

BCM88800

4.8-Tb/s Integrated Packet Processor and Traffic Manager Single-Chip Device

Overview

The Broadcom[®] BCM88800 scalable series is the industry's densest switching solution, enabling switching platforms of up to 6000 ports of 100G or 400G Ethernet.

The BCM88800 is the seventh generation of the DNX scalable switching product line and processes up to 4.8 Tb/s of line card traffic, supporting up to twenty-four 100GbE ports, or a mix of front panel ports from 1GbE to 400GbE, operating at Layer 2 through Layer 4.

The BCM88800 series, together with the BCM88790 fabric element (FE) device, enables system vendors to build a scalable product line based on a unified architecture that addresses any density or application, such as the following:

- Multi-terabit core and edge routers for data center, packet transport, or carrier network applications.
- Multiple interconnected modular systems to create a scalable core platform for switching and routing capacity multiplication.

The BCM88800 flexible packet classification engine is microcode programmable, with built-in support for data center and carrier networking applications. The large on-chip classification databases can be extended using an external knowledge-based processor (KBP) from Broadcom.

The BCM88800 traffic manager integrates deep-packet buffers with a distributed scheduling scheme that allows state-of-the-art hierarchical quality-of-service (QoS), transmission scheduling per-customer, per-service, as well as tunneling and overlay networks. Flexible flow control mechanisms support priority-based flow control (PFC), enhanced transmission selection (ETS), and explicit congestion notification (ECN).

The integrated Layer 1 (L1) switch delivers latency-sensitive services, such as optical transport network (OTN), time division multiplex (TDM), or mobile traffic.

Features

- Seventh-generation Dune scalable FAP product line.
- High-performance 2.4-Tb/s full-duplex switching.
- Fabric interface:
 - SerDes interface to the Broadcom BCM88790 fabric element.
 - Fabric-less (without the fabric element) configurations of up to three devices.
- Flexible network interface:
 - 1GbE, 10GbE, 25GbE, 40GbE, 50GbE, 100GbE, 200GbE, and 400GbE interfaces.
 - Support for a mix of OTN (ODU) and ETH ports on the same device.
- Traffic manager:
 - 4-GB on-chip deep buffering.
 - Hierarchical memory management.
 - Programmable, hierarchical scheduling.
 - Compliant with scheduling and shaping standards, including MEF and DSL-FORUM.
- Flexible and programmable packet processor:
 - Bridging, routing, MPLS, VPLS, L2VPNs, L3VPNs, and OAM.
 - Data center tunneling encapsulation including VXLAN, NV-GRE, and GENEVE.
 - Built-in support for data center, carrier and metro Ethernet, and transport applications.
 - Large modular on-chip databases, application oriented with off-chip expandability.
 - OAM accelerator engine.
 - PEM (flexible pipe).
- L1 switch:
 - Delivery of latency-sensitive services, such as OTN, TDM, or mobile traffic.
 - Unified fabric for OTN and Ethernet.
- PTP IEEE 1588, SyncE.
- Time-sensitive network (TSN) support.
- In-band management.

Figure 1: Functional Block Diagram

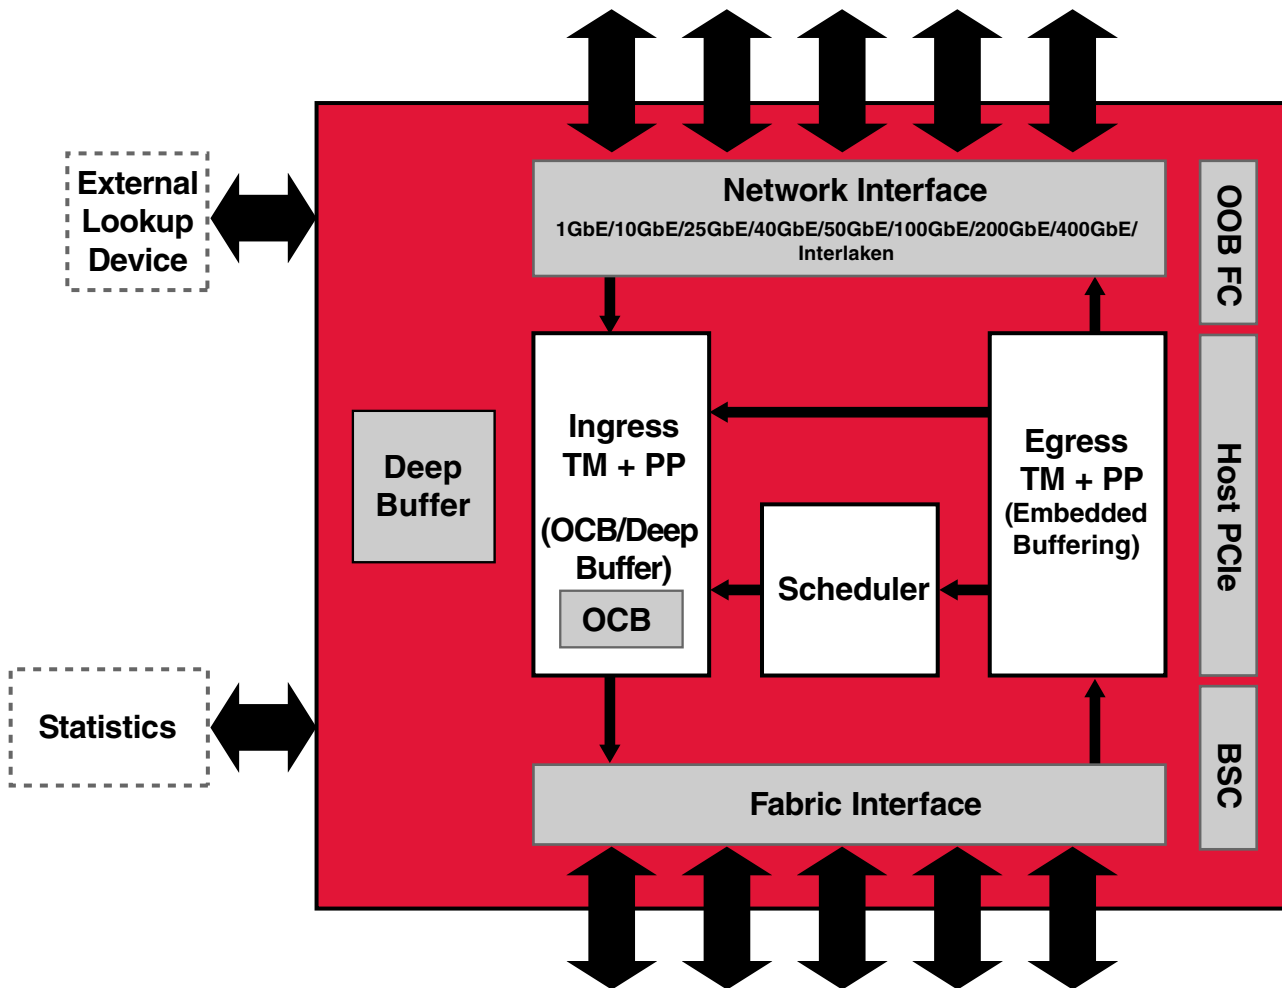


Table of Contents

| | |
|---|-----------|
| Chapter 1: Introduction | 7 |
| 1.1 Features | 7 |
| 1.2 Device Overview | 9 |
| Chapter 2: System Configurations | 12 |
| 2.1 Broadcom StrataDNX Components | 12 |
| 2.1.1 Interoperation with StrataDNX Fabric Element Devices | 12 |
| 2.1.2 Interoperation with StrataDNX Fabric Access Processor Devices | 12 |
| 2.2 Stand-Alone Configurations | 13 |
| 2.3 Mesh Fabric Configurations | 13 |
| 2.4 StrataDNX CLOS Fabric Configurations | 14 |
| Chapter 3: System Interfaces | 15 |
| 3.1 Network Interface | 15 |
| 3.1.1 Ethernet Ports | 18 |
| 3.1.1.1 400GbE Port | 22 |
| 3.1.1.2 200GbE Port | 22 |
| 3.1.1.3 100GbE Port | 22 |
| 3.1.1.4 50GbE Port | 23 |
| 3.1.1.5 40GbE Port | 24 |
| 3.1.1.6 25GbE Port | 24 |
| 3.1.1.7 12GbE Port | 25 |
| 3.1.1.8 10GbE Port | 25 |
| 3.1.1.9 1GbE Port | 25 |
| 3.1.1.10 Mixed Port Types on PM50 | 26 |
| 3.1.1.11 Mixed Port Types on PM25 | 28 |
| 3.1.2 Interlaken | 29 |
| 3.1.2.1 Interlaken Lane Selection | 29 |
| 3.1.2.2 Interlaken over Fabric SerDes | 29 |
| 3.1.2.3 Interlaken SerDes Supported Rates and Electrical Standards | 30 |
| 3.1.2.4 Interlaken FEC | 30 |
| 3.1.2.5 Interlaken Receive Burst Rules | 31 |
| 3.1.2.6 Interlaken Transmit Burst Rules | 31 |
| 3.1.2.7 Interlaken In-Band Flow Control | 32 |
| 3.2 External Lookup Engine | 33 |
| 3.3 Out-of-Band Flow Control | 33 |
| 3.4 Synchronous Ethernet | 35 |
| 3.4.1 Transmit Clock | 35 |
| 3.4.2 Recovered Clock | 35 |

- 3.5 IEEE 1588** 36
 - 3.5.1 BroadSync External Interfaces 37
- 3.6 Fabric Interface** 39
 - 3.6.1 PCS Layer Line Coding 40
 - 3.6.1.1 64b/66b Encoding 40
 - 3.6.1.2 25G Reed-Solomon FEC (RS-FEC) 40
 - 3.6.1.3 Low Latency 25G Reed-Solomon FEC (LL 25G RS-FEC) 41
 - 3.6.1.4 50G Reed-Solomon FEC (RS-FEC) 41
 - 3.6.1.5 Low Latency 50G Reed-Solomon FEC (LL 50G RS-FEC) 41
 - 3.6.1.6 Line Coding Overheads and FEC Latency 42
 - 3.6.2 Fabric SerDes Connectivity 43
- 3.7 Statistics Interface** 44
- 3.8 HBM Packet Buffers** 45
- 3.9 CPU Interface** 45
 - 3.9.1 Remote CPU Support 46
 - 3.9.2 Remote Packet Operations 47
 - 3.9.3 PCIe Interface 47
 - 3.9.4 MIIM 47
 - 3.9.5 UART 47
 - 3.9.6 LED Interface 48
- Chapter 4: Pin Signal Description** **49**
 - 4.1 Pin List and Pin Map 49
 - 4.2 Pin I/O Type Description 49
 - 4.3 Pin Description – Grouped by Function 50
- Chapter 5: Electrical Specifications** **62**
 - 5.1 Absolute Maximum Ratings 62
 - 5.2 Recommended Operating Conditions 63
 - 5.2.1 Recommended Operating Voltage 63
 - 5.2.2 Recommended Operating Voltage Range for DC Supplies 64
 - 5.3 Device Power Consumption 65
 - 5.4 Power Supply Filtering 67
 - 5.5 Power-Up and Reset Sequence 68
 - 5.5.1 Power-Up Sequence 68
 - 5.5.2 Power-Down Sequence 71
 - 5.5.3 Fail-Safe Considerations 72
 - 5.5.4 Warm Reset 73
 - 5.5.5 HBM-Only Power Down and Power Up 73
 - 5.6 Power-Up Configuration Word 74
 - 5.7 DC Electrical Specifications 75

| | | |
|--|---|-----------|
| 5.7.1 | 1.8V Digital I/Os | 75 |
| 5.7.2 | BSC/I ² C | 75 |
| 5.8 | AC Electrical Specifications | 76 |
| 5.8.1 | BSC/I2C Interface Timing | 76 |
| 5.8.2 | Management Interface Timing | 77 |
| 5.8.2.1 | MDIO AC Characteristics | 77 |
| 5.8.3 | SyncE Recovered Clocks | 78 |
| 5.8.4 | LED Timing | 79 |
| 5.8.5 | Out-of-Band Flow Control Timing | 80 |
| 5.8.5.1 | SPI4.2 Flow-Control Mode | 80 |
| 5.8.5.2 | Interlaken Flow-Control Mode | 81 |
| 5.8.6 | BroadSync and Time Sync Timing | 82 |
| 5.8.7 | PCIe Interface | 83 |
| 5.8.7.1 | PCIe Receiver | 83 |
| 5.8.7.2 | PCIe Transmitter | 83 |
| 5.8.8 | QSPI Flash Interface | 84 |
| 5.9 | Reference Clocks | 85 |
| 5.9.1 | CLOCK25 Reference Clock | 85 |
| 5.9.2 | Core and Microcontroller PLL Reference Clocks | 85 |
| 5.9.3 | PCIe PLL Reference Clock | 86 |
| 5.9.4 | SerDes Reference Clock | 86 |
| 5.9.5 | TS_PLL Reference Clock | 87 |
| 5.9.6 | DRAM PHY PLL Reference Clock | 87 |
| 5.10 | Blackhawk SerDes Operating Conditions | 88 |
| 5.10.1 | Blackhawk SerDes Features | 88 |
| 5.10.1.1 | General Features | 88 |
| 5.10.1.2 | Debug Features | 88 |
| 5.10.2 | Blackhawk SerDes Supported Rates | 89 |
| 5.10.3 | Blackhawk SerDes Receiver | 89 |
| 5.10.4 | Blackhawk SerDes Transmitter | 90 |
| 5.11 | Falcon16 SerDes Operating Conditions | 91 |
| 5.11.1 | Falcon16 SerDes Features | 91 |
| 5.11.1.1 | General Features | 91 |
| 5.11.1.2 | Debug Features | 91 |
| 5.11.2 | Falcon16 SerDes Supported Rates | 91 |
| 5.11.3 | Falcon16 SerDes Receiver | 92 |
| 5.11.4 | Falcon SerDes Transmitter | 92 |
| Chapter 6: Thermal Specifications | | 93 |
| 6.1 | Absolute and Operational Thermal Specifications | 93 |
| 6.2 | Package Thermal Specifications | 95 |

| | |
|--|------------|
| 6.2.1 Block Thermal Model | 95 |
| 6.3 Temperature Monitoring | 95 |
| 6.3.1 VTMON | 95 |
| 6.3.2 Thermal Diodes..... | 96 |
| 6.3.3 On-Die Thermal Monitor Locations | 97 |
| 6.4 Reflow Temperature..... | 98 |
| 6.5 Heat Sink Considerations..... | 98 |
| Chapter 7: Packaging | 99 |
| Related Documents | 100 |

Chapter 1: Introduction

1.1 Features

The Broadcom® BCM88800 is an integrated packet processor, traffic manager, and fabric interface single-chip switch. The following features are available with the BCM88800:

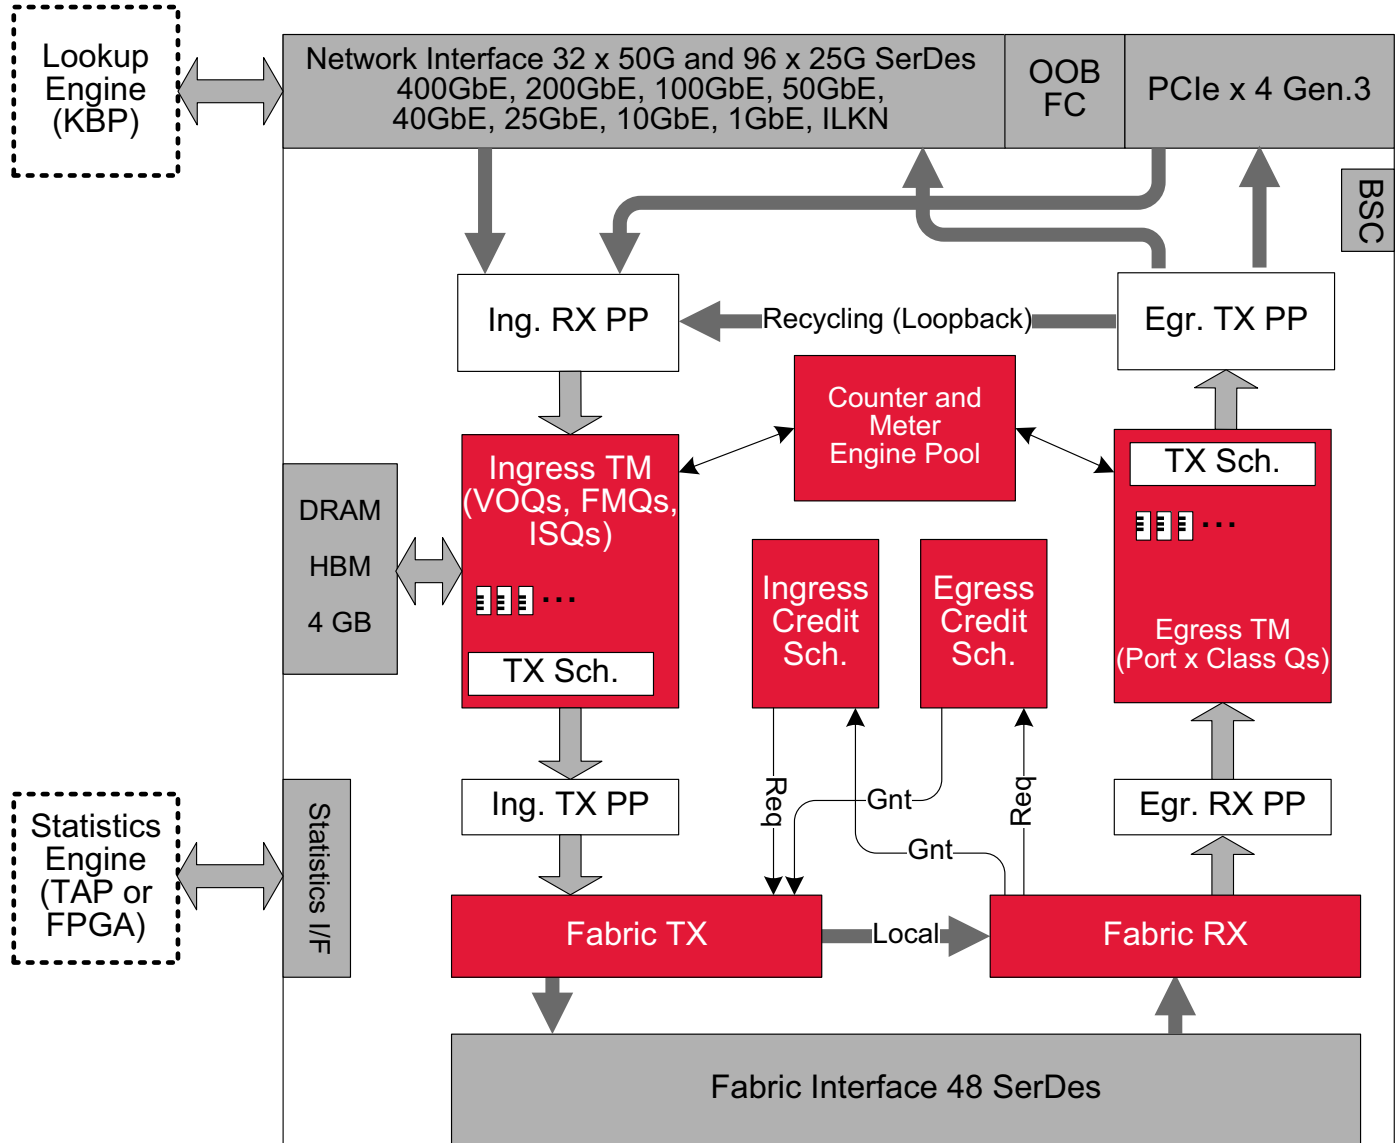
- High performance.
 - 4.8-Tb/s full-duplex, integrated fabric interface, traffic manager, and packet processor.
 - 1000 Mp/s processing rate.
- StrataDNX™ fabric interface.
 - 48 SerDes for up to 53.125-Gb/s serial links (PAM4 with symbol rate of 26.5625 Gbaud).
 - Mesh configurations of up to two devices.
 - Clos fabric configurations up to 4800 Tb/s (2K devices).
 - Dynamically variable-sized cells for highly efficient segmentation.
 - Multiple line coding options with Reed-Solomon Forward Error Correction (RS-FEC), which optimizes performance according to system characteristics.
 - Dynamic routing and load balancing over all fabric links.
 - 1 + 1, N – x, and N + x redundancy schemes.
 - Automatic fault detection and recovery with no software intervention.
- Flexible network interface.
 - 32 SerDes with rates up to 53.125 Gb/s, supporting the following port configurations:
 - 32 × 10GbE/25GbE/50GbE over one lane.
 - 16 × 40GbE/50GbE/100GbE over two lanes.
 - 8 × 40GbE/100GbE/200GbE over four lanes.
 - 4 × 400GbE over eight lanes.
 - 96 SerDes with rates up to 25.78125G. The SerDes support the following port configurations:
 - 96 × 1GbE/10GbE/25GbE over one lane.
 - 48 × 40GbE/50GbE over two lanes.
 - 24 × 40GbE/100GbE over four lanes.
 - Up to 16 Interlaken interfaces.
- Packet lengths supported in the range 64B to 10240B.
- Traffic Manager.
 - 128K programmable wire-rate queues.
 - Deep packet buffering.
 - Integrated high bandwidth memory (HBM) Gen2 cube (4-Hi stack, for a total of 4 GB).
 - Congestion management.
 - Hierarchical WRED and tail-drop policies.
 - Congestion notification – CNM generation and CNM reception (proxy).
 - Flow control generation – Fully programmable, in-band and out-of-band.
 - Flow control reception-any level – Interface, port, class, flow, traffic type-in-band and out-of-band.
 - Priority flow control (PFC) – Eight levels.
 - Congestion tracking statistics.

- Up to 288K meters.
- Three ingress meter operations per packet.
- Two egress meter operations per packet.
- Hierarchical scheduling and shaping.
 - Fully programmable to any depth.
 - Support for priority propagation.
- MEF, DSL-FORUM TR-059-compliant scheduling and shaping.
- Packet processor.
 - Bridging, routing, MPLS, VPLS, L2VPNs, L3VPNs, and OAM.
 - Data center tunneling encapsulations including VxLAN, NV-GRE, and GENEVE.
 - Built-in support for data center, carrier and metro Ethernet, and transport applications.
 - Large modular on-chip databases, application-oriented with off-chip expandability.
 - OAM accelerator engine.
 - PEM (flexible pipe).
- Counters, meters, and statistics.
 - On-chip counter pool up to 384K counters.
 - On-chip meter pool up to 288K meters.
 - Statistics interface for expandable, off-chip statistics gathering:
 - The SerDes used for the statistics interface is shared with NIF SerDes.
 - Efficient packet-based protocol based on Ethernet simplifies connectivity to KBP BCM16K, TAP BCM5235, or FPGAs.
 - Seamless connection to KBP BCM16K and TAP BCM5235 TAP devices.
- Multicast – Pointer-based ingress and/or egress multicast replication.
- IEEE 1588 support with improved time-stamping accuracy of nanosecond scale.
- In-band management.
- PCIe × four-lane Gen3 host interface with DMA.
- Hardware linkscan engine.
- LED processor.
- The BCM88800 device is fully interoperable with the previous and current generation of Broadcom StrataDNX devices. BCM88690, BCM88680, BCM88670, BCM88790, and BCM88770 devices may be used in a system simultaneously.

1.2 Device Overview

The following figure is a high-level functional block diagram of the BCM88800.

Figure 2: BCM88800 Block Diagram



As shown in [Figure 2](#), the BCM88800 includes the following functional blocks.

Traffic Manager

- Ingress traffic manager (TM):
 - Manages a pool of queues in on-chip SRAM and in HBM DRAM
 - Replicates packets for multicast, snooping, and mirroring
- Ingress and egress end-to-end credit scheduler – Schedules packets out of the ingress TM
- Egress traffic manager:
 - Manages a pool of egress queues in on-chip memory
 - Schedules traffic toward packet interfaces – Network, PCIe, internal hosts, and recycling (loopback)
 - Replicates multicast packets
- Fabric transmit and fabric receive:
 - Segments packet to cells
 - Transmit – Load-balances cells across fabric SerDes
 - Receive – Reassembles cells into packets
- Counter and meter engine pool:
 - Provides a general-purpose pool of counters and meters
 - Used for ingress and egress counting and metering
 - Includes configurable counting modes and criteria
 - Applied according to the packet processor command

Packet Processing

- Ingress receive packet processor:
 - Handles the main packet processing stage
 - Identifies incoming interface link layer, tunnel, PWE, and AC
 - Determines where to forward the packet based on packet forwarding header (L2, L3, MPLS, and so on)
 - Appends the packet processor (PP) header
 - Uses optionally expandable databases using an external lookup (ELK) interface
- Ingress transmit packet processor:
 - Edits the packet (or packet copy) before transmitting to egress PP
- Egress receive packet processor:
 - Filters packets according to various criteria
- Egress transmit packet processor:
 - Edits packets according to PP header (from ingress)

Interfaces

- Network interface:
 - 32 SerDes up to 53.125 Gb/s (PAM4)
 - 96 SerDes up to 25.78125 Gb/s
 - Port types supported include 1GbE, 10GbE, 25GbE, 40GbE, 50GbE, 100GbE, 200GbE, and 400GbE
 - Multiple Interlaken (ILKN) interfaces
- Statistics interface:
 - Sharing network interface SerDes
 - Use Ethernet ports, up to 200GbE per statistics interface
 - Generate statistics records over packets
- Fabric interface:
 - Up to 48 SerDes
 - Up to 53.125 Gb/s (PAM4) per link
 - Reed-Solomon Forward Error Correction (RS-FEC)
- External lookup interface:
 - Running over Interlaken interface, using NIF or fabric SerDes
- Out-of-band flow control (OOBFC):
 - Transmit and receive flow control
 - SPI 4.2 or ILKN protocol
- PCIe x4 lane Gen3 host interface:
 - Configuration and status register access
 - Packet transfer to and from host memory by using DMA
- BSC (NXP I2C-compatible) 2-line interface:
 - Basic device debug and register access (PCIe debug)
 - PCIe QSPI flash programming
 - Loading code used for heating when an industrial device is powered-up at a low ambient temperature

Chapter 2: System Configurations

2.1 Broadcom StrataDNX Components

The StrataDNX architecture includes two types of components: Fabric elements (FEs) and fabric access processors (FAPs):

- Fabric elements:
 - Are used to build a Clos fabric.
 - Provide self-routing, cell-based switching.
 - Support both single-stage and multistage fabric configurations.
- Fabric access processors:
 - Provide integrated fabric access and traffic management functions.
 - Provide an integrated packet processor.
 - May be used as a stand-alone TM (ingress and/or egress) or with a fabric.

2.1.1 Interoperation with StrataDNX Fabric Element Devices

The BCM88800 can interoperate in a system with the following FE devices:

- BCM88790
- BCM88770

2.1.2 Interoperation with StrataDNX Fabric Access Processor Devices


The BCM88800 can interoperate in a system with the following FAP devices:

- BCM88800
- BCM88690
- BCM88680
- BCM88670
- BCM88476

2.2 Stand-Alone Configurations

The BCM88800 may be used as a stand-alone switch device with ingress and egress intelligent TM. The integrated TM enables intelligent oversubscription with granular, per-flow (or per-customer-and-traffic-class) scheduling and shaping in the upstream direction, downstream direction, or both.

Figure 3: BCM88800 Switch with Integrated Traffic Manager



As shown in the preceding figure, the BCM88800 offers a flexible set of user interfaces, supporting a mix of 10GbE, 25GbE, 40GbE, 50GbE, 100GbE, 200GbE, and 400GbE Ethernet ports, and ILKN interfaces.

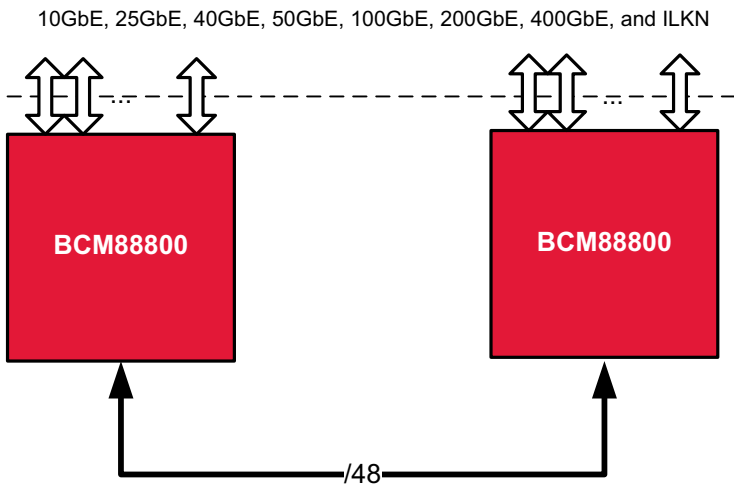
In the upstream direction from the network interfaces into the system, the integrated TM intelligently manages the oversubscription. In the downstream direction, the integrated TM can schedule and shape traffic, possibly per customer and application.

2.3 Mesh Fabric Configurations

The BCM88800 devices may be interconnected in a mesh configuration through the fabric links without the use of StrataDNX fabric elements. The mesh topology is used to create large-scale centralized (fabric-less) systems.

In Figure 4, two BCM88800 devices are connected in a mesh topology, with 48 links between the pair of devices. This configuration gives a net bandwidth of 2.16 Tb/s between the pair.

Figure 4: Two BCM88800 Devices in a Mesh Configuration



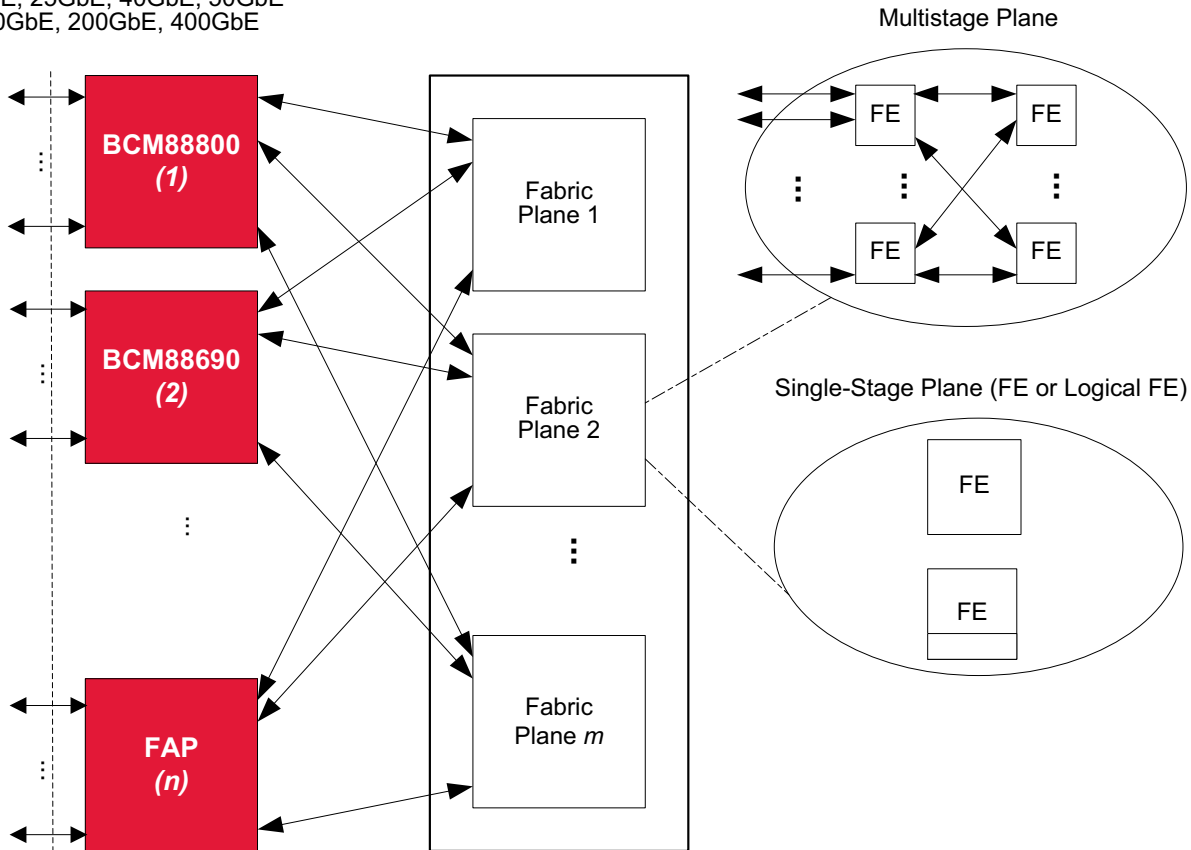
Only fabric SerDes are used for the mesh connectivity, keeping the full network interface available for front-panel connectivity, either in non-blocking or oversubscribed configurations.

2.4 StrataDNX CLOS Fabric Configurations

For medium to large nonblocking configurations, a StrataDNX fabric with fabric element devices (for example, the BCM88790) is used to interconnect the FAP devices. With the BCM88800, the StrataDNX fabric scales up to 4800 Tb/s.

Figure 5: FAPs Interconnected by StrataDNX Fabric

10GbE, 25GbE, 40GbE, 50GbE
100GbE, 200GbE, 400GbE



In the preceding figure, “n” FAP devices are interconnected through a Clos fabric composed of fabric elements. The number of interconnected FAP devices can range from 1 to 2048. A StrataDNX fabric based on the BCM88790 or BCM88770 accommodates a mix of BCM88800, BCM88690, BCM88680, and BCM88670 FAP devices, all of which are interconnected and interoperate simultaneously.

Chapter 3: System Interfaces

3.1 Network Interface

The BCM88800 network interface (NIF) includes four PM50, 24 PM25, and eight Interlaken (ILKN) cores.

NOTE: The BCM88800 can be connected to network interfaces with a total bandwidth of up to 4 Tb/s; however, the maximum device performance cannot exceed 3 Tb/s at any packet size.

Ethernet Port Macros (PMs)

PM50:

- Four Ethernet port macros supporting SerDes rates up to 53.125 Gb/s
- Each PM50 includes an octal SerDes (Blackhawk) supporting up to 53.125 Gb/s
- Each PM50 supports the following Ethernet configurations:
 - 1 × 400GbE port over eight lanes (PAM4)
 - 2 × 200GbE ports over four lanes (PAM4)
 - 4 × 100GbE ports over two lanes (PAM4)
 - 2 × 100GbE ports over four lanes
 - 8 × 50GbE ports over one lane (PAM4)
 - 4 × 50GbE ports over two lanes
 - 4 × 40GbE ports over two lanes
 - 2 × 40GbE ports over four lanes
 - 8 × 25GbE ports over one lane
 - 8 × 12GbE ports over one lane
 - 8 × 10GbE ports over one lane

PM25:

- 24 Ethernet Port Macros supporting SerDes rates up to 25.78125 Gb/s
- Each PM25 includes a quad SerDes (Falcon16) supporting up to 25.78125 Gb/s
- Each PM25 supports the following Ethernet configurations:
 - 1 × 100GbE port over four lanes
 - 2 × 50GbE ports over two lanes
 - 2 × 40GbE ports over two lanes
 - 1 × 40GbE port over four lanes
 - 4 × 25GbE ports over one lane
 - 4 × 12GbE ports over one lane
 - 4 × 10GbE ports over one lane
 - 4 × 1GbE ports over one lane

Interlaken Cores

The BCM88800 contains eight Interlaken cores, and each core can support one or two Interlaken interfaces. Two cores are connected to PM50 SerDes, and six cores are connected to PM25 SerDes.

When a single interface is in a core, the interface can support the following:

- Up to 24 lanes for rates up to 25.78125 Gb/s
- Up to 12 lanes over PM50 at 53.125 Gb/s

When two interfaces are in a core, each interface can support the following:

- Up to 12 lanes for rates up to 25.78125 Gb/s
- Up to 6 lanes over PM50 at 53.125 Gb/s (only one interface in a core can use PAM4 + FEC)

Interlaken interfaces that are connected to PM50 SerDes support optional ILKN-FEC.

The following table describes the mapping of PMs and Interlaken cores to SerDes and provides the interface naming conventions.

Table 1: BCM88800 NIF SerDes Interface Mapping

| PM Type and Number | SRD Number | PAM4 | ILKN Core | ILKN IF | Available for ELK | Available for STAT | BCM88802 BCM88804 (No HBM, No ILKN) BCM88806 (No ILKN) BCM88820 (Mesh) BCM88821 (No ELK, No FAB) BCM88823 (No FAB, No ILKN) 128 SRD | BCM88803 Supported Interfaces |
|--------------------|------------|------|-------------------|------------------|-------------------|--------------------|---|-------------------------------|
| PM50_00 | 000–007 | + | 6 | 12, 13 | + | — | Active | ELK 12 |
| PM50_01 | 008–015 | + | 6, 7 | 12–15 | + | — | Active | ELK 12 |
| PM50_02 | 016–023 | + | 6, 7 | 12–15 | + | — | Active | Blocked |
| PM50_03 | 024–031 | + | 7 | 14, 15 | — | + | Active | STAT |
| PM25_04 | 032–035 | — | 0 | 0, 1 | — | + | Active | MNG (ETH) |
| PM25_05 | 036–039 | — | 0 | 0, 1 | — | + | Active | Blocked |
| PM25_06 | 040–043 | — | 0, 1 ^a | 0–3 ^a | + | + | Active | ILKN 2,3 |
| PM25_07 | 044–047 | — | 0, 1 ^a | 0–3 ^a | + | + | Active | ILKN 2,3 |
| PM25_08 | 048–051 | — | 0, 1 | 0–3 | + | — | Active | ILKN 2,3 |
| PM25_09 | 052–055 | — | 0, 1 | 0–3 | + | — | Active | ILKN 2,3 |
| PM25_10 | 056–059 | — | 1, 2 ^b | 2–5 ^b | + | — | Active | ILKN 2,3 |
| PM25_11 | 060–063 | — | 1, 2 ^b | 2–5 ^b | + | — | Active | ILKN 2,3 |
| PM25_12 | 064–067 | — | 1, 2 ^c | 2–5 ^c | + | — | Active | Blocked |
| PM25_13 | 068–071 | — | 1, 2 ^c | 2–5 ^c | + | — | Active | Blocked |
| PM25_14 | 072–075 | — | 2, 3 | 4–7 | — | — | Active | ILKN 6,7 |
| PM25_15 | 076–079 | — | 2, 3 | 4–7 | — | — | Active | ILKN 6,7 |
| PM25_16 | 080–083 | — | 2, 3 ^d | 4–7 ^d | — | — | Active | ILKN 6,7 |
| PM25_17 | 084–087 | — | 2, 3 ^d | 4–7 ^d | — | — | Active | ILKN 6,7 |
| PM25_18 | 088–091 | — | 3 | 6, 7 | — | — | Active | ILKN 6,7 |
| PM25_19 | 092–095 | — | 3 | 6, 7 | — | — | Active | ILKN 6,7 |
| PM25_20 | 096–099 | — | 4 | 8, 9 | — | — | Active | Blocked |

Table 1: BCM88800 NIF SerDes Interface Mapping (Continued)

| PM Type and Number | SRD Number | PAM4 | ILKN Core | ILKN IF | Available for ELK | Available for STAT | BCM88802 BCM88804 (No HBM, No ILKN) BCM88806 (No ILKN) BCM88820 (Mesh) BCM88821 (No ELK, No FAB) BCM88823 (No FAB, No ILKN) 128 SRD | BCM88803 Supported Interfaces |
|--------------------|------------|------|-----------|---------|-------------------|--------------------|---|-------------------------------|
| PM25_21 | 100–103 | — | 4 | 8, 9 | — | — | Active | Blocked |
| PM25_22 | 104–107 | — | 4, 5 | 8–11 | — | — | Active | Blocked |
| PM25_23 | 108–111 | — | 4, 5 | 8–11 | — | — | Active | Blocked |
| PM25_24 | 112–115 | — | 4, 5 | 8–11 | — | — | Active | Blocked |
| PM25_25 | 116–119 | — | 4, 5 | 8–11 | — | — | Active | Blocked |
| PM25_26 | 120–123 | — | 5 | 10, 11 | — | — | Active | Blocked |
| PM25_27 | 124–127 | — | 5 | 10, 11 | — | — | Active | Blocked |

- a. If ILKN core_1 (interfaces 2 and 3) use SerDes from PM25_06 or PM25_07. This core cannot use SerDes from PM25_12 or PM25_13.
- b. If ILKN core_2 (interfaces 4 and 5) use SerDes from PM25_10 or PM25_11. This core cannot use SerDes from PM25_16 or PM25_17.
- c. If ILKN core_1 (interfaces 2 and 3) use SerDes from PM25_12 or PM25_13. This core cannot use SerDes from PM25_06 or PM25_07.
- d. If ILKN core_2 (interfaces 4 and 5) use SerDes from PM25_16 or PM25_17. This core cannot use SerDes from PM25_10 or PM25_11.

3.1.1 Ethernet Ports

Ethernet ports are implemented by the Ethernet port macro (PM). The PMs include the MAC and PCS layers of Ethernet ports. The BCM88800 includes two types of port macros. PM50 is based on an octal SerDes block (Blackhawk) that supports SerDes rates up to 53.125 Gb/s. PM25 is based on a quad SerDes block (Falcon16) that supports SerDes rates of up to 25.78125 Gb/s for standard Ethernet.

The following table depicts the Ethernet port modes supported by PM50 and PM25. Rates above 25.78125 Gb/s are supported only by PM50.

NOTE:

- RS272 FEC is a Broadcom-proprietary protocol to achieve lower latency.
- For 100G and 50G ports using RS-272 FEC mode, the FEC engine serves four SerDes lanes. If using RS-272, no other FEC mode can be used in the same FEC engine.
- BASE-R FEC indicates IEEE 802.3, Clause-74 compliant FEC.

Table 2: Supported Port Modes

| Port Speed | No. of Lanes | PM Type | Port Mode | SerDes Rate (Gb/s) | VCO Rate (GHz) | SerDes Mode | FEC | Related Standard |
|------------|--------------|-----------|---------------|--------------------|----------------|-------------|--------|--|
| 400GbE | 8 | PM50 only | 400GAUI-8 C2C | 53.125 | 26.5625 | PAM4 | RS-544 | IEEE 802.3bs Annex 120D |
| | | | 400GAUI-8 C2M | 53.125 | 26.5625 | PAM4 | RS-544 | IEEE 802.3bs Annex 120E |
| | | | 400GBASE-KR8 | 53.125 | 26.5625 | PAM4 | RS-544 | This is 400GAUI-8 with KR interface as defined in IEEE 802.3cd CL137 |
| | | | 400GBASE-CR8 | 53.125 | 26.5625 | PAM4 | RS-544 | This is 400GAUI-8 with CR interface as defined in IEEE 802.3cd CL136 |
| 200GbE | 4 | PM50 only | 200GAUI-4 C2C | 53.125 | 26.5625 | PAM4 | RS-544 | IEEE 802.3bs Annex 120D |
| | | | 200GAUI-4 C2M | 53.125 | 26.5625 | PAM4 | RS-544 | IEEE 802.3bs Annex 120E |
| | | | 200GBASE-KR4 | 53.125 | 26.5625 | PAM4 | RS-544 | IEEE 802.3cd CL137 |
| | | | 200GBASE-CR4 | 53.125 | 26.5625 | PAM4 | RS-544 | IEEE 802.3cd CL136 |

Table 2: Supported Port Modes (Continued)

| Port Speed | No. of Lanes | PM Type | Port Mode | SerDes Rate (Gb/s) | VCO Rate (GHz) | SerDes Mode | FEC | Related Standard |
|------------|--------------|---------------|---------------|--------------------|----------------|-------------|----------------|--|
| 100GbE | 2 | PM50 only | 100GAUI-2 C2C | 53.125 | 26.5625 | PAM4 | RS-544, RS-272 | IEEE 802.3cd Annex 135F |
| | | | 100GAUI-2 C2M | 53.125 | 26.5625 | PAM4 | RS-544, RS-272 | IEEE 802.3cd Annex 135G |
| | | | 100GBASE-KR2 | 53.125 | 26.5625 | PAM4 | RS-544, RS-272 | IEEE 802.3cd CL137 |
| | | | 100GBASE-CR2 | 53.125 | 26.5625 | PAM4 | RS-544, RS-272 | IEEE 802.3cd CL136 |
| | | | CAUI-2 C2C | 51.5625 | 25.78125 | PAM4 | RS-528 | IEEE 802.3.bj (CAUI-4) port, bit muxed to two lanes with C2C interface as defined in IEEE 802.3cd Annex 135F |
| | | | CAUI-2 C2M | 51.5625 | 25.78125 | PAM4 | RS-528 | IEEE 802.3.bj (CAUI-4) port, bit muxed to two lanes with C2M interface as defined in IEEE 802.3cd Annex 135G |
| 100GbE | 4 | PM50 and PM25 | CAUI-4 C2C | 25.78125 | 25.78125 | NRZ | RS-528, no FEC | IEEE 802.3bm Annex 83D |
| | | | CAUI-4 C2M | 25.78125 | 25.78125 | NRZ | RS-528, no FEC | IEEE 802.3bm Annex 83E |
| | | | 100GBASE-KR4 | 25.78125 | 25.78125 | NRZ | RS-528, no FEC | IEEE 802.3bj CL93 |
| | | | 100GBASE-CR4 | 25.78125 | 25.78125 | NRZ | RS-528, no FEC | IEEE 802.3bj CL92 |
| | | PM50 only | 100GAUI-4 C2C | 26.5625 | 26.5625 | NRZ | RS-544 | IEEE 802.3cd Annex 135D |
| | | | 100GAUI-4 C2M | 26.5625 | 26.5625 | NRZ | RS-544 | IEEE 802.3cd Annex 135E |
| 50GbE | 1 | PM50 only | 50GAUI-1 C2C | 53.125 | 26.5625 | PAM4 | RS-544, RS-272 | IEEE 802.3cd Annex 135F |
| | | | 50GAUI-1 C2M | 53.125 | 26.5625 | PAM4 | RS-544, RS-272 | IEEE 802.3cd Annex 135G |
| | | | 50GBASE-KR | 53.125 | 26.5625 | PAM4 | RS-544, RS-272 | IEEE 802.3cd CL137 |
| | | | 50GBASE-CR | 53.125 | 26.5625 | PAM4 | RS-544, RS-272 | IEEE 802.3cd CL136 |

Table 2: Supported Port Modes (Continued)

| Port Speed | No. of Lanes | PM Type | Port Mode | SerDes Rate (Gb/s) | VCO Rate (GHz) | SerDes Mode | FEC | Related Standard |
|------------|--------------|---------------|-----------------------------|--------------------|----------------|-------------|------------------------|---|
| 50GbE | 2 | PM50 and PM25 | Consortium 50G C2C | 25.78125 | 25.78125 | NRZ | RS-528, no FEC | Consortium 50GE with C2C interface as defined in IEEE 802.3bm Annex 83D |
| | | | Consortium 50G C2M | 25.78125 | 25.78125 | NRZ | RS-528, no FEC | Consortium 50GE with C2M interface as defined in IEEE 802.3bm Annex 83E |
| | | | LAUI-2 C2C | 25.78125 | 25.78125 | NRZ | RS-528, no FEC | IEEE 802.3cd Annex 135B |
| | | | LAUI-2 C2M | 25.78125 | 25.78125 | NRZ | RS-528, no FEC | IEEE 802.3cd Annex 135C |
| | | | Consortium 50G KR2 | 25.78125 | 25.78125 | NRZ | RS-528, no FEC | Consortium 50GbE with KR interface as defined in IEEE 802.3bj CL93 |
| | | | Consortium 50G CR2 | 25.78125 | 25.78125 | NRZ | RS-528, no FEC | Consortium 50GbE with CR interface as defined in IEEE 802.3bj CL92 |
| | | PM50 only | 50GAUI-2 C2C | 26.5625 | 26.5625 | NRZ | RS-544 | IEEE 802.3cd Annex 135D |
| | | | 50GAUI-2 C2M | 26.5625 | 26.5625 | NRZ | RS-544 | IEEE 802.3cd Annex 135E |
| 40GbE | 2 | PM50 and PM25 | XLAUI-2 KR | 20.625 | 20.625 | NRZ | No FEC | XLAUI-4, bit muxed to two lanes, with KR interface as defined in 802.3bj CL93 (masks scaled to 20.625G) |
| | | | XLAUI-2 CR | 20.625 | 20.625 | NRZ | No FEC | XLAUI-4, bit muxed to two lanes, with CR interface as defined in 802.3bj CL92 (masks scaled to 20.625G) |
| | | | XLAUI-2 C2C | 20.625 | 20.625 | NRZ | No FEC | XLAUI-4, bit muxed to two lanes, with C2C interface as defined in 802.3bm Annex 83D (masks scaled to 20.625G) |
| | | | XLAUI-2 C2M | 20.625 | 20.625 | NRZ | No FEC | XLAUI-4, bit muxed to two lanes, with C2M interface as defined in 802.3bm Annex 83E (masks scaled to 20.625G) |
| 40GbE | 4 | PM50 and PM25 | XLAUI | 10.3125 | 20.625 | NRZ | BASE-R, no FEC | IEEE 802.3 Annex 83A |
| | | | 40GBASE-KR4 | 10.3125 | 20.625 | NRZ | BASE-R, no FEC | IEEE 802.3 CL84 |
| | | | 40GBASE-CR4 | 10.3125 | 20.625 | NRZ | BASE-R, no FEC | IEEE 802.3 CL85 |
| | | | XLPP1 | 10.3125 | 20.625 | NRZ | BASE-R, no FEC | IEEE 802.3 Annex 86A |
| 25GbE | 1 | PM50 and PM25 | 25GAUI C2C | 25.78125 | 25.78125 | NRZ | RS-528, BASE-R, no FEC | IEEE 802.3by Annex 109A |
| | | | 25GAUI C2M | 25.78125 | 25.78125 | NRZ | RS-528, BASE-R, no FEC | IEEE 802.3by Annex 109B |
| | | | 25GBASE-KR/ 25GBASE-KR-S | 25.78125 | 25.78125 | NRZ | RS-528, BASE-R, no FEC | IEEE 802.3by CL111 |
| | | | 25GBASE-CR/ 25GBASE-CR-S | 25.78125 | 25.78125 | NRZ | RS-528, BASE-R, no FEC | IEEE 802.3by CL110 |

Table 2: Supported Port Modes (Continued)

| Port Speed | No. of Lanes | PM Type | Port Mode | SerDes Rate (Gb/s) | VCO Rate (GHz) | SerDes Mode | FEC | Related Standard |
|------------|--------------|---------------|-------------------------------|--------------------|--------------------------------------|-------------|----------------|---------------------------------|
| 12GbE | 1 | PM50 and PM25 | XFI (Scaled) | 12.5 | 25 | NRZ | No FEC | XFI+ (FC-PI-3), scaled to 12.5G |
| 10GbE | 1 | PM50 and PM25 | 10GBASE-KR | 10.3125 | PM50: 20.625 PM25: 25.78125 | NRZ | BASE-R, no FEC | IEEE 802.3 CL72 |
| | | | XFI | 10.3125 | PM50: 20.625 PM25: 25.78125 | NRZ | BASE-R, no FEC | XFI+ (FC-PI-3) |
| | | | SFI | 10.3125 | PM50: 20.625 PM25: 25.78125 | NRZ | BASE-R, no FEC | SFF-8431 |
| | | | Direct Attached Cable (DAC) | 10.3125 | PM50: 20.625 PM25: 25.78125 | NRZ | BASE-R, no FEC | — |
| 1GbE | 1 | PM25 only | 1000BASE-X, SGMII (1GbE only) | 1.25 | 25.78125 | NRZ | No FEC | — |

3.1.1.1 400GbE Port

The BCM88800 device supports standard 400GbE according to IEEE 802.3bs. This port type is supported on PM50 only (not supported on PM25).

The 400GbE port supports the following interfaces:

- 400GAUI-8 (chip-to-chip IEEE 802.3bs Annex 120D and chip-to-module IEEE 802.3bs Annex 120E)
- 400GBASE-KR8 (This is 400GAUI-8 with a KR interface as defined in IEEE 802.3cd CL137.)
- 400GBASE-CR8 (This is 400GAUI-8 with a CR interface as defined in IEEE 802.3cd CL136.)

When using the 400GbE port, Reed-Solomon FEC is supported. This FEC is compliant with IEEE 802.3bs Clause 119, RS(544,514).

3.1.1.2 200GbE Port

The BCM88800 device supports standard 200GbE according to IEEE 802.3bs and IEEE 802.3cd. This port type is supported on PM50 only (not supported on PM25).

The 200GbE port supports the following interfaces:

- 200GAUI-4 (chip-to-chip IEEE 802.3bs Annex 120D and chip-to-module IEEE 802.3bs Annex 120E)
- 200GBASE-KR4 (IEEE 802.3cd CL137)
- 200GBASE-CR4 (IEEE 802.3cd CL136)

When using the 200GbE port, an optional Reed-Solomon FEC is supported. This FEC is compliant with IEEE 802.3bs Clause 119, RS(544,514).

3.1.1.3 100GbE Port

3.1.1.3.1 100GbE over Two Lanes

The BCM88800 device supports standard 100GbE according to IEEE 802.3cd. This port type is supported on PM50 only (not supported on PM25).

The 100GbE port supports the following interfaces:

- 100GAUI-2 (chip-to-chip IEEE 802.3cd Annex 135F and chip-to-module IEEE 802.3cd Annex 135G)
- 100GBASE-KR2 (IEEE 802.3cd CL137)
- 100GBASE-CR2 (IEEE 802.3cd CL136)
- CAUI-2, which is the IEEE 802.3.bj (CAUI-4) port bit muxed to two lanes (chip-to-chip and chip-to-module).

When using the 100GbE port, an optional Reed-Solomon FEC is supported. This FEC is compliant with IEEE 802.3 Clause 91, both RS(528,514) and RS(544,514).

3.1.1.3.2 100GbE over Four Lanes

The BCM88800 device supports standard 100GbE according to IEEE 802.3. This port type is supported on both PM50 and PM25.

The 100GbE port supports the following interfaces:

- CAUI-4 (chip-to-chip IEEE 802.3bm Annex 83D and chip-to-module IEEE 802.3bm Annex 83E)
- 100GBASE-KR4 (IEEE 802.3bj CL93)
- 100GBASE-CR4 (IEEE 802.3bj CL92)
- 100GAUI-4 (chip-to-chip IEEE 802.3cd Annex 135D and chip-to-module IEEE 802.3cd Annex 135E)

When using the 100GbE port, an optional Reed-Solomon FEC is supported. This FEC is compliant with IEEE 802.3 Clause 91, both RS(528,514) and RS(544,514).

3.1.1.4 50GbE Port

3.1.1.4.1 50GbE over One Lane

The BCM88800 supports 50GbE according to IEEE 8023.cd. This port type is supported on PM50 only (not supported on PM25).

The 50GbE port supports the following interfaces:

- 50GAUI-1 (chip-to-chip IEEE 802.3cd Annex 135F and chip-to-module IEEE 802.3cd Annex 135G)
- 50GBASE-KR (IEEE 802.3cd CL137)
- 50GBASE-CR (IEEE 802.3cd CL136)

When using the 50GbE port, an optional Reed-Solomon FEC is supported. This FEC is compliant with IEEE 802.3 Clause 134, RS(544,514).

3.1.1.4.2 50GbE over Two Lanes

The BCM88800 supports 50GbE according to IEEE 802.3cd and the 25G/50G Ethernet consortium. This port type is supported on both PM50 and PM25.

The 50GbE port supports the following interfaces:

- 50GbE over two lanes according to the 25G/50G Ethernet consortium
- LAUI-2 (chip-to-chip IEEE 802.3cd Annex 135B and chip-to-module IEEE 802.3cd Annex 135C)
- 50GBASE-KR2 (This is Ethernet consortium 50GbE with a KR interface as defined in IEEE 802.3bj CL93.)
- 50GBASE-CR2 (This is Ethernet consortium 50GbE with a CR interface as defined in IEEE 802.3bj CL92.)
- 50GAUI-2 (chip-to-chip IEEE 802.3cd Annex 135D and chip-to-module IEEE 802.3cd Annex 135E)

When using the 50GbE port, an optional Reed-Solomon FEC is supported. This FEC is compliant with IEEE 802.3 Clause 134, RS(544,514), or Clause 91, RS(528,514). PM25 supports only Clause 91, RS(528,514).

3.1.1.5 40GbE Port

3.1.1.5.1 40GbE over Two Lanes

The BCM88800 supports standard 40GbE according to IEEE 802.3ba. This port type is supported on both PM50 and PM25.

The 40GbE ports support the following interfaces:

- XLAUI-2 – This is XLAUI-4, bit-muxed to two lanes, with a KR interface as defined in 802.3bj CL93 (masks scaled to 20.625G).
- XLAUI-2 – This is XLAUI-4, bit-muxed to two lanes, with a CR interface as defined in 802.3bj CL92 (masks scaled to 20.625G).
- XLAUI-2 – This is XLAUI-4, bit-muxed to two lanes, with a chip-to-chip interface as defined in 802.3bm Annex 83D (masks scaled to 20.625G).
- XLAUI-2 – This is XLAUI-4, bit-muxed to two lanes, with a chip-to-module interface as defined in 802.3bm Annex 83E (masks scaled to 20.625G).

No FEC is supported for the 40GbE port over two lanes.

3.1.1.5.2 40GbE over Four Lanes

The BCM88800 supports standard 40GbE according to IEEE 802.3. This port type is supported on both PM50 and PM25.

The 40GbE ports support the following interfaces:

- XLAUI (IEEE 802.3 Annex 83A)
- 40GBASE-KR4 (IEEE 802.3 CL84)
- 40GBASE-CR4 (IEEE 802.3 CL85)
- XLPPI (IEEE 802.3 Annex 86A)

When using the 40GbE port in four-lane mode, an optional FEC is supported. The FEC is compliant with IEEE 802.3 Clause 74.

3.1.1.6 25GbE Port

The BCM88800 supports 25GbE according to IEEE 802.3by, and the 25G/50G Ethernet consortium. This port type is supported on both PM50 and PM25.

The 25GbE ports support the following interfaces:

- 25GAUI (chip-to-chip IEEE 802.3by Annex 109A and chip-to-module IEEE 802.3by Annex 109B)
- 25GBASE-KR/25GBASE-KR-S (IEEE 802.3by CL111)
- 25GBASE-CR/25GBASE-CR-S (IEEE 802.3by CL110)

When using the 25GbE port, an optional FEC is supported. The FEC can be either IEEE 802.3 Clause 108, RS(528,514) or IEEE 802.3 Clause 74 FEC.

3.1.1.7 12GbE Port

The BCM88800 supports 12GbE. The 12GbE port is similar to the 10GbE port, with a SerDes rate of 12.5G (scaled up from 10.3125G). This port type is supported on both PM50 and PM25.

The 12GbE port supports XFI+ (FC-PI-3), and the SerDes rate is scaled up to 12.5G from 10.3125G.

The 12GbE port does not support FEC.

3.1.1.8 10GbE Port

The BCM88800 supports 10GbE according to IEEE 802.3. This port type is supported on both PM50 and PM25.

The 10GbE ports support the following interfaces:

- XFI or SFI for direct connect to optical module¹
- 10GBASE-KR (IEEE 802.3 CL72)
- Direct-attached cable (DAC)

When using the 10GbE port with the KR interface, an optional FEC is supported. The FEC is compliant with IEEE 802.3 Clause 74.

3.1.1.9 1GbE Port

The BCM88800 supports 1GbE over an SGMII or 1000BASE-X interface.

- The GbE port is supported over the PM25 port macro only.
- The GbE port supports 1GbE only (no support for 10-Mb/s and 100-Mb/s speeds).
- The GbE port does not support auto-negotiation.
- The GbE port does not support clock recovery for SyncE.
- Half-duplex mode is not supported.
- No support is available for a 2.5GbE port.

1. Supported: SR (limiting), LR (limiting), ER (limiting).
Not supported: ZR (limiting), ZR (linear), LRM (linear), DWDM (linear).

3.1.1.10 Mixed Port Types on PM50

Different port types can be supported on the same Blackhawk port macro (PM50). Mixing port types is allowed when all ports in the same PM50 are derived from one or two PLL rates. The following table defines the PLL (VCO) combinations and available ETH ports on each combination.

NOTE: For Ethernet ports, TVCO (PLL1) and OCVO (PLL0) do not have the same functionality. (See [Section 3.1.1.10.1, PM50 Mixed Port-Type Limitations.](#))

Table 3: PM50 PLL VCO Combinations and Supported Port Types

| | | | | TVCO | 25.78125G | 25.78125G | 25.78125G | 26.5625G | 26.5625G | 26.5625G |
|-----------------|----------------------|-------|----------|------|-----------|-----------|-----------|----------|----------|-----------|
| | | | | OVCO | 20.625G | 25.000G | 26.5625G | 20.625G | 25.000G | 25.78125G |
| VCO Rate (Gb/s) | Port BW | Lanes | SRD Rate | | | | | | | |
| 26.5625 | 400GbE | 8 | 53.1250 | | | | + | + | + | + |
| 26.5625 | 200GbE | 4 | 53.1250 | | | | + | + | + | + |
| 26.5625 | 100GbE | 2 | 53.1250 | | | | + | + | + | + |
| 26.5625 | 100GbE | 4 | 26.5625 | | | | + | + | + | + |
| 26.5625 | 50GbE | 1 | 53.1250 | | | | + | + | + | + |
| 26.5625 | 50GbE | 2 | 26.5625 | | | | + | + | + | + |
| 25.78125 | 100GbE | 2 | 51.5625 | | + | + | + | | | + |
| 25.78125 | 100GbE | 4 | 25.78125 | | + | + | + | | | + |
| 25.78125 | 50GbE | 1 | 51.5625 | | + | + | + | | | + |
| 25.78125 | 50GbE | 2 | 25.78125 | | + | + | + | | | + |
| 25.78125 | 25GbE | 1 | 25.78125 | | + | + | + | | | + |
| 25.000 | 12.5GbE ^a | 1 | 12.5000 | | | + | | | + | |
| 20.6250 | 40GbE | 2 | 20.6250 | | + | | | + | | |
| 20.6250 | 40GbE | 4 | 10.3125 | | + | | | + | | |
| 20.6250 | 10GbE | 1 | 10.3125 | | + | | | + | | |

a. 12.5GbE data bandwidth is 12.12G. This is a non-standard ETH rate.

3.1.1.10.1 PM50 Mixed Port-Type Limitations

The main limitations of the PM50 are as follows:

- Up to eight lanes can be supported.
- Up to two VCO rates can be supported. An ETH application on PM50 requires a TVCO rate of 25.78125 Gb/s or 26.5625 Gb/s. This requirement prevents using a combination of 20.625 Gb/s and 25.0000 Gb/s PLL rates if ETH is used.
- Any modification of TVCO causes ETH ports based on OVCO to be reset as well.

3.1.1.10.2 Software Sequence and Rules for Port Allocation

This section describes the software sequence and rules for port allocation. The goal is to provide a better understanding of the device options. For more information and updates regarding PM50 mixed ports, refer to the *BCM88690 Traffic Manager Programming Guide* (88690-PG2xx).

Ports can be allocated to the PM50 by a single port assignment or by a multi-port assignment using the following APIs:

- `bcm_port_resource_set()`
- `bcm_port_resource_multi_set()`

Ports are allocated only if the new configuration meets the TVCO and OVCO combinations in [Table 3](#) and the following software limitations are met.

For the initial allocation (there are no active ports on the PM), VCO rates are defined according to the following table.

Table 4: Initialization Options Table

| Requested PLL Rate | Port Requirement for PM50 Initialization | TVCO | OVCO |
|--------------------------------------|--|------------------------|-----------|
| Initialization with One Rate | | | |
| 20.625G | At least one ETH port | 25.78125G ^a | 20.625G |
| 20.625G | ILKN only (no ETH) | Shut-down | 20.625G |
| 25.000G | At least one ETH port | 25.78125G ^a | 25.000G |
| 25.000G | ILKN only (no ETH) | Shut-down | 25.000G |
| 25.78125G | No restrictions | 25.78125G | Shut-down |
| 26.5625G | No restrictions | 26.5625G | Shut-down |
| Initialization with Two Rates | | | |
| 25.000G + 20.625G ^b | ILKN only (no ETH) | 20.625G | 25.000G |
| 25.78125G + 20.625G | No restrictions | 25.78125G | 20.625G |
| 25.78125G + 25.000G | No restrictions | 25.78125G | 25.000G |
| 26.5625G + 20.625G | No restrictions | 26.5625G | 20.625G |
| 26.5625G + 25.000G | No restrictions | 26.5625G | 25.000G |
| 26.5625G + 25.78125G | No restrictions | 26.5625G | 25.78125G |

a. The default rate is 25.78125G for the dynamic addition of 25.78125G ports. A software flag can configure the TVCO to 26.5625G if that rate is expected.

b. This combination cannot work with ETH. It works only with ILKN.

For an allocation update (there are active ports on the PM), the following information applies:

- If all ports in the PM are included in the update command, software treats the update as if this is the initial allocation.
- If only some of the PM ports are included in the update command, software tries to perform the PLL update without impacting used resources.
- To update TVCO, all ETH ports using TVCO and OVCO, as well as all ILKN ports using TVCO, should be included in the update command. All of these ports will go through the port-down and port-up sequence. (Only ILKN ports using OVCO are not affected.)
- To update OVCO, all ETH ports using OVCO and all ILKN ports using OVCO should be included in the update command. All of these ports will go through the port-down and port-up sequence. (Ports using TVCO are not affected.)

3.1.1.10.3 PM50 Mixed Port Type Allocation Restrictions

The scenarios in this section show examples of configuration restrictions.

Scenario 1

If a PM is configured using a 40GbE port (four lanes, VCO = 20.6250G) and a 50GbE port (two lanes, VCO = 25.78125G), a dynamic configuration adding a 100GbE port (four lanes, VCO = 25.78125G) results in an error indication because more than eight lanes are required.

Scenario 2

If a PM is configured using a 40GbE port (four lanes, VCO = 20.6250G) and a 50GbE port (two lanes, VCO = 26.5625G), a dynamic configuration that adds a 50GbE port (two lanes, VCO = 25.78125G) results in an error indication because three VCO rates are required.

Scenario 3

Static configuration of a 1 × 200GbE port (four lanes, VCO = 26.5625G), and 4 × 10GbE ports (four lanes, VCO = 20.6250G) results in {PLL1 = 26.5625G, PLL0 = 20.6250G}.

A dynamic configuration that replaces the 200GbE port with a 100GbE port (four lanes, VCO = 25.78125G) results in an error indication because PLL1 should be updated, causing a port-down event on the 10GbE ports as well.

If it is acceptable to allow the 10GbE ports to go down and back up, they should be added to the allocation update (by using `bcm_port_resource_multi_set`).

3.1.1.11 Mixed Port Types on PM25

The 25GbE (25.78125G), 10GbE (10.3125G), and 1GbE (1.25G) port types can coexist simultaneously on the same PM25. Any combination of other rates is not allowed.

For example, when using 12GbE (12.5G), other rates can not be used on the same PM25.

NOTE: If PM25 is used for ETH, it cannot be used for ILKN.

3.1.2 Interlaken

The BCM88800 has eight Interlaken cores. Each Interlaken core can be configured to operate as either a single interface or as two interfaces.

- ILKN cores 0 through 5 are connected to PM25 and support a set of defined rates in the range 10.3125 Gb/s to 25.78125 Gb/s.
- ILKN cores 6 and 7 are connected to PM50 and support a set of defined rates in the range 10.3125 Gb/s to 53.125 Gb/s.

For the defined ILKN rates, see [Section 3.1.2.3, Interlaken SerDes Supported Rates and Electrical Standards](#).

ILKN core 6 and ILKN core 7 are connected to 50G SerDes and support optional RS-FEC as defined in “Interlaken Reed-Solomon Forward Error Correction Extension” protocol definition by the Interlaken Alliance.

For information about FEC usage, see [Section 3.1.2.4, Interlaken FEC](#).

Each Interlaken core can be configured to operate as either a single interface (up to 24 lanes) or as two interfaces (each up to 12 lanes).

When only the even-numbered interface (0, 2, and so on) in the core is used, the interface can support the following:

- Up to 24 lanes for rates up to 25.78125 Gb/s
- Up to 12 lanes over PM50 at 53.125 Gb/s

When the odd-numbered interface in a core is used, each interface (the even or the odd) can support the following:

- Up to 12 lanes for rates up to 25.78125 Gb/s
- Up to 6 lanes over PM50 at 53.125 Gb/s (only one interface in a core can use PAM4 + FEC)

3.1.2.1 Interlaken Lane Selection

Each ILKN core can use SerDes from a predefined group, as described in [Table 1](#). In some cases, the same SerDes can be mapped to one of two ILKN cores. This flexible approach implies some limitations, which include the following:

- For limitations of PM assignment to ILKN cores, see [Table 1](#).
- If the RX (or TX) of a specific physical SerDes is used for ILKN, the TX (or RX) of the same SerDes should go to the same ILKN interface.
- For a logical ILKN lane, RX SerDes and TX SerDes should be on the same PM.
- The Interlaken specification indicates that an interface has one FEC instance for every two adjacent lanes. If an ILKN interface uses FEC, follow the information in *Hardware Design Guidelines for StrataDNX 16-nm Devices* (DNX16-AN1xx), Section 2.4.4: Blackhawk Lane Mapping – BCM88800 ILKN, and Section 3.4.2: Falcon16 Lane Mapping – BCM88800 ILKN, to fully understand the Interlaken specification requirements and options.

NOTE: If PM25 is used for ILKN, it cannot be used for ETH or STAT.

3.1.2.2 Interlaken over Fabric SerDes

ILKN core 6 supports the option to use fabric lanes as its SerDes.

Fabric SerDes 0 to 15 can be mapped for use by ILKN core 6. When mapped this way, these SerDes can be used for a single ILKN interface or for two interfaces. Usage of one interface from fabric, and one interface from NIF is supported. However, a mix of SerDes from fabric and SerDes from NIF for the same interface is not supported.

NOTE: When ILKN core 6 uses fabric SRD[i], NIF SRD[i] cannot be used for ILKN.

3.1.2.3 Interlaken SerDes Supported Rates and Electrical Standards

The ILKN interface is targeted to comply with the electrical specifications of the standards as listed in the following table.

Table 5: ILKN Supported Rates and Electrical Specifications Standards

| SerDes Rate (Gb/s) | SerDes Mode | Standard Electrical Specifications |
|--------------------|-------------|---|
| 53.125 | PAM4 | 50GBASE-KR IEEE 802.3cd CL137 |
| | | 200GAUI-4 C2C: IEEE 802.3bs Annex 120D |
| | | 200GAUI-4 C2M: IEEE 802.3bs Annex 120E |
| 25.78125 | NRZ | 25GBASE-KR-S IEEE 802.3by CL111 |
| | | CAUI-4 C2C: IEEE 802.3bm Annex 83D |
| | | CAUI-4 C2M: IEEE 802.3bm Annex 83E |
| 25.0 | NRZ | 25GBASE-KR-S IEEE 802.3by CL111 ^a |
| | | CAUI-4 C2C: IEEE 802.3bm Annex 83D ^a |
| 20.625 | NRZ | 25GBASE-KR-S IEEE 802.3by CL111 ^a |
| 12.5 | NRZ | 10G XFI: XFI+ (FC-PI-3) ^a |
| 10.3125 | NRZ | 10GBASE-KR IEEE 802.3 CL72 |
| | | 10G XFI: XFI+ (FC-PI-3) |

a. Scaled to the appropriate rate.

3.1.2.4 Interlaken FEC

ILKN cores 6 and 7, which support PAM4 (50G), also support Interlaken Reed-Solomon Forward Error Correction (RS-FEC), which is RS-FEC(544,514).

The FEC complies with the FEC extension of the *Interlaken Reed-Solomon Forward Error Correction Extension Protocol Definition* from the Interlaken Alliance.

The following ILKN FEC rules apply:

- For an ILKN interface using SerDes in NRZ mode, FEC is optional.
- For an ILKN interface using SerDes in PAM4 mode, FEC is mandatory.

The following guidelines apply to FEC usage and SerDes rates:

- Do not use FEC for 25.78G.
- Use only a single interface per core with the 53G SerDes rate.

NOTE: Only one ILKN interface in a core can support FEC. For more information, refer to EID#8025 in the *BCM88800 Device Errata* (88800-ER1xx).

3.1.2.5 Interlaken Receive Burst Rules

Packets are transmitted across the Interlaken interface in bursts. The BCM88800 requires that burst sizes and burst intervals (Interlaken BurstShort) conform to certain rules.

- The BCM88800 Interlaken RX supports both burst interleaving and full-packet mode. If the peer device is configured to burst interleaving, assign a separate ingress reassembly context for each active channel.
- If the peer device BurstMax is 256B and burst interleaving is enabled, BCM88800 performance degradation may occur due to sending data words shorter than 512B to the ingress pipe.

Receive (into the BCM88800) burst interleaving rules are as follows:

- Supported BurstMax sizes are 256B or 512B. If the peer device can work with 512B BurstMax, use this configuration.
- Start-of-packet (SOP) bursts should be 192B or larger (in 64B increments).
- Bursts that are not end-of-packet (EOP) bursts may be either 128B, 192B, 256B, up to 512B in 64B increments.
- EOP bursts may be any size (up to the configured BurstMax).
- BurstShort (the minimum interval between burst control words) should be at least 32B.

3.1.2.6 Interlaken Transmit Burst Rules

BCM88800 ILKN TX supports only full-packet mode (burst-interleaving occurs only when TDM packets preempt data packets).

When transmitting across an Interlaken interface, the BCM88800 supports configurations consistent with the following:

- BurstShort (minimum interval between burst control words) – Configurable in increments of 32B from 32B to 256B.
- BurstMax options:
 - BurstMax = 256B, normal scheduling:
 - Non-EOP bursts are 256B.
 - EOP burst can be from 1B to 256B.
 - BurstMax = 256B, enhanced scheduling:
 - Configurable BurstMin of 64B or 128B.
 - Non-EOP bursts are 256B.
 - Penultimate burst when BurstMin is 128B: 128B or 256B.
 - Penultimate burst when BurstMin is 64B: 192B or 256B.
 - EOP bursts are from 64B to 256B.
 - BurstMax = 512B, normal scheduling:
 - Non-EOP bursts are 512B.
 - EOP burst can be 1B to 512B.
 - BurstMax = 512B, enhanced scheduling:
 - Configurable BurstMin of 128B or 256B.
 - Non-EOP bursts are 512B.
 - Penultimate burst when BurstMin is 256B: 256B or 512B.
 - Penultimate burst when BurstMin is 128B: 384B or 512B.
 - EOP bursts are from 64B to 512B.

3.1.2.7 Interlaken In-Band Flow Control

Each Interlaken interface supports in-band flow control according to the Interlaken protocol definition.

The flow-control (FC) information is carried over the control words sent across the interface. Each control word includes a 16-bit flow-control field and a reset calendar bit. Each flow-control bit indicates the flow-control status (ON/OFF) of a specific Interlaken calendar channel. The BCM88800 supports a configurable calendar length between 16 to 256 channels (16, 32, 64, 128, and 256).

The BCM88800 supports two options of link level flow control (LLFC) mapping:

- Map the LLFC to flow control calendar channel #0.
- Map the LLFC to the first flow-control calendar channel in every control word (that is, entries 0, 16, 32, and so on). This option is useful for faster reaction times for LLFC, at the expense of flow-control channels.

The BCM88800 is flexible in flow-control processing, and each calendar entry can be mapped to one of the following flow-control reaction points:

- Link level (mapped to an Interlaken NIF port), traffic stop at burst boundary
- Channel level (mapped to an OTM port), traffic stop on packet boundary
- Egress queue pair level

Flow-control generation (flow-control transmission): The BCM88800 can assign each calendar entry with flow control information or indications that may represent the following:

- Status of Interlaken port receive buffer (useful for link level).
- Status of the global resources (in other words, DRAM buffers).
- Status of a virtual-statistics queue: May be programmed to generate flow control per Interlaken channel but enables more fine-grain flow-control indications; for example, per port and traffic class, per flow, and so on.

3.2 External Lookup Engine

The packet processing engine uses various types of databases. To address the requirement of some applications for greater database capacity, the BCM88800 supports expansion of some databases with the use of an external lookup (ELK) interface.

NOTE: The BCM88800 does not support MAC table expansion over the ELK interface.

The BCM88800 ELK interface connects directly to Broadcom knowledge-based processor (KBP) devices.

The ELK interface has the following guidelines and characteristics:

- The ELK interface can use either ILKN_2 from core 1 or ILKN_12 from core 6 (when using core 6, either NIF SerDes or fabric SerDes).
 - If ILKN_2 is used for ELK, ILKN_3 cannot be used.
 - If ILKN_12 is used for ELK, ILKN_13 cannot be used.
- When using NRZ 25.78125G mode, up to 16 lanes are available for the ELK interface.
- When using PAM4 53.125G mode (supported by core 6 only), up to 12 lanes are available for the ELK interface.

3.3 Out-of-Band Flow Control

The BCM88800 supports out-of-band flow control (OOBFC) from the user to the BCM88800 (egress flow control or reception) and from the BCM88800 to the user (ingress flow control or generation).

The BCM88800 has two independent bidirectional OOBFC interfaces. Each interface can work in SPI-4.2 mode or Interlaken mode.

In all modes, the OOBFC interface supports the following:

- Reception – Each calendar entry is mapped to a flow-control reaction point that is either:
 - Link-level (NIF port)
 - Channel-level (OTM port)
 - Priority-level (egress queue pair)
- Generation (flow-control transmission) – Each calendar entry represents a flow-control generation point that is either:
 - Status of a virtual-statistics queue
 - Status of the global resources, in other words, DRAM buffers
 - Status of the port (MAC) receive buffer (link level)
- XON/XOFF signaling (no concept of credits, such as SPI-4.2)

In SPI-4.2 mode, the OOBFC interface has the following characteristics:

- SPI-4.2 status channel-like framing
- Calendar length of up to 512 entries
- Three-wire signaling for each direction – Two data, one clock (support XON/XOFF indications):
 - On TX STARVING = XON and SATISFIED = XOFF
 - On RX XON = (HUNGRY or STARVING) and XOFF = SATISFIED
- SDR operation for data signals
- Transmission clock rate: Core frequency divided by 6, 8, 10, 12, 14, or 16 (maximum rate is 166 MHz)
- Reception clock rate: DC to 200 MHz

In Interlaken mode, the OOBFC interface has the following characteristics:

- Interlaken protocol
- Calendar length of up to 512 entries
- Reception and generation of Interlaken retransmit requests
- Three-wire signaling for each direction – One each for clock, data, and sync
- 4b CRC protection
- DDR operation for data and sync signals
- Transmission clock rate: Core frequency divided by 6, 8, 10, 12, 14, or 16 (maximum rate is 166 MHz)
- Reception clock rate: Up to 180 MHz

NOTE: In-band and out-of-band flow-control signaling may be used simultaneously.

3.4 Synchronous Ethernet

The BCM88800 supports Synchronous Ethernet (SyncE) applications. The support includes two functions: controlling the transmit clock of network ports, and recovering a clock from a network port.

3.4.1 Transmit Clock

The transmit clock of each SerDes is locked to one of the NIF_[4:0]_REFCLK_P/N input reference clocks (for mapping information, see the NIF_[4:0]_REFCLK_P/N description in [Section 4.3, Pin Description – Grouped by Function](#)).

By connecting the system transmit clock to the NIF_[3:0]_REFCLK_P/N inputs, it is possible to control the transmit clock of the SerDes.

3.4.2 Recovered Clock

The BCM88800 provides up to two recovered clocks that may be used as reference clocks for an external synchronization unit. The source for the recovered clocks can be any of the NIF SerDes configured as an Ethernet port. The BCM88800 drives two recovered clocks, each as a differential signal. Two additional pads are used for the valid indication, one per recovered clock.

NOTE: Support for clock recovery is as follows:

- Clock recovery is supported for a NIF interface configured as an Ethernet port.
- Clock recovery is not supported for a 1G Ethernet interface.

The valid indication is asserted when the recovered clock is synchronized to the selected SerDes RX signal and can be used for system synchronization. The valid indication is based on a PCS lock and a link-up indication.

Each of the recovered clocks can work as either a normal clock or a squelched clock. In the normal clock mode, CLK_OUT is the recovered clock of the selected SerDes. In the squelched mode, CLK_OUT will halt if the RX of the selected link is not synced. For multilane ports (for example, 40GbE or 100GbE), the recovered clock is derived from the first SerDes of the relevant port, and the valid signal is according to the link-up status of the entire port.

The recovered clock frequency is 25 MHz.

A fractional divider is used for generating an average clock of 25 MHz.

The following example describes how the average output frequency of 25 MHz is generated for a 100GbE port (4 × 25.78125 Gb/s).

The 25.78125 GHz is divided by 40, and the fractional divider, over 32 clock cycles, uses 7 clock cycles of a 25 divider [(25.78125 ÷ 40) ÷ 25] and 25 clock cycles of a 26 divider [(25.78125 ÷ 40) ÷ 26], as shown in the following equation.

$$\frac{25.78125 \text{ GHz}}{40} \div \frac{(7 \times 25) + (25 \times 26)}{32} = 25 \text{ MHz}$$

Similarly:

- A fractional divider of 26.5625 GHz uses 14 clock cycles of a 26 divider and 18 clock cycles of a 27 divider.
- A fractional divider of 20.625 GHz uses 12 clock cycles of a 20 divider and 20 clock cycles of a 21 divider.

NOTE: An external jitter attenuator is required to clean the recovered clock before using it as a reference for the rest of the system.

3.5 IEEE 1588

The BCM88800 is a highly integrated device with many hardware hooks for designs that require network time synchronization with improved time-stamping accuracy on a nanosecond scale. The following features make the device ideally suited for time synchronization applications that comply with IEEE 1588:

- Supported modes:
 - E2E and P2P transparent clock (TC).
 - Boundary clock (BC).
 - TC + OC slave, BC + OC slave.
- One-step clock features:
 - On-the-fly egress packet modification including UDP checksum update and CRC update.
 - All modifications to the correction field are handled in hardware.
 - Very short residence time.
 - All packets timestamped on ingress.
 - Switch-packet processing engines identify IEEE 1588 packets.
- Two-step features:
 - Egress timestamps are stored in per-port FIFO, along with IEEE 1588 sequence number.
 - The CPU can indicate which packets should generate a timestamp on egress.
 - All packets are timestamped on ingress.
 - Uses switch packet processing engines to identify IEEE 1588 packets and trap to CPU.
- Synchronizable timestamp counter:
 - Can be phase-locked to external source.
 - BroadSync[®] (timecode + event clock) interface.
 - Broadcom PHY sync interface.
 - Timestamped GPIOs.
- Frequency synthesizer:
 - Additional clock divider: 10 MHz + 1 pps.
 - Low jitter.

NOTE: PTP/IEEE 1588 functionality is supported over network (NIF) Ethernet ports only. Timestamping is not supported over fabric and Interlaken interfaces.

3.5.1 BroadSync External Interfaces

The BroadSync block provides the following external interfaces for providing timing information to off-chip devices or for retrieving timing information from external devices:


- Reference clock (input)
- External Sync 1 (input or output)
- External Sync 2 (input or output)
- Three-pin BroadSync interface (input or output):
 - TS_BIT_CLK
 - TS_SYNC
 - TS_TIME_VAL

The reference clock input is used for the clocking of all of the logic in the BroadSync block. The BroadSync block operates on a separate clock domain than the rest of the switch logic. The reference clock input should be driven from a low-jitter source to ensure that all time-related functions are accurate.

The External Sync 1 and 2 signals can be used as either inputs or outputs. When the signals are being used as inputs, a rising edge on the signal can be used to sample the current value of the internal timer, which can then be retrieved by the CPU. When these signals are being used as outputs, the signal can be programmed to toggle based on a configurable interval.

The BroadSync interface is configured through the CMIC_BS_CONFIG register and is composed of three signals: the clock, the heartbeat, and the time code. The BroadSync interface can be configured to operate as either a master (output) or slave (input). The interface clocks out or takes in serial data as shown in the following figure.

Figure 6: BroadSync Interface I/O



When the BroadSync interface is operating in slave mode, an external device clocks a time code into the BCM88800. External hardware provides the TS_BIT_CLK, TS_SYNC, and TS_TIME_VAL signals. During each heartbeat period, the external hardware shifts in a time code value, which consists of the following:

- Start bit
- Lock bit
- 16-bit epoch
- 32-bit seconds
- 2-bit zero
- 30-bit nanoseconds
- 8-bit accuracy
- CRC8 (covers all bits from LOCK to ACCURACY[0])

The time value shifted in corresponds to the time of the most recent rising edge of the heartbeat signal. All of the bits clocked in from the time code are stored in the CMIC_BS_INPUT_TIME[0:2] registers. Additionally, the received CRC8 is compared against the computed CRC8, and the result of the comparison is present in the CHECKSUM_ERROR field in the CMIC_BS_INPUT_TIME_2 register.

The internal time value is calibrated to the external signals through the following process:

1. The rising edge of TS_SYNC is used to sample the device's internal free-running clock value.
2. The sampled free-running clock value is compared to the time value subsequently shifted in using the TS_TIME_VAL signal.
3. These pairs of values (shifted-in time and sampled free-running time) are provided to the CPU on an occasional basis.
4. The differences and rates of change of the differences of the two time bases are used to derive a drift value.
5. The computed drift value is used to correct the internal time counter.

When the BroadSync interface is operating in master mode, the interface drives a time code to external devices. How frequently a new time code is clocked out depends on how often the heartbeat signal goes high. The heartbeat signal toggling frequency can be configured through the CMIC_BS_HEARTBEAT_CTRL register. When the heartbeat goes high, the contents of the CMIC_BS_OUTPUT_TIME memory is clocked out. The BCM88800 automatically computes and appends the correct CRC8 value immediately after the 8-bit accuracy field is clocked out. The TS_BIT_CLK used to clock out the time code is synthesized from the BroadSync Clock domain.

The BCM88800 can optionally synthesize the TS_BIT_CLK from a highly accurate internal PLL. The PLL generates an extremely low-jitter clock that is ITU-T G.824 and ITU-T G.823 compliant.

3.6 Fabric Interface

The fabric block integrates 48 SerDes links. The SerDes core is an octal (x8) lane 53.125 Gb/s PAM4 or 25.78125 Gb/s NRZ and is suitable for optical and backplane applications. The SerDes core has two PLLs that can be configured independently, and each lane can select between the two VCO clocks. The SerDes core supports data rates from 23.0 Gb/s to 25.78125 Gb/s in NRZ (backward-compatible with the BCM88770) and a data rate of up to 53.125 Gb/s in PAM4 mode. Each lane can be configured independently to run in PAM4 or NRZ modes with one of the two PLLs' configured rates.

SerDes supports a built-in PRBS functionality. PRBS is also supported on the fabric MAC layer per lane.

The following table lists the supported SerDes rates as well as the associated electrical specification standards.

Table 6: Fabric Supported Rates and the Electrical Specifications Standards

| SerDes Rate (Gb/s) | SerDes Mode | Standard Electrical Specifications |
|--------------------|-------------|--|
| 53.125 | PAM4 | 50GBASE-CR: IEEE 802.3cd CL136 |
| | | 50GBASE-KR: IEEE 802.3cd CL137 |
| | | 200GAUI-4 C2C: IEEE 802.3bs Annex 120D |
| | | 200GAUI-4 C2M: IEEE 802.3bs Annex 120E |
| 50.0 | PAM4 | 50GBASE-KR: IEEE 802.3cd CL137 ^a |
| 25.78125 | NRZ | 100GBASE-CR4: IEEE 802.3bj CL92 |
| | | 100GBASE-KR4: IEEE 802.3bj CL93 |
| | | CAUI-4 C2C: IEEE 802.3bm Annex 83D |
| | | CAUI-4 C2M: IEEE 802.3bm Annex 83E |
| 25.0 | NRZ | 100GBASE-KR4: IEEE 802.3bj CL93 ^a |
| 23.0 | NRZ | 100GBASE-KR4: IEEE 802.3bj CL93 ^a |

a. Scaled to the appropriate rate.

NOTE: For more information about the supported rates, see [Table 41, Fabric SerDes Supported Rates](#).

3.6.1 PCS Layer Line Coding

Each fabric link supports one of the following line coding options:

- 64b/66b encoding
- 25G Reed-Solomon FEC
- Low Latency 25G Reed-Solomon FEC
- 50G Reed-Solomon FEC
- Low Latency 50G Reed-Solomon FEC

3.6.1.1 64b/66b Encoding

The 64b/66b line coding is based on the standard 10GbE 64b/66b coding (IEEE 802.3 Clause 49).

The KR-FEC block supports error indication features. Error indication enables the RX-MAC to obtain an indication about an uncorrectable FEC frame, and discard cells transmitted in these frames. Enabling the error indication improves the Mean Time to False Packet Acceptance (MTTFPA) of the fabric link. The error indication adds one additional KR-FEC frame time delay (two KR-FEC frames time in total).

The effective bandwidth of 64b/66b is ~97% (64b/66b) of the SerDes rate.

3.6.1.2 25G Reed-Solomon FEC (RS-FEC)

The RS-FEC is based on the Reed-Solomon error correction code and can be configured on a per-link basis.

For 25-Gb/s links, the BCM88800 supports RS (206, 196, t = 5) coding using 10-bit symbols. This means that each 2060-bit frame consists of 1960 bits of data and 100 bits of RS Syndrome, which is used by the decoder to correct errors.

The selected code can correct any five symbols in the frame. This RS code can correct any single-error burst of 41 bits per RS frame, and up to 50 bits of error burst if it spans no more than five symbols.

Out of the 1960 data bits of the RS-FEC frame, the BCM88800 uses only 1950 bits for real data and 10 bits as overhead. The effective bandwidth of a link running RS-FEC is ~93.2% of the link rate ($64/65 \times 1950/2060$).

The 25G RS-FEC decoder on the receive side adds a delay equivalent to ~1.5 RS-FEC frame time (the computation logic takes ~0.5 FEC frame time). For example, RS-FEC adds a latency of ~130 ns on 23 Gb/s links (~116 ns on 25.78125-Gb/s links).

The RS-FEC supports the option to indicate uncorrectable RS-FEC frames. When enabled, the RX MAC drops all cells that are part of an uncorrectable RS-FEC frame. Enabling this indication adds one additional RS-FEC frame time delay (total of ~2.5 RS-FEC frame time). For example, RS-FEC with error indication adds a latency of ~220 ns on 23 Gb/s links (~196 ns on 25.78125-Gb/s links).

3.6.1.3 Low Latency 25G Reed-Solomon FEC (LL 25G RS-FEC)

The BCM88800 supports a lower-latency version of the RS-FEC (for 25 Gb/s links only). The LL 25G RS-FEC is based on the same 25G RS (206, 196, $t = 5$) coding using 10-bit symbols; however, only half the data is transmitted. The LL 25G RS-FEC code can correct any five symbols in the frame, which is the same as 25G RS-FEC.

The LL 25G RS-FEC transmitted frame is 1080 bits (975 data bits, 5 overhead bits and 100 RS Syndrome bits). The effective bandwidth of a link running LL 25G RS-FEC is $\sim 88.9\%$ of the link rate $[(64 \div 65) \times (975 \div 1080)]$.

The LL 25G RS-FEC has a delay equivalent to ~ 2 LL 25G RS-FEC frame time (computation logic takes ~ 1 LL 25G RS-FEC frame time) and ~ 3 LL 25G RS-FEC frame time when the uncorrectable error indication is enabled. For example, LL 25G RS-FEC adds a latency of ~ 84 ns on 23-Gb/s links without uncorrectable error indication (~ 76 ns on 25.78125-Gb/s links) and ~ 131 ns on 23-Gb/s links when uncorrectable error indication is enabled (~ 118 ns on 25.78125-Gb/s links).

3.6.1.4 50G Reed-Solomon FEC (RS-FEC)

For 50 Gb/s-links, the BCM88800 supports RS (545, 515, $t = 15$) coding using 10-bit symbols. This means that each 5450-bit frame consists of 5150 bits of data and 300 bits RS Syndrome, which is used by the decoder to correct errors.

This code can correct any 15 symbols in the frame. This RS code can correct any single error burst of 141 bits per RS frame, and up to 150 bits of error burst if it spans no more than 15 symbols.

Out of the 5150 data bits of the RS-FEC frame, the BCM88800 uses only 5056 bits for real data. The effective bandwidth of a link running RS-FEC is $\sim 92.7\%$ of the link rate ($5056 \div 5450$).

The 50 Gb/s RS-FEC decoder on the receive side adds a delay equivalent to ~ 1.8 RS-FEC frame time (the computation logic takes ~ 0.8 FEC frame time). For example, RS-FEC adds a latency of ~ 185 ns on 53.125-Gb/s links.

3.6.1.5 Low Latency 50G Reed-Solomon FEC (LL 50G RS-FEC)

The BCM88800 supports a lower-latency version of the 50G RS-FEC. The LL 50G RS-FEC is based on the same 50G RS (545, 515, $t = 15$) coding using 10-bit symbols; however, only half the data is transmitted. The LL 50G RS-FEC code is able to correct any 15 symbols in the frame, which is the same as 50G RS-FEC.

The LL 50G RS-FEC transmitted frame is 3040 bits (2688 data bits). The effective bandwidth of a link running LL 50G RS-FEC is $\sim 88.4\%$ of the link rate. The LL 50G RS-FEC has a delay equivalent to ~ 2 LL 50G RS-FEC frame time (computation logic takes ~ 1 LL 50G RS-FEC frame time) and ~ 3 LL 50G RS-FEC frame time when the uncorrectable error indication is enabled.

3.6.1.6 Line Coding Overheads and FEC Latency

The following two tables show technical specifications for the line coding options.

Table 7: Line Coding Overheads

| Code | Block Length (Bits) | Data Length (Bits) | Utilization | SerDes Rate (Gb/s) | Encoded SerDes Rate (Gb/s) |
|------------------------|---------------------|--------------------|-------------|--------------------|----------------------------|
| 64b/66b | 66 | 64 | 96.97% | 23 | 22.3 |
| 25G RS-FEC | 2060 | 1920 | 93.20% | 23 | 21.43 |
| Low latency 25G RS FEC | 1080 | 975 | 88.9% | 23 | 20.45 |
| 64b/66b | 66 | 64 | 96.97% | 25.78125 | 25 |
| 25G RS-FEC | 2060 | 1920 | 93.20% | 25.78125 | 24 |
| Low latency 25G RS FEC | 1080 | 975 | 88.9% | 25.78125 | 22.92 |
| 50G RS-FEC | 5450 | 5150 | 92.7% | 50 | 46.35 |
| 50G RS-FEC | 5450 | 5150 | 92.7% | 53.125 | 49.24 |
| Low latency 50G RS FEC | 3040 | 2688 | 88.4% | 50 | 44.21 |
| Low latency 50G RS FEC | 3040 | 2688 | 88.4% | 53.125 | 46.97 |

Table 8: FEC Latency

| Code | Frame Size (Bits) | SerDes Rate (Gb/s) | Frame Transmit (ns) | Total Without Marking (ns) | Total With Marking (ns) |
|------------------------|-------------------|--------------------|---------------------|----------------------------|-------------------------|
| 25G RS-FEC | 2060 | 23 | 90 | 130 | 220 |
| Low latency 25G RS FEC | 1080 | 23 | 47 | 84 | 131 |
| 25G RS-FEC | 2060 | 25.78125 | 80 | 116 | 196 |
| Low latency 25G RS FEC | 1080 | 25.78125 | 42 | 76 | 118 |
| 50G RS-FEC | 5450 | 50 | 109 | 198 | 307 |
| 50G RS-FEC | 5450 | 53.125 | 102.6 | 184.6 | 287.2 |

3.6.2 Fabric SerDes Connectivity

When connecting devices using fabric SerDes, use the following guidelines:

- Minimize the maximum cable length. Cable length above 100 meters requires special approval.
- Minimize the cable length variance. Use similar cable lengths for all links.
- To be able to support link level flow control (LLFC), the device has further requirements. When a link input FIFO is building up, LLFC is asserted to prevent input FIFO overflow. When the LLFC is deasserted, the input FIFO should have enough data to prevent the FIFO from getting empty, resulting in potential performance loss. The LLFC assertion and deassertion threshold must be able to handle the LLFC delivery time, as well as data in-flight, which depends on various parameters. The significant parameters are the PCS line encoding, the cell size, and the cable length. The LLFC settings depend on the size of the input-link FIFO:
 - In single-pipe mode, the input-link FIFO size is 512 entries of 128 bytes per link (using the memory from the other pipes).
 - In two-pipe mode in the BCM88800, the input-link FIFO size can be set to 256 entries of 128 bytes.
 - In three-pipe mode, the input-link FIFO size is 170 entries of 128 bytes per link, per pipe.

The following table describes the cable length limitations when the BCM88800 requires generating LLFC indication, which is calculated according to the BCM88800 input link FIFO size. The LLFC cable length limitations of a link partner device must be calculated separately and may be greater or lesser than the data presented in this table.

Table 9: LLFC Fiber-Optic Cable Length Summary

| Number of Fabric Pipes | MAX Link Rate (Gb/s) | Uncorrectable Error Marking? | Line Coding | Cable Length (Meters) | Remark |
|------------------------|----------------------|------------------------------|-------------|-----------------------|--|
| Single | 53.125 | Any | Any | 100 | — |
| Single | 50 | Any | Any | 100 | — |
| Single | 25.78125 | Any | Any | 100 | — |
| Two | 53.125 | No | Any | 100 | — |
| Two | 53.125 | Yes | 50G RS-FEC | 100 | — |
| Two | 50 | No | Any | 100 | — |
| Two | 50 | Yes | 50G RS-FEC | 100 | — |
| Two | 25.78125 | Any | 25G RS-FEC | 100 | — |
| Three | 53.125 | No | 50G RS-FEC | 73/63 | 73 if connected with BCM88790 63 if connected in mesh |
| Three | 53.125 | Yes | 50G RS-FEC | 51/42 | 51 if connected with BCM88790 42 if connected in mesh |
| Three | 50 | No | 50G RS-FEC | 80/70 | 80 if connected with BCM88790 70 if connected in mesh |
| Three | 50 | Yes | 50G RS-FEC | 57/47 | 57 if connected with BCM88790 47 if connected in mesh |
| Three | 25.78125 | Any | Any | 100 | — |

3.7 Statistics Interface

The statistics interface is an event-driven interface through which statistics are pushed out. Statistics records are continuously reported at a maximum rate of one record per clock in ingress and egress, for a maximum rate of two records per clock. The BCM88800 supports one of the following global configurable options:

- Queue-Size – The statistics interface reports the queue size of ingress enqueue actions and ingress dequeue actions. This mode can be used to build an image of queue sizes, infer congestion, and enable congestion management by an ingress PP.
- Billing – The statistics interface reports ingress received packets and egress transmitted packets. Packets are tagged with information that maps into counters at the external statistic processing device.
- Ingress-Enqueue/Dequeue – In this mode, the statistics interface reports per each ingress packet enqueued and dequeued. Packets are tagged with information that map into counters at the external statistic processing device.

The statistics interface has two modes:

- Single – The statistics interface is a single port up to 200GbE. Each packet contains records from ingress and egress. This mode allows different record sizes for the ingress and the egress. Thus, it is possible to mix larger ingress records with shorter egress records without compromising the rate of the statistics interface.
- Dual – The statistics interface is made of two 100GbE interfaces. Each interface caters to the ingress or egress. In this mode, mixing larger ingress record with shorter egress records would compromise the statistics interface bandwidth.

The external statistics processor can be Broadcom KBP or TAP devices, or a custom FPGA.

The BCM88800 statistics interface can use NIF ports from PM50_03 (SerDes 24 to 31) or PM25_04, PM25_05, PM25_06, or PM25_07 (SerDes 32 to 47) as described in [Table 1](#). In single mode, each one of the two 200GbE ports of PM50_03 can be used. In dual mode, any two 100GbE ports from PM50_03 (SerDes 24 to 31) or PM25_04, PM25_05, PM25_06, or PM25_07 (SerDes 32 to 47) can be used.

NOTE: The following notes apply to STAT_IF usage:

- When STAT_IF is using PM50_03, other lanes in PM50_03 cannot be used for ETH ports (can be used for ILKN).
- When STAT_IF is using PM25_04, PM25_05, PM25_06, or PM25_07, other PM25 in this group (PM25_04, PM25_05, PM25_06, and PM25_07) cannot be used for ETH ports. For example, if STAT_IF is using PM25_06, then PM25_04, PM25_05, and PM25_07 cannot be used for ETH ports (can be for ILKN).

NOTE: Only the TX direction of the Ethernet port is used for the statistics interface. The RX direction is not used for data (statistics records) transfer and can be left unconnected.

However, connecting the RX direction is recommended because this enables the use of link-training and allows more debug capabilities over this interface.

An Ethernet-like MAC is used to send statistics packets:

- Standard Ethernet 8-byte preamble.
- 32b FCS, which is the same as standard Ethernet.
- Average IPG of 12B.

Statistics Packets format definitions, refer to the Statistics Records section in the *BCM88800 Counters and Statistic Interface* application note (88800-AN1xx-R). See [Related Documents](#).

3.8 HBM Packet Buffers

The BCM88800 includes a single HBM Gen2 module. The HBM technology integration enables low-power and high-performance memory access. The integrated HBM module is 4-Hi, supporting 4 GB of deep buffering.

3.9 CPU Interface

The Broadcom iProc block provides an interface between the host CPU and the internal registers and tables within the switch device, enabling complete management of the switch.

The iProc block is made up of the following components:

- PCIe x4 lane Gen3 interface at up to 8 Gb/s:
 - Compatible with x1 or x2 lane PCIe.
 - Compatible with PCIe Gen1 at 2.5 Gb/s or Gen2 at 5 Gb/s.
- BSC (I2C-compatible) 2-line interface.
 - Basic device debug and register access (PCIe debug)
 - PCIe QSPI flash programming
 - Loading code used for heating when an industrial device is powered-up at a low ambient temperature
- Microcontroller subsystem:
 - Two Arm Cortex-R5 microcontrollers, running at 875 MHz (independent of core clock).
 - 32-KB I-Cache, 32-KB D-Cache for each microcontroller core.
 - 128-KB I-Tightly coupled memory, 128-KB D-Tightly coupled memory for each microcontroller core.
 - 1-MB internal system memory.
- MIIM interface:
 - MDIO compatible interface.
- LED interface.
- SBUS DMA: Enables the BCM88800 to read from host memory and update the BCM88800 tables, and to read BCM88800 tables and write to the host memory.
- FIFO DMA.
- Packet TX/RX DMA.
- Remote CPU over network or fabric interface.
- Miscellaneous (endian order, reset controls).

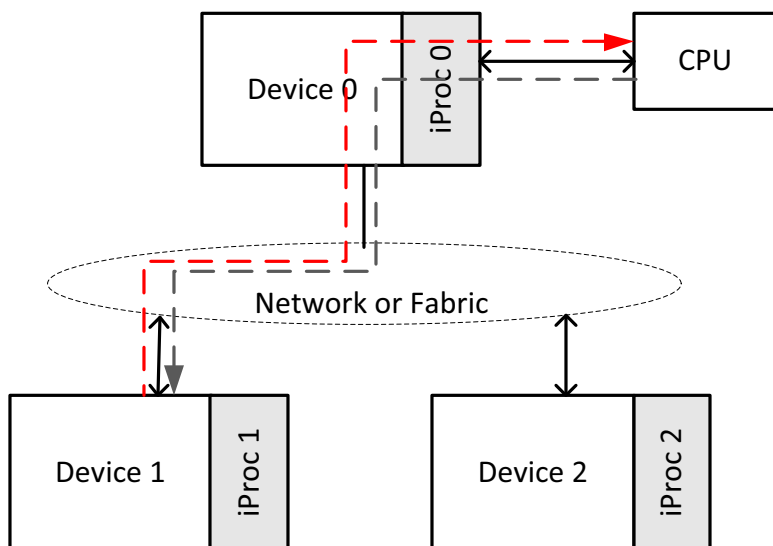
3.9.1 Remote CPU Support

The BCM88800 can be programmed from a remote CPU after initialization. Upon power-up, the device must be initialized using the PCIe interface. After the device has been initialized, the following functions are supported:

- iProc register access
- Generating RCPU packets when interrupts are triggered
- SCHAN register access
- SCHAN table access

The BCM88800 can operate in a system that is managed by a remote CPU. The CMIC communicates with the remote CPU through Ethernet packets with a special EtherType. These packets are referred to as remote CPU packets. The following figure represents a system managed by a remote CPU.

Figure 7: System Managed by Remote CPU



For the CPU to generate an S-channel operation in Device 1, the following events must occur:

- The CPU generates an Ethernet packet with MACDA addressed to iProc1 with a special EtherType value reserved for remote CPU packets. The packet is an SCHAN_REQUEST.
- iProc0 injects this packet into the ingress pipeline of Device 0. The packet is forwarded based on the MACDA.
- The packet is received at Device 1 and is forwarded to iProc1.
- iProc1 interprets the packet and performs the S-channel operation.
- iProc1 creates a new SCHAN_REPLY packet based on the result of the S-channel operation. This is also an Ethernet packet with MACDA addressed to the CPU with a special EtherType. This packet is injected into the ingress pipeline of Device 1.
- Device 1 forwards the SCHAN_REPLY packet based on the MACDA and the packet is sent to Device 0.
- Device 0 receives the packet and forwards it to iProc0.
- iProc0 sends the packet to the CPU.

3.9.2 Remote Packet Operations

Remote packets are those that are sent or received by the iProc without a local CPU (whether internal or external) being involved. The iProc receives a remote packet from the switch egress pipe. The iProc matches this packet and performs some operations based on the packet data. The iProc may then send a packet back to the sender of the original packet.

The remote CPU performs certain SCHAN operations without the local CPU's intervention. This is especially useful in stacks where the master CPU in the stack may want to perform L2 insert or delete operations. The means for having the CMIC match incoming packets and perform an SCHAN operation is provided. The CMIC may then send a reply packet back to the requesting remote CPU with the SCHAN operation's status and result data.

Although remote CPUs can send arbitrary SCHAN control packets to the device using this mechanism, it does not remove the requirement for a local CPU to configure the switch (either internal or external). The remote CPU SCHAN packets contain control information to match a reply to a request, but higher layer software must be provided to deal with lost packets, whether they are request or reply packets. Some SCHAN operations are potentially destructive in that they cannot easily be replayed if a reply is lost. No mechanism exists in this device to handle such situations.

3.9.3 PCIe Interface

The PCIe interface of the BCM88800 switch conforms to PCIe 3.1 specifications. The BCM88800 supports four lanes of Gen3 PCIe (8G in each direction). No external glue logic is required to support this interface. The protocols and electrical requirements of the PCIe specifications are strictly implemented.

3.9.4 MIIM

The iProc supports the IEEE 802.3 standard MII Management (MIIM) interface. This is a two-wire serial bus controlled by the iProc that allows register access to external PHYs in the system. The two signals for MIIM are MDC (clock) and MDIO (bidirectional data).

The CPU programs the external PHY registers using this interface. The MIIM interface can be configured to support Clause 22 or Clause 45.

The BCM88800 supports six MIIM interfaces.

3.9.5 UART

The BCM88800 has two UART interfaces. UART interfaces are used for debugging software running on the microcontrollers (one UART for each microcontroller). The UART interfaces can be used for time-of-day (ToD) synchronization

The UART interface includes only the following two data lines:

- One RxData line
- One TxData line


3.9.6 LED Interface

The device provides five serial LED output interfaces. An Arm Cortex-M0 microcontroller has control of all five interfaces, allowing the user to select which interfaces are used to provide serial LED bitstreams. A user can write code for the microcontroller that collects status for Ethernet ports, forms streams of status bits, and then shifts them out using any one of the LED interfaces.

The output frequency and refresh rate are user-configurable. These parameters are common across all five of the LED status interface outputs.


A two-wire (clock and data) LED interface controls system LEDs. Both signals are held low during periods of inactivity. A single LED refresh cycle consists of clocking out a programmable number of LED data bits. The LED data signal is pulsed high at the start of each LED refresh cycle (see the following figure).

Figure 8: Single LED Refresh Cycle



The LED refresh cycle is repeated periodically to refresh the LEDs (see the following figure).

Figure 9: LED Refresh Timing



Chapter 4: Pin Signal Description

4.1 Pin List and Pin Map

The BCM88800 pin list and pin map are provided in spreadsheet format on the Broadcom Customer Service Portal (docSAFE) for collateral distribution. The spreadsheet serves as the official document containing the device's signal mapping. Refer to the *BCM88800 PinList* file (see [Related Documents](#)).

4.2 Pin I/O Type Description

The following table lists the conventions that are used to describe the I/O nature of the pins.

Table 10: Signal I/O Type Description

| I/O | Description |
|-----------------|--|
| B | Bidirectional signal |
| B _{OD} | Open drain bidirectional signal |
| B _{PD} | Bidirectional signal, with internal pull-down ^a |
| B _{PU} | Bidirectional signal, with internal pull-up ^a |
| I | Input signal |
| I _{OD} | Open drain input signal |
| I _{PD} | Input signal, with internal pull-down ^a |
| I _{PU} | Input signal, with internal pull-up ^a |
| NC | No Connect |
| O | Output signal |
| O _{OD} | Open drain output signal |
| O _{PD} | Output, with internal pull-down ^a |
| O _{PU} | Output, with internal pull-up ^a |

a. Pull-up and pull-down values are minimum = 40 kΩ, maximum = 60 kΩ.

4.3 Pin Description – Grouped by Function

The following table provides an overview of the pins on the BCM88800.

NOTE: Information about connectivity and filters for some of the signals is available in the *Hardware Design Guidelines for StrataDNX 16-nm Devices* (DNX16-AN1xx). In the following table, this document is referred to as the HWDG.

Table 11: Pin List by Function

| Signal/Bus Name | Qty. | Type | Tech | Description |
|-----------------------|------|-----------------|---------------------|--|
| PCIe Interface | | | | PCIe interface supporting Gen1, Gen2, and Gen3. According to connectivity (x1, x2, or x4), use lanes [0], [1:0], or [3:0]. |
| PCIE_TX_[3:0]_P/N | 2×4 | O | Differential | PCIe differential TX pairs. The lanes should be AC coupled. When not in use, leave open. |
| PCIE_RX_[3:0]_P/N | 2×4 | I | Differential | PCIe differential RX pairs. The lanes should be AC coupled. RX is internally terminated. When not in use, leave open. |
| PCIE_REFCLK_P/N | 2 | I | Differential CML | PCIe reference clock inputs. 100 MHz. External 100Ω termination is required between P and N pins. For connectivity, refer to the HWDG. |
| PCIE_TESTOUT_P/N | 2 | O | Differential CML | Factory test only. Leave open. |
| PCIE_RST_N | 1 | I _{PD} | CMOS 1.8V | PCIe reset, active low. Follow the functionality described in Section 5.5, Power-Up and Reset Sequence . (This function is required for PCIe Gen1, Gen2, and Gen3.) Use external pull-down to force 0 while control logic is not initiated. |
| PCIE_PVDD0P8 | 1 | PWR | 0.8V | PCIe PLL power supply. For the recommended filter, refer to the HWDG. |
| PCIE_RTVDD0P8 | 2 | PWR | 0.8V | PCIe SerDes analog power supply. For the recommended filter, refer to the HWDG. |
| QSPI | | | | The QSPI interface accesses a serial flash memory device. The flash memory holds the PCIe SerDes firmware and configuration that is required in all PCIe modes (Gen1, Gen2, and Gen3). Output pins are driven even when the device is in reset. |
| QSPI_CS_N | 1 | O _{PD} | CMOS 1.8V | Chip select (active low) from device to flash. |
| QSPI_HOLD_N | 1 | O _{PD} | CMOS 1.8V | Hold (active low) from device to flash. Can be used to pause the serial communication with the master device without resetting the serial sequence. |
| QSPI_MISO | 1 | I _{PD} | CMOS 1.8V | Serial data from flash (SO) to device (MI). |
| QSPI_MOSI | 1 | O _{PD} | CMOS 1.8V | Serial data from device (MO) to flash (SI). |
| QSPI_SCK | 1 | O _{PD} | CMOS 1.8V | Serial clock from device to flash. |

Table 11: Pin List by Function (Continued)

| Signal/Bus Name | Qty. | Type | Tech | Description |
|------------------------------------|------|-----------------|--------------|--|
| QSPI_WP_N | 1 | O _{PD} | CMOS 1.8V | The write protect (WP_N) allows normal read/write operations when held high. When the WP_N is brought low, all write operations are blocked. |
| BSC/I2C Interface | | | | The BSC/I2C is an alternate management interface that can be used for PCIe debugging and QSPI image programming. |
| I2C_SCL | 1 | I _{OD} | CMOS 1.8V | BSC/I2C CPU interface (slave only). Clock. Open drain. Must be pulled up externally to 1.8V. |
| I2C_SDA | 1 | B _{OD} | CMOS 1.8V | BSC/I2C CPU interface (slave only). Data in/out. Open drain. Must be pulled up externally to 1.8V. |
| Miscellaneous Signals | | | | |
| INT_N | 1 | O | CMOS 1.8V | CPU interrupt output. Pseudo open drain, active low. Must be pulled up externally to 1.8V. |
| SYS_RST_N | 1 | I _{PD} | CMOS 1.8V | Device reset input. Active low. Use external pull-down to force 0 when control logic is not initiated. |
| DISCONNECT | 1 | I _{PD} | CMOS 1.8V | Disconnect control input. When 1, the device indicates to fabric link partner to stop sending data to it. This pin can be used for graceful shutdown by external logic that implements early detection of power-down or card removal. Normal operation: 0 Graceful shutdown: 1 |
| MIIM Interface | | | | |
| MDC_[7:0] | 8 | O _{PU} | CMOS 1.8V | Clock output of MIIM interface chains. Enables controlling external PHY (master mode only). Supports Clause 22/45 protocol formats with CMOS 1.8V signaling. When not in use, leave open. MDC_0 and MDC_5 are for factory test only. Leave open. |
| MDIO_[7:0] | 8 | B _{PU} | CMOS 1.8V | Data in and data out of MIIM interface chains. Enables controlling external PHY (master mode only). Supports Clause 22/45 protocol formats with CMOS 1.8V signaling. When not in use, leave open. MDIO_0 and MDIO_5 are for factory test only. Leave open. |
| Power-Up Configuration Word | | | | |
| PUC_[29:0] | 30 | I | CMOS 1.8V | Power-up configuration (PUC) word. For PUC bus information see Section 5.5, Power-Up and Reset Sequence and Section 5.6, Power-Up Configuration Word . |

Table 11: Pin List by Function (Continued)

| Signal/Bus Name | Qty. | Type | Tech | Description |
|--|------|-----------------|---|--|
| Recommended Operating Voltage | | | | |
| ROV_[2:0] | 3 | O | CMOS 1.8V | Recommended operating voltage. Pull to GND, R < 5 kΩ. These pins define the required VDDC voltage levels with which the specific device should work. Working with a VDDC level that is different from the value required by the ROV may cause the device to fail normal operation or exceed power limits. For ROV information, see Section 5.2.1, Recommended Operating Voltage . |
| LED Controller | | | | |
| LED_CLK[4:0] | 5 | O _{PD} | CMOS 1.8V | LED clock of the five LED buses. Used to latch the LED output data. When not in use, leave open. |
| LED_DATA[4:0] | 5 | O _{PD} | CMOS 1.8V | LED data of the five LED buses. Serially indicates port status. When not in use, leave open. |
| UART | | | | |
| UART[1:0]_RX | 2 | I _{PD} | CMOS 1.8V | UART data receive, one for each UART interface. Might be used for on-chip Arm debug. Usage options: <ul style="list-style-type: none"> ■ When using the Broadcom IEEE 1588 stack, both UART[1:0]_RX can be used as input of Time of Day (ToD) for Grand Master (GM) clock. ■ When not in use, leave open. |
| UART[1:0]_TX | 2 | O _{PU} | CMOS 1.8V | UART data transmit, one for each UART interface. Might be used for on-chip Arm debug. Usage options: <ul style="list-style-type: none"> ■ When using Broadcom IEEE 1588 stack, both UART[1:0]_TX can be used as output of Time of Day (ToD) for slave-clock. ■ When not in use, leave open. |
| BroadSync and Broadcom IEEE 1588 Stack Interfaces | | | For BroadSync implementation, refer to the <i>Broadcom BroadSync Solution: Description and Implementation</i> (1588-AN2xx). | |
| TS_GPIO_[5:0] | 6 | B _{PU} | CMOS 1.8V | User-programmable general-purpose I/Os. Each pin can be individually configured to act as input or output. Usage options: <ul style="list-style-type: none"> ■ When not in use, leave open. ■ When using BroadSync: <ul style="list-style-type: none"> – Available as general-purpose I/Os. – TS_GPIO_1 can be used as 1 pps for testing. ■ When using Broadcom IEEE 1588 stack: <ul style="list-style-type: none"> – Available as general-purpose I/Os. – TS_GPIO_[5:0] can be used as 1 PPS input or 1 PPS output |
| TS_PLL_BYP | 1 | I _{PD} | CMOS 1.8V | Factory test only. Pull to GND, R ≤ 1 kΩ. |
| BS_PLL_BYP | 1 | I _{PD} | CMOS 1.8V | Factory test only. Pull to GND, R ≤ 1 kΩ. |

Table 11: Pin List by Function (Continued)

| Signal/Bus Name | Qty. | Type | Tech | Description |
|--------------------|------|-----------------|------------------|--|
| TS_PLL_REFCLK_P/N | 2 | I | Differential CML | TSPLL (time stamp PLL) and BSPLL (BroadSync PLL) reference clock differential input. Clock rate is 25 MHz. For connectivity, refer to the HWDG. Usage options: <ul style="list-style-type: none"> When not in use, leave open. When using BroadSync, can be sourced from a simple free-running local oscillator. When using Broadcom IEEE 1588 stack, must be sourced from TDPLL, which is sourced from OCXO. |
| TS_PLL_TESTOUT_P/N | 2 | O | Differential CML | Factory test only. Leave open. |
| BS_PLL_TESTOUT_P/N | 2 | O | Differential CML | Factory test only. Leave open. |
| TS_PLL_LOCK | 1 | O | CMOS 1.8V | TS PLL lock indication. High indicates PLL is locked. Lock indication is available only after the software initializes and enables the PLL. |
| BS_PLL_LOCK | 1 | O | CMOS 1.8V | BS PLL lock indication. High indicates PLL is locked. Lock indication is available only after the software initializes and enables the PLL. |
| TS_PLL_AVDD1P8 | 1 | PWR | 1.8V | PLL analog power 1.8V. For the recommended filter, refer to the HWDG. |
| TS[1:0]_BIT_CLK | 2 | B _{PD} | CMOS 1.8V | Usage options: <ul style="list-style-type: none"> When not in use, leave open. When using BroadSync: <ul style="list-style-type: none"> TS0_BIT_CLK is used as BroadSync bit clock, usually 10 MHz. This signal can be configured as input for BroadSync Slave or output for BroadSync Master. TS1_BIT_CLK is not used. When using Broadcom IEEE 1588 stack, both TS[1:0]_BIT_CLK are optional 10-MHz output. |
| TS[1:0]_SYNC | 2 | B _{PD} | CMOS 1.8V | Usage options: <ul style="list-style-type: none"> When not in use, leave open. When using BroadSync: <ul style="list-style-type: none"> TS0_SYNC is used as BroadSync heartbeat pulse, 4 kHz. Marks the start of the transmission of the synchronize time value. This signal can be configured as input for BroadSync Slave or output for BroadSync Master. TS1_SYNC is not used. When using Broadcom IEEE 1588 stack, both TS[1:0]_SYNC are optional 4-kHz input or output. |

Table 11: Pin List by Function (Continued)

| Signal/Bus Name | Qty. | Type | Tech | Description |
|--------------------------------------|------|-----------------|---------------------|--|
| TS[1:0]_TIME_VAL | 2 | B _{PD} | CMOS 1.8V | Usage options: <ul style="list-style-type: none"> When not in use, leave open. When using BroadSync: <ul style="list-style-type: none"> TS0_TIME_VAL is used as BroadSync synchronized time value and serially shifts the time value one bit per rising edge of the TS0_BIT_CLK. This signal can be configured as input for BroadSync-Slave or output for BroadSync-Master. TS1_TIME_VAL is not used. When using Broadcom IEEE 1588 stack, both TS[1:0]_TIME_VAL are not used. |
| Synchronized Ethernet (SyncE) | | | | |
| SYNCE[1:0]_PLL_REFCLK_P/N | 2×2 | I | Differential CML | SYNCE PLL reference clock inputs. Factory test only. Leave open. |
| SYNCE[1:0]_PLL_LOCK | 2 | O | CMOS 1.8V | SYNCE PLL lock indication. High indicates PLL is locked. These two PLLs are not in use. Factory test only. Leave open. |
| SYNCE[1:0]_CLK_OUT_P/N | 2×2 | O | Differential CML | SyncE recovered clock. With an additional external circuit, these signals can be used for system-level clock synchronization for SyncE applications. See Section 3.4.2, Recovered Clock . |
| SYNCE[1:0]_CLK_OUT_VALID | 2 | O | CMOS 1.8V | SyncE valid indication. With an additional external circuit, these signals can be used for system-level clock synchronization for SyncE applications. See Section 3.4.2, Recovered Clock . When using this output, connect an external pull-down. |
| SYNCE_PLL_AVDD1P8 | 1 | PWR | 1.8V | PLL analog power 1.8V. For the recommended filter, refer to the HWDG. |
| Out-of-Band Flow Control | | | | |
| FC_A_RX_CLK | 1 | I _{PD} | CMOS 1.8V | Flow control interface A, RX clock. When not in use, leave open or place a pull-down resistor. |
| FC_A_RX_STAT | 1 | I _{PD} | CMOS 1.8V | Flow control interface A, RX status information. Usage options: <ul style="list-style-type: none"> When operating in ILKN mode, the functionality is xx_STAT. When operating in SPI mode, the functionality is xx_STAT[0]. When not in use, leave open. |
| FC_A_RX_SYNC | 1 | I _{PD} | CMOS 1.8V | Flow control interface A, RX synchronization signal. Usage options: <ul style="list-style-type: none"> When operating in ILKN mode, the functionality is xx_Sync. When operating in SPI mode, the functionality is xx_STAT[1]. When not in use, leave open. |
| FC_A_TX_CLK | 1 | O | CMOS 1.8V | Flow control interface A, TX clock. When not in use, leave open. |
| FC_A_TX_STAT | 1 | O | CMOS 1.8V | Flow control interface A, TX status information. Usage options: <ul style="list-style-type: none"> When operating in ILKN mode, the functionality is xx_STAT. When operating in SPI mode, the functionality is xx_STAT[0]. When not in use, leave open. |

Table 11: Pin List by Function (Continued)

| Signal/Bus Name | Qty. | Type | Tech | Description |
|----------------------|------|-----------------|---------------------|---|
| FC_A_TX_SYNC | 1 | O | CMOS 1.8V | Flow control interface A, TX synchronization signal. Usage options: <ul style="list-style-type: none"> When operating in ILKN mode, the functionality is xx_Sync. When operating in SPI mode, the functionality is xx_STAT[1]. When not in use, leave open. |
| FC_B_RX_CLK | 1 | I _{PD} | CMOS 1.8V | Flow control interface B, RX clock. When not in use, leave open or place a pull-down resistor. |
| FC_B_RX_STAT | 1 | I _{PD} | CMOS 1.8V | Flow control interface B, RX status information. Usage options: <ul style="list-style-type: none"> When operating in ILKN mode, the functionality is xx_STAT. When operating in SPI mode, the functionality is xx_STAT[0]. When not in use, leave open. |
| FC_B_RX_SYNC | 1 | I _{PD} | CMOS 1.8V | Flow control interface B, RX synchronization signal. Usage options: <ul style="list-style-type: none"> When operating in ILKN mode, the functionality is xx_Sync. When operating in SPI mode, the functionality is xx_STAT[1]. When not in use, leave open. |
| FC_B_TX_CLK | 1 | O | CMOS 1.8V | Flow control interface B, TX clock. When not in use, leave open. |
| FC_B_TX_STAT | 1 | O | CMOS 1.8V | Flow control interface B, TX status information. Usage options: <ul style="list-style-type: none"> When operating in ILKN mode, the functionality is xx_STAT. When operating in SPI mode, the functionality is xx_STAT[0]. When not in use, leave open. |
| FC_B_TX_SYNC | 1 | O | CMOS 1.8V | Flow control interface B, TX synchronization signal. Usage options: <ul style="list-style-type: none"> When operating in ILKN mode, the functionality is xx_Sync. When operating in SPI mode, the functionality is xx_STAT[1]. When not in use, leave open. |
| DRAM PHY | | | | |
| DRAM_PHY_REFCLK_P/N | 2 | I | Differential CML | DRAM PHY reference clock inputs. 100 MHz. For connectivity, refer to the HWDG. |
| DRAM_PHY_TESTOUT_P/N | 2 | O | Differential CML | Factory test only. Leave open. |
| DRAM_PHY_PVDD1P8 | 1 | PWR | 1.8V | PLL analog power 1.8V. For the recommended filter, refer to the HWDG. |
| DRAM_PHY_AGND | 1 | GND | GND | PLL analog ground. On the board, it should share the same common GND plane as the VSS pins. |
| DRAM_PHY_VDDC | 7 | PWR | 0.88V | DRAM PHY core power supply. |
| DRAM_PHY_VDDC_SENSE | 1 | O | Analog | Factory test only. Leave open. |
| DRAM_PHY_VSS_SENSE | 1 | O | Analog | Factory test only. Leave open. |
| DRAM_PHY_VDDO_SENSE | 1 | O | Analog | Factory test only. Leave open. |
| DRAM_PHY_VSSO_SENSE | 1 | O | Analog | Factory test only. Leave open. |

Table 11: Pin List by Function (Continued)

| Signal/Bus Name | Qty. | Type | Tech | Description |
|-------------------------------------|------|-----------------|--------------|--|
| HBM | | | | |
| HBM_THERM_ALARM | 1 | O _{PD} | CMOS 1.8V | An indication that a CATTRIP signal from the HBM die was asserted. Indicates that the HBM die temperature has reached a certain level where catastrophic damage may occur unless power is reduced. The CATTRIP output is sticky in that to clear a CATTRIP, power-off of the device is required to return the CATTRIP output to 0. When HBM_THERM_ALARM is asserted, power-down the following HBM power supplies: <ul style="list-style-type: none"> ■ HBM_VDDC ■ HBM_VDDO ■ HBM_VPP2P5 |
| HBM_VDDC | 17 | PWR | 1.2V | HBM core supply voltage. For the recommended filter, refer to the HWDG. |
| HBM_VDDC_SENSE | 1 | O | Analog | Factory test only. Leave open. |
| HBM_VSSC_SENSE | 1 | O | Analog | Factory test only. Leave open. |
| HBM_VDDO | 23 | PWR | 1.2V | Supply voltage for I/O, for the HBM, and the controller DRAM PHY. For the recommended filter, refer to the HWDG. |
| HBM_VDDO_SENSE | 1 | O | Analog | Factory test only. Leave open. |
| HBM_VSSO_SENSE | 1 | O | Analog | Factory test only. Leave open. |
| HBM_VPP2P5 | 6 | PWR | 2.5V | HBM VPP supply. For the recommended filter, refer to the HWDG. |
| Direct Access (DA) Test Port | | | | |
| HBM_DA_## | 38 | I/O | Not defined | HBM direct access port for vendor-specific test implementations. Factory test only. Leave open. |
| Calibration Resistors | | | | |
| FAB_RESCAL | 1 | Analog | Analog | Resistor calibration terminal for SerDes. Connect each pin through an external resistor of 4.53 kΩ (1%) to SRD_AGND. The resistor should be located between this signal via and the nearest SRD_AGND via. |
| FAB_RESCAL_AVDD0P8 | 1 | PWR | 0.8V | RESCAL analog power 0.8V. For the recommended filter, refer to the HWDG. |
| NIF_RESCAL_[2:0] | 3 | Analog | Analog | Resistor calibration terminal for SerDes. Connect each pin through an external resistor of 4.53 kΩ (1%) to SRD_AGND. The resistor should be located between this signal via and the nearest SRD_AGND via. <p>NIF_RESCAL_0: Calibrate SerDes NIF50_[15:0], NIF25_[127:96].</p> <p>NIF_RESCAL_1: Calibrate SerDes NIF50_[31:16], NIF25_[31:0].</p> <p>NIF_RESCAL_2: Calibrate SerDes NIF25_[95:32].</p> |
| NIF_RESCAL_[2:0]_AVDD0P8 | 3 | PWR | 0.8V | RESCAL analog power 0.8V. For the recommended filter, refer to the HWDG. |

Table 11: Pin List by Function (Continued)

| Signal/Bus Name | Qty. | Type | Tech | Description |
|------------------------|------|--------|---------------------|---|
| Thermal Diode | | | | |
| THERM_DIODE_N | 1 | Analog | Analog | Cathode (N) pin for the thermal diode. Use an external thermal diode reader. When not in use, connect to GND. |
| THERM_DIODE_P | 1 | Analog | Analog | Anode (P) pin for the thermal diode. Use an external thermal diode reader. When not in use, connect to GND. |
| THERM_DIODE_VDD3P3 | 1 | PWR | 3.3V | Thermal diode 3.3V supply. |
| PLLs and Clocks | | | | |
| CLOCK25 | 1 | I | CMOS 1.8V | 25-MHz clock. |
| C_PLL_REFCLK_P/N | 2 | I | Differential CML | Core PLL reference clock inputs. 25 MHz. For connectivity, refer to the HWDG. |
| C_PLL_TESTOUT_P/N | 2 | O | Differential CML | Factory test only. Leave open. |
| C_PLL_AVDD1P8 | 1 | PWR | 1.8V | PLL analog power 1.8V. For the recommended filter, refer to the HWDG. |
| C_PLL_AGND | 3 | GND | GND | PLL analog ground. On the board, it should share the same common GND plane as the VSS pins. |
| C_PLL_LOCK | 1 | O | CMOS 1.8V | Core PLL lock indication. High indicates PLL is locked. |
| U_PLL_REFCLK_P/N | 2 | I | Differential CML | Microcontroller PLL reference clock inputs. 25 MHz. For connectivity, refer to the HWDG. |
| U_PLL_TESTOUT_P/N | 2 | O | Differential CML | Factory test only. Leave open. |
| U_PLL_AVDD1P8 | 1 | PWR | 1.8V | PLL analog power 1.8V. For the recommended filter, refer to the HWDG. |
| U_PLL_LOCK | 1 | O | CMOS 1.8V | Microcontroller PLL lock indication. High indicates PLL is locked. |
| FAB_PLL_REFCLK_P/N | 2 | I | Differential CML | Fabric PLL reference clock inputs. Used for Fabric SerDes digital section. 156.25 MHz. For connectivity, refer to the HWDG. |
| FAB_PLL_TESTOUT_P/N | 2 | O | Differential CML | Factory test only. Leave open. |
| FAB_PLL_LOCK | 1 | O | CMOS 1.8V | FAB_PLL lock indication. High indicates PLL is locked. Lock indication is available only after the software initializes and enables the PLL. |

Table 11: Pin List by Function (Continued)

| Signal/Bus Name | Qty. | Type | Tech | Description |
|--|-------|-----------------|------------------|---|
| FAB_[1:0]_REFCLK_P/N | 2×2 | I | Differential CML | Fabric analog reference clock inputs. 156.25 MHz. 0: FAB50 cores (0 to 3). 1: FAB50 cores (4 to 5). These clocks should share the same source. For connectivity, refer to the HWDG. |
| NIF_[4:0]_REFCLK_P/N | 5×2 | I | Differential CML | NIF analog reference clock inputs. 156.25 MHz. 0: PM50 0 and 1. 1: PM50 2 and 3. 2: PM25 4 to 11. 3: PM25 12 to 19. 4: PM25 20 to 27. These clocks should share the same source. For connectivity, refer to the HWDG. |
| FAB_PVDD1P8 | 4 | PWR | 1.8V | PLL analog power 1.8V. For the recommended filter, refer to the HWDG. |
| NIF_[4:0]_PVDD1P8 | 5 × 2 | PWR | 1.8V | PLL analog power 1.8V. For the recommended filter, refer to the HWDG. |
| C_PLL_BYP | 1 | I _{PD} | CMOS 1.8V | Factory test only. Pull to GND, R ≤ 1 kΩ. |
| U_PLL_BYP | 1 | I _{PD} | CMOS 1.8V | Factory test only. Pull to GND, R ≤ 1 kΩ. |
| FAB_PLL_BYP | 1 | I _{PD} | CMOS 1.8V | Factory test only. Pull to GND, R ≤ 1 kΩ. |
| C_PLL_BYP_CLK | 1 | I _{PD} | CMOS 0.8V | Factory test only. Leave open. |
| U_PLL_BYP_CLK[1:0] | 2 | I _{PD} | CMOS 0.8V | Factory test only. Leave open. |
| FAB50 SerDes Links, Power, and Test Outputs | | | | |
| FAB50_TX_[47:0]_P/N | 2×48 | O | Differential | Fabric SerDes (50G) differential TX pairs. The link should be AC coupled, and have a termination on the link partner receiver. Verify that the link partner includes this circuitry. |
| FAB50_RX_[47:0]_P/N | 2×48 | I | Differential | Fabric SerDes (50G) differential RX pairs. The SerDes receiver includes internal serial AC coupling and termination. If external AC coupling is required, see Chapter 5, Electrical Specifications . |
| FAB50_[5:0]_PLL[1:0]_PVDD0P8 | 6×2 | PWR | 0.8V | SerDes octet PLL power 0.8V. For the recommended filter, refer to the HWDG. |
| FAB50_[5:0]_RVDD0P8 | 6×4 | PWR | 0.8V | SerDes octet receiver power 0.8V. For the recommended filter, refer to the HWDG. |
| FAB50_[5:0]_TVDD0P8 | 6×3 | PWR | 0.8V | SerDes octet transmitter power 0.8V. For the recommended filter, refer to the HWDG. |

Table 11: Pin List by Function (Continued)

| Signal/Bus Name | Qty. | Type | Tech | Description |
|--|------|------|---------------------|---|
| FAB50_[5:0]_TVDD1P2 | 6×2 | PWR | 1.2V | SerDes octet transmitter power 1.2V. For the recommended filter, refer to the HWDG. |
| FAB50_#_PLL#_TESTOUT_P/N | 4×2 | O | Differential CML | Factory test only. Leave open. |
| FAB50_#_PLL0_LOCK | 3 | O | CMOS 1.8V | Factory test only. Leave open. |
| NIF50 SerDes Links, Power, and Test Outputs | | | | |
| NIF50_TX_[31:0]_P/N | 2×32 | O | Differential | NIF SerDes (50G) differential TX pairs. The link should be AC coupled and have a termination on the link partner receiver. Verify that the link partner includes this circuitry. |
| NIF50_RX_[31:0]_P/N | 2×32 | I | Differential | NIF SerDes (50G) differential RX pairs. The SerDes receiver includes internal serial AC coupling and termination. If external AC coupling is required, see Chapter 5, Electrical Specifications . |
| NIF50_[3:0]_PLL[1:0]_PVDD0P8 | 8 | PWR | 0.8V | SerDes octet PLL power 0.8V. For the recommended filter, refer to the HWDG. |
| NIF50_[3:0]_RVDD0P8 | 4×4 | PWR | 0.8V | SerDes octet receiver power 0.8V. For the recommended filter, refer to the HWDG. |
| NIF50_[3:0]_TVDD0P8 | 4×3 | PWR | 0.8V | SerDes octet transmitter power 0.8V. For the recommended filter, refer to the HWDG. |
| NIF50_[3:0]_TVDD1P2 | 4×2 | PWR | 1.2V | SerDes octet transmitter power 1.2V. For the recommended filter, refer to the HWDG. |
| NIF50_##_PLL#_TESTOUT_P/N | 3×2 | O | Differential CML | Factory test only. Leave open. |
| NIF50_#_PLL0_LOCK | 2 | O | CMOS 1.8V | Factory test only. Leave open. |
| NIF25 SerDes Links, Power, and Test Outputs | | | | |
| NIF25_TX_[127:32]_P/N | 2×96 | O | Differential | NIF SerDes (25G) differential TX pairs. The link should be AC coupled, and have a termination on the link partner receiver. Verify that the link partner includes this circuitry. |
| NIF25_RX_[127:32]_P/N | 2×96 | I | Differential | NIF SerDes (25G) differential RX pairs. The SerDes receiver includes internal serial AC coupling and termination. If external AC coupling is required, see Chapter 5, Electrical Specifications . |
| NIF25_[27:4]_PVDD0P8 | 24 | PWR | 0.8V | SerDes quad PLL power 0.8V. For the recommended filter, refer to the HWDG. |
| NIF25_[27:4]_RVDD0P8 | 24×2 | PWR | 0.8V | SerDes quad receiver power 0.8V. For the recommended filter, refer to the HWDG. |
| NIF25_[27:4]_TVDD0P8 | 24 | PWR | 0.8V | SerDes quad transmitter power 0.8V. For the recommended filter, refer to the HWDG. |
| NIF25_[27:4]_TVDD1P2 | 24 | PWR | 1.2V | SerDes quad transmitter power 1.2V. For the recommended filter, refer to the HWDG. |

Table 11: Pin List by Function (Continued)

| Signal/Bus Name | Qty. | Type | Tech | Description |
|-------------------------------------|------|-----------------|------------------|--|
| NIF25_##_PLL_TESTOUT_P/N | 9×2 | O | Differential CML | Factory test only. Leave open. |
| NIF25_##_PLL_LOCK | 6 | O | CMOS 1.8V | Factory test only. Leave open. |
| BRCM Internal Test and Debug | | | | |
| SCAN_MODE | 1 | I _{PD} | CMOS 1.8V | Factory test only. Pull to GND, R ≤ 1 kΩ. |
| TEST_[4:0] | 5 | I _{PD} | CMOS 1.8V | Factory test only. Pull to GND, R ≤ 1 kΩ. |
| FTEST_[169:0] | 170 | B | CMOS 1.8V | Factory test only. Required connectivity: <ul style="list-style-type: none"> ■ FTEST_[14:13]: Connect to GND. ■ FTEST_[22:17]: Connect to GND. ■ Other FTEST pins: Leave open. |
| AVS_PVTMON_AIO | 1 | B | Analog | Factory test only. Pull to GND, R = 0Ω. |
| C_VTMON_AIO | 1 | B | Analog | Factory test only. Pull to GND, R = 0Ω. |
| FAB_VTMON_AIO | 1 | B | Analog | Factory test only. Pull to GND, R = 0Ω. |
| NIF_VTMON[2:0]_AIO | 3 | B | Analog | Factory test only. Pull to GND, R = 0Ω. |
| FAB_VTMON_AVDD1P8 | 1 | PWR | 1.8V | VTMON analog supply 1.8V. For the recommended filter, refer to the HWDG. |
| NIF_VTMON[2:0]_AVDD1P8 | 3 | PWR | 1.8V | VTMON analog supply 1.8V. For the recommended filter, refer to the HWDG. |
| SPI | | | | |
| SPI_MISO | 1 | O _{PU} | CMOS 1.8V | Factory test only. Leave open. |
| SPI_MOSI | 1 | I _{PU} | CMOS 1.8V | Factory test only. Leave open. |
| SPI_SCK | 1 | I _{PU} | CMOS 1.8V | Factory test only. Leave open. |
| SPI_SSN | 1 | I _{PU} | CMOS 1.8V | Factory test only. Leave open. |
| JTAG | | | | |
| JTAG_TCE | 1 | I _{PD} | CMOS 1.8V | JTAG test enable. 0 during normal device operation. 1 to enable JTAG functionality. An option is to connect with JTAG_TRST_N. |
| JTAG_TCK | 1 | I | CMOS 1.8V | JTAG, clock input. |
| JTAG_TDI | 1 | I | CMOS 1.8V | JTAG, input data. |

Table 11: Pin List by Function (Continued)

| Signal/Bus Name | Qty. | Type | Tech | Description |
|-----------------|------|-----------------|--------------|--|
| JTAG_TDO | 1 | O | CMOS 1.8V | JTAG, output data. |
| JTAG_TMS | 1 | I | CMOS 1.8V | JTAG, TMS test mode input. |
| JTAG_TRST_N | 1 | I _{PD} | CMOS 1.8V | JTAG TAP controller reset. Pull down to GND. (R= < 10 kΩ.) |
| Power | | | | |
| VDDC | 236 | PWR | VDDC | Core power supply. Must be adjusted according to ROV. Working with a VDDC level different from the value required by the ROV may cause the device to fail normal operation or exceed power limits. |
| VDDO | 17 | PWR | 1.8V | IO 1.8V power supply. |
| VSS | 262 | GND | GND | Connect to ground. |
| SRD_AGND | 1627 | GND | GND | Analog ground (return path) for Blackhawk, Falcon, and PCIe SerDes and their supplies (PVDD, RTVDD, RVDD, and TVDD). On the board, it should share the same common GND plane as with VSS pins. |
| HBM_VSS | 73 | GND | GND | Connect to ground. |
| VDDC_SENSE | 1 | O | Analog | Core VDDC sense, from the VDDC supply grid. Used by the system as a feedback to the voltage supply monitor. |
| VSS_SENSE | 1 | O | Analog | Core VSS sense, from the VSS ground grid. Used by the system as a feedback to the voltage supply monitor. |
| NC_OCHK | 1 | NC | None | Orientation check for BRCM testing systems. Leave not connected on customer board. |
| NB | 12 | None | None | No ball. Each corner of the package has three locations with no ball. |

Chapter 5: Electrical Specifications

5.1 Absolute Maximum Ratings

The specifications shown in the following table indicate levels where permanent damage to the device can occur. Functional operation is not guaranteed under these conditions. Operation at absolute maximum conditions for extended periods can adversely affect the long-term reliability of the device.

Table 12: Absolute Maximum Ratings

| Parameter | Min. | Max. | Unit |
|---|-------|-------|------|
| Supply voltage 1.8V | -0.25 | +2.0 | V |
| DDR and HBM VDDO supply voltage ^a | -0.3 | +1.5 | V |
| HBM VDDC supply voltage ^a | -0.3 | +1.5 | V |
| DDR PHY VDDC supply voltage (0.88V) | -0.25 | +1.04 | V |
| HBM VDDP supply voltage ^a | -0.3 | +3.0 | V |
| Supply voltage 1.2V (SerDes TX driver) | -0.25 | +1.38 | V |
| Supply voltage 0.8V, digital core | -0.25 | +1.04 | V |
| Supply voltage 0.8V, SerDes analog | -0.25 | +0.92 | V |
| Thermal diode 3P3 supply voltage | -0.25 | +3.8 | V |
| Storage temperature | -40 | +100 | °C |
| Main die maximum junction temperature (T _J) | — | +110 | °C |
| HBM die maximum junction temperature (T _J) | — | +95 | °C |

a. HBM VDDC, VDDO, and VDDP are from the HBM specification.

5.2 Recommended Operating Conditions

5.2.1 Recommended Operating Voltage

The BCM88800 is equipped with a preprogrammed recommended operating voltage (ROV) stamp indicating the nominal voltage at which the core (VDDC) of the specific BCM88800 device must be operated. The ROV stamp indicates the silicon process of this specific device. Working with a VDDC level that is different from the value required by the ROV may cause the device to fail normal operation or exceed power limits.

Power-up the BCM88800 with the initial state voltage level (according to the following table) for the VDDC. Next, read the ROV stamp from the ROV[2:0] pins. Then, trim the VDDC power supply to the ROV stamp voltage. Perform VDDC trimming before the BCM88800 is initialized. ROV does not change from power-up to power-up.

NOTE: If more than one device is on the card, make sure each device has its own VDDC power supply. This VDDC power rail cannot be shared with any other supply rail.

Table 13: VDDC Voltage Level According to ROV Stamp

| ROV[2:0] | VDDC Voltage Level | Comment |
|----------|--------------------|--|
| 000 | 0.82 | Initial state, detected before device power up |
| 001 | 0.82 | — |
| 010 | 0.90 | — |
| 011 | 0.78 | — |
| 100 | 0.80 | — |
| 101 | 0.84 | — |
| 110 | 0.86 | — |
| 111 | 0.88 | — |

5.2.2 Recommended Operating Voltage Range for DC Supplies

The following table shows the operating voltage range of the DC supplies.

Table 14: Supply Voltage Range

| Parameter | Symbol | Min. (V) | Typ. (V) | Max. (V) |
|---|--|-------------|----------|-------------|
| Core 0.8V ROV ^a digital supply | VDDC | ROV × 0.995 | ROV | ROV × 1.005 |
| SerDes 0.8V PLL supply | xxx_PLL#_PVDD0P8 ^b | 0.80 | 0.825 | 0.85 |
| SerDes 0.8V RX supply | xxx_##_RVDD0P8 ^b | 0.82 | 0.835 | 0.85 |
| SerDes 0.8V TX supply | xxx_##_TVDD0P8 ^b | 0.82 | 0.835 | 0.85 |
| SerDes 1.2V TX driver supply | xxx_##_TVDD1P2 ^b | 1.25 | 1.275 | 1.30 |
| PCIe SerDes 0.8V PLL supply | PCIe_PVDD0P8 | 0.80 | 0.825 | 0.85 |
| PCIe SerDes 0.8V RX/TX supply | PCIe_RTVDD0P8 | 0.80 | 0.825 | 0.85 |
| SerDes RESCAL 0.8V supply | FAB_RESCAL_AVDD0P8 NIF_RESCAL_[2:0]_AVDD0P8 | 0.80 | 0.825 | 0.85 |
| xPLL 1.8V supply ^c RX/TX supply | xPLL_xVDD1P8 | 1.75 | 1.8 | 1.85 |
| I/O 1.8V digital supply | VDDO | 1.71 | 1.8 | 1.89 |
| DDR PHY VDDO and HBM VDDO supply voltage | HBM_VDDO | 1.19 | 1.2 | 1.21 |
| HBM VDDC supply voltage | HBM_VDDC | 1.19 | 1.2 | 1.21 |
| HBM VPP supply voltage | HBM_VPP2P5 | 2.48 | 2.5 | 2.52 |
| DDR_PHY_VDDC (Fixed; not related to ROV VDDC) | DRAM_PHY_VDDC | 0.87 | 0.88 | 0.89 |

a. Recommended operating voltage (ROV).

b. "xxx_" is used for "FAB50_", "NIF50_", and "NIF25_".

c. xPLL_xVDD1P8 is defined for the following signals:

- C_PLL_AVDD1P8
- U_PLL_AVDD1P8
- FAB_PVDD1P8
- NIF_x_PVDD1P8
- TS_PLL_AVDD1P8
- SYNCE_PLL_AVDD1P8
- DRAM_PHY_PVDD1P8

5.3 Device Power Consumption

NOTE: Under the worst-case process, voltage, and temperature (main die at 110°C junction and HBM die at 95°C junction²), the target power consumption for an application based on the BCM88802 device does not exceed 230W (including the HBM), where 218W are consumed by the main die, and 12W are consumed by the HBM. This target is for an application that includes full Ethernet NIF (32 × 50G, 96 × 25G), full fabric 2.4 Tb/s, HBM, and no ILKN.

Broadcom does not provide a separate value for thermal power. Use the maximum power as the thermal power as well.

[Table 15](#) shows the current drawn by the different supply rails.

2. For more information, see the HBM Excursion information in [Footnote b of Table 47, Absolute Thermal Limit Specifications](#).

Table 15: Supply Rails and Their Drawn Current

| Power Rail | Related Ball | Nominal Voltage (V) | Number of Instances ^a | Max. Current Per Instance (mA) | Max. Current from Rail (mA) |
|-------------------------------|--|---------------------|----------------------------------|--------------------------------|-----------------------------|
| Device core supply | Device core supply | 0.82 | 1 | 235,000 | 235,000 |
| General CMOS I/O supply | VDDO | 1.8 | 1 | 1000 | 1000 |
| System PLLs and analog | C_PLL_AVDD1P8 U_PLL_AVDD1P8 | 1.8 | 2 | 70 | 230 |
| | DRAM_PHY_PVDD1P8 | 1.8 | 1 | 10 | |
| | SYNCE_PLL_AVDD1P8 TS_PLL_AVDD1P8 (TS_PLL_AVDD1P8 also drives two other PLLs: BS_PLL and FAB_PLL) | 1.8 | 4 | 20 | |
| | | | | | |
| DRAM-Related | | | | | |
| DRAM_PHY_VDDC | DRAM_PHY_VDDC | 0.88 | 1 | 2000 | 2000 |
| HBM_VDDO and PHY_VDDO | HBM_VDDO ^{b,c} | 1.2 | 1 | 5500 | 5500 |
| HBM_VDDC | HBM_VDDC ^{b,c} | 1.2 | 1 | 6000 | 6000 |
| HBM_VPP | HBM_VPP2P5 ^{b,c} | 2.5 | 1 | 500 | 500 |
| PCIe SerDes Rails | | | | | |
| PCIe core PVDD0P8 (PLL) | PCIE_PVDD0P8 | 0.825 | 1 | 32.5 | 32.5 |
| PCIe lane RVDD0P8 and TVDD0P8 | PCIE_RTVDD0P8 | 0.825 | 4 | 38.75 | 155 |
| Blackhawk SerDes Rails | | | | | |
| Blackhawk core PVDD1P8 | FAB_PVDD1P8 NIF_0_PVDD1P8 | 1.8 | 10 | 1 | 10 |
| Blackhawk core PVDD0P8 | FAB50_[5:0]_PLL[1:0]_PVDD0P8 NIF50_[3:0]_PLL[1:0]_PVDD0P8 | 0.825 | 10 | 282 | 2820 |
| Blackhawk lane RVDD0P8 | FAB50_[5:0]_RVDD0P8 NIF50_[3:0]_RVDD0P8 | 0.835 | 80 | 250 | 20000 |
| Blackhawk lane TVDD0P8 | FAB50_[5:0]_TVDD0P8 NIF50_[3:0]_TVDD0P8 | 0.835 | 80 | 43 | 3440 |
| Blackhawk lane TVDD1P2 | FAB50_[5:0]_TVDD1P2 NIF50_[3:0]_TVDD1P2 | 1.275 | 80 | 20.5 | 1640 |
| Falcon SerDes | | | | | |
| Falcon core PVDD1P8 | NIF_[4:1]_PVDD1P8 | 1.8 | 24 | 0.625 | 15 |
| Falcon core PVDD0P8 | NIF25_[27:4]_PVDD0P8 | 0.8 | 24 | 122.5 | 2940 |
| Falcon lane RVDD0P8 | NIF25_[27:4]_RVDD0P8 | 0.8 | 96 | 140.625 | 13500 |
| Falcon lane TVDD0P8 | NIF25_[27:4]_TVDD0P8 | 0.8 | 96 | 25 | 2400 |
| Falcon lane TVDD1P2 | NIF25_[27:4]_TVDD1P2 | 1.2 | 96 | 18.75 | 1800 |

- The number of instances represents the number of modules, PLLs, SerDes cores, or SerDes lanes in the BCM88800 device. This number is not related to the number of power pins used for the module, PLL, or SerDes core.
- The HBM uses 12W of the total power consumed by HBM_VDDO, HBM_VDDC, and HBM_VPP2P5. The main die uses the rest of the power (because the HBM shares HBM_VDDO with the PHYs, which are part of the main die).
- For more information, see [Footnote b of Table 47, Absolute Thermal Limit Specifications](#).

5.4 Power Supply Filtering

The following table lists the magnitude of supply noise allowed on the different supply rails.

Table 16: Supply Noise Specifications (AC)

| Description | Symbol | Condition | Max. | Unit |
|---|----------------------|----------------------------|------|------|
| Core 0.8V digital supply | VDDC | < 0.3 mΩ; 10 kHz to 10 MHz | 30 | mVpp |
| SerDes 0.8V PLL supply ^a | xxx_PLL[1:0]_PVDD0P8 | 100 kHz to 20 MHz | 3 | mVpp |
| SerDes 0.8V RX supply ^a | xxx_RVDD0P8 | 100 kHz to 20 MHz | 10 | mVpp |
| SerDes 0.8V TX supply ^a | xxx_TVDD0P8 | 100 kHz to 20 MHz | 10 | mVpp |
| SerDes 1.2V TX driver supply ^a | xxx_TVDD1P2 | 100 kHz to 20 MHz | 10 | mVpp |
| PCIe SerDes 0.8V PLL supply | PCIe_PVDD0P8 | 100 kHz to 20 MHz | 3 | mVpp |
| PCIe SerDes 0.8V supply | PCIe_RTVDD0P8 | 100 kHz to 20 MHz | 10 | mVpp |
| PLL 1.8V supply ^b | xPLL_xVDD1P8 | 100 kHz to 20 MHz | 3 | mVpp |
| 1.8V I/O supply | VDDO | 100 kHz to 20 MHz | 100 | mVpp |
| DRAM_PHY VDDO and HBM VDDO | HBM_VDDO | 100 kHz to 20 MHz | 72 | mVpp |
| HBM VDDC | HBM_VDDC | 100 kHz to 20 MHz | 72 | mVpp |
| HBM VPP (2.5V) | HBM_VPP2P5 | 100 kHz to 20 MHz | 150 | mVpp |
| DRAM_PHY VDDC (0.88V) | DRAM_PHY_VDDC | 100 kHz to 20 MHz | 53 | mVpp |

a. "xxx_" is used for FAB50_##, NIF50_##, and NIF25_##.

b. xPLL_xVDD1P8 is defined for the following signals:

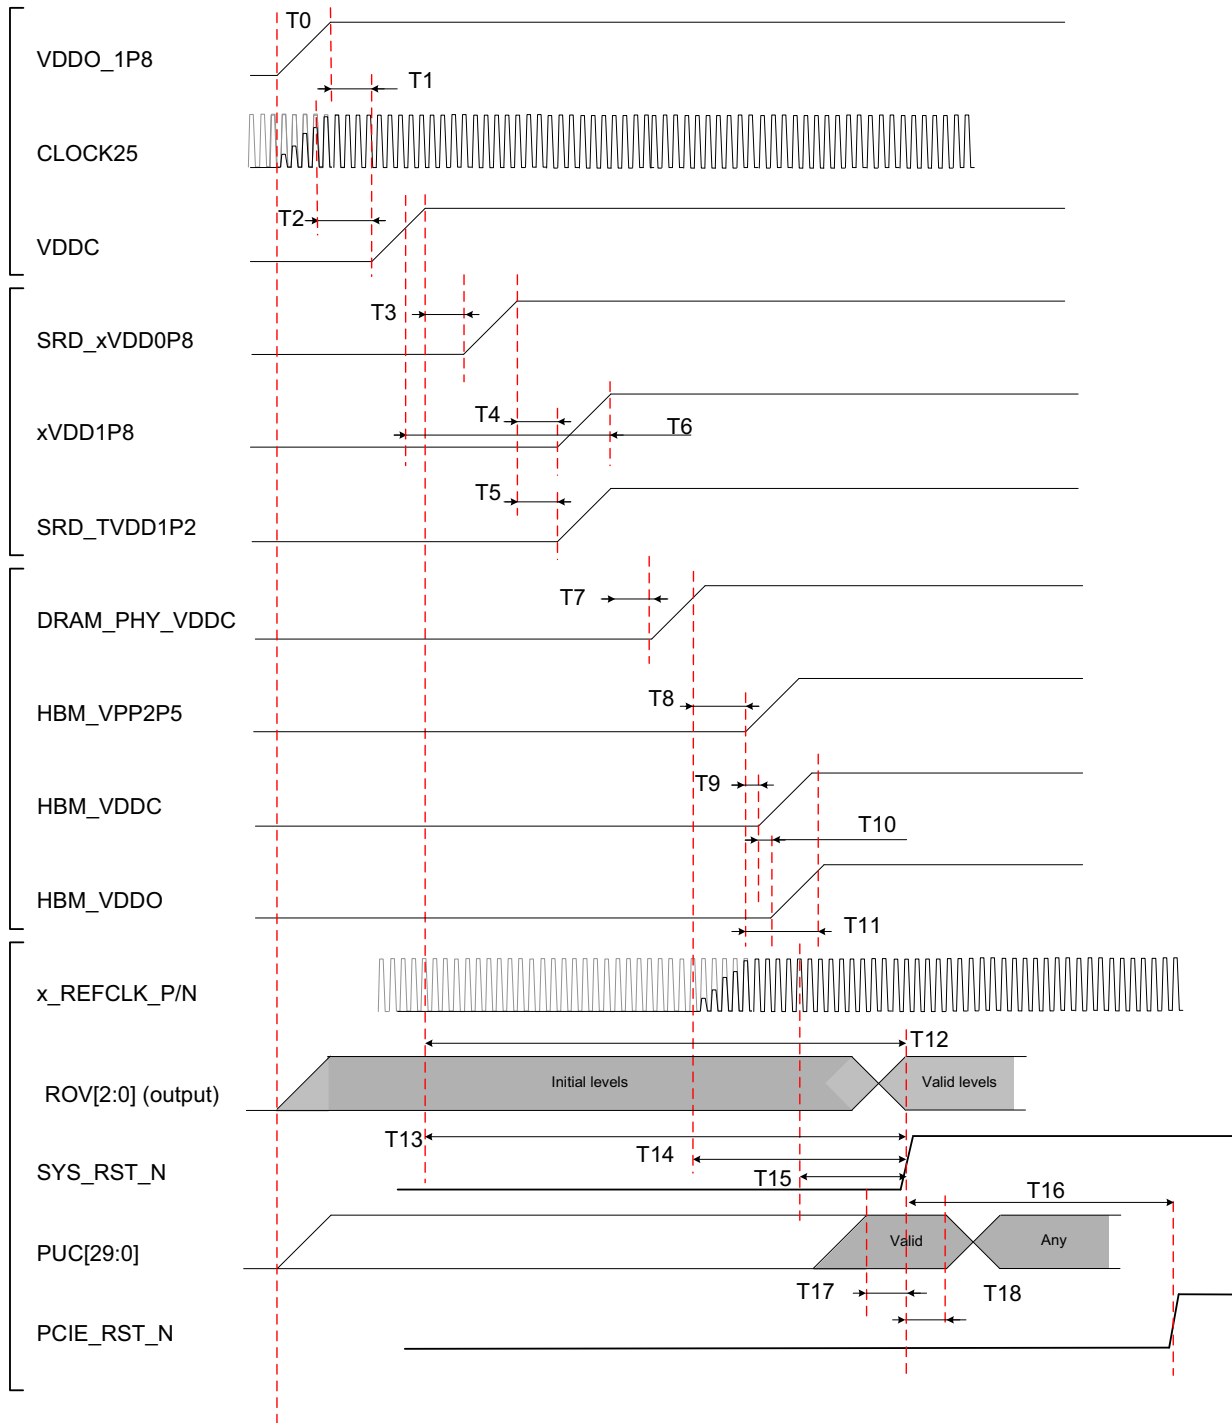
- C_PLL_AVDD1P8
- U_PLL_AVDD1P8
- FAB_PVDD1P8
- NIF_x_PVDD1P8
- TS_PLL_AVDD1P8
- SYNCE_PLL_AVDD1P8
- DRAM_PHY_PVDD1P8

5.5 Power-Up and Reset Sequence

5.5.1 Power-Up Sequence

The following figure illustrates the power-up sequence.

Figure 10: Power-Up Sequence



SRD_xVDD0P8 represents the following rails:

- PCI_PVDD0P8
- PCI_TRVDD0P8
- FAB50_[5:0]_PLL[1:0]_PVDD0P8, FAB50_[5:0]_RVDD0P8, and FAB50_[5:0]_TVDD0P8
- NIF50_[3:0]_PLL[1:0]_PVDD0P8, NIF50_[3:0]_RVDD0P8, and NIF50_[3:0]_TVDD0P8
- NIF25_[27:4]_PLL[1:0]_PVDD0P8, NIF25_[27:4]_RVDD0P8, and NIF25_[27:4]_TVDD0P8

SRD_TVDD1P2 represents the following rails:

- FAB50_[5:0]_TVDD1P2, NIF50_[3:0]_TVDD1P2, and NIF25_[27:4]_TVDD1P2

xVDD1P8 represents the following rails:

- TS_PLL_AVDD1P8
- SYNCE_PLL_AVDD1P8
- DRAM_PHY_PVDD1P8
- C_PLL_AVDD1P8
- U_PLL_AVDD1P8
- FAB_PVDD1P8
- NIF_[4:0]_PVDD1P8
- FAB_VTMON_AVDD1P8³
- NIF_VTMON[2:0]_AVDD1P8³

x_REFCLK_P/N represents the differential reference clocks:

- PCIE_REFCLK_P/N
- TS_PLL_REFCLK_P/N
- DRAM_PHY_REFCLK_P/N
- C_PLL_REFCLK_P/N
- U_PLL_REFCLK_P/N
- FAB_PLL_REFCLK_P/N
- FAB_xx_REFCLK_P/N
- NIF_xx_REFCLK_P/N

NOTE: The power-up sequence can be easily met for voltage ramp-ups as slow as 1V/3 ms.

3. This rail can also be connected to VDDO_1P8.

Table 17: Power-Up Sequence Timing

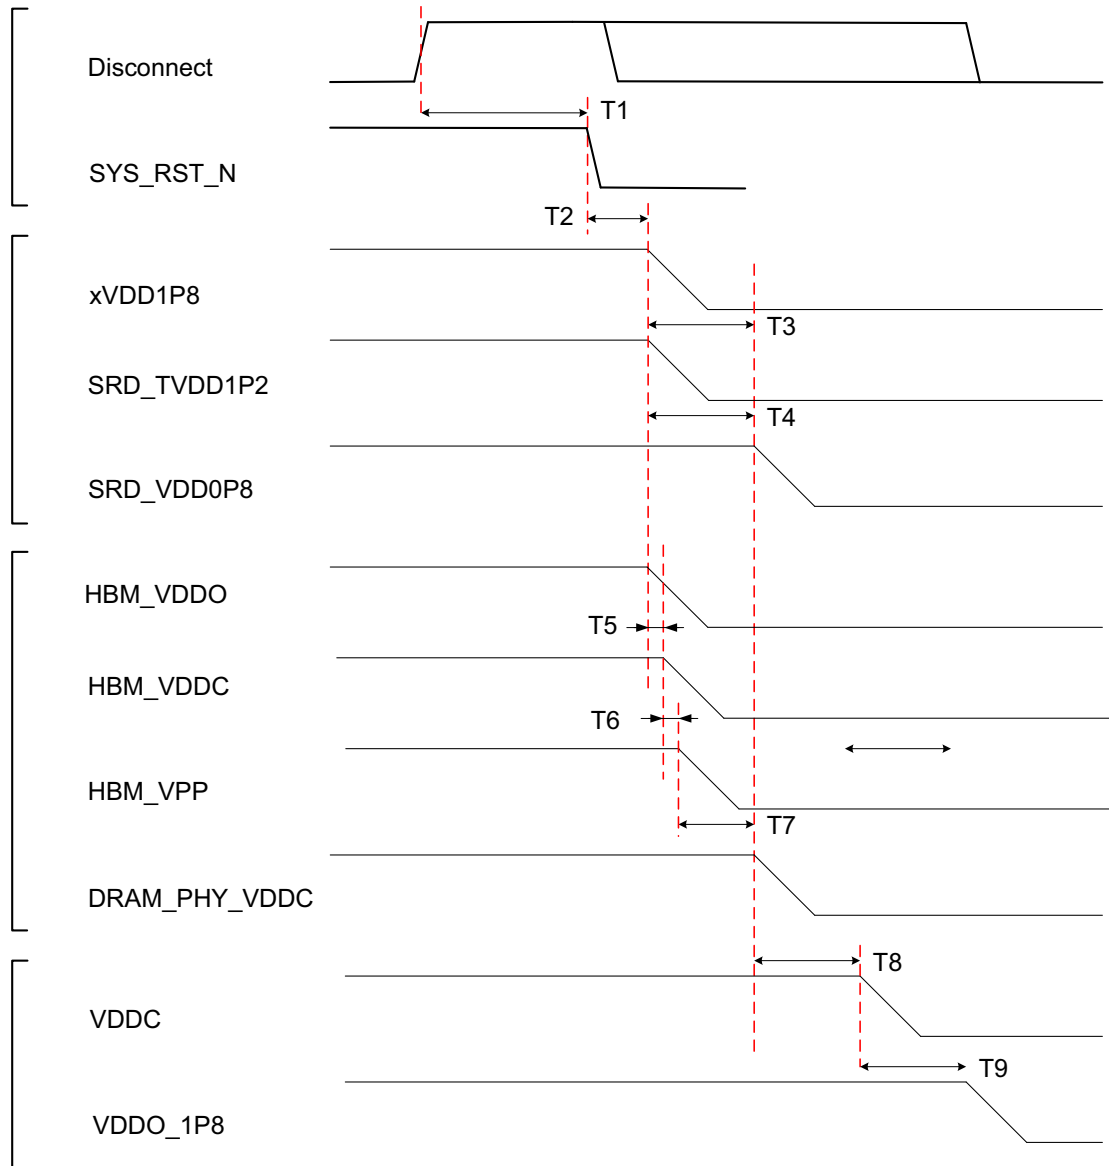
| T Number | Description | Min. | Max. | Unit |
|--------------------|--|------|------------------|---------|
| T0 | Initial condition: <ul style="list-style-type: none"> ■ All voltage rails are below 100 mV. ■ SCAN_MODE, TEST[4:0], JTAG_TCE are 0. ■ SYS_RST_N, PCIE_RST_N are 0. | — | — | — |
| Vramp | Voltage ramp-up time for 1V The rise time of 1V should be longer than 50 μ s, and shorter than 5 ms. | 50 | 5000 | μ s |
| T1 | (VDDO_1P8 > 95%) to (VDDC start) | 0 | 15 | ms |
| T2 | (CLOCK25 valid) to (VDDC start) | 0 | | ms |
| T3 | (VDDC > 95%) to (SRD_xVDD0P8 start) | 0 | 3 | ms |
| T4 ^{a, b} | (SRD_xVDD0P8 > 95%) to (xVDD1P8 start) | 0 | 3 | ms |
| T5 ^{a, c} | (SRD_xVDD0P8 > 95%) to (SRD_TVDD1P2 start) | 0 | 3 | ms |
| T6 | (VDDC = 0.55V) to (xVDD1P8 > 95%) | — | 10 | ms |
| T7 | (xVDD1P8 > 95%) to (DRAM_PHY_VDDC start) | 0 | 3 | ms |
| T8 | (DRAM_PHY_VDDC > 95%) to (HBM_VPP2P5 start) | 0 | 3 ^d | ms |
| T9 ^e | (HBM_VPP2P5 start) to (HBM_VDDC start) | 0 | 2 | ms |
| T10 ^f | (HBM_VDDC start) to (HBM_VDDO start) | 0 | See ^f | ms |
| T11 | (HBM_VPP2P5 start) to (HBM_VDDC, HBM_VDDO > 95%) | 0.01 | 12.5 | ms |
| T12 | (VDDC > 95%) to ROV valid | — | 60 | ms |
| T13 | (VDDC > 95%) to SYS_RST_N de-assertion (rising from 0 to 1) | 60 | — | ms |
| T14 | (DRAM_PHY_VDDC to reach 95%) to SYS_RST_N de-assertion (rising from 0 to 1) | 10 | — | ms |
| T15 | x_REFCLK_P/N valid to SYS_RST_N de-assertion | 10 | — | ms |
| T16 | SYS_RST_N =1 to PCIE_RST_N de-assertion (rising from 0 to 1) | 100 | — | ms |
| T17 | PUC set-up time to SYS_RST_N de-assertion | 160 | — | ns |
| T18 | PUC hold time after SYS_RST_N de-assertion | 160 | — | ns |

- a. There is no required timing between xVDD1P8 and SRD_TVDD1P2.
- b. xVDD1P8 should *not* exceed SRD_xVDD0P8 by more than 1.2V ($xVDD1P8 - SRD_xVDD0P8 < 1.2V$).
- c. SRD_TVDD1P2 should *not* exceed SRD_xVDD0P8 by more than 0.6V ($SRD_TVDD1P2 - SRD_xVDD0P8 < 0.6V$).
- d. The T8 maximum value can be ignored if power-up takes place below 0°C. For powering up at temperatures between –40°C and 0°C, refer to *Hardware Design Guidelines for StrataDNX 16-nm Devices (DNX16-AN1xx)*, Chapter 10: Thermal Aspects, Industrial-Grade Devices.
- e. The HBM_VPP2P5 supply must be applied before or at the same time as HBM_VDDC and must be equal to or higher than HBM_VDDC at all times.
- f. The HBM_VDDC supply must be applied before or at the same time as HBM_VDDO. During the power ramp, HBM_VDDC must be equal to or higher than HBM_VDDO, and $(HBM_VDDC - HBM_VDDO)$ must be less than 0.3V.

5.5.2 Power-Down Sequence

The following figure illustrates the power-down sequence.

Figure 11: Power-Down Sequence



SRD_xVDD0P8 represents the following rails:

- PCI_PVDD0P8, PCI_TRVDD0P8
- FAB50_[5:0]_PLL[1:0]_PVDD0P8, FAB50_[5:0]_RVDD0P8, and FAB50_[5:0]_TVDD0P8
- NIF50_[3:0]_PLL[1:0]_PVDD0P8, NIF50_[3:0]_RVDD0P8, and NIF50_[3:0]_TVDD0P8
- NIF25_[27:4]_PLL[1:0]_PVDD0P8, NIF25_[27:4]_RVDD0P8, and NIF25_[27:4]_TVDD0P8

SRD_TVDD1P2 represents the following rails:

- FAB50_[5:0]_TVDD1P2
- NIF50_[3:0]_TVDD1P2
- NIF25_[27:4]_TVDD1P2

xVDD1P8 represents the following rails:

- TS_PLL_VDD1P8
- SYNCE_PLL_AVDD1P8
- C_PLL_AVDD1P8
- U_PLL_AVDD1P8
- FAB_VTMON_AVDD1P8
- NIF_VTMON[2:0]_AVDD1P8
- DRAM_PHY_PVDD1P8
- FAB_PVDD1P8
- NIF_[4:0]_PVDD1P8

Table 18: Power-Down Sequence Timing

| T Number | Description | Min | Max | Unit |
|--------------------|--|-----|-----|------|
| T1 | DISCONNECT pin or SW disconnect before reset (for single stage system). This is optional to keep traffic data integrity during reset and graceful shutdown. | 15 | — | µs |
| T1 | DISCONNECT pin or SW disconnect before reset (for multi stage system). This is optional to keep traffic data integrity during reset and graceful shutdown. | 10 | — | ms |
| T2 ^{a, b} | SYS_RST_N asserted to first power-off start. This is optional to minimize the number of packets with errors during reset. | 5 | — | µs |
| T3 ^{a, c} | xVDD1P8 drop start to SRD_xVDD0P8 drop start. | 0 | — | ms |
| T4 ^{a, d} | SRD_TVDD1P2 drop start to SRD_xVDD0P8 drop start. | 0 | — | ms |
| T5 | HBM_VDDO drop start to HBM_VDDC drop start. | 0 | — | ms |
| T6 | HBM_VDDC to HBM_VPP drop start. (HBM_VPP2P5 supply must be equal to or higher than HBM_VDDC.) | 0 | — | ms |
| T7 | HBM_VPP drop start to DRAM_PHY_VDDC drop start. | 0 | — | ms |
| T8 | SRD_xVDD0P8/DRAM_PHY_VDDC drop start to VDDC drop start. | 0 | — | — |
| T9 | VDDC drop start to VDDO drop start. | 0 | — | ms |

a. Timing is not required between xVDD1P8 and SRD_TVDD1P2.

b. Power rails can start the drop without delay, however [Footnote c](#) and [Footnote d](#) must be met.

c. xVDD1P8 should *not* exceed SRD_xVDD0P8 by more than 1.2V ($xVDD1P8 - SRD_xVDD0P8 < 1.2V$).

d. SRD_TVDD1P2 should *not* exceed SRD_xVDD0P8 by more than 0.6V ($SRD_TVDD1P2 - SRD_xVDD0P8 < 0.6V$).

5.5.3 Fail-Safe Considerations

The CMOS I/Os of the BCM88800 are powered by 1.8V VDDO. Input signals can be driven before VDDO (1.8V) is supplied.

5.5.4 Warm Reset

It is possible to reset the device during normal device operation and not just during power-up. To place the device in reset, assert SYS_RST_N and PCIE_RST_N to low. Keep SYS_RST_N low for at least 10 μs before setting it high (releasing the device from the reset condition). The PCIE_RST_N should go high 100 ms after the SYS_RST_N.

When taking the device from the reset condition, all the requirements for logic signals that are part of the power-up and reset sequence should be met (see Section 5.5, Power-Up and Reset Sequence).

5.5.5 HBM-Only Power Down and Power Up

If it is necessary to power down only the HBM and later power it up, note the following requirements:

- The HBM_VPP2P5 supply must be equal to or higher than HBM_VDDC at all times.
- The HBM_VDDC supply must be applied before or at the same time as HBM_VDDO.
- During the power ramp, $HBM_VDDC > HBM_VDDO$, and $(HBM_VDDC - HBM_VDDO) < 0.3V$.

The following figure illustrates these requirements.

Figure 12: HBM-Only Power Down And Power Up




Table 19: T1 Definition

| Parameter | Description | Min. | Max. | Unit |
|-----------|--|------|------|------|
| T1 | (HBM_VPP2P5 start) to (HBM_VDDC, HBM_VDDO > 95%) | 0.01 | 12.5 | ms |

5.6 Power-Up Configuration Word

The following table describes the power-up configuration word functionality for PUC[29:0]. When needed, use pull-up or pull-down resistors where $R < 5 \text{ k}\Omega$.

Table 20: Power-Up Configuration Signal Description

| PUC | Function | Description |
|-------------|---------------------|--|
| PUC_[9:0] | CORE_PLL_N_DIV[9:0] | Set to 240 (PUC[19:0] = 018F0) For A0, set to 216 (0xD8) (PUC[19:0] = 0x018D8) |
| PUC_[18:10] | CORE_PLL_M_DIV[8:0] | Set to 6 (PUC[19:0] = 0x018F0) For A0, PUC[19:0] = 0x018D8 |
| PUC_[19] | Factory test | Set to 0 |
| PUC_[21:20] | I2C_SA[1:0] | Device I2C (BSC) slave address LSB. When the I2C (BSC) interface is used, the device physical address is made up of the following 7 bits: <ul style="list-style-type: none"> ■ [A6, A5, A4, A3, A2] fixed as 0x10001. ■ [A1, A0] set according to [PUC_21, PUC_20]. Available addresses are 0x44, 0x45, 0x46, and 0x47. |
| PUC_[25:22] | Factory test | Set to 0x0. |
| PUC_[26] | PCIE_QSPI_ENABLE | Set to 1. |
| PUC_[29:27] | Factory test | Set to 0x0. |

5.7 DC Electrical Specifications

5.7.1 1.8V Digital I/Os


The following table lists the DC specifications of the CMOS 1.8V I/Os.

Table 21: DC Specification for CMOS 1.8V I/O

| Parameters | Symbol | Conditions | Min. | Max. | Unit |
|---------------------------|----------|---------------------------|-------------------------------|-------------------------------|------|
| Input low-level voltage | V_{IL} | — | See Figure 13 | $V_{DDO} \times 0.32$ | V |
| Input high-level voltage | V_{IH} | — | $V_{DDO} \times 0.65$ | See Figure 13 | V |
| Output low-level voltage | V_{OL} | $I_{OUT} = 4 \text{ mA}$ | — | 0.4 | V |
| Output high-level voltage | V_{OH} | $I_{OUT} = -4 \text{ mA}$ | $V_{DDO} - 0.4$ | — | V |

The following figure shows CMOS 1.8V overshoot definitions.

Figure 13: CMOS 1.8V Overshoot Definitions



The high overshoot can be up to $V_{DDO} + 0.3V$, and the low overshoot can be down to $GND - 0.3V$. The duration, measured on a level of half the peak, should be less of 10% of the duty cycle and less than 10 ns.

5.7.2 BSC/I²C

The following table lists the DC specifications of the BSC I/Os.

Table 22: DC Specification for CMOS 1.8V BSC/I²C

| Parameter | Symbol | Conditions | Min. | Typ. | Max. | Unit |
|--------------------------|----------|-----------------|---------------------|------|---------------------|------|
| Input low level voltage | V_{IL} | — | — | — | $0.3 \times V_{DD}$ | V |
| Input high level voltage | V_{IH} | — | $0.7 \times V_{DD}$ | — | — | V |
| Output low level current | V_{OL} | $V_{OL} = 0.4V$ | 8 | — | — | mA |

5.8 AC Electrical Specifications

5.8.1 BSC/I2C Interface Timing

The BSC (I2C-compatible) interface supports standard I2C mode and can operate at up to 100 kHz.

The BSC interface of the BCM88800 can operate in slave mode only.

The BCM88800 samples BSC_SDA during a write operation and drives BSC_SDA during a read operation.

Figure 14: BSC Timing Diagram

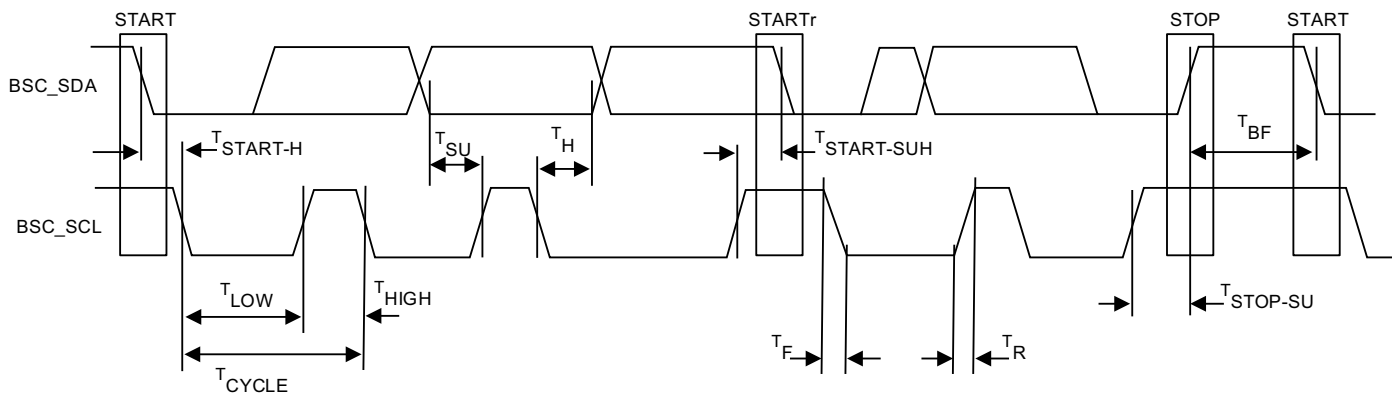


Table 23: BSC Slave Standard-Mode Timing

| Parameter | Symbol | Min. | Typ. | Max. | Unit |
|--|----------------|------|------|------|---------|
| BSC_SCL clock frequency | f_{CLK} | — | — | 100 | kHz |
| BSC_SCL cycle time | T_{CYCLE} | 10 | — | — | μs |
| BSC_SCL low time | T_{LOW} | 4.7 | — | — | μs |
| BSC_SCL high time | T_{HIGH} | 4.0 | — | — | μs |
| Data hold time | T_H | 0.0 | — | — | μs |
| Data setup time | T_{SU} | 250 | — | — | ns |
| Rise time, data ^a | T_R | — | — | 1000 | ns |
| Fall time, data | T_F | — | — | 300 | ns |
| Hold time, start, or repeated start | $T_{START-H}$ | 4.0 | — | — | μs |
| Setup time, repeated start | $T_{START-SU}$ | 4.7 | — | — | μs |
| Setup time, stop | $T_{STOP-SU}$ | 4.0 | — | — | μs |
| Bus free time (between stop and start) | T_{BF} | 4.7 | — | — | μs |

a. BSC_SCL is an open-drain input, and BSC_SDA is an open-drain input/output. The rise time is dependent on the strength of the external pull-up resistor, which must be chosen to meet the rise time requirement.

5.8.2 Management Interface Timing

5.8.2.1 MDIO AC Characteristics

Figure 15: MIIM Interface Timing Diagram




Table 24: MDC and MDIO Timing

| Parameter | Symbol | Min. | Typ. | Max. | Unit |
|---------------------|-------------|------|------|------|------|
| MDC clock frequency | f_{CLK} | — | 2.5 | 12.5 | MHz |
| MDC cycle time | T_{CYCLE} | 80 | 400 | — | ns |
| MDC duty cycle | — | 40 | — | 60 | % |
| MDIO setup time | T_{SU} | 20 | — | — | ns |
| MDIO hold time | T_H | 0 | — | — | ns |
| MDIO output delay | T_{OD} | 10 | — | 35 | ns |

NOTE:

- Output load conditions = 25 pF.
- The external device to conform to the IEEE specifications.
- The MDC rate and the MDIO output delay are configurable.

5.8.3 SyncE Recovered Clocks

The SYNCE_[1:0]_CLK_OUT_P/N recovered clocks are differential CML outputs. The following figure shows the recommended connectivity of these clocks.

Figure 16: SYNCE_[1:0]_CLK_OUT Recommended Connectivity

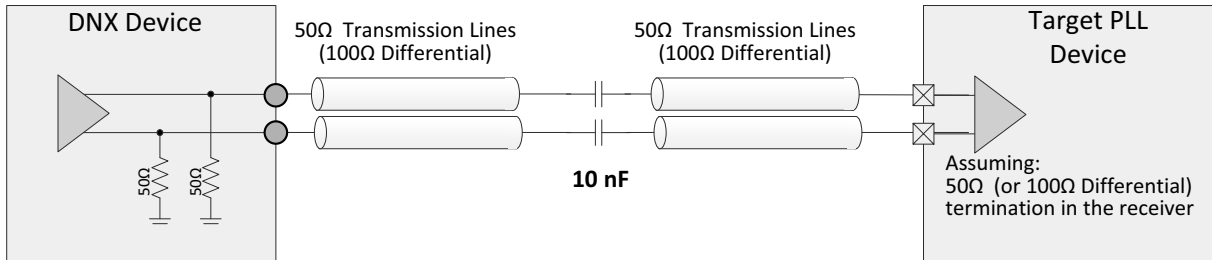


Table 25: SYNCE_[1:0]_CLK_OUT_P/N Output CML Clock Timing

| Parameters | Symbol | Min. | Typ. | Max. | Unit |
|---------------------------------|-------------------|------|------|------|------|
| Recovered clock frequency | F_{CLKOUT_AVG} | — | 25 | — | MHz |
| Output clock differential swing | $V_{DIFF-PK-PK}$ | 0.5 | 0.94 | 1.2 | Vppd |

5.8.4 LED Timing

LED[4:0]_CLK and LED[4:0]_DATA are outputs. LED[4:0]_CLK output clock period is 200 ns (5.0 MHz).

Figure 17: LED Timing Diagram




Table 26: LED Timing^a

| Parameter | Symbol | Min | Typ | Max | Unit |
|-----------------------|-------------|-----|-----|-----|------|
| LED frequency | F_{TCK} | — | 5 | 5 | MHz |
| LED period | T_{CYCLE} | 200 | — | — | ns |
| LED clock HIGH | T_{HIGH} | 70 | 100 | 130 | ns |
| LED clock LOW | T_{LOW} | 70 | 100 | 130 | ns |
| LED data output valid | T_{OV} | -15 | — | 15 | ns |

a. Timing figures are specified at the 50% crossing thresholds.

5.8.5 Out-of-Band Flow Control Timing

The out-of-band flow-control (OOBFC) interfaces support two modes of operation, which impacts both the protocol and the low-level timing.

5.8.5.1 SPI4.2 Flow-Control Mode

Figure 18: OOB Flow-Control Timing in SPI4.2 Mode




Table 27: OOB Flow-Control Timing Specifications in SPI4.2 Mode

| Parameters | Symbol | Min. | Typ. | Max. | Unit |
|--|--------|------|--------------|--------|------|
| FC_#_T_CLK, OOB output clock frequency ^a | 1/T1 | 62.5 | FCORE/n | 166.66 | MHz |
| FC_#_T_CLK, OOB output clock duty cycle | — | 45 | — | 55 | % |
| FC_#_T_STAT[1:0] output stable window size ^b | T2 | 5 | T1 – 1 | 15 | ns |
| FC_#_T_STAT[1:0] output stable before clock ^c | T3 | — | Programmable | — | ns |
| FC_#_T_STAT[1:0] output stable after clock ^c | T4 | — | Programmable | — | ns |
| FC_#_R_CLK, OOB input clock frequency | 1/T5 | — | — | 200 | MHz |
| FC_#_R_CLK, OOB input clock duty cycle | — | 40 | — | 60 | % |
| FC_#_R_STAT[1:0] input setup time ^d | T6 | 0.5 | — | — | ns |
| FC_#_R_STAT[1:0] input hold time ^d | T7 | 0.5 | — | — | ns |

a. FC_#_T_CLK output interface clock frequency is user-programmable. Supported rates are core clock divided by 2n (where n is 3 to 8). For a 1G core clock, this equals 166.66, 125, 100, 83.33, 71.43, or 62.5 MHz.

b. The given value is for FC_#_T_CLK at F(core/6). For F(core/8), add 2 ns. For F(core/16), add 10 ns.

c. The phase of FC_#_T_CLK in relation to FC_#_T_STAT[1:0] is programmable

d. The sampling clock edge of the input FC_#_R_STAT[1:0] with respect to the input clock can be programmed to either the rising or falling edge. Figure 18 illustrates a rising edge configuration.

5.8.5.2 Interlaken Flow-Control Mode

Figure 19: OOB Flow-Control Timing in Interlaken Mode




Table 28: OOB Flow-Control Timing Specifications in Interlaken Mode

| Parameters | Symbol | Min. | Typ. | Max. | Unit |
|---|--------|------|---------------------|--------|------|
| FC_#_T_CLK, OOB output clock frequency ^a | 1/T1 | 62.5 | FCORE/n | 166.66 | MHz |
| FC_#_T_CLK, OOB output clock duty cycle | — | 45 | — | 55 | % |
| FC_#_T_STAT and FC_#_T_SYNC output stable window size ^b | T2 | 2 | $0.5 \times T1 - 1$ | 7 | ns |
| FC_#_T_STAT and FC_#_T_SYNC output stable before clock ^c | T3 | — | Programmable | — | ns |
| FC_#_T_STAT and FC_#_T_SYNC output stable after clock ^c | T4 | — | Programmable | — | ns |
| FC_#_R_CLK, OOB input clock frequency | 1/T5 | — | — | 200 | MHz |
| FC_#_R_CLK, OOB input clock duty cycle | — | 45 | — | 55 | % |
| FC_#_R_STAT and FC_#_R_SYNC input setup time | T6 | 0.5 | — | — | ns |
| FC_#_R_STAT and FC_#_R_SYNC input hold time | T7 | 0.5 | — | — | ns |

a. FC_#_T_CLK output interface clock frequency is user-programmable. Supported rates are core clock divided by 6, 8, 10, 12, 14, or 16. For a 1G core clock, this equals 166.66, 125, 100, 83.33, 71.43, or 62.5 MHz.

b. ILKN FC mode is DDR. The value is for FC_#_T_CLK at F(core/6). For F(core/8), add 1 ns. For F(core/16), add 5 ns.

c. The phase of FC_#_T_CLK in relation to FC_#_T_STAT and FC_#_T_SYNC is programmable.

5.8.6 BroadSync and Time Sync Timing

The following figure and table show the slave mode input timing.

Figure 20: BroadSync Input Timing – Slave Mode




Table 29: BroadSync Input Timing – Slave Mode

| Parameters | Symbol | Min. | Typ. | Max. | Unit |
|---------------------------------------|------------|------|------|------|------|
| BS_BIT_CLK cycle time | t_{CYC} | — | 100 | — | ns |
| BS_BIT_CLK duty cycle | t_{HIGH} | 40 | — | 60 | ns |
| BS_TIME_VAL; BS_SYNC input setup time | t_{ISU} | 20 | — | — | ns |
| BS_TIME_VAL; BS_SYNC input hold time | t_{IH} | 0 | — | — | ns |

The following figure and table show the master mode input timing.

Figure 21: BroadSync Output Timing – Master Mode

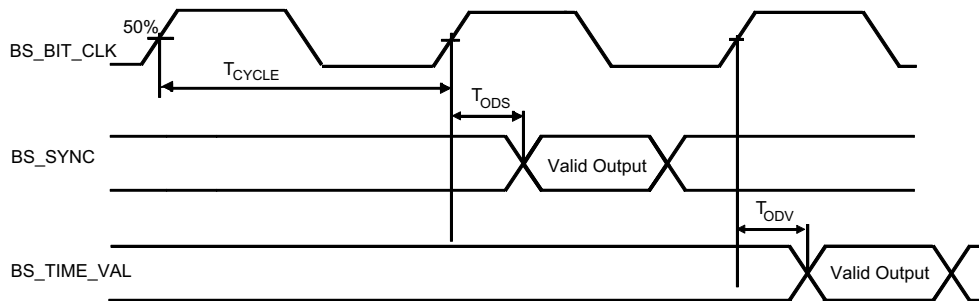


Table 30: BroadSync Output Timing – Master Mode

| Parameters | Symbol | Min. | Typ. | Max. | Unit |
|--------------------------|------------|------|------|------|------|
| BS_BIT_CLK cycle time | t_{CYC} | — | 100 | — | ns |
| BS_BIT_CLK duty cycle | t_{HIGH} | 40 | — | 60 | ns |
| BS_SYNC output delay | t_{ODS} | 0 | — | 25 | ns |
| BS_TIME_VAL output delay | t_{ODV} | 0 | — | 25 | ns |

5.8.7 PCIe Interface

The PCIe core of the BCM88800 supports PCIe Gen1 (2.5G), Gen2 (5G), and Gen3 (8G). The following sections provide basic electrical specifications.

5.8.7.1 PCIe Receiver

The following table lists the specifications of the PCIe SerDes receiver.

Table 31: PCIe SerDes RX Specifications

| Parameter | Symbol | Min. | Typ. | Max. | Unit |
|---------------------------------------|---------------|------|------|------|-------|
| Input differential swing ^a | $V_{IN-DIFF}$ | 85 | — | 1200 | mVppd |
| Input differential termination | R_{TERM} | 80 | 100 | 120 | Ohm |

a. The Receiver input should be externally AC coupled.

5.8.7.2 PCIe Transmitter

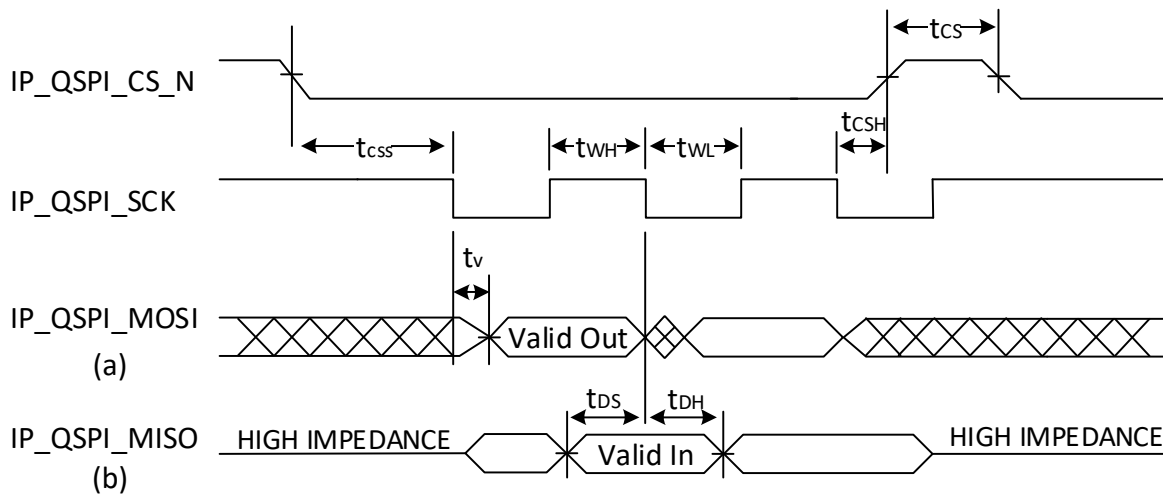
The following table lists the specifications of the PCIe SerDes transmitter.

Table 32: PCIe SerDes TX Specifications

| Parameter | Symbol | Min. | Typ. | Max. | Unit |
|--|----------------|------|------|------|-------|
| Output differential termination | R_{TERM} | 80 | 100 | 120 | Ohm |
| Output differential swing (programmable) | $V_{OUT-DIFF}$ | 400 | — | 1200 | mVppd |
| Output common mode | V_{OUT-CM} | — | 400 | — | mV |

5.8.8 QSPI Flash Interface

Figure 22: IP_QSPI Timing (Boot Read Mode Using BSPI Controller)



(a): also valid for IP_QSPI_MISO in dual/quad mode; also valid for IP_QSPI_WP_N, IP_QSPI_HOLD_N in quad mode
 (b): also valid for IP_QSPI_MOSI in dual/quad mode; also valid for IP_QSPI_WP_N, IP_QSPI_HOLD_N in quad mode

Table 33: IP_QSPI Timing (Boot Read Mode Using BSPI Controller)

| Parameter | Symbol | Min. | Typ. | Max. | Unit |
|-----------------------|-----------|---------------------|-------------------|------|------|
| SCK frequency | F_{SCK} | — | 62.5 ^a | — | MHz |
| SCK clock LOW period | t_{WL} | $0.5/F_{SCK} - 0.5$ | — | — | ns |
| SCK clock HIGH period | t_{WH} | $0.5/F_{SCK} - 0.5$ | — | — | ns |
| CS lead time | t_{CSS} | $1/F_{SCK} - 2.9$ | — | — | ns |
| CS trail time | t_{CSH} | -1.6 | — | — | ns |
| MOSI output valid | t_v | -1.6 | — | 3 | ns |
| MISO input setup | t_{SU} | 4 | — | — | ns |
| MISO input hold | t_H | 1.3 | — | — | ns |

a. The QSPI controller issues only FAST_READ commands, as opposed to READ (03h) commands. Therefore, the QSPI device's operating frequency should be based on the Fast Read commands rather than Read (03h) commands.

5.9 Reference Clocks

5.9.1 CLOCK25 Reference Clock

The CLOCK25 is a CMOS 1.8V input. This is a free-running clock (without an internal PLL) that is used for SYS_RST_N propagation, device initialization, and monitoring. The required parameters are specified in [Table 21, DC Specification for CMOS 1.8V I/O](#).

Table 34: CLOCK25 Specifications

| Parameters | Symbol | Conditions | Min. | Typ. | Max. | Unit |
|--|-------------|------------|------|------|------|------|
| Reference clock frequency | F_{REF} | — | — | 25 | — | MHz |
| Reference clock accuracy | ACC_{REF} | — | -100 | — | +100 | PPM |
| Reference duty cycle | F_{DC} | — | 40 | — | 60 | % |
| Rise/fall time | T_R/T_F | 20% to 80% | — | — | 4.0 | ns |
| Input Jitter cycle-to-cycle, peak-to-peak, 10K Samples | J_{IN} | — | — | — | 100 | ps |

5.9.2 Core and Microcontroller PLL Reference Clocks

The BCM88800 holds two PLLs that drive the internal data path and logic. These PLLs are the core clock PLL (C_PLL) and the microcontroller clock PLL (U_PLL). The C_PLL_REFCLK_P/N and U_PLL_REFCLK_P/N, respectively drive these PLLs.

The following table lists the specifications of the core PLL (C_PLL) and the microcontroller PLL (U_PLL) reference clocks.

Table 35: C_PLL and U_PLL Reference Clock Specifications

| Parameter | Symbol | Min. | Typ. | Max. | Unit |
|---|-------------|------|------|------------------|---------|
| Reference clock frequency | F_{REF} | — | 25 | — | MHz |
| Reference clock accuracy, total stability | ACC_{REF} | -32 | — | +32 | PPM |
| Duty cycle | F_{DC} | 45 | 50 | 55 | % |
| CML input reference clock swing (diff.) ^a | — | 500 | — | 2000 | mVppd |
| Input CML differential termination | — | 80 | 100 | 120 | Ohm |
| Input CML clock slew. Based on rise/fall time (10% to 90%) | T_R/T_F | — | — | 0.8 ^b | ns/Vppd |
| Input jitter (12 kHz to 5 MHz, RMS) | J_{IN} | — | — | 0.35 | ps |

a. Input should have external AC coupling. The device has an internal 100Ω differential termination.

b. There is a 0.4-ns rise/fall time for a 500-mV swing and a 1.6-ns rise/fall time for a 2000-mV swing.

5.9.3 PCIe PLL Reference Clock

The following table lists the specifications of the PCIe PLL reference clock.

Table 36: PCIe PLL Reference Clock Specifications

| Parameter | Symbol | Min. | Typ. | Max. | Unit |
|---|---------------|------|------|----------------------------|-------|
| Reference clock frequency | F_{REF} | — | 100 | — | MHz |
| Reference clock accuracy | ACC_{REF} | -300 | — | +300 | PPM |
| Reference duty cycle | F_{DC} | 40 | — | 60 | % |
| Input differential swing ^a | $V_{IN-DIFF}$ | 600 | — | 1200 | mVppd |
| Rise/fall time (20% to 80%) | T_R/T_F | 200 | — | 900 | ps |
| Reference clock jitter (RMS, up to 100 MHz) | — | — | — | Follow PCIe specifications | ps |

a. Should be AC coupled, and external 100Ω termination is required.

5.9.4 SerDes Reference Clock

The following reference clocks drive SerDes-related PLLs:

- FAB_PLL_REFCLK_P/N drives internal logic for fabric SerDes
- FAB_[1:0]_REFCLK_P/N drive fabric Blackhawk SerDes PLLs
- NIF_[4:0]_REFCLK_P/N drive NIF Blackhawk and Falcon SerDes PLLs

The following table lists the specifications of the SerDes-related reference clock.

Table 37: SerDes Reference Clock Specifications

| Parameter | Symbol | Min. | Typ. | Max. | Unit |
|---|-------------|------|--------|------|-------|
| Reference clock frequency | F_{REF} | — | 156.25 | — | MHz |
| Reference clock accuracy, total stability | ACC_{REF} | -50 | — | +50 | PPM |
| Duty cycle | F_{DC} | 40 | 50 | 60 | % |
| CML input reference clock swing (diff) ^a | — | 800 | — | 1400 | mVppd |
| Input CML differential termination | — | 80 | 100 | 120 | Ohm |
| Differential rise/fall time (20% to 80% for minimum amplitude) | T_R/T_F | 50 | — | 400 | ps |
| Input jitter (12 kHz to 20 MHz, RMS) | J_{IN} | — | — | 0.15 | ps |
| Requirement for FAB_PLL_REFCLK only ^b input jitter (12 kHz to 20 MHz, RMS) | J_{IN} | — | — | 10 | ps |

a. Input should have external AC coupling. Device has an internal 100Ω differential termination and internal bias circuit for CM.

b. The loose requirement is only for FAB_PLL_REFCLK if it is using a source or path different from other clocks referenced in the table.

5.9.5 TS_PLL Reference Clock

Table 38: TS_PLL Reference Clock Definitions

| Parameters | Symbol | Min. | Typ. | Max. | Unit |
|---|------------------|------|------|------|---------|
| Reference clock frequency | F_{REF} | — | 25 | — | MHz |
| Reference clock accuracy | ACC_{REF} | -25 | — | +25 | PPM |
| Reference clock duty cycle | F_{REF_DC} | 40 | — | 60 | % |
| CML input reference clock differential swing ^a | $V_{DIFF-PK-PK}$ | 0.5 | — | 1.8 | Vppd |
| Input reference clock rise time/fall time (10% to 90%) | T_R/T_F | — | — | 1 | ns/Vppd |
| Input jitter (12 kHz to 20 MHz, RMS) | J_{IN} | — | — | 10 | ps |

a. Input should have external AC coupling.

NOTE: For some IEEE 1588 applications, OCXO is required. For more information, contact Broadcom support.

5.9.6 DRAM PHY PLL Reference Clock

Table 39: DRAM_PHY PLL Reference Clock Definitions

| Parameters | Symbol | Min. | Typ. | Max. | Unit |
|--|------------------|------|------|------|---------|
| Reference clock frequency | F_{REF} | — | 100 | — | MHz |
| Reference clock accuracy | ACC_{REF} | -100 | — | +100 | PPM |
| Reference clock duty cycle | F_{REF_DC} | 40 | — | 60 | % |
| CML input reference clock swing ^a | $V_{DIFF-PK-PK}$ | 0.5 | — | 1.8 | Vppd |
| Input reference clock rise time/fall time (10% to 90%) | T_R/T_F | — | — | 1 | ns/Vppd |
| Input reference clock jitter (RMS, integration range 12 kHz to 20 MHz) | J_{IN} | — | — | 1 | ps |

a. Amplitude is pk-pk differential (clk_out_p to clk_out_n). The reference clock should be externally AC coupled.

5.10 Blackhawk SerDes Operating Conditions

The following table lists the Blackhawk SerDes operating conditions.

Table 40: Blackhawk SerDes Operating Conditions

| Parameter | Symbol | Min. | Typ. | Max. | Unit |
|-------------------------|-----------------|---------|------|--------|------|
| Data rate, fabric | F_{b_Fabric} | 23 | — | 53.125 | Gb/s |
| Data rate, NIF Ethernet | F_{b_NIF} | 10.3125 | — | 53.125 | Gb/s |
| Data rate, ILKN | F_{b_ELK} | 10.3125 | — | 53.125 | Gb/s |
| VCO | F_{VCO} | 15 | — | 30 | GHz |

The Blackhawk SerDes high-speed lanes are organized in cores. Each core, named Blackhawk, includes eight SerDes. Each one of the eight SerDes is an independent lane suitable for optical and backplane applications, operating in either PAM4 or NRZ line coding.

5.10.1 Blackhawk SerDes Features

The following sections describe the main features of the Blackhawk SerDes.

5.10.1.1 General Features

Blackhawk supports the following general features:

- A block of eight SerDes supporting eight serial links.
- Line rates (depending on the application) from 10.3125 Gb/s to 53.125 Gb/s per serial link (PAM4 53.125 Gb/s, NRZ 25.78125 Gb/s).
- Two independent PLLs in each Blackhawk.
- Integrated Arm micro subsystem: Monitors the signal and adaptively adjusts the gain, peaking-filter, and DFE coefficients to optimally equalize and restore the signal. During Clause72 and Clause 93 TX/RX link-training, this microcontroller is also responsible for returning feedback to the far-link partner to optimally tune its transmitter.

5.10.1.2 Debug Features

Blackhawk supports the following debugging features:

- PRBS 7, 9, 10, 11, 13 15, 20, 23, 31, 49, and 58 generator and checker with burst error length measurement.
- Digital loopback: Turns around the transmit data before it goes into the analog front-end (this is a digital data path loopback for the data coming from the PCS/MAC transmit side).
- AC-JTAG for both TX and RX.
- Full-range horizontal and vertical eye diagnostics.

5.10.2 Blackhawk SerDes Supported Rates

The following table shows the rates that the fabric SerDes supports.

Table 41: Fabric SerDes Supported Rates

| VCO Rate (GHz) | PAM4 Rate (Gb/s) | NRZ OSx1 Rate (Gb/s) | NRZ OSx2 Rate (Gb/s) |
|----------------|------------------|----------------------|----------------------|
| 20.625 | — | 20.625 | — |
| 23 | — | 23 | — |
| 25 | 50 | 25 | — |
| 25.78125 | — | 25.78125 | — |
| 26.5625 | 53.125 | — | — |

For the rates that the NIF supports, see [Table 2, Supported Port Modes](#).

For the rates that the ELK interface supports, see [Table 5, ILKN Supported Rates and Electrical Specifications Standards](#).

5.10.3 Blackhawk SerDes Receiver

The Blackhawk receiver features are as follows:

- Integrated AC coupling on RX inputs.
- Three-stage analog peaking filter (PF).
- Three-stage controlled variable gain amplifier (VGA).
- 14-tap DFE (six fixed and eight floating) with adaptive control slicer.
- Clock-phase interpolation in receiver timing recovery.
- Programmable RX polarity inversion.

The following table lists the electrical characteristics of the SerDes receiver.

Table 42: SerDes Receiver Characteristics

| Parameter | Symbol | Min. | Typ. | Max. | Unit | Notes |
|---------------------------|---------------|------|------|------|-------|--|
| Input common mode | V_{IN-CM} | — | 530 | 900 | mV | Higher than the max V_{IN-CM} requires external AC capacitor on board, 100 nF. |
| Input differential swing | $V_{IN-DIFF}$ | 85 | — | 1600 | mVppd | — |
| Absolute maximum RX input | V_{IN-Abs} | 0 | — | 1100 | mV | Measured RX_P to GND and RX_N to GND. |

5.10.4 Blackhawk SerDes Transmitter

The Blackhawk transmitter features are as follows:

- Transmitter with fully programmable six-tap FIR.
- IEEE 802.3 link-training.
- Programmable TX polarity inversion.
- TX disable.
- Controlled peak-to-peak amplitude.

The following table lists the electrical characteristics of the SerDes transmitter.

Table 43: SerDes Transmitter Characteristics

| Parameter | Symbol | Min. | Typ. | Max. | Unit | Notes |
|---------------------------|-----------------------|------|------|------|-------|--|
| Output common mode | $V_{\text{OUT-CM}}$ | — | 0.45 | — | V | When terminated with a 100Ω differential |
| Output differential swing | $V_{\text{OUT-DIFF}}$ | 0 | — | 1050 | mVppd | Programmable |

5.11 Falcon16 SerDes Operating Conditions

The following table lists the Falcon16 SerDes operating conditions.

Table 44: Falcon16 SerDes Operating Conditions

| Parameter | Symbol | Min. | Typ. | Max. | Unit |
|-------------------------|---------------------|---------|------|----------|------|
| Data rate, NIF Ethernet | F _{b_NIF} | 1.25 | — | 25.78125 | Gb/s |
| Data rate, ILKN | F _{b_ILKN} | 10.3125 | — | 25.78125 | Gb/s |

Each Falcon16 core includes four SerDes. Each one of the four SerDes is an independent lane, operating in NRZ line coding. Each core has one PLL.

Every group of Falcon16 cores receives its reference clock from the NIF_[x]_REFCLK_P/N as described in [Section 5.11.2, Falcon16 SerDes Supported Rates](#).

5.11.1 Falcon16 SerDes Features

The following sections describe the main features of the Falcon16 SerDes.

5.11.1.1 General Features

Falcon16 supports the following general features:

- A block of four SerDes supporting four serial links.
- Line rates (depending on the application) from 1.25 Gb/s to 25.78125 Gb/s per serial link.
- One PLL in each Falcon16.
- Integrated Arm micro subsystem: Monitors the signal and adaptively adjusts the gain, peaking filter, and DFE coefficients to optimally equalize and restore the signal. During Clause72 and Clause 93 TX/RX link-training, this microcontroller is also responsible for returning feedback to the far-link partner to optimally tune its transmitter.

5.11.1.2 Debug Features

Falcon16 supports the following debugging features:

- PRBS 7, 9, 11, 15, 23, 31, and 58 generator and checker.
- Digital loopback: Turns around the transmit data before it goes into the analog front-end (this is a digital data path loopback for the data coming from the PCS/MAC transmit side).
- AC-JTAG for both TX and RX.
- Full-range horizontal and vertical eye diagnostics.

5.11.2 Falcon16 SerDes Supported Rates

For the rates the NIF supports, see [Table 2, Supported Port Modes](#).

For the rates the ELK interface supports, see [Table 5, ILKN Supported Rates and Electrical Specifications Standards](#).

5.11.3 Falcon16 SerDes Receiver

The Falcon16 receiver features are as follows:

- Integrated AC coupling on RX inputs.
- Two-stage RX equalizer.
- Two-stage controlled variable gain amplifier (VGA).
- 14-tap DFE (six fixed and eight floating) with adaptive control slicer.
- Programmable RX polarity inversion.

The following table lists the electrical characteristics of the SerDes receiver.

Table 45: SerDes Receiver Characteristics

| Parameter | Symbol | Min. | Typ. | Max. | Unit | Notes |
|--------------------------|---------------|------|------|------|-------|--|
| Input common mode | V_{IN-CM} | — | 750 | 900 | mV | Higher than the maximum V_{IN-CM} requires external AC capacitor on board, 100 nF. |
| Input differential swing | $V_{IN-DIFF}$ | 85 | — | 1600 | mVppd | — |
| Absolute max RX input | V_{IN-Abs} | 0 | — | 1100 | mV | Measured RX_P to GND and RX_N to GND. |

5.11.4 Falcon SerDes Transmitter

The Falcon16 transmitter features are as follows:

- Transmitter with fully programmable five-tap FIR.
- IEEE 802.3 link-training.
- Programmable TX polarity inversion.
- TX disable.
- Controlled peak-to-peak amplitude.

The following table lists the electrical characteristics of the SerDes transmitter.

Table 46: SerDes Transmitter Characteristics

| Parameter | Symbol | Min. | Typ. | Max. | Unit | Notes |
|---------------------------|----------------|------|------|------|-------|--|
| Output common mode | V_{OUT-CM} | — | 0.45 | — | V | When terminated with a 100 Ω differential |
| Output differential swing | $V_{OUT-DIFF}$ | 0 | — | 1050 | mVppd | Programmable |

Chapter 6: Thermal Specifications

6.1 Absolute and Operational Thermal Specifications

The BCM88800 is a multi-die device. The package contains the main ASIC die from Broadcom and one HBM die from the memory vendor. Take this structure into consideration when planning the thermal definitions and thermal design.

The following table shows the specifications for the absolute thermal limits.

Table 47: Absolute Thermal Limit Specifications

| Parameter | Min. | Max. | Unit |
|---------------------------------------|------|-------------------|------|
| Storage temperature | -40 | +100 | °C |
| Main die maximum junction temperature | — | +110 ^a | °C |
| HBM die maximum junction temperature | — | +95 ^b | °C |
| HBM CATTRIP assertion temperature | — | +120 ^c | °C |

- Operating at a temperature above the maximum T_J may cause permanent damage to the device.
A maximum excursion temperature of $T_J = 125^\circ\text{C}$ for 15 days per year (fewer than 96 consecutive hours) is allowed.
Proper functionality and performance cannot be guaranteed when the device operates above the maximum junction temperature.
- Over the HBM lifetime, a maximum excursion temperature of 105°C for total of 360 hours per year is allowed. A single excursion period should not be longer than 96 hours, and the time between two consecutive excursions periods is a minimum of 24 hours. During these excursion periods, the HBM performance and reliability are not affected.
When violating the excursion conditions, functional operation failures and reliability degradation are expected.
At 105°C , the total power the HBM consumes is 15W (instead of 12W). To calculate the consumption per rail, use a factor of $\times 1.25$ for HBM_VDDO, HBM_VDDC, and HBM_VPP2P5.
- Operation above the CATTEMP (120°C) results in immediate and permanent damage to the device.

Table 48: SerDes Core Operating Temperatures (Junction)

| Parameter | Symbol | Min. (Industrial-Grade Devices) ^a | Min. (Commercial-Grade Devices) | Max | Unit |
|---------------------------------|--------------------------|--|---------------------------------|------|------|
| PCIe operating temperature | PCIE_T _N | -20 | 0 | +110 | °C |
| Blackhawk operating temperature | BLACKHAWK_T _N | -10 | 0 | +110 | °C |
| Falcon16 operating temperature | FALCON16_T _N | -10 | 0 | +110 | °C |

- Lowest temperature for operation. For powering up at temperatures between -40°C and 0°C , refer to the *Hardware Design Guidelines for StrataDNX 16-nm Devices (DNX16-AN1xx)*, Chapter 10, Thermal Aspects, Industrial-Grade Devices section.

Table 49: HBM Operating Temperature (Junction, HBM Is Powered On)

| Parameter | Symbol | Min. | Max. | Unit |
|---|--------------------|------|------|------|
| Standard operating temperature | HBM_T _N | 0 | 85 | °C |
| Extended operating temperature ^a | HBM_T _E | 85 | 95 | °C |

- The HBM extended operating temperature requires additional refresh cycles, which are inserted by the controller.

NOTE: When the HBM is powered-on, it must comply with the HBM_T_N or HBM_T_E values. The HBM must be powered-down if the die temperature is below the HBM_T_N minimum or above the HBM_T_E maximum.

Table 50: HBM Exceeding Temperature (Junction, HBM Is Powered Off)

| Parameter | Symbol | Min. | Max. | Unit |
|------------------------------------|-------------------------|------|------|------|
| Exceeding temperature ^a | HBM_T _{EXCEED} | 95 | 125 | °C |

- a. The *exceeding* temperature is the temperature that exceeds the extended temperature. For the HBM die, an exceeding junction temperature of 95°C to 125°C for 96 hours per year is allowed as long as the following conditions are met:
1. The HBM is powered-down by the board (follow the power-down sequence as described in [Section 5.5.5, HBM-Only Power Down and Power Up](#)).
 2. The total duration per year is less than 96 hours.
- During these exceeding periods, the HBM reliability is not affected. When violating the exceeding temperature conditions, functional operation failures and reliability degradation are expected.

Table 51: Ambient Temperature Conditions

| Parameter | Min. | Max. | Unit |
|------------------------------|------|------|------|
| Commercial temperature grade | 0 | +70 | °C |
| Industrial temperature grade | -40 | +85 | °C |

NOTE: For additional thermal information (including information about the temperature excursion), refer to the *Thermal Considerations for High-Power Switching Devices* application note (StrataDNX-StrataXGS-AN1xx).

6.2 Package Thermal Specifications

This section provides information about the block thermal model package thermal specifications.

6.2.1 Block Thermal Model

The block thermal model includes separate entities for the main die and the HBM die.

The HBM junction temperature can be read directly from the thermal simulation, and it is not necessary to use a formula for extracting the HBM temperature from the simulated main die temperature. In addition, it is possible to update power consumption values for the main die and HBM die without an updated thermal model. The model default settings are 215W for the main die and 12W for the HBM die.

The BCM88800 block thermal model is available on the Broadcom Customer Support Portal (docSAFE). For more information regarding block models, refer to the `Block_Thermal_Model_Brief.pdf` file, which is part of the block model package available on docSAFE.

6.3 Temperature Monitoring

6.3.1 VTMON

The BCM88800 has five internal on-die temperature sensors that can be accessed by the software API. These monitors (VTMONs) can be used only while the device is operational. They have an accuracy of $\pm 3^{\circ}\text{C}$.

6.3.2 Thermal Diodes

The BCM88800 has one on-die thermal diode that is not coupled to any of the BCM88800 power sources nor associated with any of its logic. This enables the temperature to be read while the device is powered off or in a nonoperational condition (such as during a reset).

The diode circuit has an independent supply rail, THERM_DIODE_VDD3P3.

To measure and read the on-die thermal diode indications, external circuitry is required.

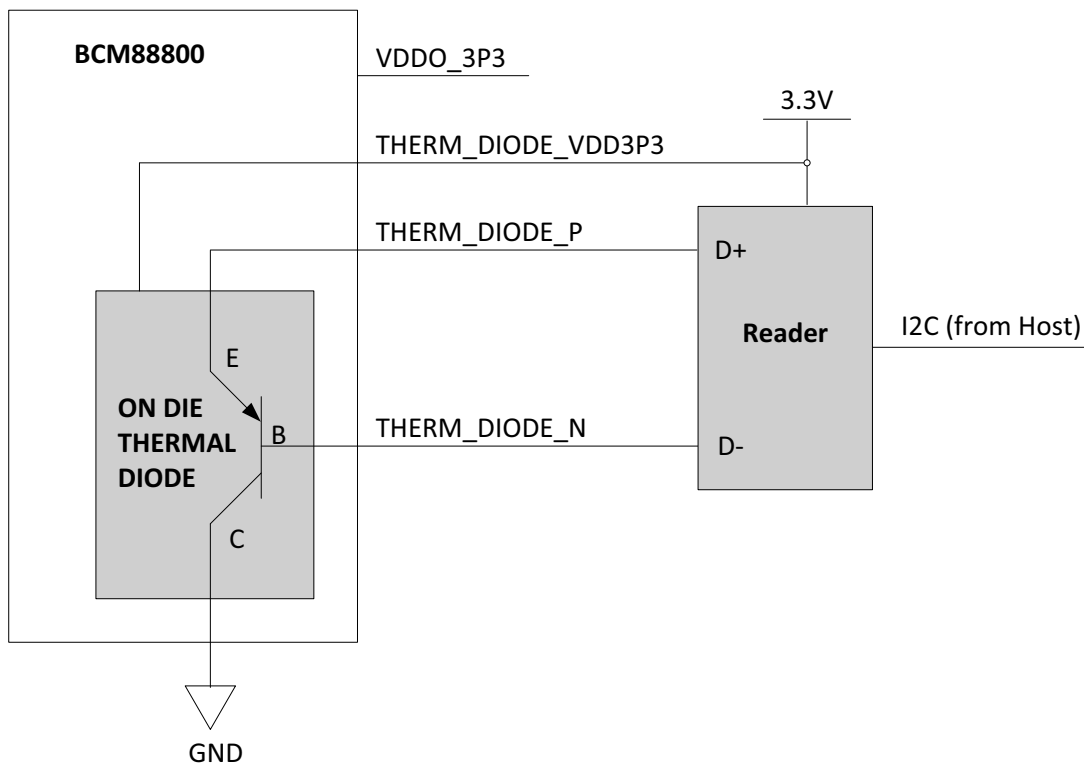
The following table shows the specifications of the on-die thermal diode.

Table 52: On-Die Thermal Diode Specifications

| Parameters | Symbol | Conditions | Min. | Typ. | Max. | Unit |
|-------------------------|-----------------|-------------------|------|-------|------|------|
| Operating temperature | T | — | -40 | — | +125 | °C |
| Force current | I _D | — | 5 | — | 120 | μA |
| Forward voltage | V _{BE} | Minimum at 5 μA | 0.25 | — | — | V |
| | | Maximum at 120 μA | — | — | 0.95 | |
| η – Ideality factor | — | — | — | 1.008 | — | — |
| Base bias (informative) | — | — | — | 0.6 | — | V |
| Accuracy ^a | — | -40°C to 110°C | -2.5 | — | +2.5 | °C |
| | | 110°C to 125°C | -1.5 | — | +1.5 | |

a. The accuracy figures are for diode currents of 120 μA and 6 μA.


Figure 23: Example Application Using the On-Die Thermal Diode



6.3.3 On-Die Thermal Monitor Locations

The locations of the thermal monitors relative to the device's package are shown in the following figure.

Figure 24: BCM88800 On-Die Thermal Monitors



6.4 Reflow Temperature

Broadcom offers a lead-free package.

NOTE: For additional reflow process recommendations, refer to the “PCB Land Pattern and Reflow Process Recommendations” section in the *Hardware Design Guidelines for StrataDNX 16-nm Devices* (DNX16-AN1xx).

The following table provides information about the solder ball composition and recommended and maximum reflow temperature.

Table 53: Solder Ball Composition and Recommended and Maximum Reflow Temperature


| Part Number | Solder Ball Composition | Recommended Reflow Peak Temperature | Maximum Allowed Reflow Peak Temperature |
|--------------------------------|--------------------------|-------------------------------------|---|
| Pb-free RoHS-compliant package | 96.5% Sn, 3% Ag, 0.5% Cu | 232°C–237°C | 245°C |

6.5 Heat Sink Considerations


For information about the heat sink, refer to the *Hardware Design Guidelines for StrataDNX 16-nm Devices* (DNX16-AN1xx).

Chapter 7: Packaging

Figure 25: Package Outline



| DIMENSIONAL REFERENCES (mm) | | | |
|-----------------------------|-----------|--------|-------|
| REF. | MIN | NOM | MAX |
| A | (5.14) | (5.39) | 5.64 |
| A1 | 0.4 | 0.5 | 0.6 |
| A2 | 2.81 | 2.86 | 2.91 |
| D | 62.40 | 62.50 | 62.60 |
| D1 | 60.00 BSC | | |
| E | 62.40 | 62.50 | 62.60 |
| E1 | 60.00 BSC | | |
| b | 0.535 | 0.635 | 0.735 |
| e | 1.00 BSC | | |
| f | 1.25 REF. | | |
| aaa | - | - | 0.20 |
| bbb | - | - | 0.25 |
| ccc | - | - | 0.20 |
| ddd | - | - | 0.25 |
| eee | - | - | 0.10 |



- PCB LAND PATTERN RECOMMENDATION: LAND PATTERN KEY E. REFER TO BROADCOM PACKAGING APPLICATION NOTE "PRINTED CIRCUIT BOARD LAND PATTERN RECOMMENDATIONS FOR BALL GRID ARRAY PACKAGES."
 - PRIMARY DATUM Z AND SEATING PLANE ARE DEFINED BY THE SPHERICAL CROWNS OF THE SOLDER BALLS.
 - DIMENSION IS MEASURED AT THE MAXIMUM SOLDER BALL DIAMETER, PARALLEL TO PRIMARY DATUM Z.
 - THIS PACKAGE CONFORMS TO THE JEDEC REGISTERED OUTLINE MS-034E.
 - ALL DIMENSIONS AND TOLERANCES CONFORM TO ASME Y14.5M-1994.
- NOTES: UNLESS OTHERWISE SPECIFIED

Related Documents

The references in this section may be used with this document. The documents are available on the Broadcom Customer Support Portal (docSAFE).

For Broadcom documents, replace the “xx” in the document number with the largest number available in the repository to ensure that you have the most current version of the document.

| Document (or Item) Name | Number | Source |
|--|--------------------------------------|---|
| Broadcom Items | | |
| <i>Hardware Design Guidelines for StrataDNX 16-nm Devices</i> | DNX16-AN1xx | Broadcom CSP |
| <i>BCM88800 Device Errata</i> | 88800-ER1xx | Broadcom CSP |
| <i>BCM88690 Traffic Manager Programming Guide</i> | 88690-PG2xx | Broadcom CSP |
| <i>BCM88800 PinList</i> | BCM88800_PinList_Revxx.zip | Broadcom CSP |
| <i>Thermal Considerations for High-Power Switching Devices</i> | StrataDNX-StrataXGS-AN1xx | Broadcom CSP |
| <i>Block Thermal Model</i> | BCM88800_Block_Thermal_Model_vxx.zip | Broadcom CSP |
| Other Items | | |
| <i>Two-Resistor Compact Thermal Model Guideline</i> | JESD15-3 | https://www.jedec.org/ |

Broadcom, the pulse logo, Connecting everything, Avago Technologies, Avago, the A logo, BroadSync, StrataXGS, and StrataDNX are among the trademarks of Broadcom and/or its affiliates in the United States, certain other countries, and/or the EU.

Copyright © 2018–2021 Broadcom. All Rights Reserved.

The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries. For more information, please visit www.broadcom.com.

Broadcom reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design. Information furnished by Broadcom is believed to be accurate and reliable. However, Broadcom does not assume any liability arising out of the application or use of this information, nor the application or use of any product or circuit described herein, neither does it convey any license under its patent rights nor the rights of others.

