

Industrial Ethernet Time Sensitive Networking Switch

FEATURES

- ▶ Ethernet MAC interfaces: 10 Mb, 100 Mb or 1 Gb per port
 - ▶ 6-port: 4x RMII/RGMII/SGMII and 2x RMII/RGMII
 - ▶ SGMII, 1000BASE-SX/1000BASE-LX/1000BASE-KX, and 100BASE-FX
- ▶ Low Latency Layer 2 Ethernet switch
 - ▶ Deterministic latency from port to port
 - ▶ Cut-through or store and forward operation
 - ▶ Traffic types and bridge delay per IEEE/IEC 60802
 - ▶ Standard bridging per IEEE 802.1Q-2018 (tailored)
 - ▶ Each port is nonblocking and Independent
 - ▶ 32 kB frame buffer per transmit port
 - ▶ 4096 VLANs
- ▶ Time synchronization
 - ▶ IEEE 802.1AS-2020
 - ▶ IEEE 1588-2019 default profile
 - ▶ IEEE C37.238-2017 energy profile
 - ▶ 8 ns timestamp resolution
 - ▶ Frames timestamped on ingress or egress
- ▶ IEEE 802.1Q time sensitive networking bridging
 - ▶ Qbv: enhancements for scheduled traffic
 - ▶ Qci: per stream filtering and policing
 - ▶ Qbu: frame preemption
 - ▶ Qch: cyclic queuing and forwarding
 - ▶ Qav: forwarding and queuing enhancements¹
 - ▶ Qcc: stream reservation protocol enhancements¹
 - ▶ Number of streams
 - ▶ 16000 TSN Layer 2 streams
 - ▶ 256 extended lookup streams (IPv4, IPv6, PCP, and so on)
- ▶ SendList control
- ▶ High availability and redundancy
 - ▶ IEEE 802.1CB Frame Replication and Elimination for Reliability
 - ▶ IEC62439-3:2016-03 HSR-/PRP-Compliant Protocol for Seamless Failover
 - ▶ IEC62439-3:2021-12 Media Redundancy Protocol
- ▶ Custom Layer 2 on two ports for PROFINET IRT, EtherNet/IP beacon-based DLR, POWERLINK (100 Mbps)¹
- ▶ Packet assist engine offloads host and manages TSN and switch features
- ▶ Portable C drivers for TSN and industrial Ethernet protocols
- ▶ NETCONF support (driver to Sysrepo translation layer)
- ▶ Hardware root of trust based security features
 - ▶ Secure boot, secure update with antirollback
 - ▶ Hardware-based cryptography
 - ▶ Secure host pairing Protocol
 - ▶ Cryptographic authenticity checking
- ▶ Interfaces to external host processor
 - ▶ RMII, RGMII, or SGMII
 - ▶ SPI, dual SPI, or Quad SPI
- ▶ Power
 - ▶ 3 external power supplies: 1.1 V, 3.3 V, and VDDIO_x (1.8 V, 2.5 V, or 3.3 V)
 - ▶ Total chip power 60 mW per port at 1 Gbps, full utilization with VDDIO_B = 1.8 V
- ▶ Package and temperature range
 - ▶ 256-ball CSP_BGA, 14 mm × 14 mm, 0.8 mm pitch
 - ▶ Fully specified for -40°C to +85 and -40°C to 105°C ambient operation

APPLICATIONS

- ▶ Factory and process automation
- ▶ Motion control, robots, and cobots
- ▶ Energy automation
- ▶ Transportation
- ▶ Instrumentation
- ▶ Building automation

GENERAL DESCRIPTION

The ADIN6310 is a 6-port Gigabit Ethernet, time sensitive networking (TSN) switch with integrated security designed for industrial Ethernet applications. Each port can be configured to operate at different speeds. This switch has uncommitted media access controller (MAC) interfaces and can be paired with Analog Devices, Inc., physical (PHY) layer devices, such as [ADIN1100](#), [ADIN1200](#), and [ADIN1300](#), to form a low power, low latency system. Four serial gigabit media independent interface (SGMII) and serializer/deserializer (serdes) interfaces allow backplane connections, connectivity to small form factor (SFP) modules and cascade switches. The switch supports the suite of IEEE 802.1Q time sensitive networking bridging features required by IEEE 60802 standard providing quality of service (QoS) for latency sensitive streams. The device also includes hardware capability to support parallel redundancy protocol (PRP) or high availability seamless redundancy (HSR) redundancy protocols, thereby offloading the host processor.

¹ Future software capability

TABLE OF CONTENTS

Features.....	1	Clocking.....	42
Applications.....	1	Power Architecture.....	43
General Description.....	1	Power Supplies and Supply Domains.....	43
Functional Block Diagram.....	4	Power-On Reset (POR).....	43
Specifications.....	5	Power Supply Sequencing.....	43
Timing Characteristics.....	10	Grounding.....	43
Power-Up, Power-Down, and Reset Timing.....	10	Power Consumption and Thermal Considerations.....	43
Management Interface Timing.....	11	Power Supply Decoupling.....	44
RGMII Transmit and Receive Timing.....	12	Time Synchronization.....	45
RMII Transmit and Receive Timing.....	13	Precision Time Protocol (PTP).....	45
SGMII Serial Interface Transmit and Receive Timing.....	14	Time Domains.....	45
Serial Interface Timing.....	15	IEEE 802.1AS 2020.....	45
Absolute Maximum Ratings.....	16	IEEE 1588-2019 Default Profile.....	45
Thermal Resistance.....	16	IEEE C37.238.2017 Energy Profile.....	45
Electromagnetic Compatibility (EMC) Performance.....	16	TSN Functionality.....	46
Electrostatic Discharge (ESD) Ratings.....	16	Scheduled Traffic (Qbv).....	46
ESD Caution.....	16	Frame Preemption (Qbu).....	46
Pin Configuration and Pin Descriptions.....	17	Per Stream Filtering and Policing (Qci).....	48
ADIN6310.....	17	Cyclic Queuing and Forwarding (Qch).....	48
Typical Performance Characteristics.....	23	Forwarding and Queuing Enhancements (Qav).....	48
Theory of Operation.....	24	SendList.....	49
Switch Overview.....	24	SendList Example.....	49
Port Overview.....	24	Redundancy.....	50
Storm Protection.....	25	High-Availability Seamless Redundancy (HSR).....	50
Switch Fabric.....	25	HSR Port Modes.....	51
Lookup Tables.....	25	HSR Statistics.....	51
Types of Lookups.....	27	Parallel Redundancy Protocol (PRP).....	52
Buffer Management.....	28	PRP Statistics.....	53
Quality of Service (QoS).....	29	Hybrid HSR and PRP Networks.....	54
Port Forwarding Mask.....	29	Frame Replication and Elimination for Reliability, 802.1CB.....	54
Ethernet Frame.....	29	Media Redundancy Protocol (MRP).....	55
Switch Latency.....	30	Other Functionality.....	57
Receive and Transmit Port Statistics.....	31	Industrial Ethernet Engine.....	57
Packet Assist Engine.....	32	Logical MAC Operation.....	57
Software Running on the Switch.....	32	Layer 2 Transmit and Receive.....	58
Firmware Update.....	32	Event Subscription.....	58
NETCONF.....	33	Link Layer Discovery Protocol (LLDP).....	59
Interfaces.....	34	IGMP Snooping.....	59
Serial Interface, SPI.....	34	Port Mirroring.....	59
PHY MDC and MDIO Interface.....	35	Multiple Spanning Tree Protocol.....	59
Ethernet Interfaces.....	36	Applications Information.....	60
MAC to MAC Interface.....	38	Application Circuit Overview.....	60
Host Interface.....	40	Application Use Cases.....	62
Host Interface Strapping Pins.....	40	Layout Guidelines.....	65
TIMER Inputs and Outputs.....	41	General Layout Guidelines.....	65
GPIO.....	41		
Hardware RESET.....	41		

TABLE OF CONTENTS

Component Placement.....	65	Ordering Guide.....	69
Silicon Anomaly.....	67	Evaluation Boards.....	69
Outline Dimensions.....	69		

REVISION HISTORY**10/2024—Revision 0: Initial Version**

FUNCTIONAL BLOCK DIAGRAM

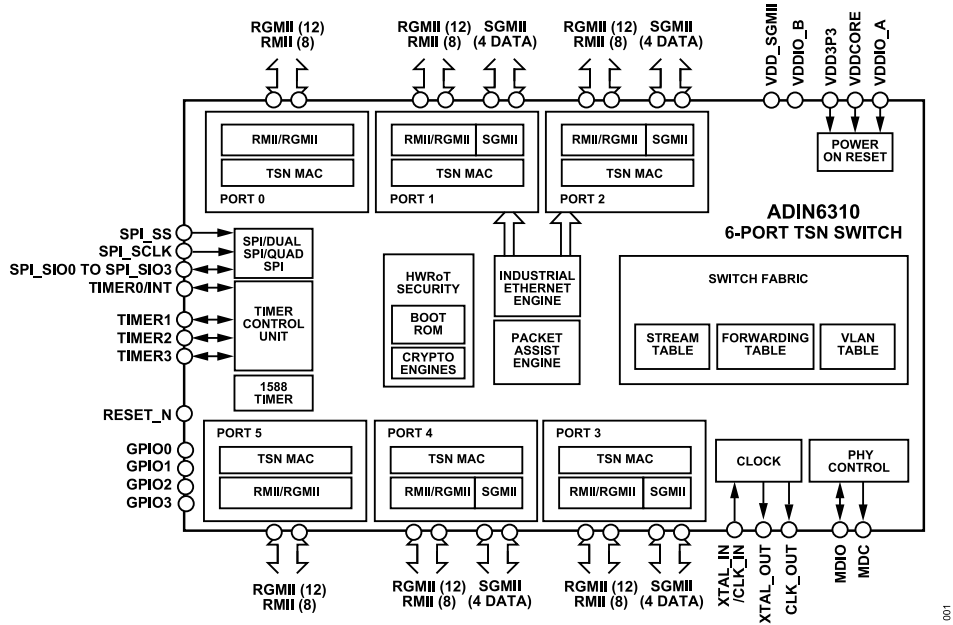


Figure 1. ADIN6310 Functional Block Diagram

SPECIFICATIONS

VDD3P3 = 3.3 V, VDDCORE = 1.1 V, VDDIO_A = 2.5 V, VDDIO_B = 1.8 V, VDD_SGMII = 3.3 V, and specifications valid over operating temperature range, T_A = -40°C to +105°C, unless otherwise noted.

Table 1. Electrical Characteristics¹

Parameter	Min	Typ	Max	Unit	Test Conditions/Comments
POWER REQUIREMENTS					
Supply Voltages					
VDD3P3	3.135	3.3	3.465	V	3.3 V ± 5%
VDDCORE	1.045	1.1	1.155	V	1.1 V ± 5%
VDDIO_A	3.135	3.3	3.465	V	3.3 V ± 5%
	2.250	2.5	2.750	V	2.5 V ± 10%
	1.71	1.8	1.89	V	1.8 V ± 5%
VDDIO_B	3.135	3.3	3.465	V	3.3 V ± 5%
	2.250	2.5	2.750	V	2.5 V ± 10%
	1.71	1.8	1.89	V	1.8 V ± 5%
VDD_SGMII	3.135	3.3	3.465	V	3.3 V ± 5%
POWER CONSUMPTION					
12 pF total capacitance on reduced gigabit media independent interface (RGMII) and reduced media independent interface (RMII) transmit lines includes package parasitics; 100% data throughput, full activity					
ADIN6310 CURRENT AND POWER					
RGMII 1000 Mbps					
VDD3P3 Current (I _{VDD3P3})		7		mA	25°C; this is for all six ports; no SGMII contribution
VDDCORE Current (I _{VDDCORE})		165		mA	
VDDIO_A Current (I _{VDDIO_A})		20		mA	VDDIO_A = 2.5 V
VDDIO_B Current (I _{VDDIO_B}) ²		93		mA	VDDIO_B = 2.5 V
		65		mA	VDDIO_B = 1.8 V
Total Power		372		mW	Including static current with VDDIO_A = 2.5 V, VDDIO_B = 1.8 V
RGMII 100 Mbps					
VDD3P3 Current (I _{VDD3P3})		7		mA	25°C; this is for all six ports; no SGMII contribution
VDDCORE Current (I _{VDDCORE})		130		mA	
VDDIO_A Current (I _{VDDIO_A})		4		mA	VDDIO_A = 2.5 V
VDDIO_B Current (I _{VDDIO_B})		23		mA	VDDIO_B = 3.3 V
		16		mA	VDDIO_B = 2.5 V
		11		mA	VDDIO_B = 1.8 V
Total Power		196		mW	VDDIO_A = 2.5 V, VDDIO_B = 1.8 V
RGMII 10 Mbps					
VDD3P3 Current (I _{VDD3P3})		7		mA	25°C; this is for all six ports; no SGMII contribution
VDDCORE Current (I _{VDDCORE})		124		mA	
VDDIO_A Current (I _{VDDIO_A})		1		mA	VDDIO_A = 2.5 V
VDDIO_B Current (I _{VDDIO_B})		8		mA	VDDIO_B = 3.3 V
		5		mA	VDDIO_B = 2.5 V
		3		mA	VDDIO_B = 1.8 V
Total Power		167		mW	VDDIO_A = 2.5 V, VDDIO_B = 1.8 V
Two SGMII and Four RGMII 1000 Mbps					
VDD3P3 Current (I _{VDD3P3})		7		mA	25°C; this is for two SGMII and four RGMII ports
VDD_SGMII (I _{VDDSGMII})		10		mA	
VDDCORE Current (I _{VDDCORE})		200		mA	
VDDIO_A Current (I _{VDDIO_A})		20		mA	VDDIO_A = 2.5 V

SPECIFICATIONS

Table 1. Electrical Characteristics¹ (Continued)

Parameter	Min	Typ	Max	Unit	Test Conditions/Comments
VDDIO_B Current (I_{VDDIO_B})		98		mA	VDDIO_B = 3.3 V
		73		mA	VDDIO_B = 2.5 V
		51		mA	VDDIO_B = 1.8 V
Total Power		418		mW	VDDIO_A = 2.5 V, VDDIO_B = 1.8 V
Four SGMII and Two RGMII 1000 Mbps					25°C; this is for four SGMII and two RGMII ports
VDD3P3 Current (I_{VDD3P3})		7		mA	
VDD_SGMII ($I_{VDDSGMII}$)		20		mA	
VDDCORE Current ($I_{VDDCORE}$)		232		mA	
VDDIO_A Current (I_{VDDIO_A})		20		mA	VDDIO_A = 2.5 V
VDDIO_B Current (I_{VDDIO_B})		71		mA	VDDIO_B = 3.3 V
		52		mA	VDDIO_B = 2.5 V
		37		mA	VDDIO_B = 1.8 V
Total Power		461		mW	VDDIO_A = 2.5 V, VDDIO_B = 1.8 V
2× 1000BASE-KX, Four RGMII 1000 Mbps					25°C; this is for two backplane connections and four RGMII ports
VDD3P3 Current (I_{VDD3P3})		7		mA	
VDD_SGMII ($I_{VDDSGMII}$)		10		mA	
VDDCORE Current ($I_{VDDCORE}$)		211		mA	
VDDIO_A Current (I_{VDDIO_A})		20		mA	VDDIO_A = 2.5 V
VDDIO_B Current (I_{VDDIO_B})		98		mA	VDDIO_B = 3.3 V
		72		mA	VDDIO_B = 2.5 V
		51		mA	VDDIO_B = 1.8 V
Total Power		430		mW	VDDIO_A = 2.5 V, VDDIO_B = 1.8 V
4× 1000BASE-KX, Two RGMII 1000 Mbps					25°C; four backplane connections and two RGMII ports
VDD3P3 Current (I_{VDD3P3})		7		mA	
VDD_SGMII ($I_{VDDSGMII}$)		20		mA	
VDDCORE Current ($I_{VDDCORE}$)		254		mA	
VDDIO_A Current (I_{VDDIO_A})		20		mA	VDDIO_A = 2.5 V
VDDIO_B Current (I_{VDDIO_B})		71		mA	VDDIO_B = 3.3 V
		54		mA	VDDIO_B = 2.5 V
		37		mA	VDDIO_B = 1.8 V
Total Power		485		mW	VDDIO_A = 2.5 V, VDDIO_B = 1.8 V
2× 1000BASE-SX/1000BASE-LX, Four RGMII 1000 Mbps					25°C
VDD3P3 Current (I_{VDD3P3})		7		mA	
VDD_SGMII ($I_{VDDSGMII}$)		10		mA	
VDDCORE Current ($I_{VDDCORE}$)		200		mA	
VDDIO_A Current (I_{VDDIO_A})		20		mA	VDDIO_A = 2.5 V
VDDIO_B Current (I_{VDDIO_B})		98		mA	VDDIO_B = 3.3 V
		72		mA	VDDIO_B = 2.5 V
		51		mA	VDDIO_B = 1.8 V
Total Power		418		mW	VDDIO_A = 2.5 V, VDDIO_B = 1.8 V
4× 1000BASE-SX/1000BASE-LX, Two RGMII 1000 Mbps					25°C
VDD3P3 Current (I_{VDD3P3})		7		mA	
VDD_SGMII ($I_{VDDSGMII}$)		20		mA	
VDDCORE Current ($I_{VDDCORE}$)		220		mA	
VDDIO_A Current (I_{VDDIO_A})		20		mA	VDDIO_A = 2.5 V

SPECIFICATIONS

Table 1. Electrical Characteristics¹ (Continued)

Parameter	Min	Typ	Max	Unit	Test Conditions/Comments
VDDIO_B Current (I_{VDDIO_B})		71		mA	VDDIO_B = 3.3 V
		54		mA	VDDIO_B = 2.5 V
		37		mA	VDDIO_B = 1.8 V
Total Power		450		mW	VDDIO_A = 2.5 V, VDDIO_B = 1.8 V
2x 100BASE-FX, Four RGMII 100 Mbps					25°C
VDD3P3 Current (I_{VDD3P3})		7		mA	
VDD_SGMII ($I_{VDDSGMII}$)		12		mA	
VDDCORE Current ($I_{VDDCORE}$)		166		mA	
VDDIO_A Current (I_{VDDIO_A})		4		mA	VDDIO_A = 2.5 V
VDDIO_B Current (I_{VDDIO_B})		23		mA	VDDIO_B = 3.3 V
		17		mA	VDDIO_B = 2.5 V
		5		mA	VDDIO_B = 1.8 V
Total Power		276		mW	VDDIO_A = 2.5 V, VDDIO_B = 1.8 V
Six RMII 100 Mbps, Internal Clock					25°C, internal 50 MHz clock used for the switch and provided out to the PHYs, using an external 50 MHz clock reduces the VDDIO_B current consumption and total power reduces by typically 20 mW
VDD3P3 Current (I_{VDD3P3})		6		mA	
VDDCORE Current ($I_{VDDCORE}$)		105		mA	
VDDIO_A Current (I_{VDDIO_A})		6		mA	VDDIO_A = 2.5 V
VDDIO_B Current (I_{VDDIO_B})		39		mA	VDDIO_B = 3.3 V
		29		mA	VDDIO_B = 2.5 V
		20		mA	VDDIO_B = 1.8 V
Total Power		186		mW	VDDIO_A = 2.5 V, VDDIO_B = 1.8 V
CUT-THROUGH SWITCH LATENCY ³					Measured in RGMII mode from ingress start of frame delimiter (SFD) to egress SFD; with a destination MAC (DA)/virtual local area network (VLAN) lookup; minByteCnt = 0x10 ⁴ ; no interfering traffic; no packet modification
RGMII Bridge Latency					Bridge latency
1000 Mbps			500	ns	
100 Mbps			2.4	µs	
10 Mbps			22	µs	
STORE AND FORWARD SWITCH LATENCY ³					Measured from ingress SFD to egress SFD; no interfering traffic; no packet modification, no gating functions enabled, with a minimum Ethernet frame of 64 bytes
1000 Mbps			0.8	µs	
100 Mbps			6	µs	
10 Mbps			58	µs	
SGMII PHY LATENCY					SGMII transmit and receive specifications represent the SGMII PHY contribution only; in SGMII mode, bridge latency must be factored into total latency ⁵
1000BASE-X, SGMII 1 Gbps Transmit		172		ns	1000BASE-SX, 1000BASE-LX, 1000BASE-KX, SGMII 1000 Mbps
1000BASE-X, SGMII 1 Gbps Receive		265		ns	
SGMII 100 Mbps Transmit		1049		ns	SGMII 100 Mbps
SGMII 100 Mbps Receive		1354		ns	
SGMII 10 Mbps Transmit		9152		ns	SGMII 10 Mbps
SGMII 10 Mbps Receive		10873		ns	

SPECIFICATIONS

Table 1. Electrical Characteristics¹ (Continued)

Parameter	Min	Typ	Max	Unit	Test Conditions/Comments
100BASE-FX Transmit		60		ns	100BASE-FX
100BASE-FX Receive		212		ns	
HOST PORT LATENCY (PORT TO SPI HOST INTERRUPT)					Measured from ingress SFD to serial peripheral interface (SPI) host interrupt, with minimum Ethernet frame of 64 bytes
1000 Mbps		112		μs	
100 Mbps		118		μs	
10 Mbps		170		μs	
DIGITAL INPUTS AND OUTPUTS					Applies to both VDDIO_A and VDDIO_B
VDDIO = 3.3 V					
Input Low Voltage (V_{IL})			0.8	V	
Input High Voltage (V_{IH})	2.0			V	
Output Low Voltage (V_{OL})			0.4	V	Output low current (I_{OL}) (minimum) = 4 mA
Output High Voltage (V_{OH})	2.4			V	Output high current (I_{OH}) (minimum) = 4 mA
VDDIO = 2.5 V					
V_{IL}			0.7	V	
V_{IH}	1.7			V	
V_{OL}			0.4	V	I_{OL} (minimum) = 4 mA
V_{OH}	2.0			V	I_{OH} (minimum) = 2 mA
	1.7			V	I_{OH} (minimum) = 4 mA
VDDIO = 1.8 V					
V_{IL}			$0.35 \times$ VDDIO	V	
V_{IH}	$0.65 \times$ VDDIO			V	
V_{OL}			0.45	V	I_{OL} (minimum) = 2 mA
V_{OH}	VDDIO – 0.45			V	I_{OH} (minimum) = 2 mA
Input Leakage Current High (I_{IH}) and Input Leakage Current Low (I_{IL})	–10		+10	μA	Except pins with internal pull-up and pull-down resistors
SGMII					AC-coupled; Px_STXP to Px_STXN
Transmit Output Differential Voltage	0.15		0.4	V	SGMII mode per Version 1.8 to copper small form-factor pluggable (SFP) or PHY
	0.8		1.6	V p-p	1000BASE-KX
	0.5		1.2	V p-p	1000BASE-SX/1000BASE-LX, 100BASE-FX
Receiver Input Differential Voltage	0.175		2.0	V p-p	AC-coupled; Px_SRXP to Px_SRXN; DC level biased in receiver
Receiver Input Impedance	80	100	120	Ω	Differential, integrated on chip
CLOCKS					
External Crystal (XTAL)					Requirements for external crystal used on XTAL_IN pin and XTAL_OUT pin
Crystal Frequency		25		MHz	
Crystal Frequency Drift	–50		+50	ppm	Includes tolerance, stability, and aging
Crystal Output Drive Level		<200		μW	
Crystal Equivalent Series Resistance (ESR)		20	100	Ω	
Crystal Load Capacitance (C_L) ⁶		10		pF	
XTAL_I Jitter			80	ps	Over the 10 kHz to 5 MHz frequency range
Crystal Oscillator RMS Jitter			1.8	ps	
Clock Input Frequency (CLK_IN)		25		MHz	Requirements for external clock applied to XTAL_IN pin; VDD3P3 voltage domain

SPECIFICATIONS

Table 1. Electrical Characteristics¹ (Continued)

Parameter	Min	Typ	Max	Unit	Test Conditions/Comments
Duty Cycle	40		60	%	
Input Voltage Range	3.135	3.3	3.465	V	
FLASH MEMORY					
Endurance	10,000			Cycles	
Data Retention	10			Years	T _J = 125°C
	20			Years	T _J = 85°C

¹ Specifications may be subject to change

² When all six ports operate at 1 Gbps, use a VDDIO_B supply voltage of 1.8 V or 2.5 V to minimize power consumption.

³ Guaranteed by design.

⁴ minByteCnt represents the number of received bytes before a cut-through frame is queued for transmission. The purpose is to remove the latency jitter due to access time in the lookup table. The default setting for minByteCnt = 0x20. For this specification, minByteCnt was reduced to 0x10 so that when the VLAN is received and the lookup table provides the routing information, the frame is immediately queued for transmission. Using lower numbers can decrease latency and increase jitter

⁵ Latency from the SGMII to SGMII port = SGMII Receive + Bridge Latency + SGMII Transmit. Latency from the SGMII to RGMII port = SGMII Receive + Bridge Latency. Latency from the RGMII to SGMII port = Bridge Latency + SGMII Transmit.

⁶ Where load capacitance (C_L) = ((C1 × C2)/(C1 + C2) + C_{STRAY}), where C_{STRAY} is the stray capacitance including routing and package parasitics.

TIMING CHARACTERISTICS

POWER-UP, POWER-DOWN, AND RESET TIMING

Table 2. Power-Up Timing^{1,2}

Parameter	Description	Min	Typ	Max	Unit
t_{RAMP}	Power supply ramp time ³	0.1		40	ms
t_1	Time from last power supply being applied to crystal settling			3	ms
t_2	Time from last supply being applied to RESET_N released	100			μ s
t_3	RESET_N rise time	100			ns
t_4	Strap input setup time to RESET_N release	5			μ s
t_5	Strap input hold time to RESET_N release	2			μ s
t_6	RESET_N low time	10			μ s
t_7	RESET_N low prior to supplies powering down	10			μ s
t_8	RESET_N release to SPI interface available ⁴		230	310	ms
	RESET_N release to the RGMII MAC interface available.		270	350	ms

¹ Specifications may be subject to change.

² No specific power sequence is required; however, the preferred sequence has VDDCORE applied last and removed first. Example ordering sequence is VDD3P3 → VDDIO_A → VDDCORE and this order is reversed when powering down.

³ Applies to all supply rails.

⁴ Includes bootloader and software configuration time.

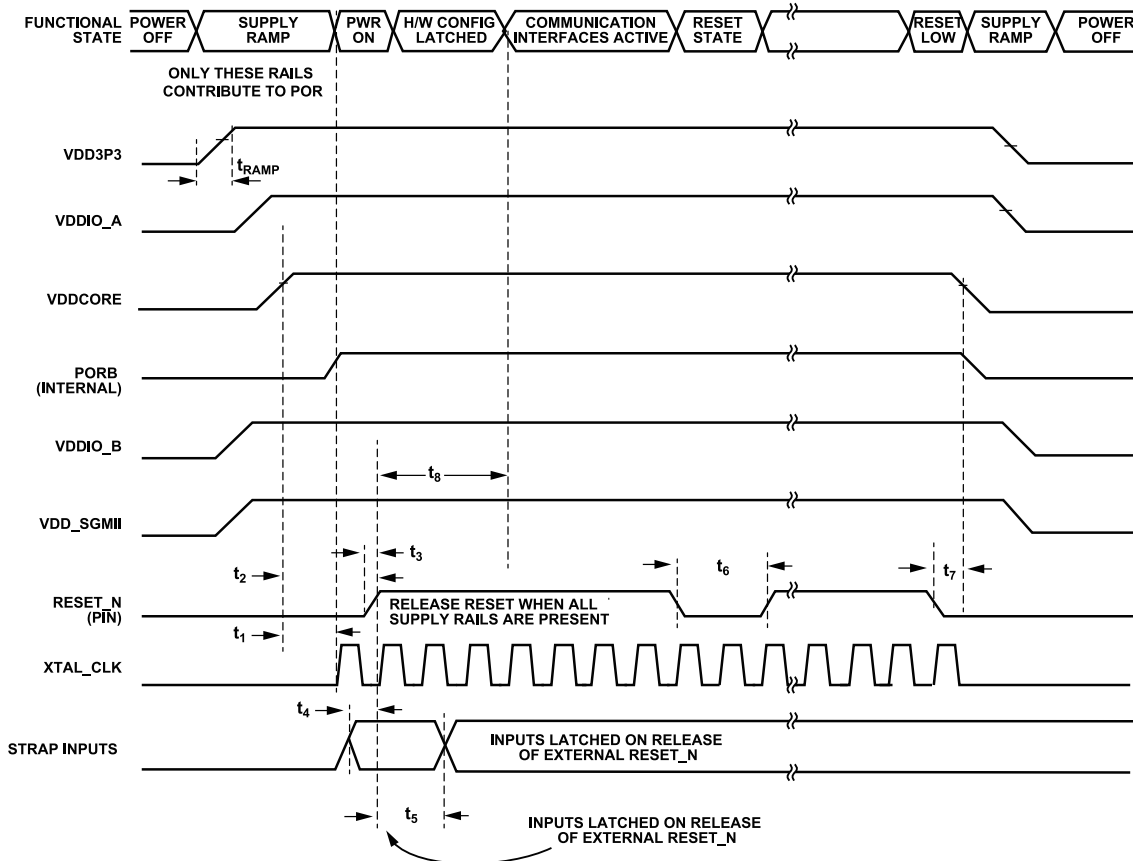


Figure 2. Power-Up, Power-Down, and Reset Timing

TIMING CHARACTERISTICS

MANAGEMENT INTERFACE TIMING

Table 3. Management Interface Timing^{1, 2}

Parameter	Description	Min	Typ	Max	Unit
	MDC duty cycle	45	50	55	ns
t_1	MDC period ³	175	400	4096	ns
t_2	MDC rise and fall time			15	ns
t_3	MDIO input setup time to MDC rising edge	21			ns
t_4	MDIO input signal hold time to MDC	0			ns
t_5	MDIO output delay to MDC	0		15	ns

¹ Specifications may be subject to change.

² Measured with a C_L of 100 pF (routing to up to six Ethernet PHYs, includes pin capacitance and trace capacitance). Timing parameters measured from 10% to 90%.

³ The typical parameters correspond to the maximum MDC frequency of 2.5 MHz as specified in the IEEE 802.3 standard.




Figure 3. Management Interface Timing

TIMING CHARACTERISTICS

RGMII TRANSMIT AND RECEIVE TIMING

Table 4. 1000M RGMII Timing^{1, 2, 3}

Parameter	Description	Min	Typ	Max	Unit
t ₁	Data to clock output skew (at transmitter) ⁴	-500	0	+500	ps
t ₂	Data to clock input skew (at receiver) ⁴	1	1.8	2.6	ns
t ₃	Data to clock output setup time (at transmitter—internal delay) ⁵	1.2	2.0		ns
t ₄	Clock to data output hold time (at transmitter— internal delay) ⁵	1.2	2.0		ns
t ₅	Data to clock input setup time (at receiver—internal delay) ⁵	1.0	2.0		ns
t ₆	Clock to data input hold time (at receiver— internal delay) ⁵	1.0	2.0		ns
t _{CYC}	Clock cycle duration ⁶	7.2	8	8.8	ns
Duty_G	Duty cycle for gigabit	45	50	55	%
	Duty cycle for gigabit (internal delay)	43	50	57	%
Duty_T	Duty cycle for 10 Mbps and 100 Mbps	40	50	60	%
t _R and t _F	Rise and fall time (20% to 80%)			0.75	ns

¹ Specifications may be subject to change.

² With C_L of 5 pF.

³ When operating at 1 Gbps, transmit and receive data is clocked on the rising and falling edge of the clocks; therefore, setup and hold times apply for both clock edges.

⁴ When operating without the RGMII internal delay, the printed circuit board (PCB) design requires clocks to be routed such that an additional trace delay of greater than 1.5 ns and less than 2.0 ns is added to the associated clock signal. For 10 Mbps and 100 Mbps, the maximum value is unspecified.

⁵ Hardware and software programmable internal delay can be enabled or disabled.

⁶ For 10 Mbps and 100 Mbps, t_{CYC} scales to 400 ns ± 40 ns and 40 ns ± 4 ns, respectively.

Table 5. 10M/100M RGMII Timing¹

Parameter	Description	Min	Typ	Max	Unit
t ₁	Data to clock output skew (at transmitter)	-500	0	+500	ps
t ₂	Data to clock input skew (at receiver)	1	1.8	2.6	ns
t ₃	Data to clock output setup time (at transmitter—internal delay)	1.2	2.0		ns
t ₄	Clock to data output hold time (at transmitter— internal delay)	1.2	2.0		ns
t ₅	Data to clock input setup time (at receiver—internal delay)	1.0	2.0		ns
t ₆	Clock to data input hold time (at receiver— internal delay)	1.0	2.0		ns
t _{CYC}	Clock cycle duration ¹	7.2	8	8.8	ns
Duty_T	Duty cycle for 10 Mbps and 100 Mbps	40	50	60	%
t _R and t _F	Rise and fall time (20% to 80%)			0.75	ns

¹ For 10 Mbps and 100 Mbps, t_{CYC} scales to 400 ns ± 40 ns and 40 ns ± 4 ns, respectively.

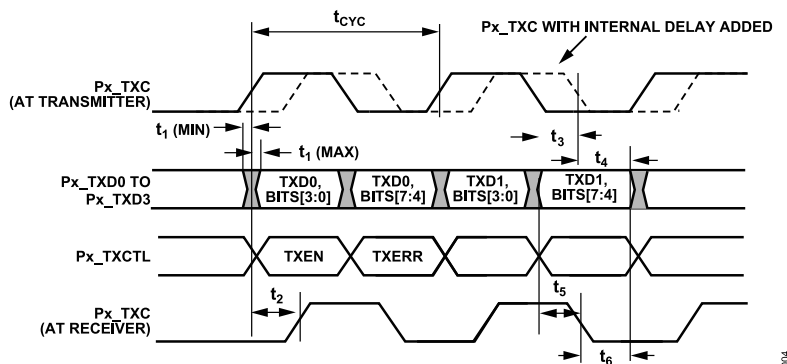


Figure 4. RGMII Transmit Timing

TIMING CHARACTERISTICS

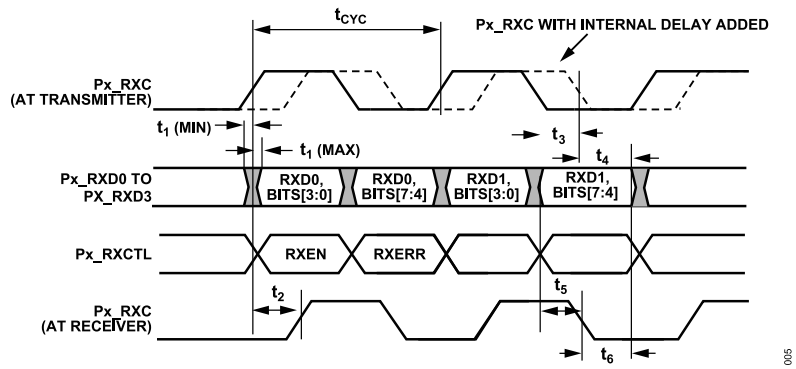


Figure 5. RGMII Receive Timing

RMII TRANSMIT AND RECEIVE TIMING

Table 6. RMII Timing¹

Parameter	Description	Min	Typ	Max	Unit
REF_CLK	Frequency of the reference clock as sourced from the Px_TXC or from an external clock source		50		MHz
	Duty cycle of the clock	35		65	%
t ₁	Data setup to REF_CLK rising edge	4			ns
t ₂	Data hold from REF_CLK rising edge	2			ns
t ₃	Output rise and fall time	1		5	ns

¹ Specifications may be subject to change




Figure 6. RMII Timing

TIMING CHARACTERISTICS

SGMII SERIAL INTERFACE TRANSMIT AND RECEIVE TIMING

Table 7. Serial Interface Transmit and Receive Timing¹

Parameter	Description	Min	Typ	Max	Unit
t_1	SGMII, 1000BASE-KX, 1000BASE-LX, and 1000BASE-SX signaling speed		1.25		GBd
	100BASE-FX signaling speed		0.125		GBd
t_2	SGMII rise time (20% to 80%)			200	ps
	1000BASE-KX rise time			320	ps
t_3	SGMII fall time (20% to 80%)			200	ps
	1000BASE-KX fall time			320	ps
t_4	STX output differential skew (STXP vs. STXN) ²			20	ps
t_5	STX total jitter			0.24	UI
	SRX total jitter tolerated			0.749	UI

¹ Specifications may be subject to change.

² Guaranteed by design.




Figure 7. SGMII Transmit Timing




Figure 8. SGMII Receive Timing

TIMING CHARACTERISTICS

SERIAL INTERFACE TIMING

Table 8. SPI, Dual SPI, and Quad SPI Timing¹

Parameter	Description	VDDIO_A = 1.8 V ± 5%			VDDIO_A = 3.3 V ± 5%			Unit
		Min	Typ	Max	Min	Typ	Max	
f _{reqSPICLK}	SPI clock frequency	5		30	5		37.5	MHz
t _{SPICLK}	SPI clock period	33.3		200	26.6		200	ns
t _{SPICLS}	SPI clock low period	15			12			ns
t _{SPICHS}	SPI clock high period	15			12			ns
t _{SDSCI}	SPI_SS assertion to first SPI_SCLK edge	15			12			ns
t _{SSPID}	Data input valid to SPI clock rising edge (data input setup)	4			4			ns
t _{HSPID}	SPI clock rising edge to data input invalid (data input hold)	2			2			ns
t _{OEEN}	SPI clock falling edge to output enabled	t _{SPICLK} + 2			t _{SPICLK} + 1			ns
t _{DSOE}	Time from clock falling edge at which SPI_SIOx is switched from input to output to output valid			t _{SPICLK} + 10			t _{SPICLK} + 7	ns
t _{DDSPID}	SPI falling clock edge to data output valid (data output delay)			10			7	ns
t _{HDS}	Last SPI clock edge to SPI_SS not asserted	15			12			ns
t _{SDSHI}	SPI_SS deassertion to data high impedance			8			6	ns
t _{TDS}	SPI_SS high time between transfers; sequential transfer delay	33.3			26.6			ns

¹ Specifications may be subject to change

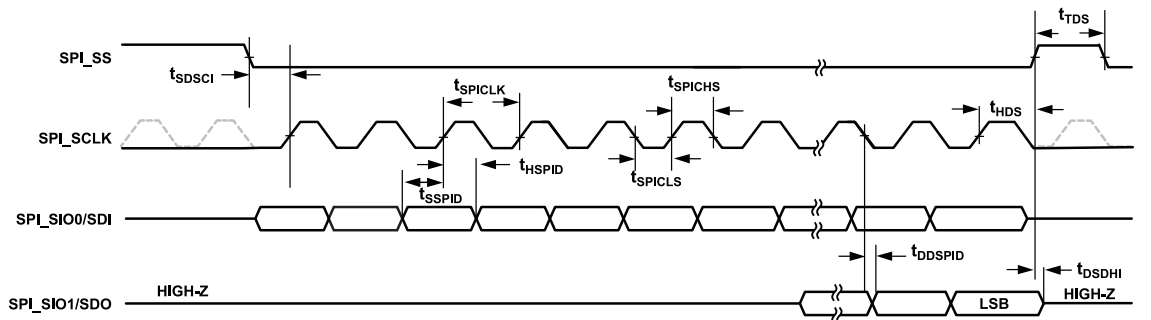


Figure 9. SPI Write and Read Timing




Figure 10. Dual and Quad SPI Write and Read Timing

ABSOLUTE MAXIMUM RATINGS

Table 9. Absolute Maximum Ratings

Parameter	Rating
VDD3P3 to VSS3P3	-0.3 V to +3.63 V
VDDCORE to DGND	-0.3 V to +1.26 V
VDDCORE_PLL to VSSCORE_DLL	-0.3 V to +1.26 V
VDDCORE_DLL to VSSCORE_DLL	-0.3 V to +1.26 V
VDD_SGMII to VSS_SGMII	-0.3 V to +3.63 V
VDDIO_A to DGND	-0.3 V to +3.63 V
VDDIO_B to DGND	-0.3 V to +3.63 V
VSS3P3 to DGND	-0.3 V to +0.3 V
VSSCORE_x to DGND	-0.3 V to +0.3 V
VSS_SGMII to VSS3P3	-0.3 V to +0.3 V
Port 0 RMII/RGMII MAC Interface to GND (P0_TXxx and P0_RXxx) ¹	-0.3 V to VDDIO_A + 0.3 V
Port 1 to Port 5 RMII/RGMII MAC Interface to GND (Px_TXxx and Px_RXxx) ²	-0.3 V to VDDIO_B + 0.3 V
SGMII MAC Interface to GND (Px_STXN, Px_STNPxx, Px_SRXN, and Px_SRPxx)	-0.3 V to VDDCORE + 0.3 V
XTAL_IN, CLK_IN, XTAL_OUT, and CLK_OUT	-0.3 V to VDD3P3 + 0.3 V
RESET_N, TIMERx, GPIOx, TDO, TDI, TCK, TMS, and TEST_EN	-0.3 V to VDDIO_A + 0.3 V
SPI (SPI_xx) to DGND	-0.3 V to VDDIO_A + 0.3 V
Management Interface (MDC and MDIO)	-0.3 V to +3.63 V
Temperature	
Industrial	-40°C to +105°C
Storage Range	-65°C to +150°C
T _J Maximum	125°C
Lead	JEDEC industry-standard
Power Dissipation	(T _J maximum - T _A)/θ _{JA}
Soldering	J-STD-020

¹ All Port 0 pins, including P0_LINK and P0_RSTN.

² All Port 1 to Port 5 pins including Px_LINK and Px_RSTN.

Stresses at or above those listed under Absolute Maximum Ratings may cause permanent damage to the product. This is a stress rating only; functional operation of the product at these or any other conditions above those indicated in the operational section of this specification is not implied. Operation beyond the maximum operating conditions for extended periods may affect product reliability.

THERMAL RESISTANCE

Thermal performance is directly linked to PCB design and operating environment. Close attention to PCB thermal design is required.

θ_{JA} is the natural convection junction to ambient thermal resistance measured in a one cubic foot sealed enclosure.

θ_{JC} is the junction to case thermal resistance.

Table 10. Thermal Resistance

Package Type	θ _{JA}	θ _{JC}	Unit
256-Ball CSP_BGA	20.84 ¹	5.96 ¹	°C/W
	20 ²	Not applicable	°C/W

¹ Based on simulated data using a JEDEC 2s2p thermal test board (with enhanced cu density) in a JEDEC natural convection environment. See JEDEC specification JESD-51 for details.

² Test Condition 1: thermal impedance measured on Analog Devices, Inc., hardware, 4S4P test board.

ELECTROMAGNETIC COMPATIBILITY (EMC) PERFORMANCE

EMC testing was performed using the EVAL-ADIN6310EBZ evaluation hardware against the following EMC test standards:

- ▶ IEC 61000-4-5 surge (Class ±4 kV)
- ▶ IEC 61000-4-4 electrical fast transient (EFT) (±4 kV)
- ▶ IEC 61000-4-2 ESD (±8 kV contact discharge)
- ▶ IEC 61000-4-3 radiated immunity
 - ▶ 80 MHz to 1 GHz (Class A: 10 V/m)
 - ▶ 1 GHz to 3.2 GHz (Class A: 3 V/m)
 - ▶ 3.2 Hz to 6 GHz (Class A: 1 V/m)
- ▶ IEC 61000-4-6 conducted immunity (Class A: 10 V)
- ▶ EN 55032 radiated emissions (Class A)
- ▶ EN 55032 conducted emissions (Class B)

ELECTROSTATIC DISCHARGE (ESD) RATINGS

The following ESD information is provided for handling of ESD-sensitive devices in an ESD protected area only.

Human body model (HBM) per ANSI/ESDA/JEDEC JS-001.

Field induced charged device model (FICDM) per ANSI/ESDA/JEDEC JS-002.

ESD Ratings for ADIN6310

Table 11. ADIN6310, 256-Ball CSP_BGA

ESD Model	Withstand Threshold (kV)	Class
HBM	2	2
FICDM	1250	C3

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.

PIN CONFIGURATION AND PIN DESCRIPTIONS

ADIN6310

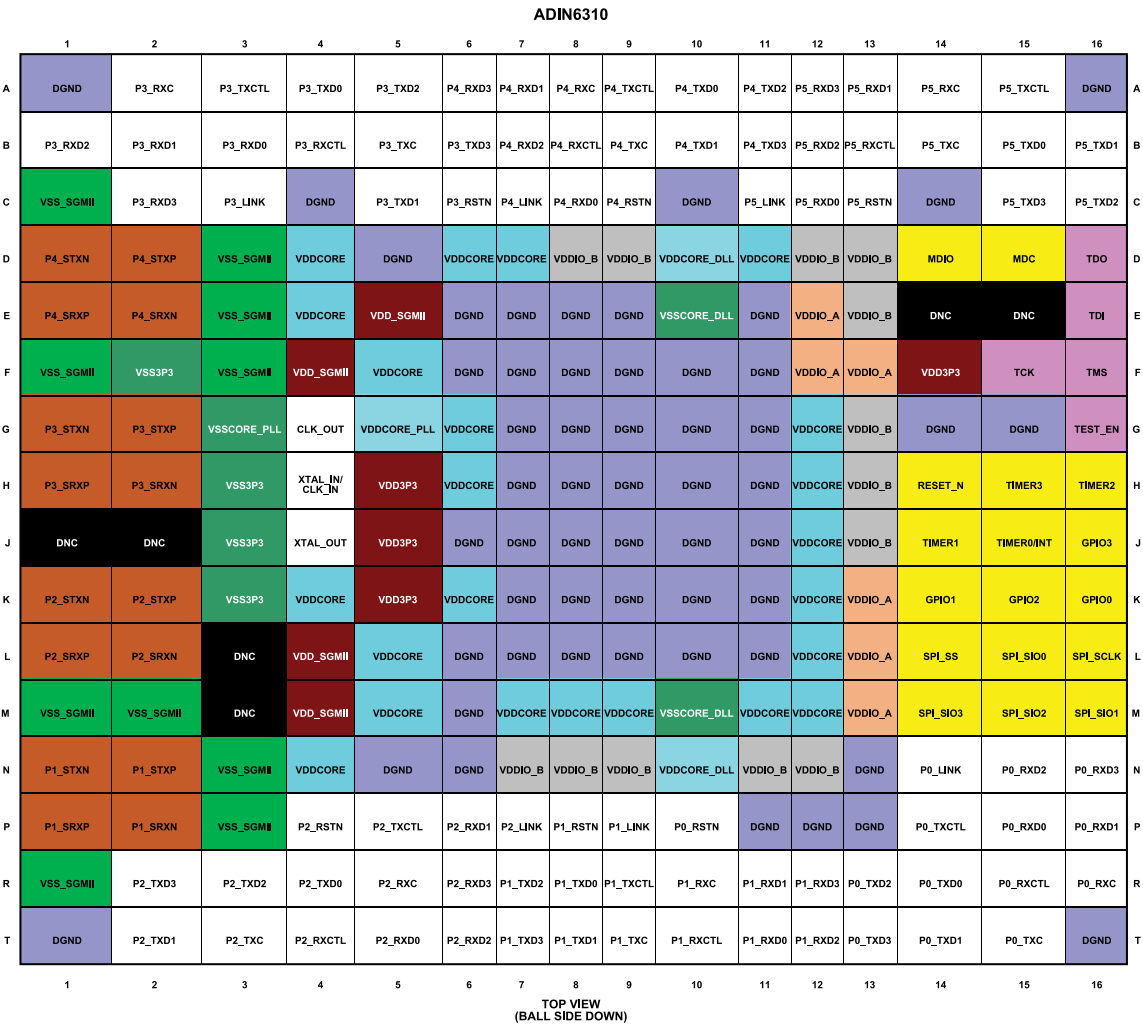


Figure 11. ADIN6310 Pin Configuration (Top View, Looking Through the Device)

Table 12. Pin Function Descriptions

ADIN6310 Ball No.	Mnemonic	Type ¹	Internal Termination at Reset ²	Power Domain ²	Description
Clock Interface					
H4	XTAL_IN/CLK_IN	A	High-Z	VDD3P3	Input for Crystal (XTAL_IN). Single-Ended 25 MHz Reference Clock (CLK_IN).
J4	XTAL_OUT	A	High-Z	VDD3P3	Second Terminal for Crystal Connection. If using a single-ended reference clock on XTAL_IN/CLK_IN, leave XTAL_OUT open circuit.
G4	CLK_OUT	A	High-Z	VDD3P3	Analog Reference Clock Output. The 25 MHz, 3.3 V reference clock from the crystal oscillator.
Digital					
H14	RESET_N	I	Pull-down	VDDIO_A	Active Low, Reset Input. RESET_N requires an external pull-down. Hold low for >10 μs.
J15	TIMER0/INT	I/O	Pull-up	VDDIO_A	Timer or General-Purpose Input and Output (GPIO). When using the SPI, TIMER0 acts as an interrupt output to the host. The boot strapping pin is used to define the host interface origin and operation.

PIN CONFIGURATION AND PIN DESCRIPTIONS

Table 12. Pin Function Descriptions (Continued)

ADIN6310 Ball No.	Mnemonic	Type ¹	Internal Termination at Reset ²	Power Domain ²	Description
J14	TIMER1	I/O	Pull-up	VDDIO_A	Timer or GPIO. The boot strapping pin is used to define the host interface origin and operation.
H16	TIMER2	I/O	Pull-down	VDDIO_A	Timer or GPIO. The boot strapping pin is used to define the host interface origin and operation.
H15	TIMER3	I/O	Pull-down	VDDIO_A	Timer or GPIO. The boot strapping pin is used to define the host interface origin and operation.
K16	GPIO0	I/O	Pull-down	VDDIO_A	GPIO.
K14	GPIO1	I/O	Pull-down	VDDIO_A	GPIO.
K15	GPIO2	I/O	Pull-down	VDDIO_A	GPIO and Activity on Industrial Ethernet Engine Port 1.
J16	GPIO3	I/O	Pull-down	VDDIO_A	GPIO and Activity on Industrial Ethernet Engine Port 2.
Management Interface					
D15	MDC	O	High-Z	VDDIO_B	Management Data Clock Output up to 5.5 MHz.
D14	MDIO	I/O	High-Z	VDDIO_B	Management Data Bidirectional Line Synchronous to the MDC Clock. The MDIO pin requires a 1.5 kΩ pull-up resistor to VDDIO_B.
Serial Interface					
L14	SPI_SS	I	Pull-up	VDDIO_A	This interface supports the SPI, dual SPI and quad SPI communication. Chip Select Input, Active Low. The boot strapping pin is used to define the host interface origin and operation.
L16	SPI_SCLK	I/O	Pull-down	VDDIO_A	Clock Input.
L15	SPI_SIO0	I/O	Pull-down	VDDIO_A	Serial Input and Output 0 for the Quad SPI. The boot strapping pin is used to define the host interface origin and operation. This pin is the Serial Input and Output 0 for the dual SPI, and the serial data input (SDI) for the SPI.
M16	SPI_SIO1	I/O	Pull-down	VDDIO_A	Serial Input and Output 1 for the Quad SPI. The boot strapping pin is used to define the host interface origin and operation. This pin is the Serial Input and Output 1 for the dual SPI, and the serial data output (SDO) for the SPI.
M15	SPI_SIO2	I/O	Pull-down	VDDIO_A	Serial Input and Output 2 for the Quad SPI. The boot strapping pin is used to define the host interface origin and operation.
M14	SPI_SIO3	I/O	Pull-down	VDDIO_A	Serial Input and Output 3 for the Quad SPI.
Port 0 Interface					
P10	P0_RSTN	O	Pull-down	VDDIO_A	Port 0 PHY Reset. Active low output.
T15	P0_TXC	O	Pull-down	VDDIO_A	Port 0 Transmit Clock. In RGMII mode, P0_TXC provides the clock output of 125 MHz for 1 Gbps, of 25 MHz for 100 Mbps, and of 2.5 MHz for 10 Mbps. In RMII mode, the 50 MHz REF_CLK is the input or output.
R14	P0_TXD0	O	Pull-down	VDDIO_A	Port 0 RGMII and RMII Transmit Data 0 Output.
T14	P0_TXD1	O	Pull-down	VDDIO_A	Port 0 RGMII and RMII Transmit Data 1 Output.
R13	P0_TXD2	O	Pull-down	VDDIO_A	Port 0 RGMII Transmit Data 2 Output.
T13	P0_TXD3	O	Pull-down	VDDIO_A	Port 0 RGMII Transmit Data 3 Output.
P14	P0_TXCTL	O	Pull-down	VDDIO_A	Port 0 RGMII Transmit Control Signal. In RGMII mode, this pin is a combination of the TX_EN and TX_ER signals using both edges of TXC. In RMII mode, this is the transmit enable input from the MAC to the PHY (TX_EN), which indicates that transmission data is available on the data lines.
R16	P0_RXC	I	Pull-down	VDDIO_A	Port 0 Receive Clock Input. In RGMII mode, expect the clock input of 125 MHz for 1 Gbps, of 25 MHz for 100 Mbps, and of 2.5 MHz for 10 Mbps.
P15	P0_RXD0	I	Pull-down	VDDIO_A	Port 0 RGMII and RMII Receive Data 0 Input.
P16	P0_RXD1	I	Pull-down	VDDIO_A	Port 0 RGMII and RMII Receive Data 1 Input.
N15	P0_RXD2	I	Pull-down	VDDIO_A	Port 0 RGMII Receive Data 2 Input.
N16	P0_RXD3	I	Pull-down	VDDIO_A	Port 0 RGMII Receive Data 3 Input. In RMII mode, this pin is RX_ER. When asserted high, it indicates that the PHY has detected a receive error.

PIN CONFIGURATION AND PIN DESCRIPTIONS

Table 12. Pin Function Descriptions (Continued)

ADIN6310 Ball No.	Mnemonic	Type ¹	Internal Termination at Reset ²	Power Domain ²	Description
R15	P0_RXCTL	I	Pull-down	VDDIO_A	Port 0 Receive Control Signal. In RGMII mode, this pin is a combination of the RX_DV and RX_ER signals using both edges of RXC. In RMII mode, this pin is CRS_DV, which is a combination of the CRS and RX_DV signals and is asserted while the receive medium is not idle.
N14	P0_LINK	I	Pull-down	VDDIO_A	General-Purpose Input. Intended to be connected to the PHY output link status pin.
Port 1 Interface					
P1	P1_SRXP	I	High-Z	VDDCORE	SGMII Receive Differential Pair Positive.
P2	P1_SRXN	I	High-Z	VDDCORE	SGMII Receive Differential Pair Negative.
N2	P1_STXP	O	High-Z	VDDCORE	SGMII Transmit Differential Pair Positive.
N1	P1_STXN	O	High-Z	VDDCORE	SGMII Transmit Differential Pair Negative.
P8	P1_RSTN	O	Pull-down	VDDIO_B	See Port 0.
T9	P1_TXC	O	Pull-down	VDDIO_B	See Port 0.
R8	P1_TXD0	O	Pull-down	VDDIO_B	See Port 0.
T8	P1_TXD1	O	Pull-down	VDDIO_B	See Port 0.
R7	P1_TXD2	O	Pull-down	VDDIO_B	See Port 0.
T7	P1_TXD3	O	Pull-down	VDDIO_B	See Port 0.
R9	P1_TXCTL	O	Pull-down	VDDIO_B	See Port 0.
R10	P1_RXC	I	Pull-down	VDDIO_B	See Port 0.
T11	P1_RXD0	I	Pull-up	VDDIO_B	See Port 0.
R11	P1_RXD1	I	Pull-down	VDDIO_B	See Port 0.
T12	P1_RXD2	I	Pull-down	VDDIO_B	See Port 0.
R12	P1_RXD3	I	Pull-down	VDDIO_B	See Port 0.
T10	P1_RXCTL	I	Pull-down	VDDIO_B	See Port 0.
P9	P1_LINK	I	Pull-down	VDDIO_B	See Port 0.
Port 2 Interface					
L1	P2_SRXP	I	High-Z	VDDCORE	See Port 1.
L2	P2_SRXN	I	High-Z	VDDCORE	See Port 1.
K2	P2_STXP	O	High-Z	VDDCORE	See Port 1.
K1	P2_STXN	O	High-Z	VDDCORE	See Port 1.
P4	P2_RSTN	O	Pull-down	VDDIO_B	See Port 0.
T3	P2_TXC	O	Pull-down	VDDIO_B	See Port 0.
R4	P2_TXD0	O	Pull-down	VDDIO_B	See Port 0.
T2	P2_TXD1	O	Pull-down	VDDIO_B	See Port 0.
R3	P2_TXD2	O	Pull-down	VDDIO_B	See Port 0.
R2	P2_TXD3	O	Pull-down	VDDIO_B	See Port 0.
P5	P2_TXCTL	O	Pull-down	VDDIO_B	See Port 0.
R5	P2_RXC	I	Pull-down	VDDIO_B	See Port 0.
T5	P2_RXD0	I	Pull-down	VDDIO_B	See Port 0.
P6	P2_RXD1	I	Pull-up	VDDIO_B	See Port 0.
T6	P2_RXD2	I	Pull-down	VDDIO_B	See Port 0.
R6	P2_RXD3	I	Pull-down	VDDIO_B	See Port 0.
T4	P2_RXCTL	I	Pull-down	VDDIO_B	See Port 0.
P7	P2_LINK	I	Pull-down	VDDIO_B	See Port 0.

PIN CONFIGURATION AND PIN DESCRIPTIONS

Table 12. Pin Function Descriptions (Continued)

ADIN6310 Ball No.	Mnemonic	Type ¹	Internal Termination at Reset ²	Power Domain ²	Description
Port 3 Interface					
H1	P3_SRXP	I	High-Z	VDDCORE	See Port 1.
H2	P3_SRXN	I	High-Z	VDDCORE	See Port 1.
G2	P3_STXP	O	High-Z	VDDCORE	See Port 1.
G1	P3_STXN	O	High-Z	VDDCORE	See Port 1.
C6	P3_RSTN	O	Pull-down	VDDIO_B	See Port 0.
B5	P3_TXC	O	Pull-down	VDDIO_B	See Port 0.
A4	P3_TXD0	O	Pull-down	VDDIO_B	See Port 0.
C5	P3_TXD1	O	Pull-down	VDDIO_B	See Port 0.
A5	P3_TXD2	O	Pull-down	VDDIO_B	See Port 0.
B6	P3_TXD3	O	Pull-down	VDDIO_B	See Port 0.
A3	P3_TXCTL	O	Pull-down	VDDIO_B	See Port 0.
A2	P3_RXC	I	Pull-down	VDDIO_B	See Port 0.
B3	P3_RXD0	I	Pull-down	VDDIO_B	See Port 0.
B2	P3_RXD1	I	Pull-down	VDDIO_B	See Port 0.
B1	P3_RXD2	I	Pull-up	VDDIO_B	See Port 0.
C2	P3_RXD3	I	Pull-down	VDDIO_B	See Port 0.
B4	P3_RXCTL	I	Pull-down	VDDIO_B	See Port 0.
C3	P3_LINK	I	Pull-down	VDDIO_B	See Port 0.
Port 4 Interface					
E1	P4_SRXP	I	High-Z	VDDCORE	See Port 1.
E2	P4_SRXN	I	High-Z	VDDCORE	See Port 1.
D2	P4_STXP	O	High-Z	VDDCORE	See Port 1.
D1	P4_STXN	O	High-Z	VDDCORE	See Port 1.
C9	P4_RSTN	O	Pull-down	VDDIO_B	See Port 0.
B9	P4_TXC	O	Pull-down	VDDIO_B	See Port 0.
A10	P4_TXD0	O	Pull-down	VDDIO_B	See Port 0.
B10	P4_TXD1	O	Pull-down	VDDIO_B	See Port 0.
A11	P4_TXD2	O	Pull-down	VDDIO_B	See Port 0.
B11	P4_TXD3	O	Pull-down	VDDIO_B	See Port 0.
A9	P4_TXCTL	O	Pull-down	VDDIO_B	See Port 0.
A8	P4_RXC	I	Pull-down	VDDIO_B	See Port 0.
C8	P4_RXD0	I	Pull-down	VDDIO_B	See Port 0.
A7	P4_RXD1	I	Pull-down	VDDIO_B	See Port 0.
B7	P4_RXD2	I	Pull-down	VDDIO_B	See Port 0.
A6	P4_RXD3	I	Pull-up	VDDIO_B	See Port 0.
B8	P4_RXCTL	I	Pull-down	VDDIO_B	See Port 0.
C7	P4_LINK	I	Pull-down	VDDIO_B	See Port 0.
Port 5 Interface					
C13	P5_RSTN	O	Pull-down	VDDIO_B	See Port 0.
B14	P5_TXC	O	Pull-down	VDDIO_B	See Port 0.
B15	P5_TXD0	O	Pull-down	VDDIO_B	See Port 0.
B16	P5_TXD1	O	Pull-down	VDDIO_B	See Port 0.
C16	P5_TXD2	O	Pull-down	VDDIO_B	See Port 0.
C15	P5_TXD3	O	Pull-down	VDDIO_B	See Port 0.
A15	P5_TXCTL	O	Pull-down	VDDIO_B	See Port 0.
A14	P5_RXC	I	Pull-down	VDDIO_B	See Port 0.

PIN CONFIGURATION AND PIN DESCRIPTIONS

Table 12. Pin Function Descriptions (Continued)

ADIN6310 Ball No.	Mnemonic	Type ¹	Internal Termination at Reset ²	Power Domain ²	Description
C12	P5_RXD0	I	Pull-up	VDDIO_B	See Port 0.
A13	P5_RXD1	I	Pull-down	VDDIO_B	See Port 0.
B12	P5_RXD2	I	Pull-down	VDDIO_B	See Port 0.
A12	P5_RXD3	I	Pull-up	VDDIO_B	See Port 0.
B13	P5_RXCTL	I	Pull-down	VDDIO_B	See Port 0.
C11	P5_LINK	I	Pull-down	VDDIO_B	See Port 0.
Other					Joint Test Action Group (JTAG) Access.
D16	TDO	O	N/A	VDDIO_A	JTAG Test Data Output.
E16	TDI	I	N/A	VDDIO_A	JTAG Test Data Input,
F15	TCK	I	N/A	VDDIO_A	JTAG Test Clock.
F16	TMS	I	N/A	VDDIO_A	JTAG Test Mode Select.
G16	TEST_EN	I	N/A	VDDIO_A	Test Enable. JTAG mode enabled when TEST_EN is pulled high. For normal operation, pull TEST_EN low.
Power ³					
D4, D6, D7, D11, E4, F5, G6, G12, H6, H12, J12, K4, K6, K12, L5, L12, M5, M7 to M9, M11, M12, N4	VDDCORE	S	N/A	N/A	Supply to the Digital and Analog Core, 1.1 V ± 5%. Place decoupling capacitors with respect to DGND.
G5	VDDCORE_PLL	S	N/A	N/A	Supply for the Analog Core, 1.1 V ± 5%. Place decoupling capacitors with respect to VSSCORE_PLL.
D10, N10	VDDCORE_DLL	S	N/A	N/A	Supply for the Analog Core, 1.1 V ± 5%. Place decoupling capacitors with respect to VSSCORE_DLL.
E12, F12, F13, K13, L13, M13	VDDIO_A	S	N/A	N/A	Supply for the Host Port (RGMII and RMII Port 0, SPI, JTAG).
D8, D9, D12, D13, E13, G13, H13, J13, N7, N8, N9, N11, N12	VDDIO_B	S	N/A	N/A	Supply for the RGMII and RMII Port 1 to Port 5, MDIO Interface.
H5, J5, K5, F14	VDD3P3	S	N/A	N/A	Analog Supply for the Analog, Crystal, and Clocking Circuitry. Place decoupling capacitors with respect to VSS3P3.
E5, F4, L4, M4	VDD_SGMII	S	N/A	N/A	Supply for the SGMII on Port 1 to Port 4. Where an application does not use the SGMII, connect the VDD_SGMII pins directly to VSS_SGMII ground on the board.
Ground					
F2, H3, J3, K3	VSS3P3	S	N/A	N/A	Ground for the Analog, Crystal, and Clock Circuit.
M10, E10	VSSCORE_DLL	S	N/A	N/A	Ground for the Analog Core.
G3	VSSCORE_PLL	S	N/A	N/A	Ground for the Analog Core.
F1, C1, D3, E3, F3, M1, M2, N3, P3, R1	VSS_SGMII	S	N/A	N/A	Ground for the SGMII on Port 1 to Port 4.
A1, A16, C4, C10, C14, D5, E6, E7 to E9, E11, F6 to F11, G7 to G11, G14, G15, H7 to H11, J6 to J11, K7 to K11, L6 to L11, M6, N5, N6, N13, P11 to P13, T1, T16	DGND	S	N/A	N/A	Digital Ground.

PIN CONFIGURATION AND PIN DESCRIPTIONS

Table 12. Pin Function Descriptions (Continued)

ADIN6310 Ball No.	Mnemonic	Type ¹	Internal Termination at Reset ²	Power Domain ²	Description
E14, E15, J1, J2, L3, M3	DNC		N/A	N/A	Do Not Connect. Leave these balls open circuit.

¹ A = analog, I = input, O = output, I/O = input and output, and S = supply.

² N/A means not applicable.

³ There are multiple balls for some of the supply rails. It is critical for proper operation of the device to ensure that each ball is directly connected to the appropriate power rail.

TYPICAL PERFORMANCE CHARACTERISTICS




Figure 12. Power Consumption vs. Temperature for VDDIO Supply Voltages, 6 Ports RGMII 1 Gbps




Figure 15. Supply Current vs. Temperature for Various VDDIO Supply Voltages, 6 Ports RGMII 100 Mbps




Figure 13. Supply Current vs. Temperature for Various VDDIO Supply Voltages, 6 Ports RGMII 1 Gbps




Figure 16. Power Consumption vs. Temperature for Various VDDIO Supply Voltages, 2x RGMII 1 Gbps, 4x 1000BASE-KX




Figure 14. Power Consumption vs. Temperature for Various VDDIO Supply Voltages, 6 Ports RGMII 100 Mbps




Figure 17. Supply Current vs. Temperature for Various VDDIO Supply Voltages, 2x RGMII 1 Gbps, 4x 1000BASE-KX

THEORY OF OPERATION

SWITCH OVERVIEW

The ADIN6310 is a 6-port Gigabit Ethernet TSN switch with integrated security primarily designed for industrial Ethernet applications. Each port can be configured to operate at different speeds. This switch is designed to pair with Analog Devices PHY devices, such as [ADIN1100](#), [ADIN1200](#), and [ADIN1300](#), to form a low power, low latency system over RMI and RGMII. Optionally four of the ports support SGMII and serdes interfaces allowing backplane connections and connectivity to SFP modules.

The switch supports the suite of IEEE 802.1 time sensitive networking bridging features required by the IEEE 60802 standard providing QoS for latency sensitive streams.

The device also includes hardware capabilities to support PRP or HSR redundancy protocols, thereby offloading the host processor.

For legacy support, two of the ports (Port 1 and Port 2) integrate Analog Devices proprietary real-time Ethernet multiprotocol switch (REMS) technology that allows simple interfacing to classical 100 Mbps industrial Ethernet protocols, such as the following:

- ▶ PROFINET RT (Class B) and PROFINET IRT (Class C)
- ▶ EtherNet/IP with device level ring (DLR)
- ▶ POWERLINK
- ▶ Modbus TCP

Software drivers are provided for the host to work in conjunction with the switch to implement the industrial Ethernet protocol stacks.

Two host interface options are supported: standard, dual, or quad SPI and Ethernet. If an Ethernet host interface is selected, this port can operate as a control plane and a data plane to the host, leaving five Ethernet ports available for industrial Ethernet applications.

The host interface may be on a different voltage domain than the connections to the external PHY layer devices, VDDIO_A vs. VDDIO_B, allowing flexibility in choosing these voltages to allow simple interfacing while conserving power.

PORT OVERVIEW

[Figure 18](#) shows a simplified overview of a port transmit and receive. The port front end supports RGMII and RMI on all six ports and SGMII on four ports (Port 1 to Port 4), and the ports are independent; therefore, they can operate different speeds or interfaces. Interface details are discussed further in the [Ethernet Interfaces](#) section.

During initial configuration, the host configures the switch ports for the desired connectivity using the `SES_InitializePorts()` application programming interface (API) (see the `SES_switch.h` file for further details, which is available as part of the software drivers from the [ADIN6310](#) product page). As part of this configuration, the user has the ability to enable all ports or keep unused ports disabled. The host identifies how each port must be configured, what MAC interface, port speed, and whether there is an Ethernet PHY connected

to the port. The initialization can also perform any PHY-related configuration needed.

The receive MAC processes the incoming frame, checks if the frame size is within the minimum or maximum range, checks the interframe gap (IFG), the preamble, the SFD, and if the incoming frame is preemptable or has been preempted, and it assesses the start packet delimiter (SMD) value, order, and fragment count. The destination address, source address, and VLAN tag are extracted for use in filtering and address lookup, and the MAC calculates the cyclic redundancy check (CRC) of the frame and compares it to the frame check sequence (FSC). Any frames found to be in error are discarded. Frames that pass the various checks continue to be processed through the receive path.

While the various checks are happening, the frame is being forwarded onto the receive path to be passed to the transmit paths of the other ports, the transmit ports buffer the frame until all appropriate information is gathered and routing decisions are complete.

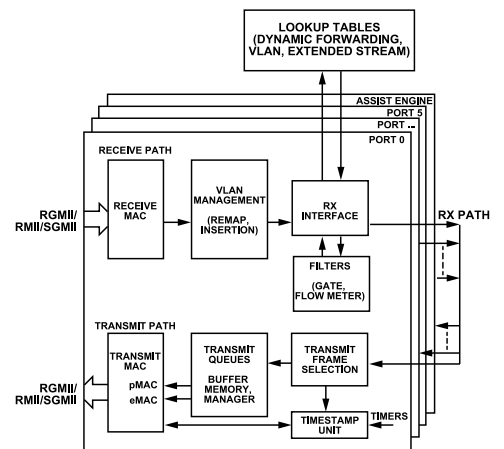


Figure 18. Simplified Port Transmit and Receive

The transmit block for each port supports eight priority queues and implements strict priority scheduling by default.

MAC, eMAC, and pMAC

The MAC consists of an express MAC (eMAC) and a preemptable MAC (pMAC) to support frame preemption (Qbu). The eMAC unit handles packets from the express queues and is capable of sending cut-through or store and forward packets. The pMAC handles packets from the preemptable queues and also monitors express queues to determine whether to fragment an active preemptable packet. The MAC verifies proper assembly of preemptable traffic by verifying proper order and combination of SMDs, continuation SMDs (SMD-Cs), and fragmentation counts, it also identifies any errors specific to preemption. The MAC front ends format the data to be driven onto the wire, maintain various statistics, such as transmit and receive port statistics, and generates mCRC and CRC for transmitted packets. Port statistics can be read on a per port basis using the `SES_GetStatistics()` API or cleared using the `SES_ClearStatistics` API (both included in `SES_switch.h`

THEORY OF OPERATION

file, which is available as part of the software drivers from the [ADIN6310](#) product page).

When transmitting, the switch ensures a minimum IFG of 11 byte-times.

STORM PROTECTION

The switch has the capability to protect the port against broadcast storms. This capability is provided as part of the per stream filtering and policing (Qci) feature (see the [Per Stream Filtering and Policing \(Qci\)](#) section).

SWITCH FABRIC

The ADIN6310 switch fabric is a high-performance engine that manages and moves data between ports and packet buffers. The switch fabric supports two modes of operation, store and forward and cut-through, with the latter providing the lowest latency. The switch has a packet storage capacity of 32 kBytes per port, which cannot be shared across ports.

Cut-Through Operation

Cut-through operation means that a frame has not been completely received before it starts to egress. In this case, not all information about the frame is known before transmission begins, such as length or if there are errors. Queuing of the frame for transmission may take place after the necessary forwarding data of the received frame has been received and processed through the forwarding database. Once transmission of a cut-through frame has begun, it always finishes.

Frames received with errors can only be discarded if they are store and forward. For a frame that has been cut-through with errors detected, the switch corrupts the CRC on egress.

Cut-through operation is the default mode of the device. Cut-through operation can be configured per egress port, per queue, see the `SES_SetStoreAndForwardMask ()` API and `SES_GetStoreAndForwardMask` API included in the `SES_switch.h` file, which is available as part of the software drivers from the [ADIN6310](#) product page. It is possible to add a static table entry that defines whether a particular frame should egress in cut-through or store and forward mode.

Cut-through operation is only possible if the egress port is not currently transmitting, and the speed of the egress port must be equal to or less than the speed of the ingress port; otherwise, the frame is forwarded in store and forward mode. Only express frames can be designated for cut-through operation. Note that preemptable frames are always store and forward.

A per-traffic class queue, `QueueMaxSDU`, setting defines the maximum data unit size for each queue, with the frames that exceeds the programmed value being discarded. These `QueueMaxSDU` settings do not apply to cut-through frames because the switch has already started forwarding before it knows the size of the frame.

Store and Forward

The switch acts as store and forward when it is not possible to cut-through a frame. The user can configure the device to operate in store and forward mode using the `SES_SetStoreAndForwardMask ()` API and `SES_GetStoreAndForwardMask` API included in `SES_switch.h` file, which is available as part of the software drivers from the [ADIN6310](#) product page. In this mode, the entire frame is stored in memory to be sent at a later time to its destination. The entire frame can be checked; therefore, errored frames are identified and discarded. `QueueMaxSDU` settings apply to store and forward frames, and frames in excess of the programmed setting are dropped. `QueueMaxSDU` settings are programmable on a per port and per queue basis.

LOOKUP TABLES

Dynamic and Static Forwarding Table

The forwarding table stores MAC addresses and associated information for dynamic and static entries. The table size is 2048 entries, the current use of the dynamic table can be read using the `SES_GetDynTblEntryUsage()` API (see `SES_switch.h` file, which is available as part of the software drivers from the [ADIN6310](#) product page). Dynamic entries get created automatically in the hardware in response to port learning from incoming frames, while static entries are typically created by management software. Each received frame generates a lookup request to the forwarding table. If an entry already exists (returns a hit), the table returns information about where to send that frame and any transforms or filters to apply. If no entry exists (for example, it returns a miss), the frame is forwarded with the default behavior. Users can configure the default miss behavior via the `SES_SetUnicastMissReturn()` API and `SES_SetMulticastMissReturn()` API included in `SES_switch.h` file. For VLAN tagged traffic, the ports must be configured to forward the VLAN IDs of interest for traffic to egress.

Learning

The learning process is a function executed for all frames received on a port unless learning is disabled. Learning is an automated mechanism for populating the dynamic forwarding table. The table is updated with a new dynamic entry if an entry does not already exist, if the received packet has no errors, if the received packet has a unicast source address, and if the packet size is valid. Assuming the frame is valid, a new entry is created in the forwarding table with the MAC address, VLAN ID, ingress port, and the age is initialized. As further frames are received and lookups performed, if an entry is matched, the age field on the entry is reset. The `SES_GetLearnMode()` API (see `SES_switch.h` file, which is available as part of the software drivers from the [ADIN6310](#) product page) allows users to check the learning mode of the switch. Users can also interrogate and read back the entries in the dynamic table.

THEORY OF OPERATION

Insert or Remove Entries

Management software can install a static entry into or remove an entry from the forwarding table.

The API to add basic static entries is **SES_AddStaticTableEntry()**, which requires the MAC address, VLAN ID, and the destination ports. It is not possible to read back these basic static entries. The **SES_RmStaticTableEntry()** API can be used to remove an entry.

More detailed static table entries can be added to the table using the **SES_AddStaticTableEntryEx()** API, which provides more configuration for the entry; this extended entry can be read back using the **SES_ReadStaticTableEntry()** API or removed using the **SES_RmStaticTableEntryByIndex ()** API. These APIs are all part of **SES_switch.h** file, which is available as part of the software drivers from the [ADIN6310](#) product page.

Aging

Each dynamic table entry has an age value associated with it. The age represents the elapsed time since the entry was last updated and/or last seen. The lookup state machine updates the age count information of a dynamic entry in the table whenever the corresponding source address and VLAN combination appears. If a record is not updated for a time and the age expires, the state machine removes the record from the table. Static entries are exempt from the aging process. The default aging value is 300 seconds and is programmable through the API. Note that the range of programmability extends from 2 seconds to 4 hours.

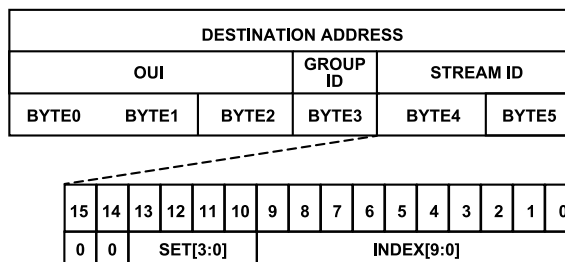
Aging time can be read or changed using the **SES_GetDynamicTblAgeOutPeriod()** API and **SES_SetDynamicTblAgeOutPeriod()** API which are part of **SES_switch.h** file, which is available as part of the software drivers from the [ADIN6310](#) product page.

Extended Lookup Table

The extended lookup table has 256 entries. The purpose of this lookup table is to allow a deeper look into the frame to identify different types of traffic to perform filtering and transformations.

Stream Lookup Table

The stream table can support up to 16k entries. It works by using a portion of the destination address along with the VLAN ID to distinguish streams. The first three bytes of the address are fixed and the lower 14 bits are used to distinguish the streams. The switch supports 16 blocks of 1024 entries to provide a total of 16k entries. The stream lookup management table is 16 entries deep (Bits[3:0]), and each entry defines the start of region from the base address + 1024 entries.



OUI: ORGANIZATIONALLY UNIQUE IDENTIFIER.

013

Figure 19. Destination Address to Stream Index

THEORY OF OPERATION

TYPES OF LOOKUPS

Figure 20 provides an overview of the types of lookups possible in the switch.




Figure 20. Lookup Flowchart

014

THEORY OF OPERATION

Standard Lookup

The standard lookup is the default type of lookup. A standard destination address and VLAN lookup is performed on each ingressing frame. The frame information is searched against all entries in the static and dynamic table, returning the mapping and transform fields associated with a hit, or otherwise, it returns a miss. The return information provided by a hit can prompt further lookups, as required, how to handle the traffic, what types of receive filtering must be applied (such as a stream gate or flow meter), and what types of transforms must be applied to the frame, such as the following:

- ▶ Inserting or removing PRP, IEEE 802.1CB, or HSR tags
- ▶ Timestamp-related changes (correction fields, source addresses, and source port identity)
- ▶ Destination address changes
- ▶ Inserting, removing, or modifying VLAN tags

A miss indicates there was no entry in the table, and the default processing is applied. The default behavior is configurable for unicast and multicast and broadcast frames.

Once the destination and VLAN lookup is complete, if additional lookups are required, these lookups can be performed in parallel.

Source Lookup

If the destination address and VLAN ID lookup prompts a source lookup, this is a lookup of the source address and VLAN ID. The return information from the lookup directs how to route the frame and any filtering or transforms to apply. A source lookup can be prompted based on how a static entry was installed. Alternatively, a port can be configured to perform source lookups on all traffic.

Extended Table Lookup

When enabled, the extended table lookup mode is performed in addition to the standard lookup or the source stream lookup. The additional analysis performed by the extended lookup results in longer processing of the frame. This mode identifies the EtherType in the frame, which is then looked up in an EtherType transform table. The EtherType transform table is a 16 entry table that indicates the search pattern for each supported EtherType. The return values are eight pairs of offset and count. The offset value indicates the data bytes past the EtherType to start populating the extended search with, while the count indicates the number of bytes starting from this offset to populate the extended search with.

The data extracted from the packet represents the total data pattern to be matched for that EtherType.

An extended lookup can be prompted based on how a static entry was installed. Alternatively, a port can be configured to perform extended lookups on all traffic.

Stream ID Lookup

The Stream ID lookup mode can be performed in addition to the standard lookup or the source stream lookup. Stream configuration APIs are part of **SES_stream.h** file, which is available as part of the software drivers from the [ADIN6310](#) product page.

VLAN Table and Management

A VLAN is used to partition the network into multiple virtual networks where traffic is directed to specific subsets of the larger network. The VLAN protocol uses a 4 byte tag included in the header of the Ethernet frame. The switch supports a VLAN table with 4096 entries, each entry in the table sets the state per port in that VLAN, such as disabled, learning, or forwarding.

The VLAN management block controls the functions applied at the receive port, remapping VLAN IDs or priorities on ingress, replacing VLAN IDs, adding VLAN tags if there is not any assigned and restricting traffic ingressing a port to particular sets of VLANs.

The default VLAN behavior is no learn and no forward on all VLAN IDs, with the exception of untagged and priority (VID 0) for all ports. Therefore, the user must configure the VLAN table directly or configure the ports as trunk or access type during the initial configuration of the device. Trunk ports can support multiple VLAN IDs or ranges of VLAN IDs, whereas access ports support only one VLAN ID.

When using a VLAN trunk or access port configuration, the switch handles the insertion and removal of VLAN tags, where required, when traffic is crossing between ports.

The VLAN configuration APIs are part of **SES_vlan.h** header file, which is available as part of the software drivers from the [ADIN6310](#) product page. The switch does not support double tagging. While the VLAN table supports 4096 entries, in practice, systems only use a small number of VLANs. When configuring the switch VLAN table, the number of active VLANs must be limited to a maximum of 64.

BUFFER MANAGEMENT

Each transmit port has a dedicated frame buffer of 32 Kbytes. In the event this buffer fills, then the oldest, lowest priority traffic gets discarded. Each traffic class queue can hold 32 entries. In the event a queue is full, the oldest entry in that queue gets removed (irrespective of traffic class). Incoming frames are stored in buffer memory as the packet is ingressed. If the egress port is available to transmit, the frame can cut-through once sufficient information is available to make a forwarding decision. If the port is busy, or the return information indicates that the frame must be treated as store and forward, the full frame is buffered and added to the egress ports transmit queue.

THEORY OF OPERATION

QUALITY OF SERVICE (QOS)

The switch QoS provides eight internal queues per port that support eight traffic priorities. Incoming frames are assigned to egress transmit queues depending on the priority in the VLAN tag or untagged. The VLAN priority of 0 traffic goes into the lowest priority queue of 0. Traffic with a VLAN priority of 7 is routed to the highest priority queue of 7. The switch uses strict priority to schedule each transmit queue; therefore, the egress port transmits frames from each queue with the higher priority queues transmitting before lower priority queues. The switch supports configuration on a per port basis to reprioritize VLAN priorities.

PORT FORWARDING MASK

The switch supports a per port forwarding port mask, which can be used to further segregate how traffic can be forwarded between ports. The host controller configures the port forwarding mask required per port as part of the device initialization routine. The API to configure this feature is **SES_ApplyForwardMask()**, which is part of **SES_logical_mac.h** header file, which is available as part of the software drivers from the **ADIN6310** product page.

ETHERNET FRAME

The Ethernet frame is shown in **Figure 21**. It consists of a preamble which is a series of 8 bytes of 0x55 followed by the SFD or SMD. The receive front end can handle preambles with fewer bytes down to a minimum of one 0x55 followed by the SMD or SFD.

If the frame preemption feature is used, the SMD indicates what type of frame follows, either an express frame, the initial portion of a preemptable packet or a continuation frame of a fragmented preemptable packet. The types of SMDs expected are as follows:

- ▶ SMD-E, the express frame delimiter = 0x5D (standard SFD)
- ▶ SMD-V, the verify packet = 0x07
- ▶ SMD-R, the respond packet = 0x19
- ▶ SMD-S0, through SMD-S3 the preemptable packet start
- ▶ SMD-C0, through SMD-C3 the continuation fragment

The Layer 2 header follows the SMD, with the destination address and the source address. A standard Ethernet frame may not have a VLAN tag; however, a tag can be added by the switch port management. TSN streams typically have a VLAN tag, which contains a priority (PCP) and a 12-bit VLAN ID. The VLAN ID is used with the destination address to identify the TSN stream in the lookup tables.

In industrial applications, typical frame sizes are less than 1522 bytes (including the VLAN tag). The switch can support reception and transmission of frames in excess of 1522 bytes; however, jumbo packets consume larger blocks of buffer memory and Ethernet CRC robustness degrades significantly for longer frame sizes, which can result in undetected errors in these frames. Frames in excess of the default maximum length (1522 bytes) are recorded in the receiving port statistics as a large frame error.

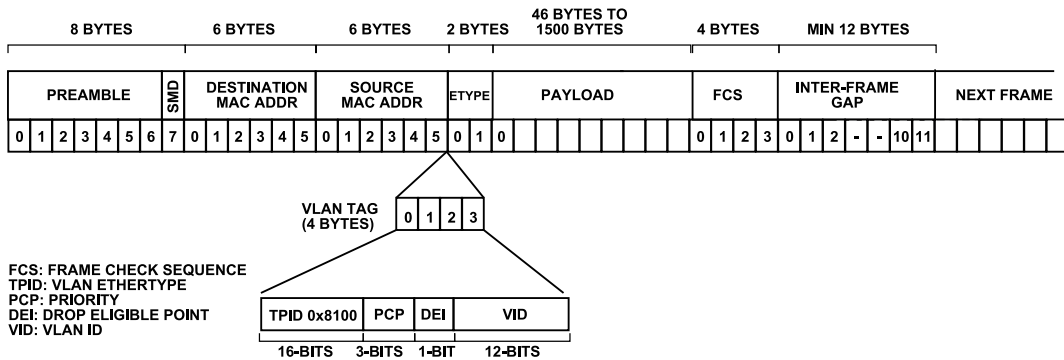


Figure 21. Ethernet Frame

THEORY OF OPERATION

SWITCH LATENCY

The switch operation (cut-through vs. store and forward) and the types of lookups have a direct effect on the latency experienced by

a frame. [Table 13](#) captures some of the key latency specifications and calculations for common use cases and lookups.

Table 13. Switch Latency vs. Lookup^{1, 2}

Type of Lookup	Latency Specification and Calculation	Conditions
Cut-Through		minByteCnt = 0x10 (Ethernet header; VLAN tag)
VLAN Insertion and Modification on the Receive Port	(DA + VID) + 0 ns	Same as DA + VID lookup
Source Address Lookup	(DA + VID) + 40 ns	Latency increase in RGMII 1000 Mbps due to additional lookup; no change at 100 Mbps or 10 Mbps
Stream Table Lookup	(DA + VID) + 0 ns	Same as DA + VID lookup
Modifications and Insertions on Transmit Port	(DA + VID) + 40 ns	Add, strip VLAN/IEEE 802.1CB header, insert timestamp, remap, and change priorities; latency increase applies to all speeds
Extended Lookup	(Inspection_Depth ³ (bytes) + 8) × Symbol_Time + 320 ns	
Store and Forward		
Port to Port	(Frame Size ⁴ (bytes) + 8) × Symbol_Time + 200 ns	No packet modification
Port to SPI Host ⁵	(Frame Size (bytes) + 8) × Symbol_Time + 110 μs	

¹ Measured from ingress SFD to egress SFD or egress host interrupt (for the host port). No interfering traffic, same frequency on ingress and egress.

² Symbol_Time = 800 ns, 80 ns, 8 ns for 10 M, 100 M, and 1000 Mb, respectively.

³ How many bytes into the frame the extended lookup analyzes.

⁴ Frame size is in bytes to transmit.

⁵ Communication from ports to the SPI host is always store and forward, which is a function of firmware.

THEORY OF OPERATION

RECEIVE AND TRANSMIT PORT STATISTICS

The switch records information is based on the traffic crossing its ports, capturing statistics on frames received, transmitted, and available to readback by the host. The receive statistics are detailed in [Table 14](#).

Table 14. Per Port Receive Statistics

Parameter	Description
Bytes Received	Number of bytes received on a port. Only applies to the frame and not the packet (that is, preamble, SFD or SMD, and CRC or mCRC)
Unicast Packets Received	Number of unicast addresses received on a port, and counter is incremented for errored packets.
Broadcast Packets Received	Number of broadcast packets received on a port, and counter is incremented for errored packets.
Multicast Packets Received	Number of multicast packets received on a port, and counter is incremented for errored packets.
Frames Received with Alignment Errors	Number of packets with alignment errors.
Frames Received with CRC and FSC Errors	Number of packets with CRC, mCRC, or FSC errors.
Frames Received with Large Frame Errors	Number of packets received on a port greater than the configured maximum length in bytes.
Frames with Receive MAC Errors	Number of frames with receive MAC errors.

The transmit statistics are detailed in [Table 15](#).

Table 15. Per Port Transmit Statistics

Parameter	Description
Bytes Transmitted	Number of bytes transmitted.
Unicast Packets Transmitted	Number of unicast packets transmitted.
Broadcast packets Transmitted	Number of broadcast packets transmitted.
Multicast Packets Transmitted	Number of multicast packets transmitted.
Frames Transmitted After Single Collision	Number of frames that experienced a single collision.
Frames Transmitted After Multiple Collision	Number of frames that experienced multiple collisions.
Frames Dropped After Excessive Collisions	Number of frames dropped because of multiple collisions.
Frames with a Collision After 512 Bits	Number of frames that experienced a late collision.
Frames Delayed by Traffic	Number of frames delayed by traffic. Carrier busy at first attempt.
Frames with Transmit MAC Errors	Number of frames with a transmit MAC errors.

PACKET ASSIST ENGINE

The packet assist engine is an on-chip microcontroller that manages the switch and external PHY devices connected to the MDIO interface. The engine controls the basic switch operation such as managing the switch configuration, installing static entries into the forwarding table, and configuring the VLAN table in response to this configuration being applied from an external host processor.

The host port connection can be through an SPI, dual SPI, or a quad SPI or alternatively through an Ethernet port.

SOFTWARE RUNNING ON THE SWITCH

Figure 22 shows an overview of the partitioning of software between the host processor and the switch packet assist engine. Blocks shown in blue are software deliverables provided by Analog Devices (by its firmware and driver libraries).

The packet assist engine runs firmware to control much of the switch-specific configuration and TSN features thereby offloading that overhead from the host processor. The firmware running on the packet assist engine is provided by Analog Devices, and this firmware can be securely updated and updates are provided as the solution progresses.

The software driver API package (which is available on the [ADIN6310](#) product page) provides the host processor with the ability to configure the features of the switch. The features supported on the ADIN6310 are as follows:

- ▶ Switch configuration management
 - ▶ Basic switch configuration, port configuration, port statistics
 - ▶ Forwarding table management for static entries, a dynamic table flushing capability, and aging control
 - ▶ VLAN table configuration, remapping, and reprioritization
 - ▶ Hardware GPIO, timer control
 - ▶ Logical MAC function
 - ▶ Port masking capability
 - ▶ Stream table configuration
 - ▶ Layer 2 transmit and receive frames
 - ▶ Event subscription
 - ▶ SendList
 - ▶ Media Redundancy Protocol (MRP) stack
 - ▶ High availability (HSR, PRP, IEEE 802.1CB)
- ▶ PHY drivers ([ADIN1100](#), [ADIN1200](#), or [ADIN1300](#) depending on the end application) communication over the MDIO interface
- ▶ Configuration and status of TSN functionality
 - ▶ Scheduled traffic, Qbv configuration for ports and hardware timer pins
 - ▶ Frame preemption, Qbu configuration, state
 - ▶ Time synchronization, IEEE 802.1AS configuration (intervals, priority, states)
 - ▶ Per stream filtering and policing, Qci configuration, statistics
 - ▶ Frame replication and elimination for reliability (FRER), IEEE 802.1CB configuration, statistics

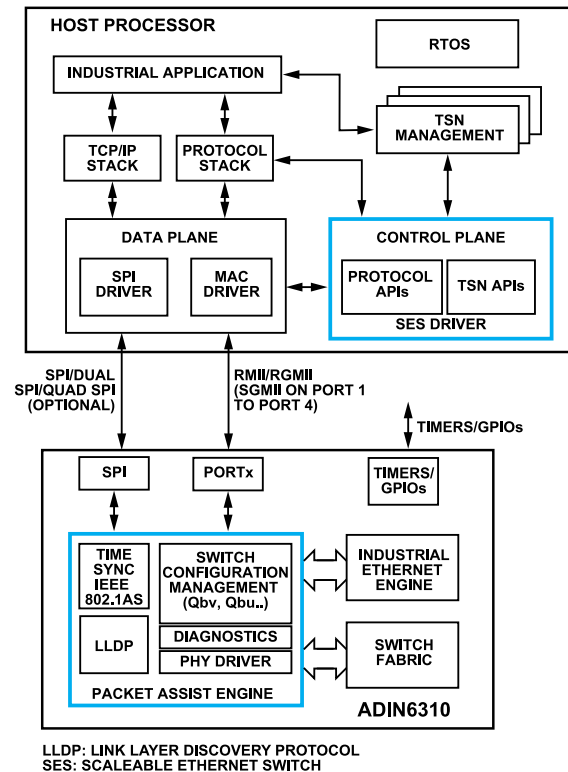


Figure 22. Overview of Software Partitioning Between Switch and Host Processor

FIRMWARE UPDATE

The switch supports secure firmware updates over the host interface via the driver (either SPI or Ethernet firmware updates are supported). The firmware is provided as part of the TSN library package. Devices are supplied unprogrammed, without firmware loaded. The first action the host must take is to load the firmware. The firmware image is approximately 380 kB (but may increase up to 500 kB). The firmware is included as part of the driver package, see [SES_firmware.c](#) file, which is available as part of the software drivers from the [ADIN6310](#) product page. The firmware image can reside in memory of the host and, therefore, requires enough space to accommodate this or have separate memory available to the host. The TSN library APIs support automatic firmware updates. When a newly installed device is powered up, its bootloader checks and requests firmware from the host. The host must respond to this request from the bootloader and support the process of loading the firmware. This process automatically starts updating the firmware, and the device resets after a successful firmware update.

Firmware can also be loaded directly using the API calls, and users can readback current firmware version running on the device.

The driver library and firmware are paired and must always match. When a new version of the driver and firmware is run, the automatic firmware update process can update the firmware running on the switch. See the [ADIN6310 Software Driver User Guide](#) for additional information and descriptions of the APIs.

PACKET ASSIST ENGINE

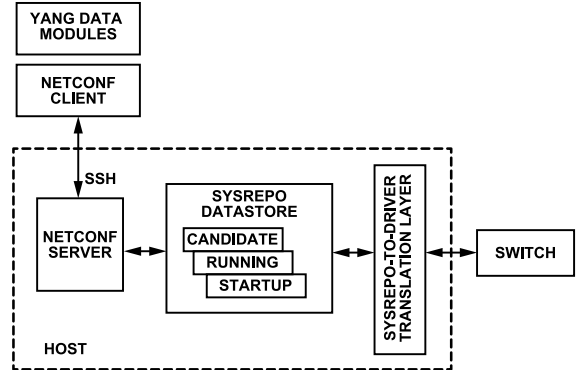
NETCONF

NETCONF is a network management configuration protocol used to configure network elements by transporting XML configuration files in the form of YANG (data modeling language) modules. NETCONF and YANG are gaining traction over older protocols, such as simple network management protocol (SNMP) and management information base (MIB) databases.

The switch supports configuration of all switch capabilities and TSN feature sets using a NETCONF server running on the host processor and loading YANG modules to the switch.

Included in the software driver library package is a translation layer between the Sysrepo datastore and the software driver.

Both TSN and other bridge functionality can be configured via NETCONF. The switch models follow the IEEE YANG models and, where necessary, include custom leaf nodes for device-specific features.



SSH: SECURE SHELL

017

Figure 23. NETCONF to Switch Overview

INTERFACES

SERIAL INTERFACE, SPI

The switch contains a synchronous SPI-compatible serial interface that can be used as the host interface. The SPI supports 32-bit data transfers, is a target interface, and can be a standard SPI or a dual or quad SPI. Table 16 lists the relevant pins for the different SPI interface modes. Configuration of the host interface mode can be made using external boot strapping (see the Table 20 section). In SPI mode, the interface allows full duplex operation and half duplex communication in dual or quad SPI.

The transactions on the interface are consumed by the packet assist engine and facilitate communications between the switch and the higher level communications and processor. The SPI provides a communication medium for low to medium performance systems with higher performance processors or field-programmable gate arrays (FPGAs) interfacing directly to the RMII, RGMII or SGMII on one of the ports.

Chip Select, SPI_SS

The active low, chip select pin (SPI_SS) enables and disables the SPI. When SPI_SS is high, the device is not selected, and the serial data outputs (SPI_SIOx) are at high impedance. When SPI_SS is brought low, the device is selected, and commands can be written to or read from the device.

Serial Clock, SPI_SCLK

The serial clock input is used to synchronize data to and from the packet assist engine of the switch. The maximum SCLK frequency is 38 MHz.

Serial Data Input and Output, SPI_SIOx

The switch supports standard SPI, dual SPI, and quad SPI communications. The device has four SPI_SIOx pins. When operating with a standard SPI, the SPI_SIO0/SDI is the unidirectional serial data input line for write operations and clocked in on the rising edge of SPI_SCLK. The SPI_SIO1/SDO is the unidirectional data output for reads and is clocked out on the falling edge of SPI_SCLK.

For dual and quad SPI, commands use the bidirectional SPI_SIOx pins to read and write data.

Data lines are only driven in response to a read command.

Table 16. SPI Pin Operation

Mnemonic	Quad SPI	Dual SPI ¹	SPI ¹
SPI_SS	SS	SS	SS
SPI_SCLK	SCLK	SCLK	SCLK
SPI_SIO0	SIO0	SIO0	SDI
SPI_SIO1	SIO1	SIO1	SDO
SPI_SIO2	SIO2	N/A	N/A
SPI_SIO3	SIO3	N/A	N/A

¹ N/A means not applicable.

Interrupt, INT

The TIMER0/INT pin provides a single interrupt output from the switch to the host and is used to coordinate SPI operations. The INT function is a shared function on the TIMER0/INT pin and is managed via the packet assist engine software, which is available on the ADIN6310 product page.

Dual and Quad SPI Interactions

Figure 24 shows an overview of the bit and byte ordering for quad SPI read communications. Reads require a two clock period turnaround before the data is available on the SPI_SIOx pins. The SPI_SIOx pins are driven low during this two clock turnaround time. The write transaction looks similar in terms of format; however, without the turnaround, the first nibble of data immediately follows the second nibble of instruction byte. Data is transmitted and received MSB bit of data on the MSB of the bus.

Dual SPI transactions are similar, SPI_SIO1 carries the most significant bit of a pair, and the most significant pair of a byte is transmitted first. Similarly, reads use a two clock period turnaround.

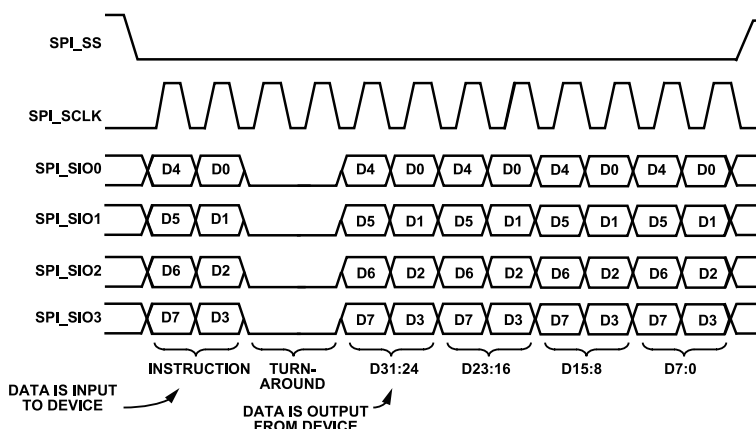


Figure 24. Quad SPI Data Read

INTERFACES

Communication over SPI

In applications using any of the SPI options as the host interface to the switch, the switch uses the `TIMER0` hardware pin as an interrupt to indicate availability of data to the host. If the host uses Ethernet as the preferred interface, then the `TIMER0` pin does not act as interrupt pin and is available as a timer pin. For all SPI modes, the host must connect the `TIMER0` pin to an available GPIO pin on the host. A rising edge event on this pin must trigger a software interrupt on the host side to prompt a read of the message from the switch. In this interrupt service routine, the hardware abstract layer (HAL) interface read function must be called to read data from the scaleable Ethernet switch (SES) and then call the **SES_ReceiveMessage()** API of the driver, which is available as part of the software drivers from the [ADIN6310](#) product page. When an interrupt from the switch is received, the host must trigger a HAL interface read of 4 bytes (`uint32_t`). The response from the SES for this interface read call results in the length of bytes available to read from the SES. The host must call another HAL interface read with this length of bytes to the SES. Finally, the response data of this call must be passed to the **SES_ReceiveMessage()** API. Ideally, the host dedicates a separate thread for this event or has a interrupt service routine (ISR).

PHY MDC AND MDIO INTERFACE

The switch has a controller MDC and MDIO interface, which is a 2-wire bus used to manage control and status of the PHY layers connected to the switch. The PHY layers act as targets on this bus, each with its own distinct address. The PHY addresses are set with external strapping resistors that are sampled at reset. Alternatively,

the switch can control each PHY address directly. This interface can have up to 32 target devices connected. The MDIO interface consists of a clock pin (MDC) and a bidirectional data line (MDIO). The clock is an output from the switch and an input to the PHY layers. The MDIO pin requires an external 1.5 k Ω resistor pulled to `VDDIO_B`. The data line is typically driven from the switch, except when a read is commanded to one of the PHY layers.

The bus is under control of the packet assist engine in switch. If an external host processor requires access to the PHY registers, it must request this access through its general interface and let the packet assist engine perform the bus transaction. The frame format is standard for this type of interface, and each instruction consists of the following:

- ▶ Preamble: establishes synchronization at beginning of the frame.
- ▶ Start of frame: 01 indicates the start of the frame.
- ▶ OP: the operation code indicates the type of frame transaction; 10 for read, and 01 for write.
- ▶ PHYAD: PHY address. MSB first, only the PHY layer with the matching PHY address responds.
- ▶ REG ADDR: register address, MSB first.
- ▶ TA: used to avoid contention during a read transition, with 2-bit time spacing between the register address field and data field.
- ▶ DATA: 16-bit field, MSB first.
- ▶ ADDRESS and DATA: 16-bit field, MSB first.
- ▶ IDLE: high-Z state, the MDIO line is pulled high by the pull-up resistor.

Table 17. Frame Format

Operation	Preamble	Start of Frame	OP	PHYAD[4:0]	REG ADDR[4:0]	TA	DATA[15:0]	IDLE
Read	32 1s	01	10	AAAAA	RRRRR	Z0	d ... d	Z
Write	32 1s	01	01	AAAAA	RRRRR	10	d ... d	Z

INTERFACES

ETHERNET INTERFACES

There are six Ethernet ports in the ADIN6310. Each port can support RMII or RGMII and four of the ports (Port 1 to Port 4) can support SGMII. The pins associated specifically with the various interfaces are listed in the following sections.

The MAC interfaces support both full duplex and half duplex for 10 Mbps and 100 Mbps speeds and full duplex for the 1000 Mbps speed. The switch does not support energy efficient Ethernet (EEE); therefore, EEE must not be enabled on the PHY layer side.

Figure 25 through Figure 27 and Figure 29 through Figure 32 show overviews of the interconnections between the switch and the ADIN1200 10 Mbps and 100 Mbps or the ADIN1300, 10 Mbps, 100 Mbps, or 1000 Mbps Ethernet PHY layers. RGMII is the preferred MAC interface when connecting to the ADIN1200 and ADIN1300 devices.

Note the ADIN1300 LINK_ST pin is active high by default but can be changed to active low through MDIO control. The Px_LINK switch defaults active low for link up. For nonhost ports, during port initialization, the user can instruct the switch to change the PHY link polarity to match what the switch port expects. Where a PHY is used with an Ethernet-based host interface, it is important to ensure that the link polarities match (drive low for link up); therefore, if using an ADIN1200 or ADIN1300 PHY on the host port, use an inverter in the path to ensure that the PHY link polarity matches what the switch expects because the switch cannot reconfigure the link polarity for a PHY layer connected on its host interface. Ensure the PHY is hardware strapped for unmanaged mode and capable of coming out of reset without any interaction from the switch. The Px_RSTN from the switch must not be connected directly to the PHY RESET_N input for the host interface port because the switch does not control this pin on the host port.

ADIN1200 and ADIN1300 PHY Addressing

For the ADIN1200 and ADIN1300 PHYs, the PHY addressing is shared with the four RXD_x data pins. The PHYs have weak internal pull-down resistors to default to a PHY address of 0. The switch includes internal pull-up and pull-down resistors to ensure each port provides a unique address for each PHY, which removes the requirement to have external pull-up and pull-down resistors for PHY address strapping purposes on the board. The default PHY addresses are shown in Table 18. In RGMII mode, all four RXD_x pins are connected between the switch and the PHY, and each port has a unique PHY address.

Table 18. PHY Addressing (RGMII)

Port Number.	PHY Address	RXD_3 to RXD_0
0	0	0000
1	1	0001
2	2	0010
3	4	0100
4	8	1000

Table 18. PHY Addressing (RGMII) (Continued)

Port Number.	PHY Address	RXD_3 to RXD_0
5	9	1001

An exception to the details in Table 18 is that the PHY address of the port assigned as the host interface always gets a PHY address of 0.

For Ethernet host connections, the expectation is that the host MAC will interface directly to the switch MAC, with no PHY in the paths. If there is a PHY in the path, the PHY will not be managed by the switch; therefore, the user must ensure that the PHY is hardware configured for the operation it requires, the link signal from the PHY provides the switch with an active low for link up, and the PHY RESET_N is pulled high.

If the use case does have a PHY layer on the host interface and chooses a port other than Port 0 as the host interface, it is recommended to use external PHY address strapping to avoid having two PHY layers in the system with Address 0.

When using RMII mode, the RXD_0 and RXD_1 of the PHY are connected to the corresponding Px_RXD0 and Px_RXD1 pins on the switch. The Px_RXD3 pin is now connected to RX_ER of the PHY layer. In this case, external pull-up resistors are required on RXD_2 and RXD_3 of the PHY for Port 3 to Port 5 to match the same port address, see Table 19.

The PHY addressing approach for the ADIN1100 PHY uses some different pins; therefore, the user must always use external PHY address strapping resistors when pairing the ADIN1100 PHY layer with the switch.

By default, the internal strapping resistors are enabled; however, these resistors can be disabled during the port initialization process if using external strapping resistors. It is important to review the PHY address strapping to ensure the expected addressing is applied. For Analog Devices ADIN product PHYs, it is recommended to hardware strap the PHYs into the software power-down mode on startup. During the switch configuration, the PHYs can be reset and brought out of software power-down mode, which ensures that the PHYs get the correct address strapping.

Table 19. PHY Addressing (RMII)

Port Number	PHY		External Pull-Up
	Address	RXD_1/RXD_0	
0	0	00	None
1	1	01	None
2	2	10	None
3	4	00	Add pull-up on RXD_2
4	8	00	Add pull-up on RXD_3
5	9	01	Add pull-up on RXD_3


INTERFACES

RMII Mode

The RMII is an eight signal interface for each port. This interface uses a 50 MHz reference clock that can be provided by the switch (on the Px_TXC pin) to the PHY layers or provided from an external clock source to both the switch and PHY layers. This interface is capable of supporting 10 Mbps and 100 Mbps data rates.

Figure 25 shows the switch connected to the 10 Mbps and 100 Mbps ADIN1200 PHY and Figure 26 shows the switch connected to the ADIN1100 10BASE-T1L PHY. The RMII connects to just two of the RXD_x pins of the PHY device. The receive data transitions synchronously to the reference clock on Px_TXC. The Px_RXCTL is a combination of the carrier sense and received data valid signal from the PHY layer. Px_RXD3 connects to RX_ER, which is synchronous to the Px_TXC clock, and is asserted when the PHY detects an error in the received frame or when a false carrier is detected in 100 Mbps.

The transmit data transitions are synchronous to the Px_TXC, and Px_TXCTL is used to indicate to the PHY TX_EN that transmission data is available on the transmit data pins.



*ADIN1200 RESET_N RESIDES IN 3.3V DOMAIN SWITCH Px_RSTN RESIDES ON VDDIO_B VOLTAGE DOMAIN 019

Figure 25. RMII Switch PHY Interface Signals (ADIN1200 PHY)

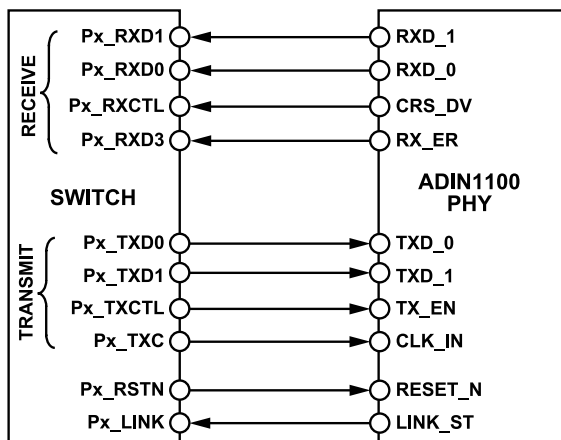



Figure 26. RMII Switch PHY Interface Signals (ADIN1100 PHY) 020

RGMII Mode

The RGMII is capable of supporting data rates of 1 Gbps, 100 Mbps, and 10 Mbps. Transmit and receive data is clocked on the rising and falling edge of the clocks when operating at 1 Gbps, thereby reducing the number of data lines for the interface.

For the receive interface, the PHY generates a 125 MHz, 25 MHz, or 2.5 MHz clock signal on RXC to synchronize the receive data in 1000 Mbps, 100 Mbps, or 10 Mbps modes, respectively.

Px_RXCTL receives a combination of data valid (RX_DV) and error information (RX_ER) and uses both edges of the clock, with data valid. The PHY transmits the RX_DV signal on the positive edge of RXC and a combination (XOR function) of RX_DV and RX_ER on the negative edge of RXC.



*ADIN1300 RESET_N RESIDES IN 3.3V DOMAIN SWITCH Px_RSTN RESIDES ON VDDIO_B VOLTAGE DOMAIN 021

Figure 27. RGMII Switch PHY Interface Signals

For the transmit interface operating in 1 Gbps mode, the switch provides the 125 MHz clock on Px_TXC and transmits data, Bits[3:0] on the positive edges and Bits[7:4] on the negative edge. In 100 Mbps and 10 Mbps modes, TXC is at 25 MHz and 2.5 MHz, respectively, and the switch transmits data, Bits[3:0] on the rising edges of TXC (but in accordance with the RGMII standard, the data may be duplicated on the falling edge of the appropriate clock). Px_TXCTL provides a combination of the TX_EN and TX_ER signals using both edges of TXC. TX_EN is transmitted on the positive edge of TXC, and TX_EN XOR TX_ER is transmitted on the negative edge of TXC. Due to the fact that data is transmitted on both edges of the clock, an accurate delay requirement of 2 ns is required on both clock edges (see Figure 28) to ensure that the delayed clock is at the center of the data window, ensuring accurate data capture. For the host port, it is possible to enable this 2 ns delay on RXC only, TXC only, on both RXC and TXC, or alternatively, no delay by using the hardware pin configuration settings. For all other ports, these delays can also be configured in software, and in addition, the driver API provides some programmability around the 2 ns delay (range of 1.69 ns to 2.55 ns). This delay is added at only one stage in the transmit and receive path; typically, it is added on the receiving side. Some MACs may not have the capability to add delay; therefore, the switch can include the delay for both

INTERFACES

sides, if necessary. Configuring the delays correctly on both sides is necessary to ensure successful communication.

In Figure 28, the 8 ns period refers to 1000 Mbps operation, which is 40 ns or 400 ns in the case of 100 Mbps and 10 Mbps, respectively. The 2 ns delay is valid for 10 Mbps, 100 Mbps, and 1000 Mbps.




Figure 28. DLL Waveform

The RGMII operates at 10 Mbps and 100 Mbps speeds exactly the same way it does at gigabit speeds with the exception that the data may be duplicated on the falling edge of the appropriate clock.

RGMII Drive Strength

The drive strength of the RGMII outputs has some programmability. When the host interface is a RGMII MAC interface, there are two choices for drive strength determined by the strapping voltage level of one of the SPI pins (SPI_SS). When the device is brought out of reset, when SPI_SS is pulled high, this results in the highest drive strength option, and the RGMII drive strength for the other ports is programmable via the SES_SetRxDelay() API and SES_SetTxDelay() API (both of which are included in SES_switch.h file, which is available as part of the software drivers from the ADIN6310 product page).

MAC TO MAC INTERFACE

Figure 29 and Figure 30 show MAC to MAC connections where a switch port is connected directly to a MAC. The Px_LINK input must be either pulled low with a pull-down resistor or alternatively driven low.

In the RMII, the Px_RXD3 acts as RX_ER. If the MAC is not providing an RX_ER signal, ensure that this pin is pulled to ground. The switch can provide a 50 MHz clock to the MAC REF_CLK, or alternatively, an external 50 MHz can be applied to both.




Figure 29. RGMII to RGMII MAC to MAC Connections

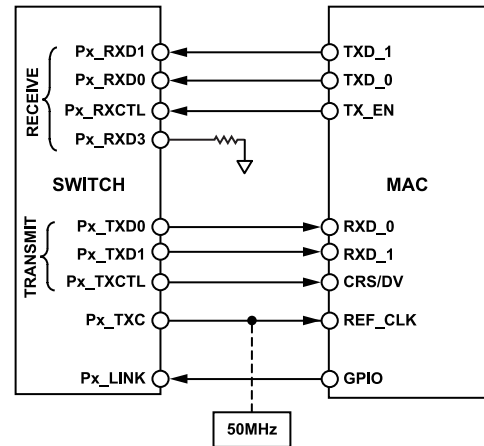


Figure 30. RMII to RMII MAC to MAC Connections

SGMII Modes

Port 1 through Port 4 support SGMII or serdes modes of operation.

SGMII Mode

The SGMII is a serial interface using two differential pairs to transfer information between the switch and PHY or host, using less pins and requiring less routing and traces compared to RMII and RGMII. SGMII transmits at 1.25 Gbps and requires external AC coupling with 0.1 μF capacitors. The SGMII is available on Port 1 to Port 4 of the ADIN6310. The interface is capable of 10 Mbps, 100 Mbps, and 1 Gbps full duplex communication with autonegotiation or forced modes. The SGMII supports autonegotiation and forced speeds for SGMII PHY link partners. When connected to another SGMII MAC link partner, autonegotiation must be disabled, and matching speeds must be configured on both sides.

INTERFACES




Figure 31. SGMII Signals

Serdes and Fiber Interface Mode

The SGMII pins are capable of functioning as SGMII or Fiber input and output pins on Port 1 to Port 4. The switch supports 1000BASE-KX for backplane applications. Figure 32 shows the switch connected to an SFP fiber transceiver. The interface supports 100BASE-FX and 1000BASE-SX and 1000BASE-LX.

When using SFPs, the fiber transceivers typically integrate the AC coupling capacitors; therefore, Figure 32 does not show external capacitors or other components required by the SFP. Refer to SFP multiple source agreement (MSA) for full details on the recommended SFP circuit. The transmit and receive paths are typically all that must be connected between the SFP and the switch port.

When a port is configured to use one of the SGMII modes, the Px_LINK pin must either be pulled to ground with a pull-down resistor, or optionally, it can be driven by the SFP loss of signal (LOS) for the port. The SFP LOS is an open-drain pin that requires a 4.7 kΩ to 10 kΩ pull-up resistor to 2 V to VDD3P3. Normal operation is indicated when LOS is driven low. When LOS is driven high, this indicates that the received optical power is less than the worst-case receiver sensitivity. If the switch VDDIO voltage is running at 1.8 V, level shifting may be required for the LOS to the Px_LINK input.

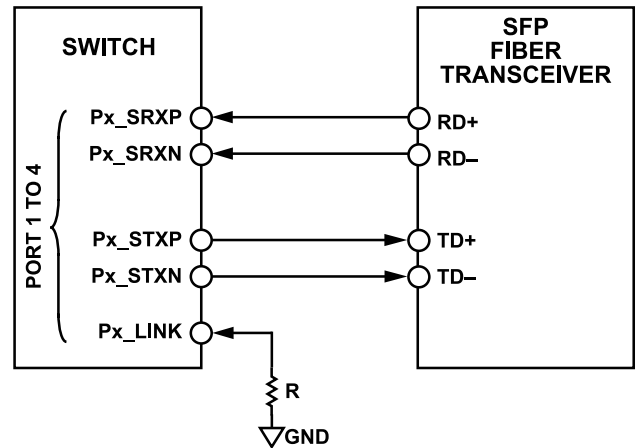
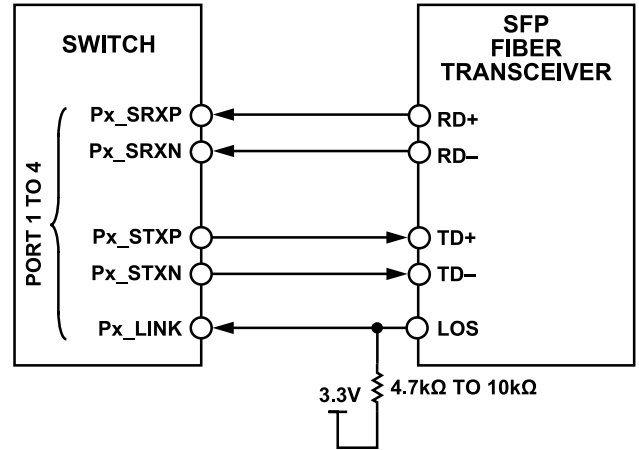


Figure 32. Connection to SFP Optical Transceiver

To understand what type of SFP module is connected, the host must access the SFP module directly because the switch does not control or interrogate the SFP.

HOST INTERFACE

The switch requires management from an external host and stack processor. The connection of the host to the switch is flexible and configured using hardware strapping pins that are sampled when the switch is brought out of reset. The switch then knows where the host resides and how to communicate. Run-time configuration to change the host interface is not supported.

The host communication port can be connected over the SPI or an Ethernet MAC interface (SGMII, RMII, or RGMII). The hardware strapping pins used are the four TIMEx pins, three of the SPI_SIOx pins, and the SPI_SS pin. The hardware strapping pins have internal pull-up and pull-downs to provide a default host interface of standard SPI. Alternatively, the user can override the default by adding external strapping resistors to select a different SPI or an Ethernet host interface. When the host interface is a MAC host, the interface can be any one of the switch ports, the voltage sensed on the SPI_SIOx pins determines which port. Port

0 is recommended as the host interface if the host is running a different input and output voltage rail to the supply used for the external physical layer devices connected to the other Ethernet ports, because the VDDIO supply rails can be partitioned, with Port 0 from the VDDIO_A rail and all other ports from the VDDIO_B rail.

There are a number of possible host interface options available as detailed in [Table 20](#).

HOST INTERFACE STRAPPING PINS

There are eight inputs and outputs used for configuration strapping listed in [Table 20](#). The voltage at the strapping pins is sampled on the rising edge of RESET_N after power on. The hardware strapping is used to let the packet assist engine know the interface location of the host and defines the primary host communication protocol as SPI or Ethernet.

Table 20. Boot Strapping Pin Configuration—Host Interface (SPI or Ethernet)

Function	SIO2 ^{1,2}	SIO1 ²	SIO0 ²	SPI_SS	TIMER3	TIMER2	TIMER1	TIMER0 ³	Notes
Default Internal	Pull-down	Pull-down	Pull-down	Pull-up	Pull-down	Pull-down	Pull-up	Pull-up	
RGMII 100 Mbps									
No Clock Delays ⁴	P2	P1	P0	ds ⁵	0	0	0	0	
Receive Clock Delay	P2	P1	P0	ds	0	0	0	1	
Transmit Clock Delay	P2	P1	P0	ds	0	0	1	0	
Receive and Transmit Clock Delays	P2	P1	P0	ds	0	1	0	1	
RGMII 1 Gbps									
No Clock Delays	P2	P1	P0	ds	1	0	0	0	
Receive Clock Delay	P2	P1	P0	ds	1	0	0	1	
Transmit Clock Delay	P2	P1	P0	ds	1	0	1	0	
Receive and Transmit Clock Delays	P2	P1	P0	ds	1	1	0	1	
RMII 10 Mbps									
RMII CLK (Internal)	P2	P1	P0	ds	0	1	0	0	Switch internal clock provided out to the PHY or host
RMII CLK (External)	P2	P1	P0	ds	1	1	0	0	External clock source provided to the PHY or host and switch
RMII 100 Mbps									
RMII CLK (Internal)	P2	P1	P0	ds	0	1	1	0	Switch internal clock provided out to the PHY or host
RMII CLK (External)	P2	P1	P0	ds	1	1	1	0	External clock source provided to the PHY or host and switch
SGMII									
100BASE-FX	P2	P1	P0	0	0	0	1	1	
SGMII 1000 Mbps	P2	P1	P0	0	0	1	1	1	
SGMII 100 Mbps	P2	P1	P0	0	1	0	1	1	
1000BASE-KX	P2	P1	P0	0	1	1	1	1	
SPI Enabled									
Single (Default)	ds	x	x	1	0	0	1	1	Default setting
Dual	ds	x	x	1	1	0	1	1	
Quad	x	x	x	1	0	1	1	1	Low drive strength
Quad	x	x	x	1	1	1	1	1	High drive strength

HOST INTERFACE

- ¹ P2 through P0 selects the port for SGMII, RMII, and RGMII host interface, see [Table 21](#).
- ² Don't care.
- ³ TIMER0, TIMER1, and SPI_SS have internal pull-ups. TIMER2, TIMER3 and SPI_SIO0 to SPI_SIO2 have internal pull-downs. The default configuration is a standard SPI when no external host strapping is used. To configure a different host interface, use the external pull-up and pull-down resistor values per [Table 22](#).
- ⁴ Ensure that the clock delays are configured correctly to interface with the host and PHY or MAC. Delay is added at only one stage in the transmit and receive paths, typically at the receiving side. For MAC interfaces that do not have the capability to add delay, the switch host strapping supports adding delays on both sides.
- ⁵ ds = Drive strength. Two drive strength options are available for the host interface (high or low). When ds = 1, this corresponds to the highest drive strength. Drive strength for other ports can be configured through the Software Driver API, which is available from the [ADIN6310](#) product page.

The default mode is the standard SPI interface based on internal pull-up and pull-down resistors. External pull-up and pull-down resistors are required to override the default values and are detailed in [Table 22](#).

After initial communication is established, software can establish the other means of communication for applications that use two interfaces between switch and host, one for control traffic (strapping interface) and one for data traffic.

Table 21. Boot Strapping Pin Configuration—Host Port Selection

Function and Host Port	Mnemonic		
	SPI_SIO2	SPI_SIO1	SPI_SIO0
SPI, Dual SPI and Quad SPI (No Ethernet Port Selected)	ds ¹	x	x
Port 0 RMII or RGMII	0	0	0
Port 1 RMII, RGMII, or SGMII	0	0	1
Port 2 RMII, RGMII, or SGMII	0	1	0
Port 3 RMII, RGMII, or SGMII	0	1	1
Port 4 RMII, RGMII, or SGMII	1	0	0
Port 5 RMII or RGMII	1	0	1
Reserved	1	1	0
Reserved	1	1	1

- ¹ ds = Drive strength. Two drive strength options are available for the host interface (high or low). When ds = 1, this corresponds to the highest drive strength. Drive strength for other ports can be configured through the `SES_SetGpio-DriveStrength()` API, which is available in the `SES_switch.h` file from the [ADIN6310](#) product page.

Because the TIMERx and SPI_SIOx pins are on the VDDIO_A power domain, external strapping resistors should be pulled to VDDIO_A.

Ensure that these pins are static and in the correct state during the hardware latch after reset. [Table 22](#) provides information on the external resistor values required for each pin to overcome the internal strapping.

Table 22. External Strapping Resistor Values

Mnemonic	Default	External Series Resistor Value
TIMER0	1	Pull-down <2.5 kΩ to latch a value 0
TIMER1	1	Pull-down <2.5 kΩ to latch a value 0
TIMER2	0	Pull-up <3.3 kΩ to latch a value 1
TIMER3	0	Pull-up <3.3 kΩ to latch a value 1
SPI_SS	1	Pull-down <400 Ω to latch a value 0

Table 22. External Strapping Resistor Values (Continued)

Mnemonic	Default	External Series Resistor Value
SPI_SIO0	0	Pull-up <10 kΩ to latch a value 1
SPI_SIO1	0	Pull-up <10 kΩ to latch a value 1
SPI_SIO2	0	Pull-up <10 kΩ to latch a value 1

TIMER INPUTS AND OUTPUTS

There are four timer pins (TIMER0 to TIMER3) that are dedicated to timer and interrupt functions, with TIMER0 providing the interrupt function for the switch SPI. The function of these pins is controlled from the packet assist engine. There are four TSN timer outputs, four legacy timer outputs, two IEEE 1588 outputs, and two input captures as detailed in [Table 23](#).

The timer pins can alternatively be configured in GPIO mode and use as GPIOs as required.

Table 23. Timer Functions

Timer	Function	GPIO Mode
TIMER0/INT	TSN timer, industrial Ethernet engine multiprotocol timer, interrupt	GPIO4
TIMER1	TSN timer, industrial Ethernet engine multiprotocol timer	GPIO5
TIMER2	TSN timer, industrial Ethernet engine multiprotocol timer, IEEE 1588 timer, timer input	GPIO6
TIMER3	TSN timer, industrial Ethernet engine multiprotocol timer, IEEE 1588 timer, timer input	GPIO7

GPIO

There are four dedicated GPIO pins available in the switch. Their function is configured by the packet assist engine.

HARDWARE RESET

The RESET_N is an active low input to the switch and resides on the VDDIO_A power supply domain. The RESET_N pin must be asserted low for a minimum of 100 μs after all power supplies are at their nominal values. After power on, when RESET_N is brought high and stays high, the external strapping pin states are latched.

During operation, to perform a hardware reset, the RESET_N pin must be driven low for a minimum 10 μs. When asserted low, all output pins are driven to a High-Z state.

HOST INTERFACE

CLOCKING

The switch requires a 25 MHz clock that can be an external crystal oscillator applied across the XTAL_IN and XTAL_OUT pins or an external clock applied to the XTAL_IN pin, as shown in Figure 33.

If using an external clock, XTAL_OUT can be left open. The crystal circuit requires a number of external components, which must have the following values populated: the feedback resistor (R_F) = 1 M Ω \pm 5%, drain resistor (R_D) = 2.1 k Ω \pm 5%, and C1 = C2 = 18 pF. Detailed specifications for suitable crystal oscillators are captured in Table 1.




Figure 33. External 25 MHz Crystal or Clock Connection

Clock Output

The CLK_OUT pin provides a 25 MHz clock generated from the 25 MHz input clock applied to the switch. It can drive a maximum C_L of 15 pF. This CLK_OUT signal is provided for the external PHY devices and requires external clock buffering. Figure 34 shows an example of a 6-channel clock buffer used in conjunction with the ADIN6310.




Figure 34. CLK_OUT Signal Buffered for External Distribution to the PHYs

POWER ARCHITECTURE

The switch has a number of power supply domains and requires a minimum of three external power supply rails. Two of supply rails (VDDIO_A and VDDIO_B) are dedicated to interfacing because the host interface may be on a different voltage domain to the external physical layer devices. The device typically requires four power supply rails to achieve optimum power consumption performance. However, it is possible to operate from a minimum of three power supplies, where the VDDIO_x domains are grouped and VDD3P3 = VDD_SGMII (if SGMII is used).

The following four power supply rails are available for the device:

- ▶ VDDCORE powers the digital logic circuitry.
- ▶ VDD3P3 powers the analog circuitry.
- ▶ VDDIO_A and VDDIO_B are supplied to the inputs and outputs involved in communication with the external host controller and PHY chips, respectively.
- ▶ VDD_SGMII is a dedicated supply for the SGMII. In applications where SGMII is not required, connect the VDD_SGMII pins directly to VSS_SGMII.

Figure 35 shows an overview of the power supply architecture.

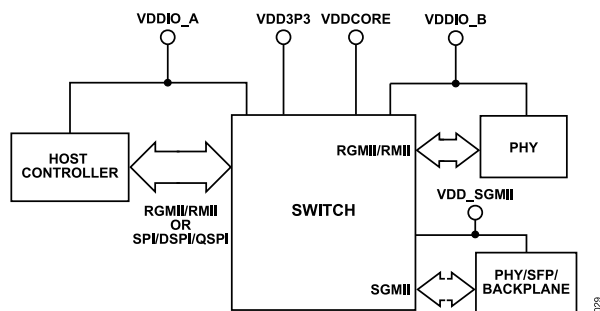


Figure 35. Power Supply Overview

POWER SUPPLIES AND SUPPLY DOMAINS

Table 24 shows the required supply rails and corresponding voltage ranges.

Table 24. Power Supplies

Power Supplies	Nominal Voltage Levels (V)
VDDIO_A	1.8, 2.5, 3.3
VDDIO_B	1.8, 2.5, 3.3
VDDCORE	1.1
VDD3P3	3.3
VDD_SGMII	3.3

The pin configuration shows that there are multiple balls for some of the supply rails (see the [Pin Configuration and Pin Descriptions](#) section). It is critical for proper operation of the device to ensure that each ball is directly connected to the appropriate power rail.

POWER-ON RESET (POR)

The device has on-chip power management to ensure the device powers up and functions reliably under normal operating conditions. The POR generation circuit operates on the main supply rails: VDD3P3, VDDIO_A, and VDDCORE. During power-up, POR remains asserted till these supplies reach stable operating ranges. If supply voltage levels on these supplies drops to less than the acceptable voltage threshold, then POR gets reasserted. The VDDIO_B and VDD_SGMII are not included in the POR scheme.

POWER SUPPLY SEQUENCING

There are no power sequence requirements for the switch; however, the preferred power sequence is to apply the VDDCORE supply last and to remove it first.

GROUNDING

There are three main ground groupings, VSS3P3, VSS_SGMII, and DGND on the ADIN6310. When designing hardware, a single ground plane (DGND) is recommended for the switch. Connect VSS_SGMII directly to the DGND plane. The remaining grounds must be star grounded together and connected to the DGND plane.

Table 25. Ground Mnemonic Grouping

Ground Mnemonic	Ground Domain	Connection
VSS3P3	VSS3P3	Star ground each VSS3P3 ball together
VSSCORE_DLL	VSSCORE_DLL	Star ground
VSSCORE_PLL, DGND, VSS3P3, VSSCORE_DLL, VSS_SGMII	DGND	Star ground together to the DGND plane

POWER CONSUMPTION AND THERMAL CONSIDERATIONS

Worst case power consumption conditions are all six ports operating at 1 Gbps. Careful consideration is required when choosing the VDDIO_B power supply rail to keep the die temperature within the operating temperature range. When operating at 1 Gbps on all six ports, use VDDIO_B at 1.8 V or 2.5 V. Operating at VDDIO_B = 3.3 V when all ports are at a 1 Gbps speed is not supported. Table 26 includes per port typical power parameters for all VDDIO supply voltages because various use cases may have the ports operating a mix of PHY technology and speeds.

Table 26. Power Consumption Per Port

Use Case	VDDIO			Unit
	1.8 V	2.5 V	3.3 V	
Base Power ¹	126.5	126.5	126.5	mW
1 Gbps RGMII	36.6	60.7	95.7	mW
100 Mbps RGMII	11	13	17.05	mW
10 Mbps RGMII	6.2	7.4	9.35	mW
1 Gbps SGMII	59.2	68.4	81.4	mW

POWER ARCHITECTURE

Table 26. Power Consumption Per Port (Continued)

Use Case	VDDIO			Unit
	1.8 V	2.5 V	3.3 V	
1000BASE-KX	65.8	75	88	mW
100BASE-FX	52	55.9	59.95	mW

¹ Base power is the background power consumption for the chip. Total Power Consumption = Base Power + Each Port Power.

POWER SUPPLY DECOUPLING

It is recommended to decouple each of the supply rails with the following capacitors detailed in [Table 27](#) to their relevant grounds. Place the smaller decoupling capacitor (0.1 μF) as close to the relevant balls as possible and ensure that the capacitor ground is routed directly into the plane.

Table 27. Recommended Decoupling

Supply Name	Capacitor Value
VDD3P3 to VSS3P3	4.7 μF per ball (0.1 μF)
VDD_SGMII to VSS_SGMII	4.7 μF per ball (0.1 μF)

Table 27. Recommended Decoupling (Continued)

Supply Name	Capacitor Value
VDDIO_A to DGND	4.7 μF per ball (0.1 μF)
VDDIO_B to DGND	4.7 μF per ball (0.1 μF)
VDDCORE to DGND	20 μF and 0.47 μF per ball (0.1 μF)
VDDCORE_PLL to VSSCORE_PLL	0.1 μF and 0.01 μF
VDDCORE_DLL to VSSCORE_DLL	Per ball (0.1 μF)

When placing the decoupling capacitors, prioritize the locations detailed in [Table 28](#) with capacitors connected across these nets.

Table 28. ADIN6310 Layout Prioritization for Bypass Capacitors

Supply	Ground	Ball Number
VDDCORE_DLL	VSSCORE_DLL	N10 to M10
	VSSCORE_DLL	D10 to E10
VDDCORE_PLL	VSSCORE_PLL	G5 to G3
VDD3P3	VSS3P3	H5 to H3
	VSS3P3	J5 to J3
	VSS3P3	K5 to K3

TIME SYNCHRONIZATION

PRECISION TIME PROTOCOL (PTP)

The purpose of time synchronization is to provide all devices on the network with an accurate, reliable easy to use time by using various profiles of IEEE 1588. The switch can support running a PTP or generalized PTP (gPTP) stack on the packet assist engine with a choice of three different time synchronization profiles, namely IEEE 802.1AS 2020, IEEE 1588-2019 (default profile), or IEEE C37.238-2017 energy profile.

The PTP stack is disabled by default and requires configuration by the host to enable the function.

TIME DOMAINS

The switch supports multiple time domains that allow a working clock to coordinate precise coordination among the elements of a machine or cell and system-wide clocks that can be used to correlate critical events, such as alarms and errors. Different synchronization approaches can be coordinated by the switch, bridging time between, for example, one set of ports operating with IEEE 802.1AS-2020 and another port connected to a network synchronized with the IEEE 1588-2019 default profile.

IEEE 802.1AS 2020

The switch supports running the IEEE 802.1AS 2020 stack on the packet assist engine with two step sync messaging (one step in future). The PTP stack supports running IEEE 802.1AS 2020's best timeTransmitter clock algorithm (BTCA) to determine the highest quality clock in the network. Alternatively, the external port configuration allows ports to have their roles configured manually.

The switch supports the IEEE 802.1AS 2020 grandmaster or master capability on all PTP ports and supports PTP end instance and PTP relay instance.

Configuration is done using the **ieee802-dot1as-ptp.yang** module, which includes the following:

- ▶ Support for all ports as PTP ports for relay instances
- ▶ Peer to peer and common peer delay mechanism for path delay
- ▶ External port configuration
- ▶ Two step

- ▶ One step capability (future software update)
- ▶ Application interfaces per Clause 9

IEEE 1588-2019 DEFAULT PROFILE

The switch can support running a PTP stack configured for the default profile per IEEE Standard 1588-2019, Annex I. The PTP stack runs on the switch packet assist engine and supports the following:

- ▶ I.3 delay request-response default profile
- ▶ I.4 peer to peer default profile (future software update)
- ▶ Ordinary clocks and boundary clocks (Section 9)
- ▶ End to end transparent clocks (Section 10.2)
- ▶ Peer to peer transparent clocks (Section 10.3)
- ▶ Multiple domains and instances (maximum of 2)
- ▶ L2 and L3 transport
- ▶ Two step PTP messaging (Section 7.5.2.5)
- ▶ One step PTP messaging (Section 7.5.2.5 and future software update)
- ▶ VLAN tagged PTP frames (not supported with VLAN trunk or access port configuration)
- ▶ Path trace mechanism, path trace TLV (Section 16.2)
- ▶ Data sets (Section 8): defaultDS, currentDS, parentDS, time-PropertiesDS, pathTraceDS, portDS, descriptionPortDS, externalPortConfigurationPortDS, commonServicesPortDs, cmlDsDefaultDs, cmlDsLinkPortDs, cmlDsLinkPortStatisticsDs, and cmlDsAsymmetryMeasurementModeDs
- ▶ Transparent clock data sets (Section 8.3 not supported)
- ▶ Backward compatibility with devices supporting Version 1588-2008
- ▶ Different profile on different instances (1588-2019 on one instance, and IEEE 802.1AS 2020 profile on another)
- ▶ Configuration is done using the **ieee1588-ptp-tt.yang** module

IEEE C37.238.2017 ENERGY PROFILE

The IEEE C37.238.2017 energy profile is used for power system applications. The switch supports configuration of the time synchronization capability to support this profile.

TSN FUNCTIONALITY

The switch supports IEEE 802.1Q time sensitive networking (TSN) bridging with the following features:

- ▶ Scheduled traffic (Qbv)
- ▶ Frame preemption (Qbu)
- ▶ Per stream filtering and policing (Qci)
- ▶ Cyclic queuing and forwarding (Qch)
- ▶ Forwarding and queuing enhancements (Qav)

SCHEDULED TRAFFIC (QBV)

Scheduled traffic is a means to provide a time protected channel to ensure priority traffic can egress on time without delays due to other traffic.

Scheduled traffic takes advantage of the time capability added to the Ethernet from 802.1AS. By using time, a time-protected channel can be used to ensure only one type of traffic is on the network at a time. The reason for this is to avoid interfering traffic. Because the Ethernet was originally specified as best-effort, certain traffic could interfere with other traffic even if that traffic was given a higher priority. Therefore, in addition to priority, these traffic classes are scheduled in time to ensure these messages are not interfered with.

Streams are placed into queues until the scheduled time arrives, and then the streams in that queue are allowed to pass onto the network until the time window for that queue closes. Because everyone on the network has an accurate time base, the schedule windows do not overlap, and therefore, the streams do not interfere with each other.

The switch supports eight queues per port with a 256-entry gate control list and seamless switchover when the gate control list is updated. The traffic classes are based on VLAN priority and can be mapped to the eight queues as required. The user can configure the network cycle time individually for each port with control of the cycle time extension and base time.

The hardware timer pins (TIMER0 to TIMER3) can also have a network cycle programmed. This feature can be used to provide signals that synchronization application hardware and software with other elements on the network.

The configuration of scheduled traffic transmit gates is based on the IEEE 802.1Q scheduled traffic YANG model. The APIs to configure Qbv are part of the `SES_scheduled_traffic.h` header file, which is available as part of the software drivers from the [ADIN6310](#) product page.

Strict Priority Scheduling

Each port supports eight queues and operates with strict priority, where the highest priority queued traffic transmits first.

Guard Bands

When using a schedule to guarantee the transmission time, it is important to guarantee that the transmitter is idle when a higher

priority gate is getting ready to open. If the transmitter is not idle, it is possible for a frame of lower priority to potentially block, or at least delay, the start of the transmission of the higher priority traffic. To guarantee that the transmitter is idle and can start immediately sending higher priority traffic when its gate is opened, a guard band is inserted ahead of the gate open time. While active, the guard band inhibits the transmitter from fetching the next frames. To be effective, the guard band must be long enough so that the longest possible message that can be queued transmits completely before the next gate open event occurs. The size of that longest frame that can be queued is defined in the parameter, maximum service data unit (**maxSDU**), in bytes. The switch supports per port and per queue maximum SDU limits and can be set and read using the `SES_QbvSetQueueMaxSduTable()` API and `SES_QbvGetQueueMaxSduTable()` API in `SES_scheduled_traffic.h` header file, which is available as part of the software drivers from the [ADIN6310](#) product page. The hardware default for per queue maximum SDU frames is 10,000 bytes.

When guard bands are enabled, the per queue, per port maximum SDU settings are used in conjunction with the established link speed to size the guard band duration. The guard band is applied for both store and forward mode and cut-through forwarding mode to protect the next gate open. Default switching behavior for all queues is to cut traffic through if conditions allow, which is configurable at the API level. The **QueueMax** SDU setting does not apply to cut-through traffic because these frames started forwarding before the hardware knows the size of the frame. However, understanding the size of the frames and sizing the **QueueMax** SDU settings accordingly protects the next gate open.

When scheduled traffic and frame preemption are used in combination, the device supports hold and release, which allows an explicit guard band to be implemented around a protected transmission window, but one that is smaller in duration than normally needed. This hold and release is used for configurations where the preemptable traffic gate is scheduled to be open at all times and could have the potential to delay the start of the schedule traffic. Using hold and release allows the express traffic protected window to be completely protected from interference while increasing bandwidth available to the preemptable traffic.

FRAME PREEMPTION (QBU)

Frame preemption can be used to further ensure high priority traffic arrives at a destination with a fixed latency by interrupting transmission of a frame designated preemptable, transmitting the high priority (express) traffic, then resuming transmission of the preemptable frame. The receiver node regenerates the preemptable frame from the fragments transmitted between bursts of priority frames. Without frame preemption, the time taken for the highest priority traffic to transit a network can vary from a minimum of the sum of the bridge and cable delays in the path to a maximum of the previous sum plus a maximum sized frame delay at each hop (for example, at 1 Gbit with a 1528 byte maximum transmission unit (MTU), this is a variance of up to more than 12 μ s per hop).

TSN FUNCTIONALITY

With preemption, this variance is reduced from 1528 byte interfering frames to 128 byte interfering fragments, or 12+ μ s to less than 1.2 μ s per hop, or an arrival time jitter reduction of 90%.

Another use for frame preemption is when frames are scheduled (scheduled traffic is enabled). Typically, guard bands are used to prevent large frames from interfering with the prompt start of a time protected channel. When there are large frames in the network, guard banding can be an inefficient use of network bandwidth because the guard band size needs to cater to the maximum frame size. Because frame preemption can fragment large frames into smaller pieces, it is possible to have much smaller guard bands and optimize network bandwidth. The preemptable traffic can start egressing before the start of the smaller guard band and get interrupted to ensure the critical traffic transits on time, with remaining fragments being allowed to finish their transmissions after the time window passes.

Traffic—Standard and Express Frames

Standard or express Ethernet frames have a preamble that ends with SFD or SMD-E, which denotes the beginning of frame data. The frame ends with a 32-bit CRC.

Traffic—Preemptable Ethernet Frames

Frames to which preemption have been applied have a start fragment, where the preamble ends with an SMD-S denoting the start of frame. The frame ends with an mCRC. If the mCRC is the same as the CRC, there is only one fragment. If the mCRC is the CRC with 16 bits XORd, this fragment does not finish the frame, and additional fragments will follow, which tells the receiving MAC that the frame has been fragmented, and it can expect a continuation fragment to follow.

The switch receiving MAC always accepts frames with SMD-S and SMD-C, frames to which preemption has been applied, because the receive side always has preemption enabled.

The switch transmitting MAC only sends frames with an SMD-S and SMD-C, frames to which preemption has been applied, only after it has been established that the link partner supports preemption, and the switch has been instructed to enable preemption for this link.

Enabling Frame Preemption with Verification

The process of discovering support for preemption is done through the exchange of additional Ethernet capability type length value (TLV). When preemption is enabled in the switch, the TLVs advertising preemption capability and status can be observed in the Link Layer Discovery Protocol (LLDP) frames originating from that port. Initially, the switch sends an LLDP frame with the TLV, indicating support of preemption; however, that preemption is not enabled or active.

If the switch does not receive an LLDP frame from the link partner with a TLV indicating support of preemption, the switch does not enable or activate preemption with the link partner.

If the switch receives an LLDP frame from the link partner with a TLV indicating support of preemption, then if verification is enabled, the switch changes the LLDP TLV to indicate supported and enabled but not active, and sends out a verify mPacket (SMD-V = 0x07). If the link partner sends a response mPacket (SMD-R = 0x19), the switch changes the LLDP TLV to indicate supported, enabled, and active, and starts to transmit frames from the selected queues with SMD-S and/or SMD-C.

In the event there was no response frame received, after a timeout, the switch sends another verify frame, after three fails (no response from link partner), it stops sending verifies and changes its LLDP TLV to indicate support only.

The verification process previously described is the expected normal operation and how the switch is intended to be configured.

There are debug options that allow users to force preemption to enable, skipping the verification process. In this case, the switch changes its LLDP TLV to indicate support, enabled, and active, and sends preemptable traffic from the defined queues. Skipping using verification should only be done as a debug mode because if the link partner does not support preemption, the preempted frames may get dropped as the receiving side may not understand the SMD values.

The switch supports programmable minimum nonfinal fragment sizes of 64 bytes, 128 bytes, 192 bytes, and 256 bytes. The APIs to configure frame preemption are included in **SES_preemption.h** header file, which is available as part of the software drivers from the [ADIN6310](#) product page.

Preemption Statistics

When using frame preemption, the transmit and receive MACs capture a number of statistics related to operation. All registers clear on read, and when full, stay at the maximum value and do not overflow.

Table 29. Per Port Frame Preemption Statistics (Transmit and Receive)

Parameter	Description
Frame Assembly Error Count	Number of frame reassembled with errors
Frame SMD Error Count	Number of frames with SMD errors
Frame Assembly OK Count	Number of frames reassembled successfully
Fragment Count Receive	Number of valid continuation fragments received
Fragment Count Transmit	Number of fragments transmitted
Hold Count	Number of times the hold capability was asserted

TSN FUNCTIONALITY

PER STREAM FILTERING AND POLICING (QCI)

Per stream filtering and policing, as defined by the IEEE 802.1Qci standard provides filtering and policing for a stream of ingress traffic.

The purpose of this feature is to prevent traffic overload conditions from affecting the receiving node. which is done by filtering traffic on a per stream basis by providing a stream filter for each stream. The stream filter serves to enforce a contract between the talker and listener.

The filtering and policing capabilities apply on a stream basis to the receive path. The following stream filters available are:

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

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

Stream Filter

The stream filter can be blocked or unblocked. The filter can be configured based on the incoming frame size. If the frame is over a specified size, it is discarded. The stream filter can be associated with a stream gate and flow meter. There are 32 size-based filters per port.

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, perform the required lookups, and handle as required. Alternatively, if the stream arrives when the gate is closed, discard the frame.

The stream gate can change the internal priority vector (IPV) of a frame.

The stream gate can allow ports to support a defined amount of traffic in a certain amount of time, octets per unit time based on the port timer control unit (TCU).

Flow Meter

The flow meter 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 are in each bucket (commit or excess). If there is enough tokens in either bucket, the frame can proceed.

CYCLIC QUEUING AND FORWARDING (QCH)

The IEEE 802.1Qch feature (which is a future software update) is a combination of two existing features, using the stream gate and scheduled traffic together.

The stream gate is used with the stream filter to change the internal priority vector of the frame, so that the gate is essentially directing what queues traffic is going in at what time, where one transmit queue is getting filled while the other queue is transmitting because it uses scheduled traffic to transmit the data.

The Qch is a popular approach for audio and video applications.




Figure 36. Cyclique Queuing and Forwarding

FORWARDING AND QUEUING ENHANCEMENTS (QAV)

The forwarding and queuing enhancements (IEEE 802.1Qav) function is a credit-based shaper (CBS), which is a future software update, and its purpose is to ensure that lower priority traffic gets some bandwidth.

The transmit function is based on traffic classes. The CBS runs on high priority queues, and the switch has two CBSs per port.

A port can only transmit if the credit is ≥ 0 . If the queue is empty, the credit is 0. When a frame shows up in the queue but traffic is already being transmitted, the frame cannot start transmission. The frame is blocked until the current frame finishes transmission and during that the time credit increases.




Figure 37. Credit-Based Shaper

SENDLIST

The switch supports the ability to develop a generic SendList function that can meet the requirements of IEC 60802 and the needs of devices operating Layer 2 data protocols like PROFINET, CC-LINK, and OPC-UA FLC.

The primary motivation is to support devices that rely on the SPI of the switch for transmission of frames and are, therefore, unable to have precise control of the timing. A secondary goal is to support transmission of a subset of traffic using this facility for devices connected over a MAC interface but lacking the timing, queueing, and direct memory access (DMA) facilities to fully implement a SendList.

A SendList is a periodic transmission of a number of sublists. Each sublist contains a number of Ethernet frames to be sent at a specific offset in the SendList period. As the SendList is periodic, if the frame data is not updated before a sublist begins the next cycle transmission, the last frame data is sent. The SendList function relies on dedicating a port in the switch to the SendList function. This port will be configured into loopback mode internally. In the 6-port device, the loopback port can be any of the six ports (just not the host interface port). When using the SendList functionality, the loopback port is not available for regular traffic, and as such, if the SendList is always active, this port does not need external circuitry (for example, PHY).

Once the port is configured in loopback mode as part of the SendList configuration, it is not available for regular ingress traffic and must be dedicated to the SendList function. The loopback configuration connects the transmit MAC with the receive MAC internally. Only one SendList instance can be supported in a device.

When creating a SendList, the following steps are the intended order of operation:

1. Create a SendList with the number of sublists and desired period.
2. Create sublists with the desired start offset from the start of the SendList period.
3. Register the SendList frames that will be added to the sublists.

4. Add a static table entry for each registered frame with the desired port map.
5. Add the registered frames to the appropriate sublists.
6. Configure the schedule for the loopback port.
7. Send registered frame data with the transmit enable parameter set for each registered frame.
8. Start the SendList with the desired start time.
9. To stop transmitting a specific frame, use the enable or disable parameter of the update registered frame or call the stop registered frame function with the frame ID of the frame to stop transmitting.
10. To stop transmitting a SendList, disable or stop all registered frames.
11. To change or update a SendList program, create the new SendList and sublists. Then, add registered frames to the sublists and start the new SendList.

The normal assist engine transmits are allowed between the end of the sublist and the start of the next. The SendList halts the assist engine transmit function the maximum frame time before the start of a sublist. The SendList starts transmitting the largest registered frame time before the production time to ensure the frame is queued before the appropriate gate opens. The Ethernet host transmits are prioritized based on the bridge configuration and priority.

SENDLIST EXAMPLE

The [ADIN6310 Software Driver User Guide](#) includes example code showing how to configure the SendList feature with four SendLists like those shown in [Figure 38](#). Sublist 0 has three frame entries, while Sublist 1 to Sublist 3 have just two frame entries.




Figure 38. SendList Example with Four Sublists

REDUNDANCY

Redundancy is used to guarantee delivery of traffic streams in the event of a break in the path. There are different methods to this, namely FRER (IEEE 802.1CB), HSR, and PRP. All these methods of redundancy provide zero recovery time or seamless redundancy in the event of a failure.

These features provide redundancy by sending two copies of a message along a maximally disjoint path to ensure delivery. The redundant paths minimize packet loss due to link or device failures or congestion. Duplicate frames get discarded on receipt.

HSR and PRP are protocols defined by the IEC 62439-3 standard and are used in applications to provide lightweight redundancy for reliable delivery of traffic, with zero recovery times in the event of link failure, device failure, or stream congestion.

The switch supports hardware-based 802.1CB, HSR and PRP implementations, thereby offloading the host the task of duplicate generation, identification, and removal.

IEEE 801.2CB is the flavor of redundancy used in TSN applications, whereas typically, HSR or PRP are used in energy applications such as smart grid substations.

The switch also supports MRP, which is a redundancy protocol for ring topologies.

HIGH-AVAILABILITY SEAMLESS REDUNDANCY (HSR)

HSR is used typically in ring applications to increase reliability of a communication channel by providing duplicate frames that travel across separate paths within the network and to provide for a seamless switchover time for a single network failure without placing any requirement on the application. HSR implements redundancy through the use of doubly attached nodes obeying HSR (DANH).

The switch provides a hardware-based approach aligned with the IEC standard and supports operation as a DANH or redundancy box (RedBox) where the hardware manages the duplication of frames, insertion of the HSR tag, consumption of the first frame, discard of the duplicate frame, and removal of the HSR tag. Because the HSR tag is just after the VLAN tag, HSR can operate in cut-through mode, resulting in shorter residence time bridge to bridge.




Figure 39. HSR Tag Location in Frame

HSR topologies are rings where all devices in such a network must be capable of interpreting HSR tags. A DANH has two HSR ports,

and a source DANH sends the same frame over each HSR port by duplication.

NonHSR devices can be connected to a HSR ring using a RedBox.

HSR can support more than one ring, where rings can be connected together using quad boxes, these are versions of RedBoxes that can connect to other RedBoxes in HSR to HSR mode for a QuadBox or connect to a PRP network when connecting between HSR and PRP.




Figure 40. Overview of Devices in HSR

Frame Forwarding

Frames sent into the ring from the DANH or singly attached node (SAN) devices get duplicated and forwarded out both ring ports. On a ring port, if an incoming frame has no HSR header, it is either consumed locally or discarded (not forwarded), such as a peer delay request message. If a HSR frame is sent out a port, the same frame (that is, same source MAC address and sequence count) is not sent out the same port again.

A unicast frame targeted at a particular device is not forwarded out the other ring port of that device. It is consumed by the listener, and the duplicate is discarded.

Multicast or broadcast frames may be of interest to multiple nodes in the ring; therefore, they get forwarded by each node around the ring. When the frame arrives at the DANH where it originated, it is not forwarded to the other HSR port or host, thereby breaking the ring and satisfying the requirement that a node not forward a frame that it injected into the ring.

If a frame is forwarded cut-through and an error is detected, the following frame with the same signature is forwarded (erroneous frames are not counted in duplicate discard).

REDUNDANCY

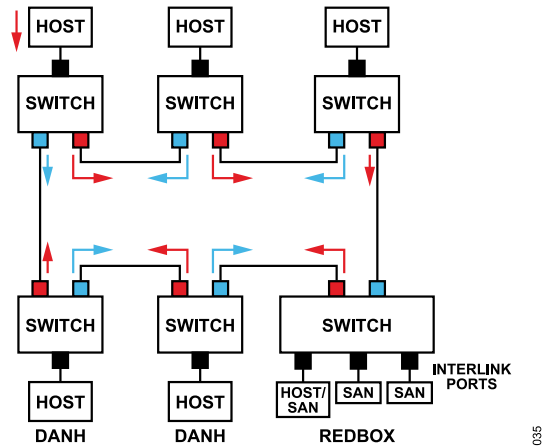


Figure 41. Simplified HSR Ring, Showing Multicast Frame Path

Forwarding Table Considerations for HSR

When in HSR mode, normal learning does not take place. Only learning for HSR is supported, where the switch learns the combination of the source MAC and sequence number. The forwarding table is shared across all ports and supports 2048 entries. Some of these entries are consumed by static entries such as PTP and LLDP, leaving a balance for HSR streams and other entries.

Size of HSR Ring

The number of devices supported by the switch depends on the traffic in the ring, how many frames, and aging time, which is driven by the depth of the discard buffer that uses the forwarding table entries for learning and discarding.

The residence time from bridge to bridge is <600 ns across the switch (including the HSR tag), but it does not include PHY transmit and receive latency. When using the [ADIN1300](#) Gigabit PHY, the transmit and receive latency in RGMII mode is on order of <300 ns for the PHY.

ADIN6310 as a DANH or RedBox

The ADIN6310 can be configured as a DANH or a RedBox device with up to eight SAN devices connected. Only one instance of HSR is supported in the ADIN6310. When the switch is in HSR mode, because normal learning is disabled, no learning is available for the nonHSR ports, and while static entries can be installed in the table, this consumes entries available for HSR learning.

Host Interface When in HSR Mode

When in HSR mode, the switch host interface can be connected over either the SPI or Ethernet MAC. In either case, the host interface can participate in the HSR network.

HSR PORT MODES

The switch supports all the port modes defined in the IEC standard. The default mode is Mode H.

The following is a list of the port modes:

- ▶ **Mode H (default mode):** HSR-tagged forwarding. The DANH inserts HSR tags on frames originating with the host and forwards ring traffic. Ring traffic that originated from the host is not forwarded. Unicast frames addressed to the host are not forwarded on the ring. Frames that have already been sent out a port are not sent a second time.
- ▶ **Mode N:** no forwarding. This mode operates like the Mode H device except that no traffic is forwarded between the two ring ports. This mode is used for special network topologies.
- ▶ **Mode T:** transparent forwarding. HSR tags are not injected into frames from the host before forwarding, and tags are stripped from all frames that traverse port to port.
- ▶ **Mode M:** mixed forwarding. HSR-tagged frames are handled like Mode H, nonHSR-tagged frames received on the ring ports are forwarded according to standard switching rules. Frames from the host either have HSR tags added and are duplicated, or these frames are treated like traffic on a standard switch according to the destination address and/or VLAN.
- ▶ **Mode U:** unicast forwarding. This node operates like in Mode H except that all unicast frames targeting the host are also duplicated out the other ring port.
- ▶ **Mode X:** no sending on counter duplicate. In this mode, the node behaves like 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.

HSR Supervision Frames

The switch sends supervision frames into both ports of the HSR network, they are generated periodically with or without a VLAN tag every **LifeCheckInterval** of 2 seconds. When configured as a RedBox, the switch sends corresponding supervision frames for the RedBox and also transmits supervision frames for any virtual DANHs/SANs connected to the RedBox at the same interval.

HSR STATISTICS

The switch captures a number of HSR-related statistics and error counters associated with each HSR port. The statistics also return how many nodes are in the network. The following HSR statistics are captured:

- ▶ **Rx Count** reflects the number of frames received by Port A or Port B that have HSR tag added.
- ▶ **Tx Count** reflects the number of frames transmitted by Port A or Port B that have HSR tag added.
- ▶ **Error Count** specifies the number of frames with errors received on the link redundancy entity (LRE) Port A or Port B.

REDUNDANCY

- ▶ **Duplicate Count** specifies the number of entries in the duplicate detection mechanism on Port A or Port B for which one single duplicate was received.
- ▶ **Multi Count** specifies the number of entries in the duplicate detection mechanism on Port A or Port B for which more than one duplicate was received.
- ▶ **Unique Count** specifies the number of entries in the duplicate detection mechanism on Port A or Port B for which no duplicate was received.
- ▶ **Node Count** returns the number of nodes detected in the system.

Nodes Table

The switch maintains a nodes table to track DANHs and SANs in the HSR network. The node table is populated based on incoming traffic and HSR supervision frames observed on HSR ports. The packet assist engine automatically inserts or deletes entries based on the node table logic. The nodes table records duplicate discard or duplicate accept based upon received HSR supervision frames for a node. This table records the last time a frame was received from a particular node, and it refreshes the nodes table every 60 seconds. If no frames have been received after this time, the node entry is removed. The user can interrogate the nodes table through the `SES_GetLreNodesEntry()` API, which is available as part of the software drivers on the [ADIN6310](#) product page. The nodes table can accommodate 1024 entries maximum.

Proxy Nodes Table

The proxy nodes table is a list of the detected SANs that are connected to the RedBox, and the last time they were seen. The proxy nodes table learns the SAN and virtual DANH MAC based on ingressing traffic on an interlink port. The user can interrogate the proxy nodes table using the `SES_GetLreProxyNodeEntry()` API, which is available as part of the software drivers on the [ADIN6310](#) product page. 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.

PARALLEL REDUNDANCY PROTOCOL (PRP)

Like HSR, PRP is used in redundant applications to increase reliability of a communication channel by providing duplicate frames traveling across separate networks and lossless communications in the event of a single network failure. PRP places no extra requirements on the network devices; however, it does require usage of a duplicate network.

PRP implements redundancy through the use of doubly attached nodes obeying PRP (DANP).

Again, only one instance of PRP can be configured in the 6-port ADIN6310 switch, the selection of ports that are part of the PRP

network is flexible and can be defined by the user during initial configuration.

A DANP is attached to two independent LANs of similar topology, named LAN A and LAN B, which operate in parallel. A source DANP takes the nonPRP frame from a SAN and duplicates the frame for egress from the two ports connected to LAN A and LAN B. These frames have a PRP redundancy control trailer (RCT or PRP tag) attached at the end of the frame. These duplicate frames are identical, except for the PRP tag (LAN ID) and FCS. A destination DANP receives the duplicate frames from both LANs, within a certain time, consumes the first frame to arrive, and discards the duplicate. The DANP can also remove the PRP tag upon egress.




Figure 42. Simplified PRP Network

Any switches in the LAN path do not care about the PRP tag because it simply looks like part of the payload. All traffic from a DANP to a DANP is accomplished with PRP; traffic from a DANP to a SAN does not require PRP.

The two LANs are identical in protocol at the MAC-LLC level, but they can differ in performance and topology. Transmission delays may also be different. The two LANs have no connection between them and are assumed to be fail independent.

The switch provides a hardware-based approach for PRP aligned with the IEC standard and supports operation as a DANP where the hardware manages the duplication of frames, insertion of the PRP tag, consumption of the first frame, discard of the duplicate frame, and removal of the PRP tag.

Because the PRP trailer is at the end of the frame, just before the FCS, PRP operates in store and forward mode in the DANP and RedBox but switches in the network can be in cut-through mode.




Figure 43. PRP Tag—Trailer at End of Frame

Host Interface When in PRP Mode

When in PRP mode, the switch host interface can be connected over either the SPI or Ethernet MAC. In either case, the host interface can participate in the PRP network.

REDUNDANCY

Maximum Reside Time

The maximum reside time is the PRP learning age out time for duplicate frames. If the duplicate has not arrived within the maximum reside time, the duplicate gets forwarded. In practice, when running at 10 Mbps or 100 Mbps speeds, the full range of maximum reside time is possible (to a maximum of 400 ms). When operating with a higher data rate over a Gigabit link, smaller maximum reside times are used to avoid consuming all of the forwarding table entries, which would result in duplicates not getting discarded. The switch forwarding table can support up to 2k entries. For the duplicate discard algorithm, the table stores the PRP sequence number and source MAC address for the programmed maximum reside time. Consider a 1532 byte frame, including Ethernet, IFG and PRP overhead, which takes 12.3 μ s to transit at Gigabit speeds, and the tables 2000 entries get utilized within 24.5 ms. A default setting of 10 ms for the maximum reside time is best suited when operating at Gigabit speeds.

The host can subscribe to the **DynamicTableLimitExceed** event to monitor the usage of the forwarding table and be alerted if the table is close to full.

PRP Supervision Frames

The switch sends supervision frames into the PRP network. These frames are generated periodically with or without a VLAN tag every **LifeCheckInterval** of 2 seconds. When configured as a RedBox, the switch sends corresponding supervision frames for RedBox and also transmits supervision frames for any virtual DANPs and SANs connected to the RedBox at the same interval.

PRP STATISTICS

The switch captures a number of PRP-related 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 following HSR statistics are captured:

- ▶ **Rx Count** reflects the number of frames received by Port A or Port B that have PRP RCT trailers added.
- ▶ **Tx Count** reflects the number of frames transmitted by Port A or Port B that have PRP RCT trailers added.
- ▶ **Error Count** specifies the number of frames with errors received on the LRE Port A or Port B.
- ▶ **Wrong LAN error count** specifies the number of frames with the wrong LAN identifiers received on LRE Port A or Port B.
- ▶ **Duplicate Count** specifies the number of entries in the duplicate detection mechanism on Port A or Port B for which one single duplicate was received.
- ▶ **Multi Count** specifies the number of entries in the duplicate detection mechanism on Port A or Port B for which more than one duplicate was received.
- ▶ **Unique Count** specifies the number of entries in the duplicate detection mechanism on Port A or Port B for which no duplicate was received.

- ▶ **Node Count** returns the number of nodes detected in the system.

Nodes Table

The switch maintains a nodes table to track DANPs and SANs in the PRP network. The node table is populated based on incoming traffic and PRP supervision frames on PRP ports. The packet assist engine automatically inserts or deletes entries based on the node table logic. The nodes table records duplicate discard or duplicate accept based upon received PRP supervision frames for a node. This table records the last time a frame was received from a particular node, and it refreshes the nodes table every 60 seconds. If no frames have been received after this time, the node entry is removed. The user can interrogate the nodes table through driver APIs in the software driver APIs, which are available as part of the software drivers on the [ADIN6310](#) product page. The nodes table can accommodate 1024 entries maximum.

Proxy Nodes Table

The proxy nodes table is a list of the detected SANs that are connected to the RedBox and the last time they were seen. The proxy nodes table learns the SAN and virtual DANP MAC based on ingress traffic on an interlink port. The user can interrogate the proxy nodes table. 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 PRP RedBox is 8.

Duplicate Discard for Wrong LAN ID

If a frame is received on a PRP port with a valid PRP RCT, but wrong LAN ID (ID 0xA on Port B or ID 0xB on Port A), in the PRP DANP and RedBox scenarios, the switch incorrectly performs a duplicate discard and strips the PRP trailer. As a result, the host DANP or SAN receives only a single copy of the frame with the trailer removed rather than receiving two copies as expected, which should not cause issue in the application layer because typically the trailer is ignored, and the duplicate frames are discarded in the corresponding stack. The error counters are incremented correctly which indicates to the host that the LANs are incorrectly connected.

The duplicate discard for wrong LAN ID only applies for the PRP DANP or PRP RedBox use cases. This behavior does not impact HSR and HSR RedBox use case.

PRP Traffic With the Wrong LSDU Size in the PRP Tag

If a peer device transmits PRP frames with the wrong link service data unit (LSDU) size in the PRP tag, the PRP DANP or PRP RedBox device drops those frames to the host processor and locally attached SANs as long as those frames arrive within the age out time that is, **lreDupListResideMaxTime**.

REDUNDANCY

HYBRID HSR AND PRP NETWORKS

The switch can support a hybrid network configuration where there is a combination of HSR and PRP. The networks are bridged by a Redbox in PRP and HSR (future software update) mode, with the switch supporting data transfer from PRP to HSR to PRP. The switch manages the identification and removal of the HSR tag and insertion of the PRP tag and similar in other direction. The sequence number is retained in either direction. Note that there is a fixed relationship between the Path ID of the HSR tag and the LAN ID of the PRP tag.

FRAME REPLICATION AND ELIMINATION FOR RELIABILITY, 802.1CB

Frame replication and elimination for reliability (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 stream 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 packet. The network ensures that no matter which of the paths the stream takes, it arrives where and when it was supposed to. The switch includes the sequence generation and recovery algorithms for FRER support.




Figure 44. 802.1CB Replication and Elimination

The redundancy tag (R-tag) includes a 2 byte sequence number that gets incremented for each frame received by the generator. The sequence number is used in the listener recovery function to identify and to eliminate duplicates of a frame. The R-tag also includes an EtherType (0xF1C1).



Figure 45. IEEE 802.1CB R-Tag

Stream Identification

The stream identification function is used to determine which ports the streams go to. The two types of stream identification are 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 30 details the different types of stream identification available, the parameters these streams examine in a packet, and the parameters that these streams overwrite.

The two available sequence recovery algorithms are match and vector algorithm. The switch can support both recovery algorithms. The match algorithm is more straightforward. The match algorithm accepts the first packet received as valid. The subsequent packets are then evaluated based on their match status with the last accepted packet. If a packet matches the last accepted one, it is discarded. If a packet does not match, it is accepted. Additionally, each accepted packet resets a timer. If this timer reaches zero, the algorithm resets, and the next received packet is accepted. The match algorithm is more suitable for scenarios of intermittent streams where no more than one packet is in flight on any given path compared to other paths.

For information on the vector algorithm, see the Sequence Recovery section.

Table 30. Stream Identification Types

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

REDUNDANCY

Sequence Recovery

The vector algorithm provides a more robust duplicate elimination. Upon packet arrival for the vector algorithm, 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. Any packets outside of this range are discarded. Additionally, duplicated packets within the history length are also discarded. Each time a packet is accepted, the timer restarts. When the timer expires, the vector algorithm resets, allowing acceptance of any sequence number in the next arriving packet. Increasing the history length of the vector algorithm makes it more suitable for scenarios of bulk streams where there can be more than one packet in flight on any given path.

Individual Recovery

The individual recovery addresses specific errors, such as a stuck transmitter that 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 the repeating sequence number within a single member stream and removes them early on, which allows early detection of errors.

MEDIA REDUNDANCY PROTOCOL (MRP)

The MRP is another redundancy protocol used to avoid single points of failure in industrial communications networks. MRP is a recovery protocol based on a ring topology and is aligned with IEC 62439-2:2021 standard. The MRP can be used in ring networks for up to 50 devices. For full details on the MRP protocol review the detailed standard. The following sections provide an overview of the MRP to help describe the functionality provided by the ADIN6310 devices and does not intend to be a full overview of MRP function:

- ▶ [MRP Stack on the Switch](#)
- ▶ [Recovery Profiles](#)
- ▶ [MRP Operation](#)
- ▶ [MRP Roles](#)
- ▶ [MRP Port and Ring States](#)
- ▶ [MRP Voting Process](#)

MRP Stack on the Switch

The MRP can be configured on start-up of the switch. The MRP stack is running on the packet assist engine, thereby offloading the MRP overhead from the host. Review the TSN driver library for full details of the APIs and configuration for the switch, which is available as part of the software drivers from the [ADIN6310](#) product page.

The switch supports operation as a media redundancy client (MRC), media redundancy manager (MRM), or media redundancy automanager (MRA).

The switch does not support interconnected rings.

One instance of MRA, MRC, or MRM is supported on a 6-port device.

Recovery Profiles

When the device 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 have the same recovery profile.

MRP Operation

There are always two kinds of nodes in the ring, one node acts as the MRM that monitors and controls ring topologies, while the other node acts as the MRC.

MRP Roles

There can be only one MRM device in any ring, and this device manages the ring state. In normal operation, the MRM blocks all traffic on one of its ring ports; however, it allows the MRP traffic through, which avoids loops in the ring. The MRM sends test frames out of one ring port, and if the ring is closed, expects to receive these frames on its other ring port. The MRP sends these frames in both directions. Failure of the ring is detected if the MRM does not receive these test frames back. The MRM reacts by opening the previously blocked ring port to normal traffic and sends the **MRP_TopologyChange** frames out to the ring.

Other nodes in the ring are MRCs. In normal operation, MRCs allow traffic through, their ring ports are forwarding, and the **MRP_Test** frames can cross these devices. An MRC reacts on received reconfiguration frames from the MRM and can detect and signal link changes on its ring ports. The MRC notifies a change by sending the **MRP_LinkChange** frames through both its ring ports.

Certain nodes or all nodes in ring can start as a MRA. MRAs select one MRM for the ring using a voting protocol, and the others transition to the MRC role. When configuring MRP device roles, the standard does not support having one or more MRA roles with a manually configured MRM in the same ring. Expected configurations for rings are one or more MRA(s) with multiple MRCs or one MRM with multiple MRCs.

MRP Port and Ring States

A switch has two dedicated ports known as ring ports. The ring ports can be any ports on the switch, which ports are defined as part of the configuration. By default, the ring ports are Port 1 and Port 2, but users can configure the device to have different ports be part of the ring. Note that only one instance of MRP can run at one time; therefore, only two ring ports are allowed in the ADIN6310 6-port switch.

Ring ports can be in one of the three following states:

- ▶ Disabled: all packets received by the port are dropped.

REDUNDANCY

- ▶ Blocking: all packets received by the port are blocked, except for the MRP packets, PTP, and LLDP frames.
- ▶ Forwarding: all packets received by the port are forwarded.

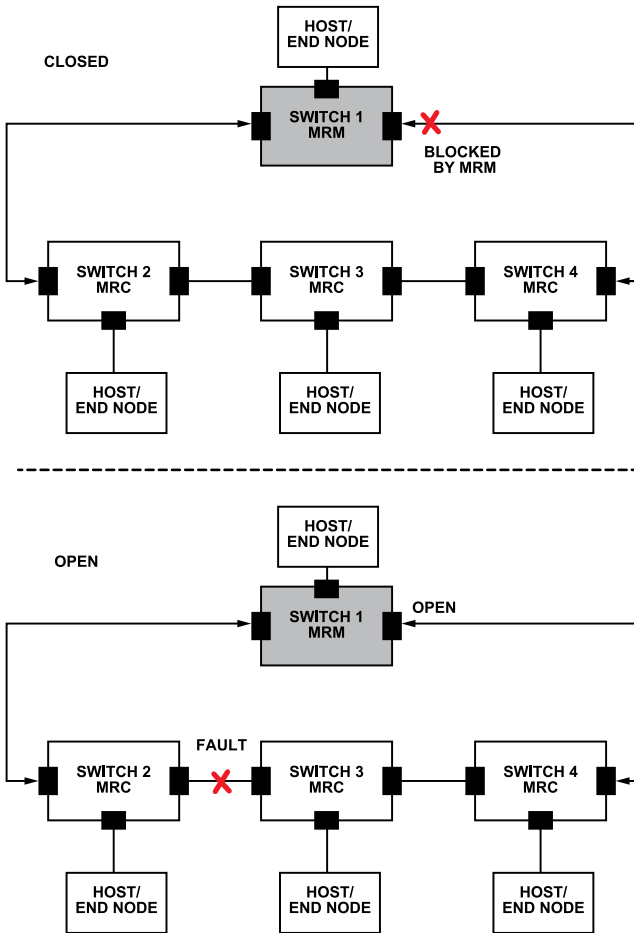


Figure 46. MRP Ring State Closed and Open

The ring itself can be in one of the following states:

- ▶ Closed: both ports on all nodes are linked up; one port on the MRM is blocking, and all other ports are forwarding.
- ▶ Open: both ports on the MRM are forwarding, either because one of the MRCs has signaled a port down, or the MRM failed to receive its own test frames.

MRP Voting Process

The MRA role is a combination of the MRM and MRC, and the MRA supports a manager voting process. Only one active MRM is allowed in the ring. All other nodes must be MRC.

If more than one node has the MRM ability (configured for the MRA), manager voting is used to decide which node should become the MRM, while the other nodes take the MRC role. Do not configure devices manually as a MRM when rings are configured with the MRA roles.

Nodes with the MRM ability have different priorities (**MRP_Prio** field of the test frame). **MRP_Prio-value** + the MAC address = the unique priority. The lower the **MRP_Prio-value** the higher the priority; therefore, if **MRP_Prio-value** is the same across nodes, the lower MAC results in that device becoming the MRM.




Figure 47. MRP Voting Before and After

MRA devices send **MRP_test** frames on both ring ports, and the **MRP_Test** frames include **MRP_Prio** information. Remote MRA devices compare received priority to their own priority. If nodes own priority is higher than the received priorities, the node sends a negative test manager acknowledgment (**MRP_TestMgrNACK**) frame.

If a negative test manager acknowledgment with its own MAC address is received, the receiving MRA transitions to a MRC role and propagates the role change via **MRP_TestPropagate**.

The MRA or MRM with the highest priority becomes the one and only MRM of the ring and starts to manage the ring.

If the MRM fails or is removed from the ring, the voting process restarts. If another MRA is inserted into a ring with an existing MRM, this MRM and new MRA exchange **MRP_Test** frames and vote to be the MRM.

OTHER FUNCTIONALITY

INDUSTRIAL ETHERNET ENGINE

The switch has the ability to optionally support legacy protocols (future software updates) on two ports while supporting the TSN capability on the remaining ports.

The industrial Ethernet engine supports Custom Layer 2 support on Port 1 and Port 2 for PROFINET IRT, EtherNet/IP beacon-based DLR, and POWERLINK protocols.

When using this feature, the TSN MAC on this pair of ports is bypassed; therefore, the TSN functions are not available on these two ports.

The industrial Ethernet protocol frames only traverse through the path between Port 1, the industrial Ethernet engine, and Port 2. Communication to the host is routed through the packet assist engine. Note that there is no direct bridge between the TSN ports and this pair of ports when this block is active.

The host controller can configure the operation of the industrial Ethernet engine through the software driver API, which is available as part of the software drivers from the ADIN6310 product page.

When using this block, the device supports operation at speeds of 100 Mbps over RGMII or SGMII. Gigabit speeds are not supported when using this feature.

Operating this block with RMII is not supported.

LOGICAL MAC OPERATION

The logical MAC feature provides the ability to logically partition the physical MAC ports into different logical MAC groups with a specific link port and one or more network ports. This feature can add tag information at the end of the frames sent to the link port, thereby allowing the stack processor to distinguish which network port the frames originated from. From the stack processor side, it can control which port or ports a frame that it originates is transmitted out of. The switch can support changing the source MAC address of the egressing packets from the switch network ports. The host can register callbacks for frames received on the desired logical MAC with the tag information indicating which port it originated from.

The following example details a configuration where there is one logical MAC group with three ports. The group contains Port 0, Port 1, and Port 2, where the link port is Port 0, and Port 1 and Port 2 are the network ports.

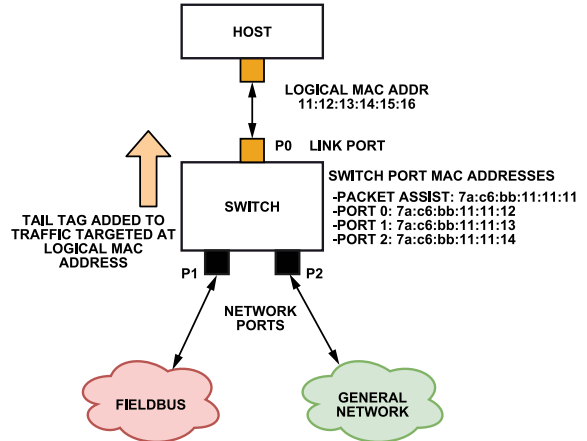


Figure 48. Example of a Logical MAC Group with Three Ports

After configuring this logical MAC group, traffic from the network ports targeted at the logical MAC address (to the host and/or link port) have a 6-byte tail tag added at the end of the frame right before the FCS. The values of interest in this tag are the upper three bits (Bits[13:15]) of this tail tag, which indicate which port the traffic originated from. The tail tag provided allows users to identify any of the six ports as shown in Table 31. An example of a tag added into the frame originating from Port 2 aimed at the logical MAC address has 60x 00xx 00 30 88 fb added at end of the frame. The 6 indicates it is from Port 2. As this feature repurposes the PRP RCT trailer, the 0x88fb at the end is the PRP suffix. Other bit fields within the tag vary depending on the frame size and sequence counter, and the sequence counter gets reset every 5 ms to ensure it does not overwrite the port ID information.

Table 31. Tail Tag Port Identification

Port Number	Tail Tag Value	Top 3 Bits (Bits[15:13])
0	0x2xxx/0x3xxx	001
1	0x4xxx/0x5xxx	010
2	0x6xxx/0x7xxx	011
3	0x8xxx/0x9xxx	100
4	0xAxxx/0xBxxx	101
5	0xCxxx/0xDxxx	110

Creating a logical MAC instance takes care of placing entries into the static table to route messages to the logical MAC address. For messages with a source MAC address as the port MAC address, the frame automatically sends out that port.

Other traffic is not forwarded automatically to the host and/or link port; therefore, for other traffic of interest to the host, static entries must be installed in the forwarding table using the **SES_AddStatic-TableEntry()** API, which is available as part of the software drivers from the ADIN6310 product page. When creating a logical MAC grouping, the API takes care of creating a port forwarding mask between the logical MAC ports to ensure traffic is constrained to the specific ports of the group. The user is not required to do any additional configuration of the port forwarding.

OTHER FUNCTIONALITY

Network ports cannot be part of two different groups, but link ports can belong to different groups.

The following example details a configuration where there are two logical MAC groups, each with two ports. In this scenario, the host has two MAC addresses associated with it. The first group is Port 0 and Port 1, where the link port is Port 0, and Port 1 is the network port, while in the second group, the link port is Port 0 and the network port is Port 2. As previously discussed, link ports can belong to more than one group; however, network ports can only belong to one group instance.




Figure 49. Example of a Logical MAC Group Where the Host MAC Has Two MAC Addresses; Therefore Two Logical MAC Groups Are Required

Configuring Logical MAC Frame (Replace Source MAC)

The `SES_ConfigureLogicalMacFrame()` API, which is available as part of the software drivers from the [ADIN6310](#) product page, allows users to enable the automatic source MAC address replacement for specified frames transmitted out a network port. Users can configure the source MAC addresses that they want replaced with a choice of replacing it with the logical MAC address or with the egressing port MAC address.

In the following example, any frames ingressing Port 0 with source MAC address 11:11:11:11:55 destined to egress on Port 1 will have the source MAC updated to either the port MAC address 7a:c6:bb:11:11:13, or alternatively, with the logical MAC address; for example, 11:12:13:14:15:16.




Figure 50. Example Where Source MAC Address Gets Replaced on Port 1

LAYER 2 TRANSMIT AND RECEIVE

The switch provides the ability to transmit and receive Layer 2 frames. The Layer 2 receive function returns the requested frame and associated data using a callback function registered by the user. Requested frames can be distinguished by destination address (first priority) or Ethertype (second priority) and can be received locally by the stack processor or remotely by the scaleable Ethernet switch. The user can request that remotely received frames (frames received at a scaleable Ethernet switch port) include ingress timestamp and port data. The Layer 2 transmit function provides a mechanism to send Layer 2 frames through the scaleable Ethernet switch. The user is able to request frame return egress timestamps via callback. Additionally, the transmit API (which is available as part of the software drivers from the [ADIN6310](#) product page) supports the scaleable Ethernet switch port features for automatically generating the FCS and priority. Finally, the transmit API supports use or override of the scaleable Ethernet switch standard frame forwarding.

EVENT SUBSCRIPTION

Events generated on the switch can be propagated to the host. The host subscribes to events of interest using the `SES_SubscribeEvent()` API, which is available as part of the software drivers from the [ADIN6310](#) product page. The types of events available include the following:

- ▶ Port link up
- ▶ Port link down
- ▶ Configuration complete event
- ▶ MAC set
- ▶ Port configuration successful
- ▶ Network sync
- ▶ Network sync ready
- ▶ LLDP events (new neighbor, LLDP change)
- ▶ Dynamic table nearly full
- ▶ Input capture event (trigger capture of timestamp)

OTHER FUNCTIONALITY

LINK LAYER DISCOVERY PROTOCOL (LLDP)

LLDP is a Layer 2 neighbor discovery protocol. It is an IEEE standard protocol that defines messages encapsulated in Ethernet frames to announce device information to neighboring devices. The switch supports an LLDP stack running on the packet assist engine. The stack is compatible with the IEEE 802.1AB – 2016 standard. The switch uses LLDP to advertise its capability for frame preemption. The LLDP stack is disabled by default and must be initialized and enabled during configuration. LLDP can be left disabled, and the host processor can run its own LLDP stack.

The default configuration is for LLDP frames to be sent out every 30 seconds with a time to live (TTL) of 121 seconds. These values are programmable, and the user can choose other intervals. APIs are provided for user control of LLDP transmission, reception, and addressing, which is available as part of the software drivers from the [ADIN6310](#) product page. The user can also add or update TLVs to the LLDP frames transmitted at a port and query received LLDP frames. The user can also subscribe to notifications for LLDP-related events, such as a new neighbor being detected, a change observed in an incoming LLDP frame, or in remote statistics or a neighbor shutting down.

IGMP SNOOPING

Internet group management protocol (IGMP) snooping is used to identify multicast groups in a network, which are groups of devices that all receive the same traffic. IGMP snooping ensures that multicast traffic is only forwarded to the relevant ports and suppresses the unnecessary forwarding of multicast data, thus saving bandwidth on the ports that do not require to receive the multicast messages.

The IGMP protocol allows for several kinds of IGMP messages, such as the following:

- ▶ Membership query (general): sent by a multicast router to discover which multicast groups have members on a particular network segment.
- ▶ Membership query (group-specific): sent by a multicast router to inquire about the status of a specific multicast group.
- ▶ Membership report: sent by a host to indicate its interest in receiving multicast traffic for a particular group, which includes the following two types:
 - ▶ Version 1 is a simple report that indicates that the host wants to receive traffic for a specific group.
 - ▶ Version 2 includes additional information about the preferences of the host for multicast traffic.
- ▶ Leave group: sent by a host to indicate that it is no longer interested in receiving traffic for a specific multicast group.

The switch examines IGMP messages, learns the location of multicast routers and group members, and routes traffic accordingly to ensure that multicast traffic is forwarded to the correct destinations

only and not out all ports. The example shown in [Figure 51](#) shows a scenario where the multicast source out on Port 2 (P2) sends IGMP queries. The devices on Port 0 (P0) and Port 5 (P5) send IGMP reports indicating interest in this multicast group, and the switch controls the flow of traffic to ensure only these two ports receive this particular multicast traffic.




Figure 51. IGMP Snooping

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.

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. If a membership report is not received from the host side against the query message within the timeout, the switch forwards the multicast packets to all the ports.

PORT MIRRORING

By default, during normal operation, the switch does not support port mirroring.

MULTIPLE SPANNING TREE PROTOCOL

The switch can support the multiple spanning tree protocol (MSTP) in accordance with the IEEE 802.1Q 2022 standard. The switch supports port roles per Section 13.12, (IEEE 802.1Q 2022's root port, master port, designated port, alternative port, and backup port) and port states per Section 13.16 (discarding, learning, and forwarding).

The configuration supports automatic determination of the multiple spanning tree region with each region mapped to a set of VLAN IDs and automatic reconfiguration of the spanning tree topology as a result of bridge failure or a breakdown in the data path. Four MST instances are supported (MSTIO and MSTIs).

MSTP has backward compatibility with rapid spanning tree protocol (RSTP) and spanning tree protocol (STP); therefore, the switch can be configured to work in RSTP or STP mode as required.

APPLICATIONS INFORMATION

APPLICATION CIRCUIT OVERVIEW

Figure 52 shows an overview of the components required around the ADIN6310 switch.




Figure 52. Typical Application Circuit for ADIN6310

APPLICATIONS INFORMATION

Interface Selection

The example in in [Figure 52](#) shows the host configured for RGMII on Port 0. See the [Host Interface](#) section for details on other options.

MDIO Bus

The MDIO bus is routed from the switch to all PHYs. This bus resides in the VDDIO_B voltage domain. In the event there is a PHY on Port 0, and VDDIO_A is at a different supply voltage than VDDIO_B, level shifting may be required on the MDIO bus between the voltage domains.

PHY Strapping

When using the [ADIN1300](#) or [ADIN1200](#) PHYs, the ADIN6310 can provide unique PHY addresses for each PHY; therefore, no external PHY address strapping resistors are necessary. If using the [ADIN1100](#) PHY, external PHY address strapping resistors must be used. The default MAC interface of the PHY is the RGMII. In this configuration, the PHY is shown strapped for auto-negotiation, all speeds and software power down (SWPD) after reset, whether the switch communicates over the MDIO bus to bring the PHYs out of SWPD. For other strapping configurations, visit the [ADIN1300](#) data sheet for further details.

SFP Usage

In the [Figure 52](#) example, one port is shown connected to an SFP module. The SFP module includes the ac coupling capacitors required. The LOS is an open-drain and/or collector output signal indicating that the received optical power is in normal operation when low or less than the worst-case receiver sensitivity when high. The LOS output expects a pull-up voltage between 2 V and the VCCT voltage level; therefore, it can be connected directly to Px_LINK if VDDIO_x is >2 V. If VDDIO_x is operating from an 1.8 V rail, the LOS must be pulled to VCCT, and a level shift must be used between the LOS and Px_LINK.

APPLICATIONS INFORMATION

APPLICATION USE CASES

6-Port Configuration

The ADIN6310 is designed for use in programmable logic controller (PLC), distributed control system (DCS), or remote input and output units where it provides scalable port counts to enable connection of external devices within the automation environment. The first use case for the ADIN6310, shown in Figure 53, illustrates a 6-port switch configuration where the host processor interfaces with the ADIN6310 through the SPI (standard SPI, dual SPI, or quad SPI) only to enable all six ports to be used for connection to external devices.

The Ethernet PHYs can be a mix of different speed PHYs using the ADIN1300 10, Mbps, 100 Mbps, and 1000 Mbps and ADIN1200 10 Mbps and 100 Mbps devices. The interface shown between the ADIN6310 and ADIN1200 is the RGMII, but it could also be the RMII; however, the RGMII is preferred because of latency, ease of PHY addressing, and simplification of the clocking requirement (because the voltage domain between the ADIN6310 and ADIN1200 can differ).

Similarly, the 6-port switch can be configured to provide four external ports connecting to factory cells, human machine interface (HMI), robots, vision systems, or other inputs, and the remaining two ports can be used for internal connection to the PLC backplane with one Ethernet connection to the host processor via the RGMII.

Figure 53 also shows the host interface residing on the optional separate VDDIO_A voltage domain, which includes Port 0, the SPI, timers, and GPIO. The PHYs on Port 2 through Port 5 reside on the VDDIO_B voltage domain and can run from a different voltage than VDDIO_A. If a PHY is used on Port 0, this also resides on the VDDIO_A voltage domain; therefore, some level shifting may be required.

The supply voltage requirements for the switch can be satisfied by using the LTM4668A module, which provides four voltage rails. There are two common voltage rails (VDDIO and AVDD3P3) between the ADIN6310 and the ADIN1300 PHYs; however, the PHY requires an additional 0.9 V rail, and the switch also requires a 1.1 V rail.

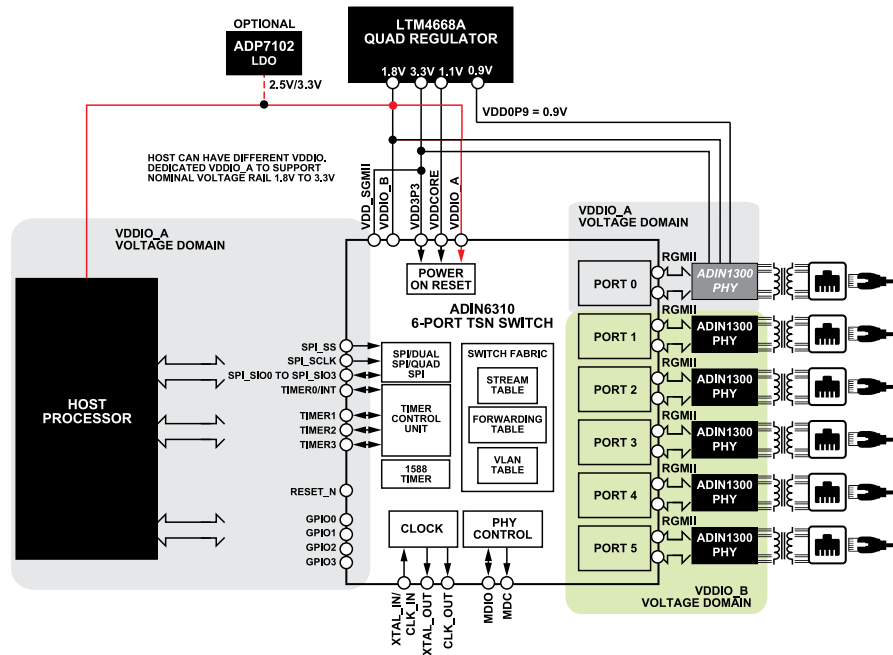


Figure 53. Application Circuit for 6-Port Switch Use Case Where the Host Processor Interface Is via SPI

APPLICATIONS INFORMATION

Field Switch Use Case

Figure 54 shows an overview of an ADIN6310 being used in a process automation application of a field switch topology that enables the connection of 10BASE-T1L, physical layer-enabled field instruments. The configuration shows two trunk ports configured for 10 Mbps communication with the ADIN1100, 10BASE-T1L Ethernet PHYs, enabling communication over 1 km of single twisted pair cabling to a power switching unit. Alternatively, the uplink port can be configured to support 10 Mbps, 100 Mbps, or 1 Gbps connectivity to a PLC or DCS via the use of the ADIN1200 and ADIN1300 Ethernet

PHYs, or fiber for longer distance requirements over the SGMII capable ports. The spur connections are all enabled to support 10BASE-T1L communication via the ADIN1100 to supporting field devices, such as temperature or pressure transmitters or flow meters that reside in intrinsically safe, Zone 0 environments. The field switch itself can be self powered or use Analog Devices power over data lines (PoDL) technology. The ADIN1100 has programmable transmit levels to suit trunk and intrinsically safe spur use cases. The interface between the ADIN6310 and ADIN1100 is RMIi or RGMII.




Figure 54. Application Circuit for an APL Field Switch with Two Trunk Ports and Four Spurs Using the ADIN1100 10BASE-T1L PHY

APPLICATIONS INFORMATION

Cascading Multiple Switches

Larger port counts can be built by cascading multiple switches as shown in Figure 55. The host processor manages the switch configuration and is connected via the SPI to the primary switch in the chain. Two ports of the primary device can be used for trunk connections to the Ethernet network, and one port as an interconnection to the chain of secondary switches. The intercon-

nect is either RGMII or SGMII, with SGMII preferred due to the reduced number of traces required and its capability to route over longer traces and across connectors. The interconnect consumes two ports on each secondary switch device, leaving four ports available to connect to the ADIN1100 10BASE-T1L PHYs per switch; therefore, each additional switch adds four spur ports to the system.

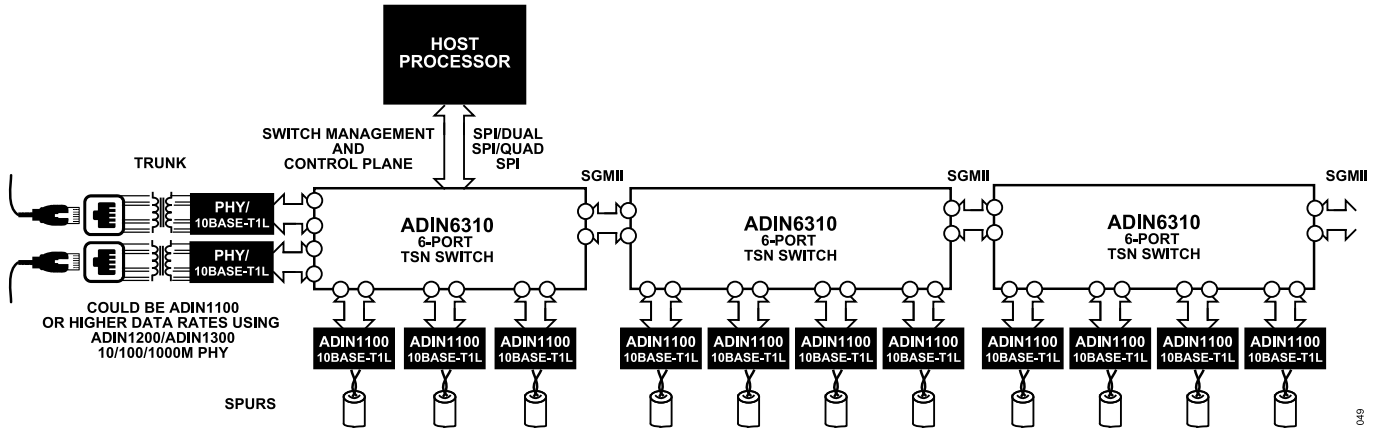


Figure 55. Cascading Multiple Switches to Build a Larger Field Switch

LAYOUT GUIDELINES

GENERAL LAYOUT GUIDELINES

The general layout guidelines include the following:

- ▶ Maintain 50 Ω characteristic impedance for all single-ended traces.
- ▶ Maintain 100 Ω characteristic impedance for all differential traces.
- ▶ Route the differential circuits away from the noisy clock and power circuit
- ▶ Provide adjacent solid GND returns for all the planes and signals.
- ▶ Provide a high decoupling connection between the digital GND return and chassis GND.
- ▶ Avoid discontinuities such as splits, slots, gaps, and cuts in the ground or return on the board.
- ▶ Any component GND pin connecting to GND fill or pour must have a via to the GND plane at the pin.
- ▶ In the case of signal transition from layer to layer, provide a GND via adjacent to the signal via.
- ▶ Locate high-speed connectors on one corner or on one edge of the board.
- ▶ Locate clock drivers near the IC if the device is used for high-frequency communication or consider providing these clock drivers locally.
- ▶ Avoid routing high-frequency signals beneath or near ringing or noisy circuits or components.
- ▶ Avoid overlapping power planes.
- ▶ Avoid routing power planes on the edge of the board.
- ▶ Avoid routing power planes like traces, instead use planes.
- ▶ Avoid stubs on all signal traces,
- ▶ Avoid vias where possible, routing signals on one layer as straight as possible or with 45° angles where traces bend.

COMPONENT PLACEMENT

Prioritization of the critical traces, power planes, and components helps simplify the routing exercises. Place and orient the critical high-speed differential, single-ended, and clock traces, power planes, and components first to ensure an effective layout with minimal turns, vias, and crossing traces.

For the switch layout, the important components are the crystal, the MAC interface traces, the SGMII traces, and all bypass capacitors local to the device. Prioritize these components and the routing to them. The following sections provide more detail for each of the areas.

For stack up, ensure to provide a solid GND return plane to all the traces and power planes adjacent to it.

Power Supply Decoupling

From a PCB layout point of view, it is preferred to locate decoupling capacitors on the same side of the IC and close to the IC power pins to help minimize the loop area inductance of the capacitor connected to IC power pins. If the capacitors are on the opposite side, it creates a higher connection inductance. If there is no room on the top layer, keep decoupling capacitors close to or under the IC. Do not put traces on capacitors for GND connection. Ensure that decoupling capacitors via connection to the power planes and GND plane are not far apart. Avoid decoupling capacitors sharing the GND pour connections with other circuitry components. Choose filtering ferrites of smaller size and appropriate impedance to the frequency of interest.

See [Table 27](#) and [Table 28](#) for recommended capacitor values and priority locations.

Crystal Oscillator

The switch requires a 25 MHz clock that can be an external crystal oscillator applied across the XTAL_IN and XTAL_OUT pins or an external clock applied to the XTAL_IN pin. To ensure minimum current consumption and to minimize stray capacitance, make connections between the crystal, capacitors, and ground as close to the device as possible. If possible, place the crystal and capacitors on the same side of PCB as the switch device as follows;

- ▶ Provide symmetrical traces.
- ▶ Provide the same GND return for the crystal oscillator, its traces, and load capacitors, and ensure that capacitor GND pads are well stitched to all the GND layers, especially adjacent to solid GND returns.

MAC Interface

Each port can support the RMII or RGMII, and four of the ports (Port 1 to Port 4) can support the SGMII.

RMII and RGMII

The RMII is an eight signal interface for each port capable of 10 Mbps and 100 Mbps speeds, while the RGMII can support data rates of 10 Mbps, 100 Mbps, and 1000 Mbps speeds requiring 12 pins. Where possible, route these interface pins on the same side as component pins, as follows:

- ▶ Keep trace lengths as short as possible.
- ▶ Route traces with an impedance of 50 Ω to ground.
- ▶ All signals within the transmit group must be length matched (to within 100 mil), similarly for all signals within the receive group. Note that, no matching is required from port to port.
- ▶ Avoid vias where possible, routing signals on one layer as straight as possible or with 45° angles where traces bend.

LAYOUT GUIDELINES

- ▶ Avoid crossover of the signals where possible. Avoid stubs on all signal traces.
- ▶ Route traces over a continuous reference plane with no interruptions to reduce inductance.
- ▶ It is recommended to route traces on the same layer.

The switch provides user programmability to adjust the drive current of the RGMII pins to help improve signal integrity and minimize ringing. Alternatively, series termination resistors can be placed on all RGMII pins if further tuning is required. When event series termination resistors are used, place these resistors close to the output pins on each side.

SGMII

The SGMII is a serial interface using two differential pairs to transfer information between the switch and SFP or host, using less pins and requiring less routing and traces compared to the RMI and RGMII. The SGMII transmits at 1.25 Gbps and requires external AC coupling with 0.1 μF capacitors (typically included in the SFP modules). The SGMII is available on Port 1 to Port 4 of the ADIN6310. The interface is capable of 10 Mbps, 100 Mbps, and 1 Gbps full duplex communication with auto-negotiation or forced modes. The SGMII pins can be connected to an SFP transceiver for fiber communication or over a backplane.

The SGMII lines are differential and must be length matched to connector to within 50 mils. Also ensure the following:

- ▶ Route the Px_STxP and Px_STxN and Px_SRxP and Px_SRxN signal traces as 100 Ω differential line route signals in homogeneous internal layers using buried and blind vias. Keep impedances constant throughout because any discontinuities can affect signal integrity.
- ▶ Co-planar routing is recommended—stitch vias on the ground shield on both sides of the traces or use a ground pour.
- ▶ Keep routing lengths short and avoid stubs.
- ▶ Each pair must be routed in parallel and on the same layer, with trace widths kept the same throughout and length matched to within 50 mil.
- ▶ Avoid any right angles on these traces (use curves in traces or 45° angles).
- ▶ Avoid crossing of differential traces with other lines and among each other. Place the components such that crossing of differential pairs is not necessary.
- ▶ Route differential pairs without vias if possible. In the case of signal transition from layer to layer, provide a symmetrical GND via adjacent to the signal via.
- ▶ Components or vias on the differential pair must be symmetrical.

- ▶ Where possible, ensure that a solid return paths are underneath all signal traces. Avoid routing signal traces across plane splits.




Figure 56. Things to Avoid when Routing Differential Pairs

MDIO

The switch has a controller MDC or MDIO interface that is a 2-wire bus used to manage the control and status of the external Ethernet PHYs connected to the switch MAC interfaces. The MDIO signal line must have a 1.5 k Ω pull-up resistor to the VDDIO_x line.

Length match the MDIO interface traces to within 100 mils.

SPI

The switch contains a synchronous SPI-compatible interface that can be optionally used as the host interface. The SPI supports 32-bit data transfers, is a target interface, and can be configured as a standard SPI, dual SPI, or quad SPI. Length match SPI traces to within 100 mils.

Power and Ground Plane

For the best high-frequency decoupling, ensure that the power and return planes are not more than 3 mils to 4 mils apart. Closer is the better for low inductance connections and interplane capacitance. Also, ensure that the following is done:

- ▶ Avoid overlapping power planes.
- ▶ Avoid routing power planes on the edge of the board.
- ▶ Avoid routing power planes like traces, and use planes.

Isolation Guidelines—RJ45 Layout

For better electromagnetic interference (EMI) performance consider using an inbuilt common-mode choke (CMC), 12-core magnetic RJ45 connector connected to the Ethernet PHY and provide a high decoupling connection between the digital GND return and the chassis GND.

SILICON ANOMALY

This anomaly list describes the known bugs, anomalies, and workarounds for the switch.

Analog Devices is committed, through future silicon revisions, to continuously improving silicon functionality. Analog Devices tries to ensure that these future silicon revisions remain compatible with your present software and systems by implementing the recommended workarounds outlined within this section.

Table 32. Functionality Issues

Silicon Revision Identifier	Silicon Status	Anomaly Sheet	Number of Reported Anomalies
B1/U3	Released	0	7

Table 33. Statistic Counters for Per Stream Filtering and Policing Function (Qci)

Background	Statistic counters gathered for debugging and diagnostic purposes are not presented per the 802.1Qci specification.
Issue	Statistic counters not provided per the IEEE 802.1Q specification. This issue does not impact operation of the filter, gate, or flow meter function or proper flow of the traffic across switch. Stream filter: statistic counters (0 to 31) are indexed by StreamID, instead of the stream filter ID. Because there can be multiple stream filters associated with a stream ID, therefore, multiple stream filters could increment the same counters. Stream gate: statistic counters are indexed by their own identifiers (0 to 15) instead of the active stream filter (0 to 31). Flow meter: statistics are indexed by their own identifiers (0 to 7) instead of the active stream filter (0 to 31). Implications: stream gate and flow meter statistics are not linked to the active filter. In addition, there are fewer counter entries available for both.
Workaround	None.
Related Issues	None.

Table 34. PRP Function Handling of Traffic in a Misconfigured Network (Wrong LAN Connections)

Background	In PRP mode, if a port receives a frame with a PRP RCT tag with a wrong LAN ID, the switch increments the counter but incorrectly performs a duplicate discard and strips the PRP trailer.
Issue	As a result, the host SAN receives only a single copy of the frame with the PRP RCT trailer removed, rather than two copies. Frames arriving with a wrong LAN ID is a misconfiguration of the network. There are error counters in place to detect this condition and allow the host to handle it accordingly.
Workaround	Error counters are available for the host to manage this situation at the application level.
Related Issues	None.

Table 35. PRP Traffic with Wrong LSDU Size in the PRP Tag Gets Dropped

Background	PRP frames with wrong LSDU size is dropped.
Issue	IEC 62439-3:2021 Section 4.2.7.5.1 and Section 4.2.7.5.3 indicate frames with the wrong LSDU size in the PRP tag are to be forwarded to the upper layers. If a peer device transmits PRP frames with the wrong LSDU size in the PRP tag, the switch configured as a PRP DANP or PRP RedBox device drops those frames to the host processor and locally attached SANs as long as those frames arrive within the age out time (that is, <code>IreDupListResideMax</code>).
Workaround	None.
Related Issues	None.

Table 36. FRER Frame Without CB Tag Is Modified by the Transmit Transform

Background	Frames matching the stream entry with no CB-tag gets corrupted.
Issue	Applying a sequence recovery to a frame that has no CB-tag results in the frame getting corrupted. This scenario is not expected in a properly configured FRER network because identified frames have an existing CB-tag. The ingressing frame must match an existing stream entry with a transmit transform to remove the CB-tag. In the event such an untagged frame arrives and matches a stream entry, the switch forwards the tagless frame, and increases the packet and encoding error counts; however, the Ethertype gets corrupted, and six bytes are removed from the end of the frame.
Workaround	None.
Related Issues	None.

SILICON ANOMALY**Table 37. SGMII Alignment Errors**

Background	Alignment error count increments on odd sized frames.
Issue	In all SGMII modes, a valid frame with an odd number of bytes increments the received alignment error count statistic, and the FCS error count does not increment unless the frame is corrupt.
Workaround	Alignment errors without corresponding FCS errors can be ignored. Errored frames can be detected by FCS errors alone.
Related Issues	None.


Table 38. Extended Table Lookup Cannot Inspect Last 10 Bytes of Frame

Background	An extended lookup cannot perform a search on the last 10 bytes of the frame at Gigabit speed.
Issue	This issue causes the extended lookup to return a miss, resulting in the frame being handled by the return information provided from the previous lookup (destination address and VLAN ID, source address and VLAN ID, or miss).
Workaround	Ensure extended lookups do not extend into the last 10 bytes of the frame. Increase the frame size to ensure lookup can be performed.
Related Issues	None.

Table 39. Transform to Remove VLAN Tag from a 64-Byte Frame Results in Corrupted Frame

Background	Application of a transform to remove a VLAN tag from a 64-byte frame results in a corrupted frame.
Issue	When a port is configured as a VLAN access port, the switch applies a transform to remove the VLAN tag as traffic egresses the port. The switch was designed with the expectation that the minimum sized VLAN tagged frame would be 68 bytes. Any runt frame (less than 64 bytes) egressing the switch gets corrupted with 0x00 data to force a CRC error. Application of a transform to a VLAN tagged frame sizes <68 bytes results in a runt frame. This behavior is not conformant with the IEEE802.1Q Annex G standard.
Workaround	None.
Related Issues	None.

OUTLINE DIMENSIONS



COMPLIANT TO JEDEC STANDARDS MO-275-JJAB-2

Figure 57. 256-Ball Chip Scale Package Ball Grid Array [CSP_BGA] (BC-256-6)
Dimensions shown in millimeters

ORDERING GUIDE

Model ¹	Temperature Range	Package Description	Package Option
ADIN6310BBCZ	-40°C to +85°C	256-Ball Chip Scale Package Ball Grid Array [CSP_BGA]	BC-256-6
ADIN6310BBCZ-RL	-40°C to +85°C	256-Ball Chip Scale Package Ball Grid Array [CSP_BGA]	BC-256-6
ADIN6310CBCZ	-40°C to +105°C	256-Ball Chip Scale Package Ball Grid Array [CSP_BGA]	BC-256-6
ADIN6310CBCZ-RL	-40°C to +105°C	256-Ball Chip Scale Package Ball Grid Array [CSP_BGA]	BC-256-6

¹ Z = RoHS Compliant Part.

EVALUATION BOARDS

Table 40. Evaluation Boards

Model ¹	Description
EVAL-ADIN6310EBZ	Evaluation Board

¹ Z = RoHS-Compliant Part.