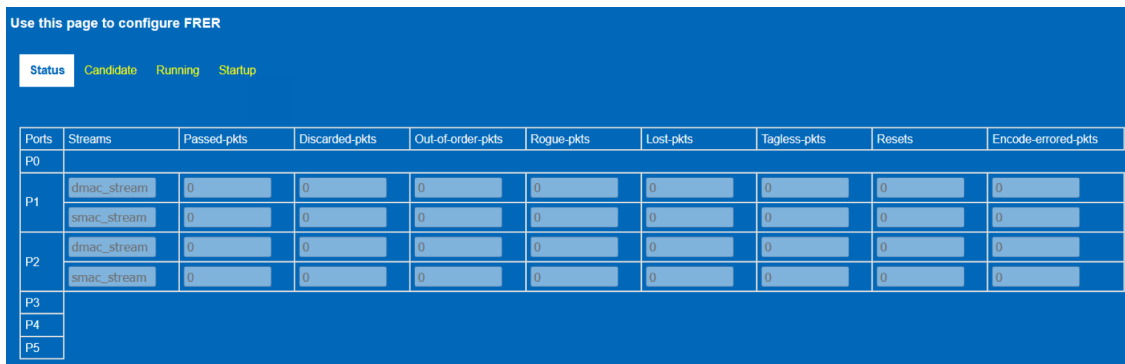


Figure 167. FRER Statistics

Table 20. Listener Match Algorithm Traffic Handling and Statistics Count

EtherType	Incoming Sequence Number	Match #	Counters				
			Passed	Discarded	Tagless	Out-of-Order	Encoder Error
0x1234	X	X	1	0	1	0	1
0xF1C1	4	X	1	0	0	0	0
0xF1C1	4	4	0	1	0	0	0
0xF1C1	8	4	1	0	0	1	0
0xF1C1	9	8	1	0	0	0	0
0xF1C1	7	9	1	0	0	1	0

FRAME REPLICATION AND ELIMINATION FOR RELIABILITY (FRER), 802.1CB

Next, you can examine how a vector-based algorithm manages incoming frames and updates its statistics counters, see [Table 21](#). In this example, the history length is 8 bits, and the algorithm maintains a history vector that tracks which recent sequence numbers have been received, see [Table 22](#). As each frame arrives, the algorithm checks the R-Tag sequence number. If it matches the current sequence number already stored in the history (bit field = 1), the frame is identified as a duplicate and discarded. If the incoming sequence number is lower than the current value, the history vector is consulted to determine whether the frame is a duplicate (and must be discarded) or a late arrival (and can still be accepted if bit field = 0). If the sequence number is higher than the current value, the algorithm updates both the current sequence number and the history vector to reflect the new state.

The first frame matches the stream identification parameters but does not contain an R-tag. As a result, it is discarded and both the encoder error and tagless counts are incremented. The second frame includes an R-Tag with a sequence number of 32 (for this example, we assume that no frames with sequence numbers prior to 32 have been received). This frame is accepted and forwarded. Sequence Number 32 is then considered the recovered sequence number, the sequence history (vector) is updated accordingly and the bit field for 32 is set to 1.

The third frame arrives with a sequence number of 30. According to the history vector, this frame has not yet been received (the bit corresponding to 30 is 0). Therefore, it is accepted and forwarded, and its bit in the history vector is set to 1. The recovered sequence number remains at 32, as this frame is older. Frame 4 arrives

with a sequence number of 35, which is higher than the current recovered sequence number. The algorithm updates the recovered sequence number to 35 and shifts the history vector to the left to accommodate the new frame. During this shift, Sequence Numbers 25, 26, and 27 fall outside the history window and are marked as lost (their bits were 0 as they were never received). The lost counter is incremented by 3 and the out of order counter is incremented because Frame 35 arrived out of order relative to Frame 30.

Frame 5 arrives with Sequence Number 32, which has already been seen. It is therefore a duplicate, discarded and the discard count is incremented. The recovered sequence number remains at 35. Frame 6 has a sequence number of 44, which is outside the acceptable history length (35 + 8 = 43). It is considered a rogue frame and discarded. The rogue counter is incremented. Frame 7 arrives with Sequence Number 41, which becomes the new recovered sequence number. The history vector shifts left, and Sequence Numbers 28 to 33 are shifted out. Since Frames 28, 29, 31, and 33 were never received (bits = 0), the lost counter is incremented by 4. The frame is accepted and the out-of-order counter is incremented.

Frame 8 arrives with Sequence Number 44, which is now within the valid window (since the recovered sequence number is 41). It is accepted and becomes the new recovered sequence number. The history vector shifts again, discarding Sequence Numbers 34, 35, and 36. Frames 34 and 36 were not received (bit = 0), so the lost counter is incremented by 2. Frame 8 is accepted and the out-of-order counter is incremented by 1.

Table 21. Listener Vector Algorithm Traffic Handling and Statistics Count

Frame Number	EtherType	Incoming Sequence Number	Action	Recovered Sequence Number	Counters							
					Passed	Discarded	Tagless	Rogue	Out-of-Order	Encoder Error	Lost	
1	0x1234	X	Discard	X	0	1	1	0	0	0	1	0
2	0xF1C1	32	Accept	X	1	0	0	0	0	0	0	0
3	0xF1C1	30	Accept	32	1	0	0	0	1	0	0	0
4	0xF1C1	35	Accept	32	1	0	0	0	1	0	0	3
5	0xF1C1	32	Duplicate	35	0	1	0	0	0	0	0	0
6	0xF1C1	44	Rogue	35	0	0	0	1	0	0	0	0
7	0xF1C1	41	Accept	35	1	0	0	0	1	0	0	4
8	0xF1C1	44	Accept	41	1	0	0	0	1	0	0	2

Table 22. History Vector and Bit Field for Vector Recovery Algorithm Example

Frame Number	Bit Field and History Vector										
1	0	0	0	0	0	0	0	0	0	0	0
	X	X	X	X	X	X	X	X	X	X	X
2	0	0	0	0	0	0	0	0	0	0	1
	25	26	27	28	29	30	31	32			
3	0	0	0	0	0	1	0	0	0	0	1
	25	26	27	28	29	30	31	32			

FRAME REPLICATION AND ELIMINATION FOR RELIABILITY (FRER), 802.1CB

Table 22. History Vector and Bit Field for Vector Recovery Algorithm Example (Continued)

Frame Number	Bit Field and History Vector							
4	0	0	1	0	1	0	0	1
	28	29	30	31	32	33	34	35
5	0	0	1	0	1	0	0	1
	28	29	30	31	32	33	34	35
6	0	0	1	0	1	0	0	1
	28	29	30	31	32	33	34	35
7	0	1	0	0	0	0	0	1
	34	35	36	37	38	39	40	41
8	0	0	0	0	1	0	0	1
	37	38	39	40	41	42	43	44

FRAME REPLICATION AND ELIMINATION FOR RELIABILITY (FRER), 802.1CB

TALKER–LISTENER CONFIGURATION EXAMPLE

The following example shows a configuration for a **Talker** and a **Listener** system. In this example, each Switch is configured independently from a separate PC/web server, but the GUI does support configuration of multiple individual switches from one PC either daisy-chained or through individual network adapters.

Two types of traffic streams are used in this example, a null stream with destination MAC address 00:00:00:11:11:11 and a source MAC/VLAN stream with source address 00:00:00:22:22:22. A VLAN ID of 100 is used for both streams. For the Talker configuration, install the stream identification entries for each type of stream. The **Port Map** is configured for Port P1 and Port P2 as egress ports.

Configure the talker system (**Talker**) in the FRER **Candidate** page as shown in [Figure 169](#). Sequence generation port is set to Port P0 for each stream. Traffic ingressing on Port P0 matching the

destination address for the null stream and the source address for the other stream will be replicated by the FRER functionality and routed out both P1 and P2 ports.

On the Listener side, install the stream identification entries in the **Stream Table**. For the listener, the **Port Map** is set to Port 0 to egress streams at this port.

For the Listener configuration, the example uses a **Vector** sequence recovery algorithm with a **History Length** of 2. The sequence recovery algorithm is applied at the port where the streams merge, in this case Port P0.

For both devices, configure the VLAN mode in the **VLAN Table** page. Port 0, Port 1, and Port 2 are set to learn and forward mode for the VID of interest (100) as shown in [Figure 170](#). The user must also review if changes are required in the MSTP configuration to exclude FRER VLAN from the CIST, see [Figure 171](#). In this case, exclude VLAN ID from MSTP.




Figure 168. Talker–Listener Configuration Example

Figure 169. Talker and Listener FRER Configuration Example

FRAME REPLICATION AND ELIMINATION FOR RELIABILITY (FRER), 802.1CB




Figure 170. VLAN Table Configuration

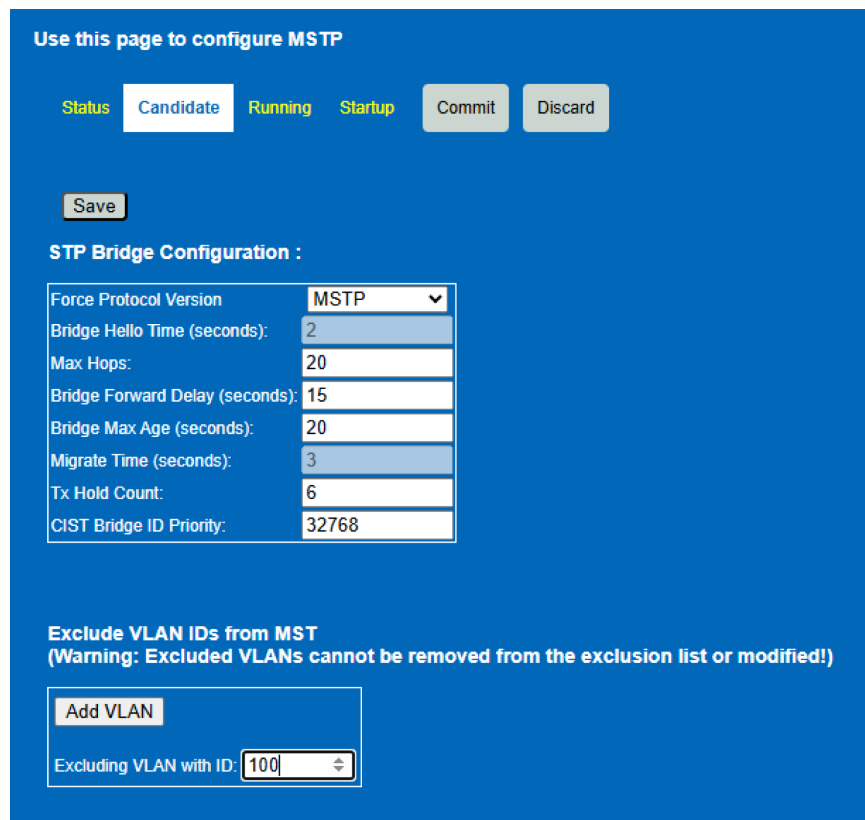


Figure 171. MSTP VLAN Exclusion Configuration

INTERNET GROUP MANAGEMENT PROTOCOL (IGMP) SNOOPING

IGMP snooping is a method network Switches use to identify multicast groups and forward packets accordingly. Multicast is a one-to-many communication method where data is sent from one source to multiple specific destinations. In a multicast setup, data packets are addressed to a specific group of devices that have expressed interest in receiving the data, which makes it more efficient than broadcast for targeted communication. As such, multicasting is highly efficient as it reduces unnecessary data transmission and processing with only the Hosts that need the data receiving it, thus conserving bandwidth and reducing processing load on uninterested Hosts.

IGMP snooping works by the Switch observing IGMP network traffic and using this information to map the ports of interest in a particular multicast group in order to control traffic flow. The Switch can support IGMP snooping (versions 1, 2, and 3).

IGMP messages are sent by devices informing their intention to join or leave a multicast group. The Switch snoops on these messages and maintains an internal map of which ports are members of which IP multicast transmission, which ensures that the multicast traffic is only sent to the Hosts that have requested it. The example shown in [Figure 172](#) is a scenario where the multicast source out on Port 2 sends IGMP queries, and devices on Port 0 and Port 5 send IGMP reports indicating interest in this multicast group. Subsequently, the Switch controls the flow of traffic to ensure only these two ports receive this particular multicast traffic.

There is no status information provided as part of this feature and no visibility into the internal mapping of ports to IP multicast.




Figure 172. Example of IGMP

IGMP snooping is disabled by default. To use IGMP snooping, enable the function in the **IGMP Snooping Configuration** page. By default, both the **Router Timeout** and the **Group Member Timeout** are programmed to 260 seconds. Click the **Save** button followed by **Commit** button for any changes. The page automatically shows the **Running** view and the Switch starts monitoring the IGMP traffic crossing the Switch and handle any multicast traffic accordingly.

ROUTER TIMEOUT

The **Router Timeout** is the duration for which a Switch considers a multicast router to be present on a particular port. When a Switch receives IGMP queries from a router on a port, it marks that port as having an active multicast router.

The purpose of the **Router Timeout** is to ensure that the Switch does not keep forwarding multicast traffic to a port where the multicast router is no longer active. This helps prevent unnecessary flooding of multicast traffic. The default timeout is 260 seconds.

INTERNET GROUP MANAGEMENT PROTOCOL (IGMP) SNOOPING

GROUP MEMBER TIMEOUT

The **Group Member Timeout** is the duration for which a Switch considers a Host to be a member of a particular multicast group. When a Switch receives IGMP membership reports from Hosts on specific ports, it marks those ports as having active members for the corresponding multicast groups.

The purpose of the **Group Member Timeout** is to ensure that the Switch does not keep forwarding multicast traffic to ports where there are no longer interested receivers for a particular multicast group. The default timeout value is 260 seconds.

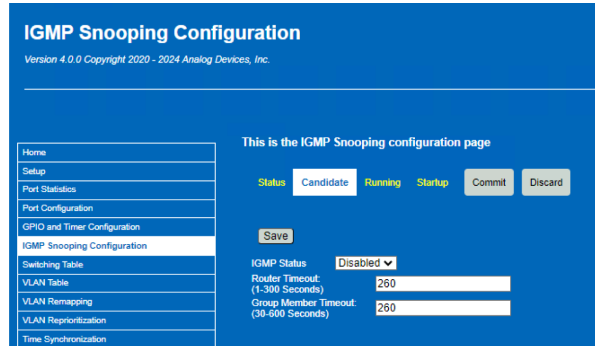


Figure 173. IGMP Snooping Web Page – Candidate View

INTERNET GROUP MANAGEMENT PROTOCOL (IGMP) SNOOPING

IGMP VERSIONS

IGMPv1:

- ▶ **Basic Operation:** Hosts send IGMP Reports to join groups. The router periodically sends IGMP Queries to verify group memberships.
- ▶ **Leave Group:** Hosts do not explicitly send Leave messages. The Switch relies on the absence of Reports to determine if a Host has left the group.

IGMPv2:

- ▶ **Enhanced Operation:** Similar to IGMPv1 but includes the ability for Hosts to send Leave Group messages, which allows a quicker leave detection.
- ▶ **Leave Group:** Hosts can send IGMP Leave messages, prompting the router to send a Group-Specific Query to verify if any other Hosts are still interested in the group.

IGMPv3:

- ▶ **Advanced Features:** Supports source-specific multicast (SSM), which allows Hosts to specify which sources they want to receive traffic from within a multicast group.
- ▶ **Leave Group:** Enhanced leave mechanisms and ability to manage memberships based on specific sources.

IGMP SNOOPING EXAMPLE

With IGMP snooping disabled, multicast traffic is visible on all ports. IGMP report messages sent from any port to join a multicast group are ignored and all ports continue to receive the multicast traffic.

When IGMP snooping is enabled, consider the following scenario (see Figure 174) where an IGMP report is sent from the device connected to Port P3 (Ethernet_5) to join a multicast group. At this point, a multicast channel is established between Port P0 (Ethernet_6) and Port P3 (Ethernet_5). All multicast traffic entering Port P0 are forwarded only to Port P3 and not to the other ports.




Figure 174. IGMP Snooping Example

This configuration remains active until the **Group Member Timeout** is reached, which is 40 seconds in this case, per configuration in Figure 175. If no new IGMP report is received within this period, the Switch deletes the multicast group, and any multicast traffic of this group sent into Port P0 is treated as broadcast and is forwarded on all ports.

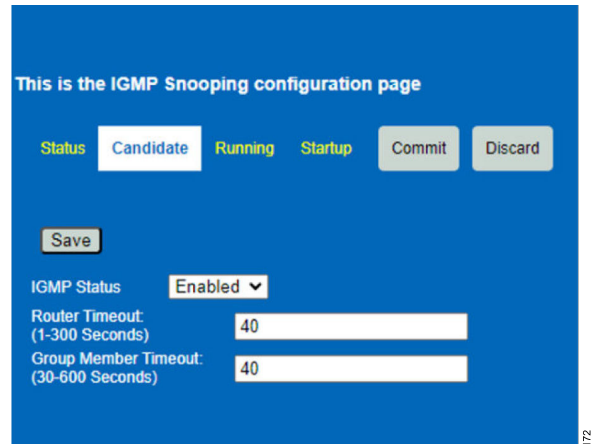


Figure 175. IGMP Snooping Configuration in Web Server

The Wireshark shown in Figure 176 is the behavior of the network traffic. The group of messages highlighted in orange occurs within the 40 second timeout period. During this time, multicast traffic sent into Port P0 is seen only by Port P3 (Ethernet_5). After the 40 second period, since no further IGMP report is received, the Switch deletes the multicast group. The subsequent multicast traffic sent into Port P0 is treated as broadcast and forwarded on all ports, as highlighted in red.

With IGMPv1, a device that is part of the multicast group continue receives multicast traffic until the **Group Member Timeout** elapses. It is not possible for a device to stop receiving before the timeout with IGMPv1. This limitation is addressed in IGMPv2, which introduces Leave Group messages allowing ports to explicitly signal when they no longer wish to receive multicast traffic.

INTERNET GROUP MANAGEMENT PROTOCOL (IGMP) SNOOPING

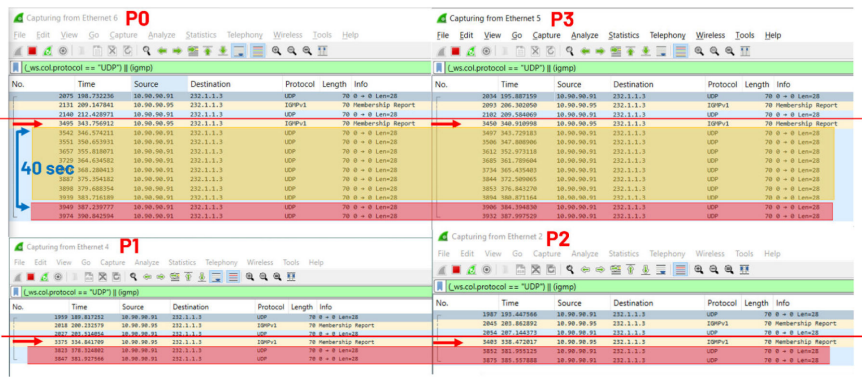


Figure 176. IGMPv1 Multicast Group Created and Group Member Timeout

INTERNET GROUP MANAGEMENT PROTOCOL (IGMP) SNOOPING

IGMPv2

The use of IGMP version 2 is demonstrated to show that it is possible for a Host to leave a multicast group before the **Group Member Timeout**. The scenario is shown in [Figure 177](#). The first multicast UDP packet (length 70) is sent into Port P0 (Ethernet_6) and gets forwarded on all other ports. A membership report is then sent into Port P3 (Ethernet_5), followed by two more UDP packets of length 70 are sent into Port P0. The result is that these multicast UDP packets are seen only on Port P3 (Ethernet_5) since it joined the multicast group with an IGMPv2 report. At this point, an IGMPv2 leave message is sent into Port P3, and another two UDP packets

(this time with a length of 80, to distinguish them in Wireshark) are sent again into Port P0 (Ethernet_6). These packets are visible on all ports because Port P3 has left the group, and if the Switch does not see any members, it treats the packets as broadcast.

This behavior shows that IGMPv2 allows a Host to dynamically manage its membership in multicast groups, which provides greater flexibility and efficiency in network traffic management. By leaving the group, the Host ensures it no longer receives unnecessary multicast traffic, thereby, optimizing network resources and performance.

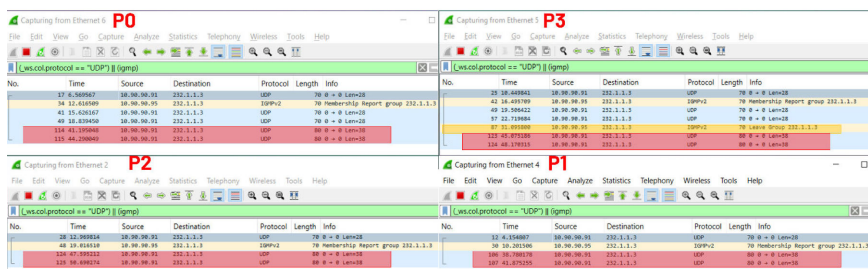


Figure 177. IGMPv2 – Leaving Message Scenario on Port P3

INTERNET GROUP MANAGEMENT PROTOCOL (IGMP) SNOOPING

IGMPv3

IGMPv3 provides the same functionality as IGMPv2 but also supports SSM, which allows a device to join a multicast group and specify from which server it wants to receive traffic. For example, if two servers (10.90.90.1 and 10.90.90.2) are sending multicast traffic to group 239.1.1.1, a device on Port P3 of the Switch can request traffic only from server 10.90.90.1 when it sends the membership report. The Switch is fully compatible with IGMPv3 queries, report, and leave messages. However, SSM selection is not supported.

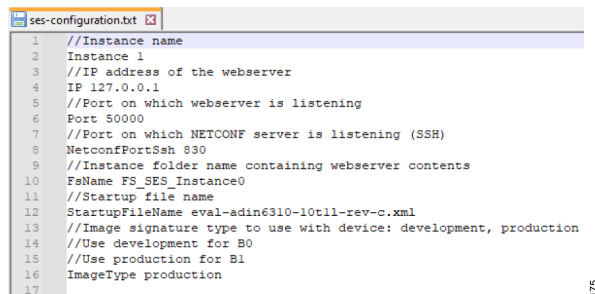
NETCONF/YANG

The **TSN Switch Evaluation** application runs a NETCONF server (netopeer2) on the Windows platform in parallel to the web server application. The server is not running on the Switch itself, instead on the PC acting as a Host. The server supports SSH.

The user can run a NETCONF client to configure the Switch functionality, the user name can be anything, the password must be blank.

The datastore repository is based on [Sysrepo](#). The repository supports four datastores: **Startup**, **Candidate**, **Running**, and **Operational**. All YANG models used by the Switch follow the models as required in IEC/IEEE 60802 with Switch specific features included in the YANG models as custom leaf nodes.

To use a NETCONF client to configure the device, first launch the GUI and **Find and Configure** the Switch or chain of Switch devices in the usual way described in the [ses-configuration File](#) section. This launches a windows-tsn-io process and NETCONF server for each device. Each server instance listens on the SSH port defined in the ses-configuration.txt file. There is a different SSH port for each Switch if there is a chain of Switches. By default, the first device is accessed through SSH port 830, the second via 50831.



```

1 //Instance name
2 Instance 1
3 //IP address of the webserver
4 IP 127.0.0.1
5 //Port on which webserver is listening
6 Port 50000
7 //Port on which NETCONF server is listening (SSH)
8 NetconfPortSsh 830
9 //Instance folder name containing webserver contents
10 FsName FS_SES_Instance0
11 //Startup file name
12 StartupFileName eval-adin6310-10c11-rev-c.xml
13 //Image signature type to use with device: development, production
14 //Use development for B0
15 //Use production for B1
16 ImageType production
17

```

Figure 178. ses-configuration.txt File Showing Netconf SSH Port

SYSREPO DATASTORE

The datastore model implemented by Sysrepo is defined in RFC 8342 Section 5 and includes four datastores defined as follows:

- ▶ **Startup:** The startup configuration datastore (<startup>) is a configuration datastore holding the configuration loaded by the device when it boots.
- ▶ **Candidate:** The candidate configuration datastore (<candidate>) is a configuration datastore that can be manipulated without impacting the device's current configuration and that can be committed to <running>.
- ▶ **Running:** The running configuration datastore (<running>) is a configuration datastore that holds the current configuration of the device.
- ▶ **Operational:** The operational state datastore (<operational>) is a read-only datastore that consists of all config true (ct) and config false (cf) nodes defined in the datastore's schema.

This means the <operational> datastore is a superset of <running> that augments configuration data (ct) with state data (cf).

As a result, retrieving data from <operational> is different from retrieving data from any of the other three datastores, because it uses device-specific code to provide the data whereas reading from <startup>/<candidate>/<running> return the values stored in the datastore itself.

Note that the RFC 8342 includes the <intended> datastore, which is a read-only version of <running>, created after transformations are applied to the configuration stored in <running> prior to applying configuration to the device. But also note that for simple implementations, <running> and <intended> are identical, which is what Sysrepo implements.

YANG MODELS

The **modules** folder in the evaluation package contains the YANG modules relevant for the switch. This includes IEEE and IETF models in addition to Switch custom leaf nodes. User can save the candidate file in XML format in the [Setup Page](#) to view and edit the YANG parameters and use as a template for models including custom leafs. To see all parameters associated with a feature, ensure to configure the relevant feature within in the web server and then save the candidate file.

CUSTOM LEAF NODES

Custom leaf nodes are used for Switch specific functionality or functionality that is not currently included in the standard YANG modules. The file ses.yang has all the custom leaf nodes. A summary are as follows:

- ▶ **Port Configuration:** MII modes, PHY types
- ▶ **VLAN operation:** Learning/Forwarding
- ▶ **Switching mode:** Cut-through or Store and forward
- ▶ **Crossover type (PHY) Auto/Man-MDIX/MDI**
- ▶ **Port Statistics**
- ▶ **Lookup-Types:** Combinations of lookups (destination, source, and extended)
- ▶ **Timer/GPIO function**
- ▶ **Time Synchronization:** PHY delays
- ▶ **Scheduled Traffic Guard-Bands**
- ▶ **Frame Preemption:** Peer-supported/enabled/active. Preemption enabled, ignore-peer, fragment size, and statistics
- ▶ **Frame replication and elimination for reliability**
- ▶ **Stream identification**

STARTUP CONFIGURATION

The **ses-default-startup.xml** in the **modules** folder contains all the **Startup** configuration details for the device. In the startup file, custom leafs can be seen with a prefix of adi:ses. This startup file is only used during the initial configuration for SES and from then on, the startup datastore contents are used to initialize the device. The datastore is contained in the **Repository** folder in the file system for each device.

NETCONF/YANG

Netconf-setup.xml file in the **modules** folder contains configuration data for the NETCONF server, it is not intended the user edits this file when using the evaluation package.

WEB SERVER USE AND NETCONF

The web server and NETCONF server both share the same connection to Sysrepo. Sysrepo manages the access to the datastore, therefore the web server is still available to use and any changes made in the web server are reflected in the datastore and vice versa.

YANG MODEL EXAMPLES

The **Setup** page allows user to save the candidate YANG file in XML format. Go to the **Setup** page and click **Save Candidate as XML**. The candidate file is saved to the **Downloads** location and captures the YANG model parameters based on the current repository candidate.

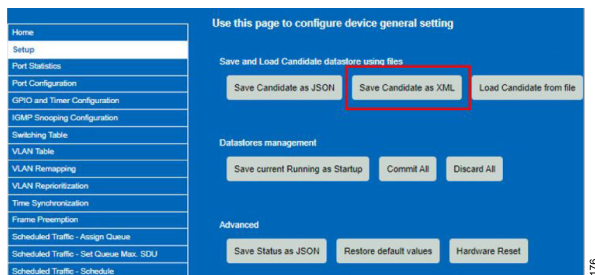


Figure 179. Save Candidate in Setup Page

Cut-Through vs. Store and Forward

The following snippet of code shows an example YANG model that a NETCONF client sends to the server to enable store and forward operation on Port 0 (P0) for queues 2, 3, and 4. Switching mode is a custom leaf, therefore has the prefix `adi:ses`.

```
<interfaces xmlns="urn:ietf:params:xml:ns:yang:ietf-interfaces"
  xmlns:sched="urn:ieee:std:802.1Q:yang:ieee802-dot1q-sched"
  xmlns:ses="urn:adi:ses">
  <interface>
    <name>P0</name>
    <bridge-port xmlns="urn:ieee:std:802.1Q:yang:ieee802-dot1q-bridge">
      <switching-mode xmlns="urn:adi:ses">
        <queue0>cut-through</queue0>
        <queue1>cut-through</queue1>
        <queue2>store-and-forward</queue2>
        <queue3>store-and-forward</queue3>
        <queue4>store-and-forward</queue4>
        <queue5>cut-through</queue5>
```

```
<queue6>cut-through</queue6>
<queue7>cut-through</queue7>
</switching-mode>
</bridge-port>
</interface>
</interfaces>
```

Scheduled Traffic Example

The following snippet of code shows an example YANG model the NETCONF client sends to the server to enable Scheduled traffic on Port 0 (P0), with a Gate Control List of 3, first entry for 250 μ s with gate 0 open, second entry another 250 μ s with gate 1 open and remainder of the cycle time (1 ms) with all gates open. Guard bands are disabled in this example.

```
<interfaces xmlns="urn:ietf:params:xml:ns:yang:ietf-interfaces"
  xmlns:sched="urn:ieee:std:802.1Q:yang:ieee802-dot1q-sched"
  xmlns:ses="urn:adi:ses">
  <interface>
    <name>P0</name>
    <bridge-port xmlns="urn:ieee:std:802.1Q:yang:ieee802-dot1q-bridge">
      <gate-parameter-table xmlns="urn:ieee:std:802.1Q:yang:ieee802-dot1q-sched-bridge">
        <gate-enabled>true</gate-enabled>
        <admin-gate-states>255</admin-gate-states>
        <admin-cycle-time>
          <numerator>1</numerator>
          <denominator>1000</denominator>
        </admin-cycle-time>
        <admin-base-time>
          <seconds>0</seconds>
          <nanoseconds>0</nanoseconds>
        </admin-base-time>
        <admin-control-list>
          <gate-control-entry>
            <index>0</index>
            <operation-name>sched:set-and-release-mac</operation-name>
            <gate-states-value>1</gate-states-value>
            <time-interval-value>250000</time-interval-value>
          </gate-control-entry>
          <gate-control-entry>
            <index>1</index>
            <operation-name>sched:set-and-re
```

NETCONF/YANG

```

lease-mac</operation-name>
  <gate-states-value>2</gate-states-
value>
  <time-interval-value>250000</time-
interval-value>
  </gate-control-entry>
</gate-control-entry>
  <index>2</index>
  <operation-name>sched:set-and-re▶
lease-mac</operation-name>
  <gate-states-value>255</gate-
states-value>
  <time-interval-value>500000</time-
interval-value>
  </gate-control-entry>
</admin-control-list>
  <admin-cycle-time-extension>0</admin-
cycle-time-extension>
  <config-change>true</config-change>
  <ses:guard-band-gate-event>>false</
ses:guard-band-gate-event>
  <ses:guard-band-hold-event>>false</
ses:guard-band-hold-event>
  </gate-parameter-table>
</bridge-port>
</interface>
</interfaces>

```

Frame Preemption Example

The following snippet of code shows the example YANG configuration the NETCONF client sends to the server to enable Frame Preemption on Port 3 with Queue 5/Priority 5 configured as express and all other queues preemptable.

```
<interfaces xmlns="urn:ietf:par▶
```

```

ams:xml:ns:yang:ietf-interfaces">
  <interface>
    <name>P3</name>
    <type
xmlns:ianaift="urn:ietf:params:xml:ns:yang:ia▶
na-if-type">ianaift:ethernetCsmacd</type>
    <enabled>>true</enabled>
    <bridge-
port xmlns="urn:ieee:std:802.1Q:yang:ieee802-
dot1q-bridge">
      <frame-preemption-param▶
ters xmlns="urn:ieee:std:802.1Q:yang:ieee802-
dot1q-preemption-bridge">
        <frame-preemption-status-table>
          <priority0>preemptable</priority0>
          <priority1>express</priority1>
          <priority2>express</priority2>
          <priority3>preemptable</priority3>
          <priority4>preemptable</priority4>
          <priority5>express</priority5>
          <priority6>preemptable</priority6>
          <priority7>preemptable</priority7>
        </frame-preemption-status-table>
        <preemption-enabled
xmlns="urn:adi:ses">true</preemption-enabled>
        <ignore-peer
xmlns="urn:adi:ses">>false</ignore-peer>
        <verify-disable
xmlns="urn:adi:ses">>false</verify-disable>
        <fragment-size xmlns="urn:adi:ses">0</
fragment-size>
        <verify-period xmlns="urn:adi:ses">10</
verify-period>
      </frame-preemption-parameters>
    </bridge-port>
  </interface>
</interfaces>

```

FIRMWARE UPDATE

The device supports firmware updates over the Host interface (SPI or Ethernet). With the **TSN Switch Evaluation** application, firmware updates are supported over the Ethernet port.

Firmware is developed by Analog Devices only and updates are provided through the product web page to add features and implement bug fixes as software development progresses.

AUTOMATIC FIRMWARE UPDATE

The software package supports automatic firmware update. Simply power on the hardware, connect to the Host port, launch the new software package, and the GUI application coordinates loading the new firmware.

PAIRED FIRMWARE AND WEB SERVER

Firmware (binary file) and web server files are paired and only function together in corresponding package pair. When new packages

come available, migrate to the newer package and continue to use that new version.

FIRMWARE DOWNGRADE

Firmware running on a device updates automatically when a new package is first run. The **Firmware Update** page is available if user wanted to roll back to a previous version of firmware or check what version is currently present. If rolling back to a previous version of firmware, then once the firmware update is complete to the older version, user must revert to use the matching GUI application that came with that firmware version. Alternatively, revert to an older version by using the previous GUI application, it downgrades the firmware on the device to match the firmware supported by that GUI version.

Figure 180. Firmware Update Page

FIRMWARE UPDATE

Perform the following steps:

1. Browse and upload the firmware image binary and click the **Submit** button.

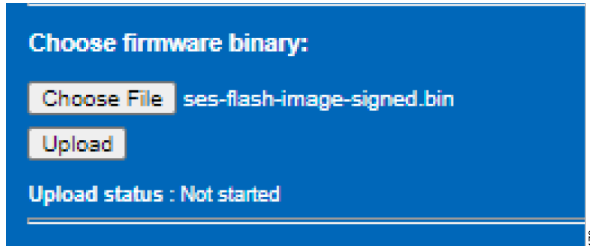


Figure 181. Firmware Update – Navigate and Select the Binary File

2. Figure 182 shows the progress of the file upload into the web server.




Figure 182. Firmware Update – Upload to Web Server

3. The web page refreshes to show progress as the upload progresses. Wait until the Progress bar reaches 100%, now the firmware upload has completed. The device automatically gets software reset once upload is complete, so the connection with the Switch is lost (MAC address gets reset to default).




Figure 183. Firmware Update – Firmware Upload Complete

4. To reestablish communication with the device, return to the GUI. Using the keyboard **Ctrl** button, click **Close All Running Processes**. All LEDs must turn off on the GUI.

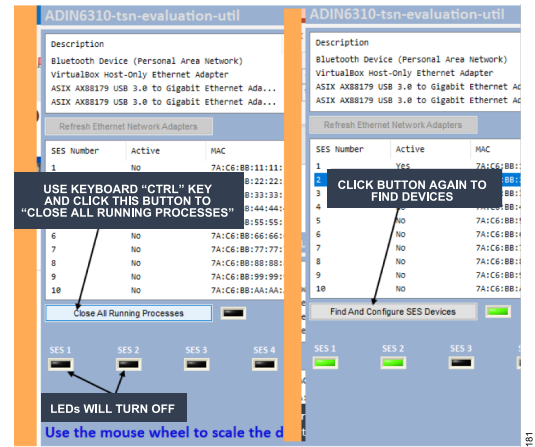


Figure 184. Firmware Update – Close Processes and Search for Devices Again

5. Click **Find And Configure SES devices** to identify and connected boards (shown with green LEDs) the devices again.
6. Firmware version and Bootloader versions appears on the **Firmware Update** page.

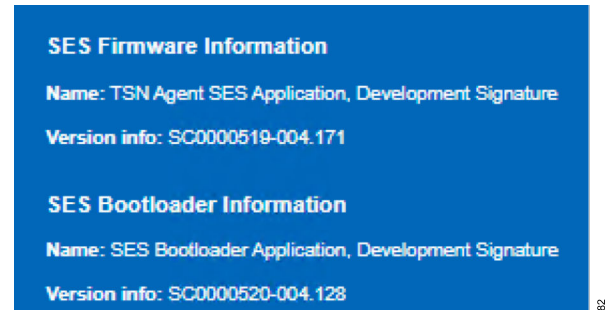


Figure 185. Firmware Update Version

TROUBLESHOOTING

GUI DOES NOT FIND ADIN6310 DEVICES

If the GUI does not find the ADIN6310 devices, do the following steps:

1. Check the board is powered on (blue power LED lights up near P2 connector).
2. Check that the correct Network Adapter is selected and the speed of the link established is 1 Gbps.
3. Check the Ethernet Cable on this Network Adapter is connected to the Host Port 0 Ethernet Port. When a device is found, the SES LED turns green.




Figure 186. No ADIN6310 Devices Found

4. Check the network adapter used and whether it has a property MAC passthrough enabled. This may cause it to pick up a different MAC address and interfere with ability of the GUI to communicate with the switch. Disable this property for the adapter, see Figure 187 and restart the GUI.

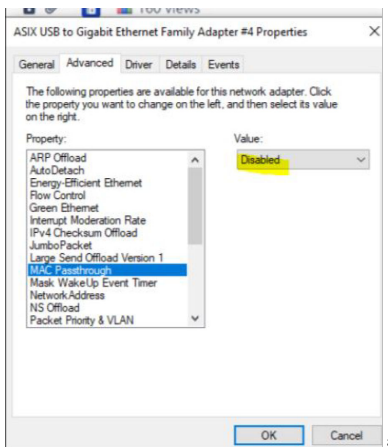


Figure 187. Modify Adapter Property for MAC Passthrough

5. Check the speed of the link established by the Network Adapter. The ADIN1300 PHY on the Host port (Port 0) autonegotiate with the Network Adapter to bring up a link of the highest common speed. The Switch port strapping (Table 23) must

match the speed of the link established. If for example, the established link is 100 Mbps and the Switch strapping configures the Switch Host port for 1 Gbps, then there is a mismatch in the Switch Host Port speed and the PHY speed. Default hardware strapping for the Switch Host port is for RGMII 1 Gbps speed.

Table 23. Host Port Strapping Link Configuration

RGMII	Timer3	Timer2	Timer1	Timer0	SPI_SS
RGMII 1 Gbps	INSERT	OPEN	INSERT	INSERT	INSERT
RGMII 100 Mbps	OPEN	OPEN	INSERT	INSERT	INSERT

GUI TABLE REMAINS BLANK

If the GUI table remains blank, even after double-clicking the selected NIC connected to the board, this may be due to the Npcap installation. Older versions of Npcap had an option for legacy loopback, if using an older version, check if it is installed with Legacy loopback support enabled. Try reinstalling Npcap with this disabled and then launch the GUI again. Double-clicking the Network Adapter connected to the Switch Port 0 should fill the GUI table.

Also verify that Npcap is installed with WinPCAP API-compatible mode enabled as discussed in Npcap Installation section.

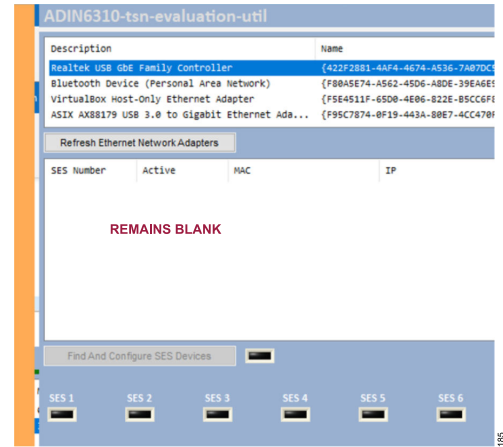


Figure 188. GUI Table Does Not Fill, Find And Configure SES Devices Button Stays Greyed Out

WEB PAGE FAILS TO LOAD

The Browser page returns to the This page isn't working. Do the following steps:

1. Check if the GUI application is still running. Aborting the application while using the web pages stops the web page from communicating with the ADIN6310 devices. Keep the GUI running while using the web pages.
2. Power cycle the board, restart the GUI and search for the Switch again.
3. Try changing the URL to 127.0.0.5:50000 or 127.0.5.1:50000.

TROUBLESHOOTING




Figure 189. No Response from Web Page

4. Navigate to **FileSystemFolders/FS_SES_InstanceX**, delete the repository folder and its contents, power cycle the device, restart the GUI and start searching again.

FIRMWARE DID NOT UPDATE

If the Firmware did not update, do the following steps:

1. Confirm correct .bin file is loaded.
2. Check installation path for software. Default path is **C:\Analog\ADIN6310EVKSW-Rel.x.x**. If installing to another location, ensure there are no spaces in the folder names. Use one-word folder name, for example, **AnalogDevices** instead of Analog Devices.

GUI INCONSISTENT AT FINDING DEVICES

If GUI is inconsistent at finding devices, do the following step:

1. Using a USB-Ethernet Network Adapter can result in some inconsistent/unstable connection between the PC Host and the Switch. Where possible, use an integrated Network Adapter for configuration of the Switch.



ESD Caution

ESD (electrostatic discharge) sensitive device. Charged devices and circuit boards can discharge without detection. Although this product features patented or proprietary protection circuitry, damage may occur on devices subjected to high energy ESD. Therefore, proper ESD precautions should be taken to avoid performance degradation or loss of functionality.

Legal Terms and Conditions

By using the evaluation board discussed herein (together with any tools, components documentation or support materials, the "Evaluation Board"), you are agreeing to be bound by the terms and conditions set forth below ("Agreement") unless you have purchased the Evaluation Board, in which case the Analog Devices Standard Terms and Conditions of Sale shall govern. Do not use the Evaluation Board until you have read and agreed to the Agreement. Your use of the Evaluation Board shall signify your acceptance of the Agreement. This Agreement is made by and between you ("Customer") and Analog Devices, Inc. ("ADI"), with its principal place of business at One Analog Way, Wilmington, MA 01887-2356, U.S.A. Subject to the terms and conditions of the Agreement, ADI hereby grants to Customer a free, limited, personal, temporary, non-exclusive, non-sublicensable, non-transferable license to use the Evaluation Board FOR EVALUATION PURPOSES ONLY. Customer understands and agrees that the Evaluation Board is provided for the sole and exclusive purpose referenced above, and agrees not to use the Evaluation Board for any other purpose. Furthermore, the license granted is expressly made subject to the following additional limitations: Customer shall not (i) rent, lease, display, sell, transfer, assign, sublicense, or distribute the Evaluation Board; and (ii) permit any Third Party to access the Evaluation Board. As used herein, the term "Third Party" includes any entity other than ADI, Customer, their employees, affiliates and in-house consultants. The Evaluation Board is NOT sold to Customer; all rights not expressly granted herein, including ownership of the Evaluation Board, are reserved by ADI. CONFIDENTIALITY. This Agreement and the Evaluation Board shall all be considered the confidential and proprietary information of ADI. Customer may not disclose or transfer any portion of the Evaluation Board to any other party for any reason. Upon discontinuation of use of the Evaluation Board or termination of this Agreement, Customer agrees to promptly return the Evaluation Board to ADI. ADDITIONAL RESTRICTIONS. Customer may not disassemble, decompile or reverse engineer chips on the Evaluation Board. Customer shall inform ADI of any occurred damages or any modifications or alterations it makes to the Evaluation Board, including but not limited to soldering or any other activity that affects the material content of the Evaluation Board. Modifications to the Evaluation Board must comply with applicable law, including but not limited to the RoHS Directive. TERMINATION. ADI may terminate this Agreement at any time upon giving written notice to Customer. Customer agrees to return to ADI the Evaluation Board at that time. LIMITATION OF LIABILITY. THE EVALUATION BOARD PROVIDED HEREUNDER IS PROVIDED "AS IS" AND ADI MAKES NO WARRANTIES OR REPRESENTATIONS OF ANY KIND WITH RESPECT TO IT. ADI SPECIFICALLY DISCLAIMS ANY REPRESENTATIONS, ENDORSEMENTS, GUARANTEES, OR WARRANTIES, EXPRESS OR IMPLIED, RELATED TO THE EVALUATION BOARD INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTY OF MERCHANTABILITY, TITLE, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT OF INTELLECTUAL PROPERTY RIGHTS. IN NO EVENT WILL ADI AND ITS LICENSORS BE LIABLE FOR ANY INCIDENTAL, SPECIAL, INDIRECT, OR CONSEQUENTIAL DAMAGES RESULTING FROM CUSTOMER'S POSSESSION OR USE OF THE EVALUATION BOARD, INCLUDING BUT NOT LIMITED TO LOST PROFITS, DELAY COSTS, LABOR COSTS OR LOSS OF GOODWILL. ADI'S TOTAL LIABILITY FROM ANY AND ALL CAUSES SHALL BE LIMITED TO THE AMOUNT OF ONE HUNDRED US DOLLARS (\$100.00). EXPORT. Customer agrees that it will not directly or indirectly export the Evaluation Board to another country, and that it will comply with all applicable United States federal laws and regulations relating to exports. GOVERNING LAW. This Agreement shall be governed by and construed in accordance with the substantive laws of the Commonwealth of Massachusetts (excluding conflict of law rules). Any legal action regarding this Agreement will be heard in the state or federal courts having jurisdiction in Suffolk County, Massachusetts, and Customer hereby submits to the personal jurisdiction and venue of such courts. The United Nations Convention on Contracts for the International Sale of Goods shall not apply to this Agreement and is expressly disclaimed. All Analog Devices products contained herein are subject to release and availability.

