
7-Port Layer-2 Carrier Ethernet Switch with MPLS/ MPLS-TP Support and VeriTime

INTRODUCTION

This document consists of descriptions and specifications for both functional and physical aspects of the VSC7416-01 device. It is intended for system designers and software developers.

In addition to the datasheet, Microchip maintains an extensive part-specific library of support and collateral materials that you may find useful in developing your own product. Depending upon the Microchip device, this library may include:

- Application notes that provide detailed descriptions of the use of the particular Microchip product to solve real-world problems
- White papers published by industry experts that provide ancillary and background information useful in developing products that take full advantage of Microchip product designs and capabilities
- User guides that describe specific techniques for interfacing to the particular Microchip products
- Reference designs showing the Microchip device built in to applications in ways intended to exploit its relative strengths
- Software Development Kits with sample commands and scripts
- Presentations highlighting the operational features and specifications of the device to assist in developing your own product road map
- Input/Output Buffer Information specification (IBIS) models to help you create and support the interfaces available on the particular Microchip product

Visit and register as a user on the Microchip Web site to keep abreast of the latest innovations from research and development teams and the most current product and application documentation. The address of the Microchip Web site is <https://www.microchip.com/>.

Register Notation

This datasheet uses the following general register notation:

`<TARGET>:<REGISTER_GROUP>:<REGISTER>.<FIELD>`

`<REGISTER_GROUP>` is not always present. In that case, the following notation is used:

`<TARGET>::<REGISTER>.<FIELD>`

When a register group does exist, it is always prepended with a target in the notation.

In sections where only one register is discussed, or the target (and register group) is known from the context, the `<TARGET>:<REGISTER_GROUP>:` may be omitted for brevity, and uses the following notation:

`<REGISTER>.<FIELD>`

Also, when a register contains only one field, the `.<FIELD>` is not included in the notation.

TO OUR VALUED CUSTOMERS

It is our intention to provide our valued customers with the best documentation possible to ensure successful use of your Microchip products. To this end, we will continue to improve our publications to better suit your needs. Our publications will be refined and enhanced as new volumes and updates are introduced.

If you have any questions or comments regarding this publication, please contact the Marketing Communications Department via E-mail at docerrors@microchip.com. We welcome your feedback.

Most Current Data Sheet

To obtain the most up-to-date version of this data sheet, please register at our Worldwide Web site at:

<http://www.microchip.com>

You can determine the version of a data sheet by examining its literature number found on the bottom outside corner of any page. The last character of the literature number is the version number, (e.g., DS30000000A is version A of document DS30000000).

Errata

An errata sheet, describing minor operational differences from the data sheet and recommended workarounds, may exist for current devices. As device/documentation issues become known to us, we will publish an errata sheet. The errata will specify the revision of silicon and revision of document to which it applies.

To determine if an errata sheet exists for a particular device, please check with one of the following:

- Microchip's Worldwide Web site; <http://www.microchip.com>
- Your local Microchip sales office (see last page)

When contacting a sales office, please specify which device, revision of silicon and data sheet (include -literature number) you are using.

Customer Notification System

Register on our web site at www.microchip.com to receive the most current information on all of our products.

Contents

Introduction	1
1.0 Product Overview	5
1.1 General Features	5
1.2 Applications	8
1.3 Functional Overview	14
2.0 Functional Descriptions	26
2.1 Port Modules	26
2.2 SERDES1G	33
2.3 SERDES6G	37
2.4 Statistics	42
2.5 MPLS-TP Classifier	49
2.6 Basic Classifier	59
2.7 VCAP	69
2.8 Analyzer	113
2.9 Policers	132
2.10 Shared Queue System	136
2.11 Scheduler and Shapers	143
2.12 Rewriter	150
2.13 Versatile OAM Processor	163
2.14 CPU Port Module	190
2.15 VRAP Engine	198
2.16 Layer-1 Timing	202
2.17 Hardware Timestamping	204
2.18 Clocking and Reset	208
3.0 VCore System and CPU Interfaces	209
3.1 VCore Configurations	210
3.2 Clocking and Reset	211
3.3 Shared Bus	212
3.4 VCore CPU	216
3.5 External CPU Support	218
3.6 PCIe Endpoint Controller	225
3.7 Frame DMA	230
3.8 VCore System Peripherals	239
4.0 Features	267
4.1 Switch Control	267
4.2 Port Module Control	267
4.3 Layer-2 Switch	271
4.4 IGMP and MLD Snooping	298
4.5 Quality of Service (QoS)	299
4.6 VCAP Applications	302
4.7 CPU Extraction and Injection	309
4.8 MPLS Operation	312
4.9 Ethernet over MPLS Service	312
5.0 Registers	320
5.1 ANA	321
5.2 DEV	367
5.3 DEVCPU_GCB	385
5.4 DEVCPU_ORG	414
5.5 DEVCPU_QS	420
5.6 HSIO	426
5.7 ICPU_CFG	445
5.8 VCAP_CORE	490
5.9 OAM_MEP	499
5.10 PCIE	545
5.11 QSYS	551

5.12	REW	569
5.13	SBA	574
5.14	SYS	576
5.15	TWI	590
5.16	UART	608
6.0	Electrical Specifications	617
6.1	DC Characteristics	617
6.2	AC Characteristics	624
6.3	Current and Power Consumption	642
6.4	Operating Conditions	643
6.5	Stress Ratings	643
7.0	Pin Descriptions	645
7.1	Pin Diagram	645
7.2	Pins by Function	645
7.3	Pins by Number	652
7.4	Pins by Name	655
8.0	Package Information	658
8.1	Package Drawing	658
8.2	Thermal Specifications	660
8.3	Moisture Sensitivity	660
9.0	Design Guidelines	661
9.1	Power Supplies	661
9.2	Power Supply Decoupling	661
9.3	Reference Clock	661
9.4	Interfaces	662
10.0	Ordering Information	667
11.0	Design Considerations	668
Appendix A: Revision History		670
The Microchip WebSite		671
Customer Change Notification Service		671
Customer Support		671

1.0 PRODUCT OVERVIEW

The VSC7416-01 Carrier Ethernet switch contains four 10/100/1000 Mbps SGMII/SerDes Ethernet ports, two 10/100/1000/2500 Mbps SGMII/SerDes ports, and one 10/100/1000 Mbps SGMII/SerDes Node Processor Interface (NPI) Ethernet port. The NPI port can operate as a 5th triple-speed Ethernet port. The device provides a rich set of Carrier Ethernet switching features such as MPLS-TP, MEF Ethernet services, hierarchical QoS scheduling, provider bridging, protection switching, OAM, 1588 precision time protocol, and synchronous Ethernet. Advanced TCAM-based service, VLAN, and QoS processing enable delivery of differentiated services with per-service SLA guarantees. Security is assured through frame processing using a TCAM-based Microchip Content Aware Processor (VCAP). In addition, the device contains a powerful 416 MHz CPU enabling full management of the switch.

1.1 General Features

- 4× 1G Ethernet ports, which are tri-speed 10/100/1000 Mbps ports
- 2× 2.5G Ethernet ports, which are quad-speed 10/100/1000/2500 Mbps ports
- 1× 1G Node Processor Interface Ethernet port, which is a tri-speed 10/100/1000 Mbps port
- All ports support both 100-BASE-FX and 1000-BASE-X-SERDES
- Eight megabits of integrated shared packet memory
- Fully nonblocking wire-speed switching performance with weighted random early detection (WRED) for all frame sizes
- Eight QoS classes with 2616 queues that can be assigned to ports, QoS classes, and services
- 330 dual leaky bucket policers, per service, per QoS class, and per port
- 64 dual leaky bucket policers, flow-based through TCAM matching
- Hierarchical QoS scheduling with up to three levels of scheduling
- 256 scheduling elements, each with a DWRR/strict priority scheduler and a dual leaky bucket shaper. All combinations of DWRR vs. strict scheduling between QoS classes, services or ingress ports are possible.
- MPLS-TP LER (Ethernet pseudo-wires) and LSR with up to 128 unidirectional label switched paths and up to 128 bidirectional Ethernet pseudowires
- Up to 1024 TCAM-based classification entries for service, Quality of Service (QoS) and VLAN membership
- Up to 1024 host identity entries for source IP guarding
- Up to 512 TCAM-based security enforcement entries
- Up to 1024 TCAM-based egress tagging and MPLS label stack entries
- OAM Up and Down-MEP support for CCM, loss measurements, delay measurements, and loopback
- Service Activation Test frames generation and checking
- L1 Synchronous Ethernet
- VeriTime—Microchip's patent-pending distributed timing technology that delivers the industry's most accurate IEEE 1588v2 timing implementation for both one-step and two-step clocks
- Energy Efficient Ethernet (IEEE 802.3az) supported by the switch core
- VCore-III CPU system with integrated 416 MHz MIPS 24KEc™ CPU with MMU and DDR3 SDRAM controller
- PCIe 1.x CPU interface

1.1.1 LAYER 2 SWITCHING

- 8,192 MAC addresses
- 4,096 VLANs (IEEE 802.1Q)
- Push/pop/translate up to two VLAN tags; translation on ingress and/or on egress
- TCAM-based VLAN classification and translation with pattern matching against Layer 2 through Layer 4 information such as MAC addresses, VLAN tag headers, EtherType, DSCP, IP addresses, and TCP/UDP ports and ranges
- Up to 1024 QoS and VLAN TCAM entries
- Up to 1024 VLAN egress tagging TCAM entries
- Link aggregation (IEEE 802.3ad)
- Link aggregation traffic distribution is programmable and based on Layer 2 through Layer 4 information
- Wire-speed hardware-based learning and CPU-based learning configurable per port
- Independent and shared VLAN learning

- Provider Bridging (VLAN Q-in-Q) support (IEEE 802.1ad)
- Rapid Spanning Tree Protocol support (IEEE 802.1w)
- Multiple Spanning Tree Protocol support (IEEE 802.1s)
- Jumbo frame support up to 10 kilobytes (kB) with programmable MTU per port

1.1.2 LAYER 2 MULTICAST

- Up to 8,192 Layer 2 multicast groups with up to 64 port masks
- Up to 8,192 IPv4/IPv6 (G) multicast groups and port masks
- Up to 8,192 IPv4/IPv6 (S, G) multicast groups and port masks with up to 512 unique sources
- Internet Group Management Protocol version 2 (IGMPv2) support
- Internet Group Management Protocol version 3 (IGMPv3) support with source specific multicast forwarding
- Multicast Listener Discovery (MLDv1) support
- Multicast Listener Discovery (MLDv2) support with source specific forwarding

1.1.3 CARRIER ETHERNET

- Provider Bridge (Q-in-Q) switch

8K MACs, 4K VLANs

- MPLS-TP switch

Label Edge Router (LER) with Ethernet pseudo-wires

Label Switch Router (LSR)

H-VPLS MTU-s (H-VPLS edge) functions

MPLS-TP forwarding in a Ethernet-based ring, where the Ethernet layer supports Drop and Continue

MPLS-TP Point-to-Multipoint (P2MP)

32 MPLS Link Layers, each defined by destination and source MAC address, and up to one VLAN tag

128 uni-directional Label Switched Paths (LSPs), each path defined by up to three labels

Extract LSP OAM frames to CPU (GAL)

128 Ethernet pseudo-wires (PWE)

Switched Provider Edge (S-PE; multisegment pseudo-wire)

Terminating Provider Edge (T-PE)

Extract PW OAM frames to CPU (GAL, VCCV1, VCCV2, VCCV3)

- 256 MEF E-LAN, E-LINE, E-TREE service points

128 bidirectional E-LINEs or 64 E-LANs/E-TREEs each with four ports

256 dual leaky bucket service policers

PCP and DSCP remarking per service point

Statistics and tagging options per service point

- Versatile OAM processor (VOP) with hardware for generating and checking CCM messages, frame delay measurements, loss measurements, and loopback

Down MEP support for 7 ports

Up-MEP and Down-MEP support for 16 paths and services

- Hardware-assisted linear protection (ITU-T G.8031) and ring protection (ITU-T G.8032)

Service, path and port protection with few milliseconds restoration time

Fully hardware-supported CCM frame generation and checking with sub-millisecond CCM intervals.

- L1 Synchronous Ethernet with up to 10 recovered clock outputs
- L2 IEEE 1588 timestamping hardware, with one-step and two-step clock support

Independent clock domains for data path forwarding (L1 Sync Ethernet) and timing (IEEE 1588)

Nanosecond precision timestamping

Asymmetry correction support in hardware

IEEE 1588 over MPLS-TP and IP/MPLS support

- Enhanced Carrier Ethernet software API

1.1.4 QUALITY OF SERVICE

- Eight QoS classes and two drop precedence levels
- Strict, deficit weighted round-robin (DWRR), or frame based round-robin (FBRR) scheduling
- TCAM-based QoS classification with pattern matching against Layer 2 through Layer 4 information
- Up to 1024 QoS and VLAN TCAM entries
- DSCP translation, both ingress and/or egress
- DSCP remarking based on QoS class and drop precedence level
- VLAN (PCP, DEI, and VID) translation, both ingress and egress
- PCP and DEI remarking based on QoS class and drop precedence level
- 392 policers, selectable per QoS class per port, per port, per service point, and per security entry through TCAM-based pattern matching, programmable in steps of 33.3 kbps
- Three-level hierarchical QoS scheduling (QoS classes, services, ingress ports, or service groups) with 256 dual leaky bucket shapers and DWRR/FBRR/strict scheduling
- Full-duplex flow control (IEEE 802.3X) and half-duplex backpressure, symmetric and asymmetric
- Priority-based full-duplex flow control (IEEE 802.1Qbb)

1.1.5 SECURITY

- Versatile Content Aware Processor (VCAP) packet filtering engine using ACLs for ingress and egress packet inspection:

Up to 512 security VCAP entries

Up to 1024 source IP guarding entries

64 shared VCAP rate policers with rate measurements in frames per second or bits per second

Eight shared range checkers supporting ranges based on TCP/UDP port numbers, DSCP values, and VLAN identifiers

VCAP match patterns supporting generic MAC, ARP, IPv4, and IPv6 protocols

VCAP actions including permit/deny, police, count, CPU-copy, and mirror

Special support for IP fragments, UDP/TCP port ranges, and ARP sanity check

Extensive CPU DoS prevention by VCAP rate policers and hit-me-once functions

Surveillance functions supported by 32-bit VCAP counters

- Generic storm controllers for flooded broadcast, flooded multicast, and flooded unicast traffic

- Selectable CPU extraction queues for segregation of CPU redirected traffic, with 8 extraction queues supported
- Per-port, per-address registration for copying/redirecting/discarding of reserved IEEE MAC addresses (BPDU, GARP, CCM/Link trace)
- Port-based and MAC-based access control (IEEE 802.1X)
- Per-port CPU-based learning with option for secure CPU-based learning
- Per-port ingress and egress mirroring
- Mirroring per VLAN and per VCAP match

1.1.6 MANAGEMENT

- MIPS 24KEc™ CPU system with memory management unit (MMU), and 32 kilobytes of instruction cache (I-cache) and 32 kilobytes of data cache (D-cache)
- PCIe 1.x enabling frame injection/extraction and register access from external CPU
- CPU frame extraction (eight queues) and injection (two queues) through DMA, which enables efficient data transfer between Ethernet ports and CPU/PCIe
- EJTAG debug interface
- Configurable 16-bit or 8-bit DDR2/3 SDRAM interface supporting up to one gigabyte (GB) of memory
- Thirty-two pin-shared general-purpose I/Os

Serial GPIO and LED controller controlling up to 32 ports with four LEDs each

Dual PHY management controller (MIIM)

Dual UART

Built-in two wire serial interface multiplexer

External interrupts

1588 synchronization I/Os

SFP loss of signal inputs


- External access to registers through PCIe, SPI, MIIM, or through an Ethernet port with inline Versatile Register Access Protocol
- Per-port 32-bit counter set with support for the RMON statistics group (RFC 2819) and SNMP interfaces group (RFC 2863)
- Support CPU models with internal CPU only, external CPU only, or dual CPU

1.2 Applications

VSC7416-01's target applications are:

- Microwave Backhaul (ODU/IDU) and all-integrated ODU
- Network Interface Devices (NID) and Ethernet Access Devices (EAD)
- Ethernet/MPLS demarcation
- Customer Premise Equipment (CPE) and Business VPN service Equipment
- Fiber Backhaul Equipment

The following illustration shows the target applications in an access and aggregation network.

FIGURE 1-1: ACCESS AND AGGREGATION NETWORK

VSC7416-01 provides the required set of User to Network Interface (UNI) and Network to Network Interface (NNI) features in a cost-effective manner:

- Maps Customer frame formats into Provider frame formats
 - Classifies frames and map to appropriate QoS profiles
 - Applies Provider Bridge (Q-in-Q) encapsulations
 - Applies MPLS pseudowire, MPLS Label stack, and MPLS Link-Layer Ethernet encapsulations
- Meters the customer traffic and ensures that the customer Service Level Agreement (SLA) is met
 - Polices using MEF-defined Dual Leaky Bucket algorithm
 - Marks frames as Committed (Green) or Discard Eligible (Yellow)
 - Provides correct QoS treatment (traffic management)
 - Provides traffic statistics per customer in a manner consistent with the SLA
 - Shapes using MEF-defined Dual Leaky Bucket shaping algorithm
- Enables end-to-end Service OAM by the customer
- Implements the service as defined by the SLA
 - E-LINE for point-point or backhaul services
 - E-LAN for multipoint/bridged services
 - E-TREE for video distribution or backhaul services
- Enables management and protection schemes as required by the Provider
 - Link aggregation or other port protection schemes if used for access

OAM at the Operator and Service Provider levels for remote management, fault diagnosis, and protection switching

- Supports network timing and synchronization requirement

Synchronous Ethernet and IEEE 1588 functionality

1.2.1 NETWORK INTERFACE DEVICE, ETHERNET ACCESS DEVICE


The following illustration is an example of how an Enterprise VPN service can be provided through an access and aggregation network. Enterprise VPN services provide connectivity between corporate headquarters and branch offices, as well as providing voice, video, internet, storage, and other network-hosted applications.

The example in the illustration shows one operator (the “Service Provider”) who sells services directly to the Enterprise, and a second operator (the “transport operator”) who sells local access to the Service Provider for a location where the Service Provider does not have direct access to the customer’s corporate office.

In general, NIDs provide demarcation between a Customer and a Network. In this example, one NID interfaces between the Service Provider (the Network) and the Enterprise customer (the Customer). A second NID interfaces between the Transport Operator (the Network for NID) and the Service Provider (the Customer for this NID).

The NIDs are located at the customer site. One NID is owned/managed by the Service Provider, the second by the Transport Operator. This two-box model is prevalent, but a one-box Virtual NID concept is emerging, as shown in red. In the Virtual NID concept, the Transport Operator owns the NID but there are separate management domains for each operator.

FIGURE 1-2: NID APPLICATION



Transport Operator Networks predominantly uses Provider Bridging (Q-in-Q) and/or MPLS-TP technology for backhauling customer traffic further into the network.

The CPE and NIDs shown in this example would likely have a relatively low port count and bandwidth. It is also possible to have CPE and NIDs with many customer-facing ports and perhaps two network-facing ports.

Very limited switching and networking functionality is present in these devices, because their purpose is simply to isolate customers from one another and from the network and allow the network operator to monitor and manage the customers’ traffic.

For NID and Virtual NID applications, VSC7416-01 provides the following features.


- Highly flexible classification, tagging and encapsulation for Provider Bridges (Q-Q), MPLS pseudowires and MPLS-TP transport.
- Fully-featured Ethernet-based OAM. Hardware OAM implementation is required for precision and scale.
- At the UNI-C and ENNI-C, MEF-compliant DLB shaping and Down-MEP functions. Down-MEP allows the Customer to monitor the service purchased from the Network.
- At the UNI-N and ENNI-N, MEF-compliant DLB policing and Up-MEP functions. Up-MEP functions allow the network to monitor the service sold to the customer. In the virtual NID model, the ENNI is an internal interface to the device.
- At the UNI-C and ENNI-C some applications require per-service queuing for separation.
- Per-service and color statistics both ingress, discarded and egress.

- Service Activation Measurements.
- IEEE-1588 network timing functions.
- Linear and Ring protection mechanisms as well as path and port protection for dual-homed NID and EAD installations.
- Integrated CPU sub-system for management and control.

1.2.2 OUTDOOR UNIT FOR MICROWAVE BACKHAUL

Both traditional ODUs for split-mount and all integrated ODUs for 4G/LTE equipment are target applications for the VSC7416-01 Carrier Ethernet switch. The following illustration shows a split-mount IDU and ODU installation to the right, and an all-integrated Outdoor Unit is shown to the left. In some larger installations, the Radio Base Station and the all-integrated ODU functionality (shown to the left) is built into a single piece of equipment. In the illustration, the three networks (fronthaul, backhaul, and pre-aggregation) are all operated by separate network providers. The backhaul network provides a transport service between the radio base station and the pre-aggregation network. Other models of network ownership are possible. VSC7416-01 supports UNI, I-NNI, and E-NNI functions as defined by Metro Ethernet Forum to support all of these models. Only one microwave link is shown in the following illustration, however, applications with multiple microwave interfaces exist supporting link aggregation and protection, as well as ring and mesh network topologies.

FIGURE 1-3: ODU AND IDU APPLICATIONS



Microwave radio links are highly susceptible to weather and use adaptive modulation techniques to deliver data reliably during bad weather. As a result, the microwave data transfer rate (link speed) can change rapidly creating jitter problems for IEEE 1588 and QoS problems for the networking gear. The link speed is also limited to a few hundred Mbps per microwave link under good weather conditions, so link bundling and load balancing are common. Next generation ODUs must support timing and synchronization such as synchronous Ethernet and IEEE 1588 packet-based timing.

Small IDUs (pizzaboxes) generally supports a carrier Ethernet feature set similar to the NID devices previously described. In addition, IDU functionality can include MPLS-TP functionality.

Integrated ODU/IDU equipment also exists, especially for 4G/LTE femtocell deployment. These units will likely have a low port count but advanced features such as MPLS-TP, protection switching mechanisms, and strong QoS.

VSC7416-01 supports the following key features for the Microwave-based Mobile Backhaul application.

- Low size, low cost, and low power, especially for the ODUs.
- Ethernet and MPLS-TP carrier features at the IDUs and all-integrated ODUs.
- UNI/NNI and Ethernet OAM as previously described for NIDs.
- IEEE-1588 functionality is critical to deliver timing to Base Stations as required for 4G/LTE networks. Various proprietary techniques are supported by VSC7416-01 for combating the jitter created by adaptive modulation on the microwave link. VSC7416-01 supports different methods to provide IEEE 1588 transparent clock over a microwave link.
- Strong QoS is required. QoS may be statically configured to handle the varying link speed. 802.1Qbb priority-based flow control may be used as in-band dynamic flow control. Versatile Register Access Protocol (VRAP) can also be used as in-band dynamic scheduler/shaper rate adjustment. Hierarchical queuing can also be helpful for customer separation in IDUs and all-integrated ODUs.
- Linear and Ring protection mechanisms


- Due to the microwave link being throttled frequently, a larger buffer memory and advanced congestion avoidance mechanisms are required to efficiently utilize bandwidth when available, while not jeopardizing low latency traffic such as voice. VSC7416-01 has a large amount of integrated memory, supports WRED and advanced scheduling and shaping. External packet memory may be required in some applications.
- Industrial temperature.

1.2.3 CARRIER ETHERNET SWITCH WITH MPLS-TP

The following illustration is an example where VSC7416-01 adds and drops Ethernet-over-MPLS traffic onto an Ethernet Ring using MPLS pseudowires (EoMPLS PW LER operation). When using VSC7416-01 the Ethernet ring can be operated at 1 Gbps or 2.5 Gbps. Two independent Ethernet layers are shown. The outermost Ethernet layer is the (MPLS) Link Layer Ethernet (LL-Ethernet). The inner layer Ethernet is the (MPLS) Service Layer Ethernet (SL-Ethernet). VSC7416-01's service concept is required to establish MPLS pseudowires for carrying SL-Ethernet traffic over LL-Ethernet.

MPLS Label Switch Router (LSR) operation is also supported by VSC7416-01's service concept. VSC7416-01 can operate as a LSR front-end to an attached routing function within a Carrier Ethernet Switch Router (CESR) application.

FIGURE 1-4: MPLS OVER ETHERNET RING



Any traffic carried over Ethernet can be forwarded within the bridge attached to the frame's arrival port without the use of VSC7416-01's service concept. That is, both LL- and SL-Ethernet bridge are fully featured provider bridges (Q-in-Q). The illustration shows examples of provider bridging within both the LL-Ethernet bridge and SL-Ethernet bridge.

The orange arrows show how an EoMPLS PW E-TREE service can be carried over an Ethernet ring using drop-and-continue forwarding at the Link-Layer Ethernet. At the ingress ring node the E-TREE service traffic arrives at the SL-Ethernet bridge and is classified to an Ethernet pseudowire. A PW label, an MPLS label stack, and a Link-layer Ethernet multicast header are pushed before the service traffic is forwarded to the Link-Layer bridge.

Whether the E-TREE service is forwarded clock-wise or counter clock-wise around the Link-Layer ring can be determined entirely by Ethernet ring protection mechanisms (G.8032) implemented at the Link-Layer of the ingress node. In the illustration, counter clock-wise direction is shown.

In the next ring node, there are no attached subscribers for this E-TREE service and the traffic is "continued" on the ring using bridging (L2 multicasting) in the LL-Ethernet Bridge.

At the next ring node (shown to the right in the illustration), a "continue" copy of the E-TREE service is again provided by L2 multicasting in the LL-Ethernet Bridge. Also, LER operation "drop" is performed as subscribers to the E-TREE service exists. The LL-Ethernet header, the MPLS label stack, and the PW label are popped in the LER block and only the original customer frame is forwarded to the Service Layer bridge. In this example, multiple copies are generated to subscribers of the L2 multicast flow within the E-TREE service.

VSC7416-01 provides the following capabilities for use in MPLS and/or Carrier Ethernet edge equipment:

- EoMPLS PW LER (Single-Segment and Multi-Segment Pseudowire)
- MPLS LSR
- Hierarchical VPLS (H-VPLS)

- Separate MAC-tables and VLAN-tables for MPLS Link-Layer Ethernet and Service-Layer Ethernet
- Service concept
 - Per Service ingress dual leaky bucket policing
 - Per-service and color statistics both ingress, discarded and egress
- Protection switching
 - 1:1 PW protection (Service Layer protection)
 - 1:1 LSP protection (Path Layer protection)
 - Facility Bypass protection per RFC-4090 (a group of services at a time)
 - Ethernet Linear protection (ITU-T G.8031)
 - Ethernet Ring protection (ITU-T G.8032)
- OAM
 - Fully-featured Ethernet-based OAM. Hardware OAM implementation is required for precision and scale.
 - MPLS-TP OAM assist functions for internal/external CPU software or external MPLS-TP OAM processor
- Service Activation Measurements
- IEEE-1588 network timing functions
- Integrated CPU sub-system for management and control

1.2.4 OAM ENGINE

VSC7416-01's OAM functionality described as part of this application is also supported in all other applications of VSC7416-01 previously described even though not detailed as part of those descriptions.

VSC7416-01 can be used as a powerful OAM off-load engine in new designs or to upgrade existing systems to become OAM aware. The following illustration shows two different usages of VSC7416-01 as an Ethernet OAM engine on a line card. Similar pizzabox models exist.


For the linecard at the top of the illustration, VSC7416-01 is side-band to the data flow. In this model, VSC7416-01 processes OAMs redirected by the NPU/ASIC and pseudo Up-MEPs and Down-MEPs can be supported fully in hardware. If used in combination with a Microchip Carrier Ethernet PHY, hardware accurate Down-MEP delay measurements and delay variation measurements (DMM/DMR) can be supported using this model. Continuity check (CCMs), synthetic loss measurements using LMM/LMR, and loopback (LBM/LBR) is also supported fully in hardware using VSC7416-01's Ethernet OAM capabilities.

The line card shown at the bottom in the illustration depicts VSC7416-01 used as an in-band OAM engine. In this model, true Up-MEPs and Down-MEPs can be supported all in hardware and with hardware accuracy and scale for continuity check, delay measurements, loss measurements, and loopback.

For both OAM models shown, OAM software can reside in VSC7416-01's integrated CPU system. This enables a distributed OAM model for chassis applications considerably reducing the OAM processing requirements to a central OAM management entity.

For the in-band OAM engine model, VSC7416-01 can also assist MPLS-TP OAM implementations. VSC7416-01 detects and redirects MPLS-TP OAM and Pseudowire OAM to internal CPU or external OAM processor, which implements actual OAM protocol handling. When extracting MPLS-TP OAM, VSC7416-01 provides arrival port information and service classification results to the internal CPU or external device. In addition, when MPLS-TP OAM is injected to VSC7416-01, the internal CPU or external device can include proprietary service mapping information. As a result, the internal CPU or MPLS-TP OAM processing device can dictate parameters such as destination port or ports, QOS and queuing, per egress port encapsulation, and statistics for each frame.

FIGURE 1-5: OAM ENGINE MODELS



The following are some of the advanced features of VSC7416-01 used in OAM engine applications.


- Fully-featured Ethernet-based OAM. All hardware OAM generation and checking functions. Hardware OAM implementation is required for precision.
- MPLS-TP OAM assist functions for internal or external CPU software or external MPLS-TP OAM processor.
- Multiple concurrent Maintenance Entity Group (MEG) Levels for port, path, and service OAM.
- OAM loopback.
- Generalized loopback functions not only applicable to LBM/LBR handling.
- Automatic timer-based hardware OAM frame generation targeted at CCMs, LMM, DMM, and LBM/TST generation.
- FDMA assisted OAM frame generation for Service Activation Measurements.
- Integrated CPU system running OAM software.

1.3 Functional Overview

This section provides an overview of all major blocks and functions involved in the forwarding operation in the same order as a frame traverses through the VSC7416-01 device. It also outlines other major functionality of the device such as the CPU port module, the CPU system, and CPU interfaces.

The following illustration shows the block diagram for the VSC7416-01 device.

FIGURE 1-6: BLOCK DIAGRAM



1.3.1 FRAME ARRIVAL

The Ethernet interfaces receive incoming frames and forward these to the port modules. The 1G SGMII and 2.5G SGMII ports support both 100BASE-X and 1000BASE-X-SERDES.

Each port module contains a Media Access Controller (MAC) that performs a full suite of checks, such as VLAN Tag aware frame size checking, Frame Check Sequence (FCS) checking, and Pause frame identification.

Each port module connects to a SerDes block and contains a Physical Coding Sublayer (PCS), which perform 8 bits/10 bits encoding, auto-negotiation of link speed and duplex mode, and monitoring of the link status.

Full-duplex is supported for all speeds, and half-duplex is supported for 10 Mbps and 100 Mbps. Symmetric and asymmetric pause flow control are both supported as well as priority-based flow control (IEEE 802.1Qbb).

All Ethernet ports support Energy Efficient Ethernet (EEE) according to IEEE 802.3az. The shared queue system is capable of controlling the operating states, active or low-power, of the PCS. The PCS understands the line signaling as required for EEE. This includes signaling of active, sleep, quiet, refresh, and wake.

1.3.2 MPLS-TP CLASSIFICATION

On ports where MPLS is enabled, incoming MPLS-encapsulated frames are sent to the MPLS-TP classifier in the ingress processing module. Non-MPLS-encapsulated frames bypass the MPLS-TP classifier and are sent directly to the basic classifier.

Based on the MPLS Link Layer and MPLS label stack, the MPLS-TP classifier determines if the frame is to be forwarded using the MPLS Link Layer Ethernet (outer Ethernet), MPLS layer, or Service Layer Ethernet (inner/customer Ethernet for Ethernet pseudowire). In some cases forwarding at two layers simultaneously is possible.

The MPLS-TP classification is based on up to three MPLS labels and 1 reserved label.

Based on the MPLS-TP classification, any Ethernet pseudo-wire (PW) can be terminated and the MPLS-TP label stack popped (LER). The frame is then forwarded based on the customer ethernet encapsulation. The forwarding function can forward the frame into a new pseudo-wire (S-PE function) or to a customer port (T-PE). When terminating a PW, VSC7416-01 can pop the MPLS link layer, 3 MPLS labels and the control word.

This implementation of customer Ethernet forwarding function supports E-LAN over MPLS (VPLS). Forwarding is limited to one pseudo-wire per physical port per E-LAN enabling a low-cost H-VPLS MTU-s edge.

Based on MPLS-TP classification, any LSP can be forwarded to an MPLS port and the MPLS label stack is modified. The MPLS label stack operation can include terminating up to 2 tunnels and initiating up to 2 new tunnels in addition to label swapping and TTL processing (LSR).

The MPLS-TP classifier supports Ethernet Link Layer drop and continue for any LSP, allowing a copy of the frame to be forwarded to a network port based on the Ethernet layer.

The MPLS-TP classifier can extract PW OAM (GAL, VCCV1, VCCV2, VCCV3), LSP OAM (GAL) and Segment OAM (GAL) to the CPU. Further, the MPLS-TP classifier can be configured to extract MPLS-TP frames with any reserved label in the MPLS label stack.

1.3.3 BASIC AND ADVANCED FRAME CLASSIFICATION


The basic and advanced frame classification in the ingress processing module receive all frames except frames for which the MPLS-TP classifier has indicated that further processing is not required. The basic and advanced classifiers determine a range of frame properties such as arrival service point, VLAN, and Quality of Service (QoS) class, and drop precedence level. This information is carried through the switch together with the frame and affects policing, drop precedence marking, statistics collecting, security enforcement, Layer-2 forwarding, and rewriting.

The classification is a combination of a basic classification using configurable logic and more advanced classification using a TCAM.

The classification engine understands up to two VLAN tags and can look for Layer-3 and Layer-4 information behind two VLAN tags. If frames are triple tagged, the higher-layer protocol information is not extracted.

Ethernet pseudo-wire LER frames are terminated by the MPLS-TP classifier and then processed at the Service Ethernet layer (inner) by the basic and advanced classifier. MPLS-encapsulated frames may also be processed purely at the MPLS Link Layer (outer Ethernet) by the MPLS-TP classifier or the basic and advanced classifier.

The following illustration shows the basic and advanced frame classification.

FIGURE 1-7: BASIC AND ADVANCED FRAME CLASSIFICATION

The basic classification classifies each frame to a VLAN, a QoS class, a drop precedence (DP) level, DSCP value, and an aggregation code. The basic classification also performs a general frame acceptance check. The output from the basic classification may be overwritten or changed by the more intelligent advanced classification using the IS1 TCAM.

Frame Acceptance The frame acceptance filter checks for valid combinations of VLAN tags against the ingress port's VLAN acceptance filter where it is possible to configure rules for accepting untagged, priority-tagged, C-, and S-tagged frames. In addition, the filter also enables discarding of frames with illegal MAC addresses (for instance null MAC address or multicast source MAC address).

VLAN Every incoming frame is classified to a VLAN by the basic VLAN classification. This is based on the VLAN in the frame, or if the frame is untagged or the ingress port is VLAN unaware, it is based on the ingress port's default VLAN. A VLAN classification includes the whole TCI (PCP, DEI, and VID) and also the TPID (C-tag or S-tag).

For double-tagged frames, it is selectable whether the inner or the outer tag is used.

The device can recognize S-tagged frames with the standard TPID (0x88A8) or S-tagged frames using a custom programmable value. One custom value is supported by the device.

QoS, DP, and DSCP Each frame is classified to a Quality of Service (QoS) class and a drop precedence level (frame color: green/yellow). The QoS class and DP level are used throughout the device for providing queuing, scheduling, and congestion control guarantees to the frame according to what is configured for that specific QoS class and color.

The QoS class and DP level in the basic classification are based on the class of service information in the frame's VLAN tags (PCP and DEI) and/or the DSCP values from the IP header. Both IPv4 and IPv6 are supported. If the frame is non-IP or untagged, the port's default QoS class and DP level are used.

The DSCP values can be remapped before being used for QoS. This is done using a common table mapping the incoming DSCP to a new value. Remapping is enabled per port. In addition, for each DSCP value, it is possible to specify whether the value is trusted for QoS purposes.

Each IP frame is also classified to an internal DSCP value. By default, this value is taken from the IP header but it may be remapped using the common DSCP mapping table or rewritten based on the assigned QoS class. The classified DSCP value may be written into the frame at egress – this is programmable in the rewriter.

Aggregation Code The basic classification calculates an aggregation code, which is used to select between ports that are member of a link aggregation group. The aggregation code is based on selected Layer-2 through Layer-4 information, such as MAC addresses, IP addresses, IPv6 flow label, ISDX, and TCP/UDP port numbers. The aggregation code ensures that frames belonging to the same conversation are using the same physical ports in a link aggregation group.

1.3.3.1 Advanced Classification

Following basic classification, Layer-2 and Layer-4 information is extracted from each frame and matched against a TCAM, IS1, with one of the following six different IS1 keys:

- **NORMAL.** Up to 512 entries with primary fields in key consisting of SMAC, outer VLAN tag, 32-bit source IP address, IP protocol, TCP/UDP source and destination port
- **NORMAL_IP6.** Up to 256 entries with primary fields in key consisting of SMAC, inner and outer VLAN tags, 128-

bit source IP address, IP protocol, TCP/UDP source and destination port

- 7TUPLE. Up to 256 entries with primary fields in key consisting of source and destination MAC addresses, inner and outer VLAN tags, 64-bit source and destination IP addresses, IP protocol, TCP/UDP source and destination port
- 5TUPLE_IP4. Up to 512 entries with primary fields in key consisting of inner and outer VLAN tags, 32-bit source and destination IP addresses, IP protocol, TCP/UDP source and destination port
- 5TUPLE_IP6. Up to 256 entries with primary fields in key consisting of inner and outer VLAN tags, 128-bit source and destination IP addresses, IP protocol, TCP/UDP source and destination port
- S1_DBL_VID. Up to 1024 entries with primary fields in key consisting of inner and outer VLAN tags

The TCAM embeds powerful protocol awareness for well-known protocols such as LLC, SNAP, IPv4, IPv6, and UDP/TCP. For each frame, three keys are generated and matched against the TCAM. The keys are selectable per ingress port per frame type (IPv4, IPv6, non-IP) per IS1 lookup.

The actions associated with each entry (programmed into the TCAM action RAM) include the ability to select a service (ISDX), overwrite or translate the classified VLAN, overwrite the priority code point (PCP) or the drop eligibility indicator (DEI), overwrite the QoS class and DP level, or overwrite the DSCP value. Each of these actions is enabled individually for each of the three lookup.


In addition, a policy association group (PAG) is assigned to the frame. The PAG identifies a security profile to which the frame belongs. The PAG is used in the succeeding security frame processor, IS2, to select which access control lists to apply to the frame. The PAG enables creating efficient ACLs that only are applicable to frames with the same PAG.

A lookup in IS1 can enable the use of a custom lookup in IS2 - the security enforcement TCAM. The custom lookup can match against selectable frame fields from the incoming frame. In total, 40 bytes from the frame are matched against the TCAM. This is a powerful future-proofing feature enabling handling of new protocols.

1.3.4 VERSATILE CONTENT AWARE PROCESSOR (VCAP)

All frames are inspected by the VCAP IS2 before they are passed on to the Layer-2 forwarding. The following illustration depicts VCAP security enforcement.

FIGURE 1-8: VCAP SECURITY ENFORCEMENT



The VCAP uses a TCAM-based frame processor enabling implementation of a rich set of security features. The flexible VCAP engine supports wire-speed frame inspection based on Layer 2-4 frame information, including the ability to perform longest prefix matching and identifying port ranges. The action associated with each VCAP entry (programmed

into the VCAP action RAM) includes the ability to do frame filtering, dual leaky bucket rate limitation (frame or byte based), snooping to CPU, redirection to CPU, mirroring, 1588 timestamping, OAM actions, and accounting. Even though the VCAP is located in the ingress path of the device, it possesses both ingress and egress capabilities.

The VCAP embeds powerful protocol awareness for well-known protocols such as LLC, SNAP, ARP, IPv4, IPv6, and UDP/TCP. IPv6 is supported with full matching against both the source and the destination IP addresses.

Each frame is looked up three times in IS2. The first lookup is a host match lookup for IPv4 and IPv6 frames enabling MAC/IP binding and IP source guarding. The key consists of information identifying the source host: ingress port number, source MAC address, and full source IP address. The output is an action informing the following two lookups in IS2 of whether the host is accepted into the network. The following two lookups in IS2 construct a key based on the frame type (LLC, SNAP, ARP, IPv4, IPv6, UDP/TCP, OAM, custom) extracting relevant information. IS2 supports up to 1024 host match entries, 512 LLC, SNAP, ARP, IPv4-TCP/UDP, or OAM entries, and 256 IPv6-TCP/UDP and custom entries.

1.3.5 ETHERNET OAM ENGINE

VSC7416-01 supports various OAM layer interactions as defined in ITU-T Y.1731, MEF E-LMI, ITU-T Y.1564 (Y.156sam), IEEE 802.3ah, and IEEE 802.1ag. The following OAM functions are supported in hardware:

- Continuity check frame generation and checking with sub-millisecond transmission periods.
- Loopback with frame generation and checking for both LBM and LBR.
- Frame loss measurements with support for both single-ended and dual-ended with hardware accurate counter readings for both LMM and LMR. Dual-ended loss measurements use CCM.
- Frame delay measurements with support for both one-way and two-way delay measurements with hardware accurate timestamps for both 1DM, DMM, and DMR.

Other OAM functions are supported but require CPU interaction.

VSC7416-01 contains 23 Versatile OAM endpoints (VOE), each handling the OAM functions of a MEG endpoint (MEP). Both port, path, and service MEP monitoring levels are supported and a MEP can be either an up or a down MEP. Each incoming OAM frames is processed by the assigned VOE, which include verification of the frame contents against configurations, forwarding based on the MEG level, and processing according to the received OAM type.

1.3.6 POLICING

Each frame is subject to a number of different policing operations. The device features 392 programmable policers. The policers are split into the followings groups:

- 256 service policers: Ingress service index (ISDX) selects which policer to use.
- 64 queue policers: Ingress port number and QoS class determine which policer to use.
- 8 port policers: Ingress port number determines which policer to use.
- 64 VCAP IS2 policers: An IS2 action can point to a policer.

The internal CPU port is the 8th port and also includes port and queue policers. The policers can measure frame rates or bit rates.

Each frame can trigger up to three policers: A service or queue policer, a port policer, and a VCAP IS2 policer.

Each policer is a MEF-compliant dual leaky bucket policer supporting both color-blind and color-aware operation. The initial frame color is derived from the drop precedence level from the frame classification. For color-aware operation, a coupling mode is configurable for each policer.

Using these policers ensures Service Level Agreement (SLA) compliance. The outcome of this policing operation is to mark each accepted frame as in-profile (green) or out-of-profile (yellow). Yellow frames are treated as excess or discard-eligible and green frames are committed. Frames that exceed the yellow/excess limits are discarded (red).

Each frame is counted in associated statistics reflecting the arrival service point (ISDX), and the frame's color (green, yellow, red). This statistics counts bytes and frames. In addition, each frame is also counted in statistics reflecting the ingress port, the QoS class, and the frame's color (green, yellow, red). This statistics can count bytes or frames.

Finally, the analyzer contains a group of storm control policers that are capable of policing various kinds of flooding traffic as well as CPU directed learn traffic. These policers are global policers working on all frames received by the switch. Storm policers measure frame rates.

1.3.7 LAYER-2 FORWARDING

After the policers, the Layer-2 forwarding block (the analyzer) handles all fundamental forwarding operations and maintains the associated MAC table, the VLAN table, the ISDX table, and the aggregation table. The device implements an 8K MAC table and a 4K VLAN table.

The main task of the analyzer is to determine the destination port set of each frame. This forwarding decision is based on various information such as the frame's ingress port, the source MAC address, the destination MAC address, the VLAN identifier, and the ISDX, as well as the frame's VCAP action, mirroring, and the destination port's link aggregation configuration.

The switch performs Layer-2 forwarding of frames. For unicast and Layer-2 multicast frames, this means forwarding based on the destination MAC address and the VLAN. For IPv4 and IPv6 multicast frames, the switch performs Layer-2 forwarding, but based on Layer-3 information, such as the source IP address. The latter enables source-specific IPv4 multicast forwarding (IGMPv3) and source-specific IPv6 multicast forwarding (MLDv2).

The following are some of the contributions to the Layer-2 forwarding.

- VLAN classification. VLAN-based forward filtering includes source port filtering, destination port filtering, VLAN mirroring, asymmetric VLANs, and so on.
- Service classification. Service-based forward filtering includes destination port filtering and forced forwarding.
- Security enforcement. The security decision made by the VCAP can, for example, redirect the frame to the CPU based on some abnormality detection filters.
- MSTP. The VLAN identifier maps to a Multiple Spanning Tree instance, which determines MSTP-based destination port filtering.
- MPLS-TP. Destination set determined by the MPLS-TP classifier.
- MAC addresses. Destination and source MAC address lookups in the MAC table determine if a frame is a learn frame, a flood frame, a multicast frame, or a unicast frame.
- Learning. By default, the device performs wire-speed learning on all ports. However, certain ports could be configured with secure learning enabled, where an incoming frame with unknown source MAC address is classified as a "learn frame" and is redirected to the CPU. The CPU performs the learning decision and also decides whether the frame is forwarded.

Learning can also be disabled. In that case, it does not matter if the source MAC address is in the MAC table.


- Link aggregation. A frame targeted at a link aggregate is further processed to determine which of the link aggregate group ports the frame must be forwarded to.
- Mirroring. Mirror probes may be set up in different places in the forwarding path for monitoring purposes. As part of a mirror a copy of the frame is sent either to the CPU or to another port.

1.3.8 SHARED QUEUE SYSTEM AND HIERARCHICAL EGRESS SCHEDULER

The analyzer provides the destination port set of a frame to the shared queue system. It is the queue system's task to control the frame forwarding to all destination ports.

The shared queue system embeds 8 megabits of memory that can be shared between all queues and ports. The queue system consists of 2616 queues with programmable use. Frames are mapped to queues through a programmable mapping function allowing ingress ports, egress ports, QoS classes, and services to be taken into account. The sharing of resources between queues and ports is controlled by an extensive set of thresholds.

Each egress port implements a default scheduler and shapers, as shown in the following illustration. Per egress port, the scheduler sees the outcome of aggregating the egress queues (one per ingress port per QoS class) into eight QoS classes. By default, aggregation is done in a deficit weighted round-robin fashion (DWRR) with equal weights to all ingress ports within a QoS class. As a result, byte-based round robin between ingress ports is performed for each QoS class. Scheduling between QoS classes for the port is by default a combination of strict and DWRR scheduling. QoS classes 7 and 6 have strict highest priority, and QoS classes 5 through 0 are scheduled using deficit weighted round-robin.

FIGURE 1-9: DEFAULT EGRESS SCHEDULER AND SHAPER CONFIGURATION

Scheduling between QoS classes within the port can use one of three methods:


- Strict. Frames with the highest QoS class are always transmitted before frames with lower QoS class.
- Combination of strict and Deficit Weighted Round Robin (DWRR) scheduling. Any split between strict and DWRR QoS classes can be configured. Each QoS class serviced using DWRR sets a DWRR weight ranging from 0 to 31.
- Combination of strict and Frame Based Round Robin scheduling (FBRR) scheduling. Any split between strict and FBRR QoS classes can be configured.

All shapers shown in [Figure 1-9](#), are dual leaky bucket shapers.

It is also possible to configure a hierarchical Quality of Service (H-QoS) egress scheduler on a few egress ports, such as network facing ports. There are a total pool of 256 scheduler elements (SE). The first 218 scheduler elements connect to the 2616 queues in the queue system. The remaining 38 SEs are highly flexible and can be used to build H-QoS scheduling. Each scheduler element has 12 inputs and 1 output. The output contains a dual leaky bucket shaper. The scheduling algorithm within each element can be strict, deficit weighted round robin (DWRR), or frame based round robin (FBRR). Combinations of strict versus DWRR/FBRR scheduling is supported as described above for the scheduling between QoS classes.

The scheduler elements are programmed to connect to each other (output of one element is input to another and so on) to achieve the overall desired scheduling hierarchy. The following illustration shows an example of a three-level hierarchy on an egress port implementing a queue system for 96 service groups with each eight QoS classes. A total of 768 egress queues is associated with the port in this example. This H-QoS scheduler utilizes 13 flexible SEs and 96 fixed SEs.

FIGURE 1-10: ADVANCED SERVICE-BASED EGRESS SCHEDULER AND SHAPER




In addition to the dual leaky bucket shaper per scheduler element, the queue system also implements in total 436 single leaky bucket shapers, which enables dedicated shaping on QoS class 6 and 7.

1.3.9 REWRITER AND FRAME DEPARTURE

Before transmitting the frame on the egress line, the rewriter can modify selected fields in the frame, such as MPLS encapsulation, VLAN tags, DSCP value, timestamping, OAM fields, and FCS.

The rewriter controls the final VLAN tagging of frames based on the classified VLAN, the VLAN pop count, and egress-determined VLAN actions from the ES0 TCAM lookup. The egress VLAN actions are by default given by the egress port settings. These include normal VLAN operations such as pushing a VLAN tag, untagging for specific VLANs, and simple translations of DEI and PCP.

FIGURE 1-11: ADVANCED FRAME REWRITING



By using the egress TCAM, ES0, much more advanced VLAN tagging operations can be achieved. ES0 enables pushing up to two VLAN tags and allows for a flexible translation of the VLAN tag header. The key into ES0 is the combination of the ingress port, the egress port, and either the classified VLAN tag header or the classified service.

The PCP and DEI bits in the VLAN tag are subject to remarking based on translating the classified tag header or by using the classified QoS value and the frame's drop precedence level from ingress.

MPLS popping and pushing is handled by the ES0 TCAM. Up to 3 MPLS labels and 1 control word can be popped and pushed as well as an Ethernet link layer header (DMAC, SMAC, VLAN tag, and MPLS EtherType).

The TC bits in the MPLS labels are subject to remarking based on the classified values, a mapping to new values based on QoS class and drop precedence level, or new programmable values. The S bit in the MPLS labels can either use the classified value or a new programmable value.

In addition, the DSCP value in IP frames can be updated using the classified DSCP value and the frame's drop precedence level from ingress. The DSCP value can be remapped at egress before writing it into the frame.

Finally, the rewriter updates the FCS if the frame was modified before the frame is transmitted.

The egress port module controls the flow control exchange of pause frames with a neighboring device when the interconnection link operates in full-duplex flow control mode. When the connected device triggers flow control through transmission of a pause frame, the MAC stops the egress scheduler's forwarding of frames out of the port. Traffic then builds up in the queue system but sufficient queuing is available to ensure wire speed loss-less operation.

In half-duplex operation, the port module's egress path responds to back pressure generation from a connected device by collision detection and frame retransmission.

1.3.10 CPU PORT MODULE

The CPU port module contains eight CPU extraction queues and two CPU injection queues. These queues provide an interface for exchanging frames between the internal CPU system and the switch core. An external CPU using the serial interface can also inject and extract frames to and from the switch core by using the CPU port module. Any Ethernet interface on the device can also be used for extracting and injecting frames.

The switch core can intercept a variety of different frame types and copy or redirect these to the CPU extraction queues. The classifier can identify a set of well-known frames such as IEEE reserved destination MAC addresses (BPDUs, GARPs, CCM/Link trace), as well as IP-specific frames (IGMP, MLD). The security TCAM, IS2, provides another very

flexible way of intercepting all kinds of frames, for instance specific OAM frames, ARP frames or explicit applications based on TCP/UDP port numbers. In addition, frames can be intercepted based on the MAC table, the VLAN table, or the learning process.

Whenever a frame is copied or redirected to the CPU, a CPU extraction queue number is associated with the frame and used by the CPU port module when enqueueing the frame into the 8 CPU extraction queues. The CPU extraction queue number is programmable for every interception option in the switch core.

1.3.11 SYNCHRONOUS ETHERNET AND PRECISION TIME PROTOCOL

The device supports Layer-1 ITU-T G.8261 Synchronous Ethernet and Layer-2 IEEE 1588 Precision Time Protocol for synchronizing network timing throughout a network.

Synchronous Ethernet allows for the transfer of network timing from one reference to all network elements. In VSC7416-01, each port can recover its ingress clock and output the recovered clock to one of up to ten recovered clock output pins. External circuitry can then generate a stable reference clock input used for egress and core logic timing in VSC7416-01.

The Precision Time Protocol (PTP) allows for the network-wide synchronization of precise time of day. It is also possible to derive network timing. PTP can operate with a one-step clock or a two-step clock. For one-step clocks, a frame's residence time is calculated and stamped into the frame at departure. For two-step clocks, a frame's residence time is simply recorded and provided to the CPU for further processing. The CPU can then initiate a follow-up message with the recorded timing.

PTP is supported for a range of encapsulations including PTP over Ethernet/IEEE 802.3 and PTP over UDP over IPv4/IPv6, and PTP over UDP over IP/MPLS.

The VSC7416-01 device contains two separate timing domains; one for Synchronous Ethernet and data path forwarding, and one for PTP timing synchronization.

1.3.12 CPU SYSTEM AND INTERFACES

The device features a VCore-III CPU system containing a powerful 416 MHz MIPS 24KEc™ CPU. It is suitable for lightly managed and fully managed applications.

VCore-III includes a general-purpose direct memory access engine (GPDMA) that also supports frame-based direct memory access (FMDA) operations. The FMDA offloads the CPU when injecting and extracting frames to and from the switch core. VCore-III boots from a serial flash and uses an 8-bit or 16-bit DDR2/3 SDRAM for memory, in addition to its built-in 32 kilobytes of instruction cache and 32 kilobytes of data cache. An external debugger can be attached to the EJTAG interface.

In addition to the integrated processor, the CPU system permits the attachment of an external CPU. For configuration of switch register, an external CPU can use either a serial interface, an MII Management interface, a PCIe interface, or an SGMII port.

When using an SGMII port, the Versatile Register Access Protocol (VRAP) allows an external CPU to access registers in VSC7416-01 using VRAP Ethernet frames. No internal CPU is required. Each VRAP request frame contains a list of read, read-modify-write, and write commands that are processed by the hardware. For each request frame, a VRAP reply frame is generated containing the results of the requests. Each frame can contain multiple commands only restricted by the MTU size.

For frame transfers, the external CPU has the option of using the serial interface, an MII Management interface, a PCIe interface, or an SGMII port.

The VSC7416-01 device includes a GPIO interface with 32 individually configurable pins. Through the GPIO pins, VSC7416-01 supports the following interfaces:

- Two-wire serial interface (two GPIO pins)
- Dual UART (four GPIO pins)
- External interrupts (two interrupt pins)
- Serial GPIO (SGPIO) and LED interface (four GPIO pins)
- MII Management slave interface for accessing switch registers from an external CPU (three GPIO pins)
- Dual MII Management master interface for controlling external PHYs (four GPIO pins)
- SFP signal detect per SerDes for fast detect and reaction to loss of signal (11 GPIO pins)
- IEEE 1588 pins with programmable synchronized 1588 clocks (two GPIO pins)
- IEEE 1588 pins with load-and-store functionality of 1588 clocks for synchronizing with external equipment (two

GPIO pins)

The serial GPIO and LED interfaces can be used specifically for driving external LEDs for external copper PHYs or for serializing external interrupts. For instance, link down events from external PHYs, before being input to the device.

2.0 FUNCTIONAL DESCRIPTIONS

This section provides information about the functional aspects of the VSC7416-01 Carrier Ethernet switch device, available configurations, operational features, and testing functionality.

2.1 Port Modules

The port modules contain the following functional blocks.

- MAC
- PCS

2.1.1 PORT MODULE NUMBERING AND MACRO CONNECTIONS

The port modules connect to the interface macros, which can be of two types:

- SERDES6G
- SERDES1G

The interface macros connect to the external interface pins. For more information about the SerDes macros, see [Section 2.2, SERDES1G](#) and [Section 2.3, SERDES6G](#).

The following table lists the mapping from the switch core port modules to the interface macros. The number next to the interface macro type (for example, “8” in cell “SERDES1G, 8”) indicates the macro number that must be used when addressing the macros.

TABLE 2-1: PORT MAPPING FROM SWITCH CORE PORT MODULE TO INTERFACE MACROS

Switch Core Port Module	VSC7416-01
0-3	SERDES1G, 0-3
8-9	SERDES6G, 0-1
10	SERDES1G, NPI
11	Internal CPU port

2.1.2 MAC

This section provides information about the high-level functionality and the configuration options of the Media Access Controller (MAC) that is used in each of the port modules.

The MAC supports the following speeds and duplex modes:

- SERDES1G ports: 10/100/1000 Mbps in full-duplex mode and 10/100 Mbps in half-duplex mode
- SERDES6G ports: 10/100/1000/2500 Mbps in full-duplex mode and 10/100 Mbps in half-duplex mode.

The following table lists the registers associated with configuring the MAC.

TABLE 2-2: MAC CONFIGURATION REGISTERS

Register	Description	Replication
DEV::CLOCK_CFG	Reset and speed configuration	Per port
DEV::MAC_ENA_CFG	Enabling of Rx and Tx data paths	Per port
DEV::MAC_MODE_CFG	Port mode configuration	Per port
DEV::MAC_MAXLEN_CFG	Maximum length configuration	Per port
DEV::MAC_TAGS_CFG	VLAN tag length configuration	Per port
DEV::MAC_ADV_CHK_CFG	Type length configuration	Per port
DEV::MAC_IFG_CFG	Interframe gap configuration	Per port
DEV::MAC_HDX_CFG	Half-duplex configuration	Per port
SYS::MAC_FC_CFG	Flow control configuration	Per port
DEV::MAC_FC_MAC_LOW_CFG	LSB of SMAC used in pause frames	Per port
DEV::MAC_FC_MAC_HIGH_CFG	MSB of SMAC used in pause frames	Per port
DEV::MAC_STICKY	Sticky bit recordings	Per port

2.1.2.1 Reset

There are a number of resets in the port module. All of the resets can be set and cleared simultaneously. By default, all blocks are in the reset state. With reference to register `CLOCK_CFG`, the resets are as follows:

- `MAC_RX_RST`: Reset of the MAC receiver
- `MAC_TX_RST`: Reset of the MAC transmitter
- `PORT_RST`: Reset of the ingress and egress queues
- `PCS_RX_RST`: Reset of the PCS decoder
- `PCS_TX_RST`: Reset of the PCS encoder

When changing the MAC configuration, the port must go through a reset cycle. This is done by writing register `CLOCK_CFG` twice. On the first write, the reset bits are set. On the second write, the reset bits are cleared. Bits that are not reset bits in `CLOCK_CFG` must keep their new value for both writes.

For more information about resetting a port, see [Section 4.2.3, Port Reset Procedure](#).

2.1.2.2 Port Mode Configuration

The MAC provides a number of handles for configuring the port mode. With reference to the `MAC_MODE_CFG`, `MAC_IFG_CFG`, and `MAC_ENA_CFG` registers, the handles are as follows:

- Duplex mode (`FDX_ENA`). Half or full duplex.
- Data sampling (`GIGA_MODE_ENA`). Must be 1 in 1 Gbps and 2.5 Gbps and 0 in 10 Mbps and 100 Mbps.
- Enabling transmission and reception of frames (`TX_ENA/RX_ENA`). Clearing `RX_ENA` stops the reception of frames and further frames are discarded. An ongoing frame reception is interrupted. Clearing `TX_ENA` stops the dequeuing of frames from the egress queues, which means that frames are held back in the egress queues. An ongoing frame transmission is completed.
- Tx-to-Tx inter-frame gap (`TX_IFG`)

The link speed is configured using `CLOCK_CFG.LINK_SPEED` with the following options.

- Link speed (`CLOCK_CFG.LINK_SPEED`)
1 Gbps (125 MHz clock)

Ports 8 and 9: 1 Gbps or 2.5 Gbps (125 MHz or 312.5 MHz clock). The actual clock frequency depends on the SerDes configuration.

100 Mbps (25 MHz clock)

10 Mbps (2.5 MHz clock)

2.1.2.3 Half Duplex

A number of special configuration options are available for half-duplex (HDX) mode:

- **Seed for back-off randomizer.** Field `MAC_HDX_CFG.SEED` seeds the randomizer used by the backoff algorithm. Use `MAC_HDX_CFG.SEED_LOAD` to load a new seed value.
- **Backoff after excessive collision.** Field `MAC_HDX_CFG.WEXC_DIS` determines whether the MAC backs off after an excessive collision has occurred. If set, backoff is disabled after excessive collisions.
- **Retransmission of frame after excessive collision.** Field `MAC_HDX_CFG.RETRY_AFTER_EXC_COL_ENA` determines whether the MAC retransmits frames after an excessive collision has occurred. If set, a frame is not dropped after excessive collisions, but the backoff sequence is restarted. This is a violation of IEEE 802.3, but is useful in non-dropping half-duplex flow control operation.
- **Late collision timing.** Field `MAC_HDX_CFG.LATE_COL_POS` adjusts the border between a collision and a late collision in steps of 1 byte. According to IEEE 802.3, section 21.3, this border is permitted to be on data byte 56 (counting frame data from 1); that is, a frame experiencing a collision on data byte 55 is always retransmitted, but it is never retransmitted when the collision is on byte 57. For each higher `LATE_COL_POS` value, the border is moved 1 byte higher.
- **Rx-to-Tx inter-frame gap.** The sum of `MAC_IFG_CFG.RX_IFG1` and `MAC_IFG_CFG.RX_IFG2` establishes the time for the Rx-to-Tx inter-frame gap. `RX_IFG1` is the first part of half-duplex Rx-to-Tx inter-frame gap. Within `RX_IFG1`, this timing is restarted if carrier sense (CRS) has multiple high-low transitions (due to noise). `RX_IFG2` is the second part of half-duplex Rx-to-Tx inter-frame gap. Within `RX_IFG2`, transitions on CRS are ignored.

When enabling a port for half-duplex mode, the switch core must also be enabled (SYS::FRONT_PORT_MODE.HDX_MODE).

2.1.2.4 Frame and Type/Length Check

The MAC supports frame lengths of up to 16 kilobytes. The maximum length accepted by the MAC is configurable in MAC_MACLEN_CFG.MAX_LEN.

The MAC allows tagged frames to be 4 bytes longer and double-tagged frames to be 8 bytes longer than the specified maximum length (MAC_TAGS_CFG.VLAN_LEN_AWR_ENA). The MAC must be configured to look for VLAN tags. By default, EtherType 0x8100 identifies a VLAN tag. In addition, a custom EtherType can be configured in MAC_TAGS_CFG.TAG_ID. The MAC can be configured to look for none, one, or two tags (MAC_TAG_CFG.VLAN_AWR_ENA, MAC_TAG_CFG.VLAN_DBL_AWR_ENA).

The type/length check (MAC_ADV_CHK_CFG.LEN_DROP_ENA) causes the MAC to discard frames with type/length errors (in-range and out-of-range errors).

2.1.2.5 Flow Control (IEEE 802.3x)

VSC7416-01 supports both standard full-duplex flow control (IEEE 802.3x) and priority-based flow control (IEEE 802.1Qbb). This section describes standard full-duplex flow control, and [Section 2.1.2.6, Priority-Based Flow Control \(IEEE 802.1Qbb\)](#) describes priority-based flow control.

In full-duplex mode, the MAC provides independent support for transmission of pause frames and reaction to incoming pause frames. This allows for asymmetric flow control configurations.

The MAC obeys received pause frames (MAC_FC_CFG.RX_FC_ENA) by pausing the egress traffic according to the timer values specified in the pause frames. In order to evaluate the pause time in the incoming pause frames, the link speed must be specified (SYS::MAC_FC_CFG.FC_LINK_SPEED).

The transmission of pause frames is triggered by assertion of a flow control condition in the ingress queues caused by a queue filling exceeding a watermark. For more information, see [Section 2.10.7, Ingress Pause Request Generation](#). The MAC handles the formatting and transmission of the pause frame. The following configuration options are available:

- Transmission of pause frames (MAC_CFG_CFG.TX_FC_ENA).
- Pause timer value used in transmitted pause frames (MAC_FC_CFG.PAUSE_VAL_CFG).
- Flow control cancellation when the ingress queues de-assert the flow control condition by transmission of a pause frame with timer value 0 (MAC_FC_CFG.ZERO_PAUSE_ENA).
- Source MAC address used in transmitted pause frames (MAC_FC_MAC_HIGH_CFG, MAC_FC_MAC_LOW_CFG).

The MAC has the option to discard incoming frames when the remote link partner is not obeying the pause frames transmitted by the MAC. The MAC discards an incoming frame if a Start-of-Frame is seen after the pause frame was transmitted. It is configurable how long reaction time is given to the link partner (MAC_FC_CFG.FC_LATENCY_CFG). The benefit of this approach is that the queue system is not risking being overloaded with frames due to a non-complying link partner.

In half-duplex mode, the MAC does not react to received pause frames. If the flow control condition is asserted by the ingress queues, the industry-standard backpressure mechanism is used. Together with the ability to retransmit frames after excessive collisions (MAC_HDX_CFG.RETRY_AFTER_EXC_COL_ENA), this enables non-dropping half-duplex flow control.

2.1.2.6 Priority-Based Flow Control (IEEE 802.1Qbb)

VSC7416-01 supports priority-based flow control on all ports for all QoS classes. The following table lists the specific registers associated with priority-based flow control.

TABLE 2-3: PRIORITY-BASED FLOW CONTROL CONFIGURATION REGISTERS

Register	Description	Replication
SYS::MAC_FC_CFG	Flow control configuration	Per port
DEV::MAC_FC_MAC_LOW_CFG	LSB of SMAC used in pause frames	Per port
DEV::MAC_FC_MAC_HIGH_CFG	MSB of SMAC used in pause frames	Per port
ANA::PFC_CFG	Configuration of Rx priority-based flow control per priority.	Per port

TABLE 2-3: PRIORITY-BASED FLOW CONTROL CONFIGURATION REGISTERS (CONTINUED)

Register	Description	Replication
QSYS::SWITCH_PORT_MODE	Configuration of Tx priority-based flow control per priority	Per port
DEV::PORT_MISC.FWD_C-TRL_ENA	Enabling forwarding of priority-based pause frames to analyzer.	Per port
ANA::CPU_FWD_BPDU_CFG	Disable forwarding of priority-based pause frames beyond the analyzer	Per port

VSC7416-01 provides independent support for transmission of pause frames and reaction to incoming pause frames, which allows asymmetric flow control configurations.

The device obeys received pause frames per priority (ANA::PFC_CFG.RX_PFC_ENA) by pausing the egress traffic according to the timer values specified in the pause frames. Transmission of frames belonging to QoS class n is paused if bit n is set in the priority_enable_vector in the incoming pause frame. The pause time for QoS class n is given by time[n] from the pause frame. The link speed must be specified in order to evaluate the pause times (ANA::PFC_CFG.FC_LINK_SPEED).

The transmission of priority-based pause frames is triggered by assertion of a flow control condition in the ingress queues caused by the memory consumption for a priority for an ingress port exceeding the BUF_Q_RSRV_I watermark. For more information about the watermark, see [Table 2-120](#). The MAC handles the formatting and transmission of the priority-based pause frame. The following configuration options are available:

- Transmission of priority-based pause frames per QoS class (QSYS::SWITCH_PORT_MODE.TX_PFC_ENA).
- Pause timer value used in transmitted priority-based pause frames (SYS::MAC_FC_CFG.PAUSE_VAL_CFG). All congested priorities use the same pause timer value. Uncongested priorities use pause timer value 0.
- Source MAC address used in transmitted priority-based pause frames (MAC_FC_MAC_HIGH_CFG, MAC_FC_MAC_LOW_CFG).
- Priority protection mode (QSYS::SWITCH_PORT_MODE.TX_PFC_MODE), which enables that when a priority congests and causes a pause frame to be sent, then the pause frame will also pause all lower priorities.

All transmitted priority-based pause frames have the priority_enable_vector set to 0xFF, independently of whether a priority is enabled for flow control. However, the pause timer value, time[n], is always 0 for disabled priorities.

The MAC generates and transmits a priority-based pause frames whenever a queue is congested. The VSC7416-01 device prevents excessive pause frame generation by waiting half the pause timer value between transmissions of pause frames.

When an ingress queue de-asserts the flow control condition, the MAC does not generate a priority-based pause frame with pause timer 0 for the priority. Instead, the timer in the link partner must expire. However, if another queue asserts flow control, then a priority-based pause frame is generated for that priority and for all uncongested queues a pause timer 0 is signaled to the link partner.

2.1.2.7 Frame Aging

The following table lists the registers associated with frame aging.

TABLE 2-4: FRAME AGING CONFIGURATION REGISTERS

Register	Description	Replication
SYS::FRM_AGING	Frame aging time	None
REW::PORT_CFG.AGE_DIS	Disable frame aging	Per port

The MAC supports frame aging where frames are discarded if a maximum transit delay through the switch is exceeded. All frames, including CPU-injected frames, are subject to aging. The transit delay is time from when a frame is fully received until that frame is scheduled for transmission through the egress MAC. The maximum allowed transit delay is configured in SYS::FRM_AGING.

Frame aging can be disabled per port (REW::PORT_CFG.AGE_DIS).

Discarded frames due to frame aging are counted in the c_tx_aged counter.

2.1.2.8 Rx and Tx Timestamps

The following table lists the registers associated with Rx and Tx timestamping.

TABLE 2-5: RX AND TX TIMESTAMPING REGISTER

Register	Description	Replication
SYS::IO_PATH_DELAY	I/O delays	Per port
ANA:PORT:PTP_CFG.PTP_BACKPLANE_MODE	Ingress backplane configuration	Per port

The MAC supports Rx and Tx timestamping where a frame's receive time and transmit time are sampled using a nanoseconds counter. The receive and transmit timestamps are shifted in time so that the resulting timestamps align with the exact reception and transmission of the first byte in the frame. This adjusts for local delays in the receive path and the transmit path. The timestamps are individually adjusted as follows:

- Rx timestamp: Sampling of the MAC's nanoseconds counter plus SYS:IO_PATH_DELAY/PORT.RX_PATH_DELAY.
- Tx timestamp: Sampling of the MAC's nanoseconds counter plus SYS:IO_PATH_DELAY/PORT.TX_PATH_DELAY.

The Rx and Tx path delays are signed so negative adjustments can be achieved by setting the most significant bit.

When the ingress port is operating in backplane mode (ANA:PORT:PTP_CFG.PTP_BACKPLANE_MODE), the Rx timestamp is set to the value of the 4-byte reserved field at offset 16 bytes in the incoming frame's PTP header instead of using the adjusted Rx timestamp from the MAC.

The Rx timestamp follows the frame to the rewriter where it can be used to calculate the frame's residence time or it can be sent to the CPU together with the frame. The Tx timestamp is used by the rewriter to calculate the frame's residence time or it can be written to a timestamp FIFO queue accessible by the CPU. For more information about hardware timestamping, see [Section 2.17, Hardware Timestamping](#).

2.1.3 PCS

This section provides information about the Physical Coding Sublayer (PCS) block, where the auto-negotiation process establishes mode of operation for a link. The PCS supports both SGMII mode and two SerDes modes, 100BASE-X and 100BASE-FX.

The following table lists the registers associated with PCS.

TABLE 2-6: PCS CONFIGURATION REGISTERS

Registers	Description	Replication
PCS1G_CFG	PCS configuration	Per PCS
PCS1G_MODE_CFG	PCS mode configuration	Per PCS
PCS1G_SD_CFG	Signal detect configuration	Per PCS
PCS1G_ANEG_CFG	Configuration of the PCS auto-negotiation process	Per PCS
PCS1G_ANEG_NP_CFG	Auto-negotiation next page configuration	Per PCS
PCS1G_LB_CFG	Loop-back configuration	Per PCS
PCS1G_ANEG_STATUS	Status signaling of the PCS auto-negotiation process	Per PCS
PCS1G_ANEG_NP_STATUS	Status signaling of the PCS auto-negotiation next page process	Per PCS
PCS1G_LINK_STATUS	Link status	Per PCS
PCS1G_LINK_DOWN_COUNTER	Link down counter	Per PCS
PCS1G_STICKY	Sticky bit register	Per PCS

The PCS is enabled in PCS1G_CFG.PCS_ENA and supports both SGMII and 100BASE-X SERDES mode (PCS_MODE_CFG.SGMII_MODE_ENA), as well as 100BASE-FX. For information about enabling 100BASE-FX, see [Section 2.1.3.7, 100BASE-FX](#).

The PCS also supports the IEEE 802.3, Clause 66 unidirectional mode, where the transmission of data is independent of the state of the receive link (PCS_MODE_CFG.UNIDIR_MODE_ENA).

2.1.3.1 Auto-Negotiation

Auto-negotiation is enabled in PCS1G_ANEG_CFG.ANEG_ENA. To restart the auto-negotiation process, PCS1G_ANEG_CFG.ANEG_RESTART_ONE_SHOT must be set.

The advertised word for the auto-negotiation process (base page) is configured in PCS1G_ANEG_CFG.ADV_ABILITY. The next page information is configured in PCS1G_ANEG_NP_CFG.NP_TX.

When the auto-negotiation state machine has exchanged base page abilities, the PCS1G_ANEG_STATUS.PAGE_RX_STICKY is asserted indicating that the link partner's abilities were received (PCS1G_ANEG_STATUS.LP_ADV_ABILITY).

If next page information is exchanged, PAGE_RX_STICKY must be cleared, next page abilities must be written to PCS1G_ANEG_NP_CFG.NP_TX, and PCS1G_ANEG_NP_CFG.NP_LOADED_ONE_SHOT must be set. When the auto-negotiation state machine has exchanged the next page abilities, the PCS1G_ANEG_STATUS.PAGE_RX_STICKY is asserted again, indicating that the link partner's next page abilities were received (PCS1G_ANEG_STATUS.LP_NP_RX). Additional exchanges of next page information are possible using the same procedure.

After the last next page is received, the auto-negotiation state machine enters the IDLE_DETECT state and the PCS1G_ANEG_STATUS.PR bit is set indicating that ability information exchange (base page and possible next pages) is finished and software can now resolve priority. Appropriate actions, such as Rx or Tx reset, or auto-negotiation restart, can then be taken, based on the negotiated abilities. The LINK_OK state is reached one link timer period later.

When the auto-negotiation process reaches the LINK_OK state, PCS1G_ANEG_STATUS.ANEG_COMPLETE is asserted.

2.1.3.2 Link Surveillance

The current link status can be observed through PCS1G_LINK_STATUS.LINK_STATUS. The LINK_STATUS is defined as either the PCS synchronization state or as bit 15 of PCS1G_ANEG_STATUS.LP_ADV_ABILITY, which carries information about the link status of the attached PHY in SGMII mode.

Link down is defined as the auto-negotiation state machine being in neither the AN_DISABLE_LINK_OK state nor the LINK_OK state for one link timer period. If a link down event occurs, PCS1G_STICKY.LINK_DOWN_STICKY is set, and PCS1G_LINK_DOWN_CNT is incremented. In SGMII mode, the link timer period is 1.6 ms; in SerDes mode, the link timer period is 10 ms.

The PCS synchronization state can be observed through PCS1G_LINK_STATUS.SYNC_STATUS. Synchronization is lost when the PCS is not able to recover and decode data received from the attached serial link.

2.1.3.3 Signal Detect

The PCS can be enabled to react to loss of signal through signal detect (PCS1G_SD_CFG.SD_ENA). At loss of signal, the PCS Rx state machine is restarted, and frame reception stops. If signal detect is disabled, no action is taken upon loss of signal. The polarity of signal detect is configurable in PCS1G_SD_CFG.SD_POL.

The source of signal detect is selected in PCS1G_SD_CFG.SD_SEL to either the SerDes PMA or the PMD receiver. If the SerDes PMA is used as source, the SerDes macro provides the signal detect. If the PMD receiver is used as source, signal detect is sampled externally through one of the GPIO pins on the VSC7416-01 device. For more information about the configuration of the GPIOs and signal detect, see [Section 3.8.7, GPIO Controller](#).

PCS1G_LINK_STATUS.SIGNAL_DETECT contains the current value of the signal detect input.

2.1.3.4 Tx Loopback

For debug purposes, the Tx data path in the PCS can be looped back into the Rx data path. This feature is enabled through PCS1G_LB_CFG.TBI_HOST_LB_ENA.

2.1.3.5 Test Patterns

The following table lists the registers associated with configuring test patterns.

TABLE 2-7: TEST PATTERN REGISTERS

Registers	Description	Replication
PCS1G_TSTPAT_MODE_CFG	Test pattern configuration	Per PSC
PCS1G_TSTPAT_MODE_STATUS	Test pattern status	Per PCS

PCS1G_TSTPAT_MODE_CFG.JTP_SEL overwrites normal operation of the PCS and enables generation of jitter test patterns for debugging. The jitter test patterns are defined in IEEE 802.3, Annex 36A, and the following patterns are supported.

- High frequency test pattern
- Low frequency test pattern
- Mixed frequency test pattern
- Continuous random test pattern with long frames
- Continuous random test pattern with short frames

PCS1G_TSTPAT_MODE_STATUS register holds information about error and lock conditions while running the jitter test patterns.

2.1.3.6 Low Power Idle

The following table lists the registers associated with low power idle (LPI).

TABLE 2-8: LOW POWER IDLE REGISTERS

Registers	Description	Replication
PCS1G_LPI_CFG	Configuration of the PCS low power idle process	Per PSC
PCS1G_LPI_WAKE_ERROR_CNT	Error counter	Per PCS
PCS1G_LPI_STATUS	Low Power Idle status	Per PCS

The PCS supports Energy Efficient Ethernet (EEE) as defined by IEEE 802.3az. The PCS converts Low Power Idle (LPI) encoding between the MAC and the serial interface transparently. In addition, the PCS provides control signals allowing to stop data transmission in the SerDes macro. During low power idles the serial transmitter in the SerDes macro can be powered down, only interrupted periodically while transmitting refresh information, which allows the receiver to notice that the link is still up but in power down mode.

For more information about powering down the serial transmitter in the SerDes macros, see [Section 2.2, SERDES1G](#) or [Section 2.3, SERDES6G](#).

It is not necessary to enable the PCS for EEE, because it is controlled indirectly by the shared queue system. It is possible, however, to manually force the PCS into the low power idle mode through PCS1G_LPI_CFG.TX_ASSERT_LPI-DLE. During LPI mode, the PCS constantly encodes low power idle with periodical refreshes. For more information about EEE, see [Section 2.10.10, Energy Efficient Ethernet](#).

The current low power idle state can be observed through PCS1G_LPI_STATUS for both receiver and transmitter:

- RX_LPI_MODE: Set if the receiver is in low power idle mode.
- RX_QUIET: Set if the receiver is in the Quiet state of the low power idle mode. If cleared while RX_LPI_MODE is set, the receiver is in the refresh state of the low power idle mode.

The same is observable for the transmitter through TX_LPI_MODE and TX_QUIET.

If an LPI symbol is received, the RX_LPI_EVENT_STICKY bit is set, and if an LPI symbol is transmitted, the TX_LPI_EVENT_STICKY bit is set. These events are sticky.

The PCS1G_LPI_WAKE_ERROR_CNT wake-up error counter increments when the receiver detects a signal and the PCS is not synchronized. This can happen when the transmitter fails to observe the wake-up time or if the receiver is not able to synchronize in time.

2.1.3.7 100BASE-FX

The following table lists the registers associated with 100BASE-FX configuration.

TABLE 2-9: 100BASE-FX REGISTERS

Registers	Description	Replication
PCS_FX100_CFG	Configuration of the PCS 100BASE-FX mode	Per PSC
PCS_FX100_STATUS	Status of the PCS 100BASE-FX mode	Per PCS

The PCS supports a 100BASE-FX mode in addition to the SGMII and 1000BASE-X SerDes modes. The 100BASE-FX mode uses 4-bit/5-bit coding as specified in IEEE 802.3 Clause 24 for fiber connections. The 100BASE-FX mode is enabled through PCS_FX100_CFG.PCS_ENA, which masks out all PCS1G related registers.

The following options are available:

Far-End Fault facility. In 100BASE-FX, the PCS supports the optional Far-End Fault facility. Both Far-End Fault generation (PCS_FX100_CFG.FEF_GEN_ENA) and Far-End Fault Detection (PCS_FX100_CFG.FEF_CHK_ENA) are supported. An Far-End Fault incident is recorded in PCS_FX100_STATUS.FEF_FOUND.

Signal Detect. 100BASE-FX has a similar signal detect scheme to the SGMII and SerDes modes. For 100BASE-FX, PCS_FX100_CFG.SD_ENA enables signal detect, PCS_FX100_CFG.SD_POL controls the polarity, and PCS_FX100_CFG.SD_SEL selects the input source. The current status of the signal detect input can be observed through PCS_FX100_STATUS.SIGNAL_DETECT. For more information about signal detect, see [Section 2.1.3.3, Signal Detect](#).

Link Surveillance. The PCS synchronization status can be observed through PCS_FX100_STATUS.SYNC_STATUS. When synchronization is lost, the link breaks and PCS_FX100_STATUS.SYNC_LOST_STICKY is set. The PCS continuously tries to recover the link.

Unidirectional mode. 100BASE-FX has a similar unidirectional mode as SGMII and SerDes modes. PCS_FX100_CFG.UNIDIR_MODE_ENA enables unidirectional mode.

2.2 SERDES1G

SERDES1G is a high-speed SerDes interface that operates at 1 Gbps (SGMII/SerDes) and 100 Mbps (100BASE-FX). The 100BASE-FX mode is supported by oversampling. Port modules 0 through 7 and 10 connect to a SERDES1G.

The following table lists the registers associated with SERDES1G.

TABLE 2-10: SERDES1G REGISTERS

Registers	Description	Replication
SERDES1G_COMMON_CFG	Common configuration	Per SerDes
SERDES1G_DES_CFG	Deserializer configuration	Per SerDes
SERDES1G_IB_CFG	Input buffer configuration	Per SerDes
SERDES1G_SER_CFG	Serializer configuration	Per SerDes
SERDES1G_OB_CFG	Output buffer configuration	Per SerDes
SERDES1G_PLL_CFG	PLL configuration	Per SerDes
SERDES1G_MISC_CFG	Miscellaneous configuration	Per SerDes

For increased performance in specific application environments, SERDES1G supports the following:

- Programmable loop-bandwidth and phase regulation of deserializer
- Input buffer signal detect/loss of signal (LOS) options
- Input buffer with equalization
- Programmable output buffer features, including:
 - De-emphasis
 - Amplitude drive levels
 - Slew rate control
 - Idle mode
- Synchronous Ethernet support
- Loopbacks for system test

2.2.1 SERDES1G BASIC CONFIGURATION

The SERDES1G is enabled in SERDES1G_COMMON_CFG.ENA_LANE. By default, the SERDES1G is held reset and must be released from reset before the interface is active. This is done through SERDES1G_COMMON_CFG.SYS_RST and SERDES1G_MISC_CFG.LANE_RST.

2.2.1.1 SERDES1G Frequency Configuration

To operate the SERDES1G block at 1.25 GHz (corresponding to 1 Gbps data rate), the internal macro PLL must be configured as follows:

1. Configure SERDES1G_PLL_CFG.PLL_FSM_CTRL_DATA to 200.
2. Set SYS_RST = 0 (active) and PLL_FSM_ENA = 0 (inactive).
3. Set SYS_RST = 1 (deactive) and PLL_FSM_ENA = 1 (active).

2.2.2 SERDES1G LOOPBACK MODES

The SERDES1G interface supports two different loopback modes for testing and debugging data paths: equipment loopback and facility loopback.


Equipment loopback (SERDES1G_COMMON_CFG.ENA_ELOOP) Data is looped back from serializer output to deserializer input, and the receive clock is recovered. The equipment loopback includes all transmit and receive functions, except for the input and output buffers. The Tx data can still be observed on the output.

Facility loopback (SERDES1G_COMMON_CFG.ENA_FLOOP) The clock and parallel data output from deserializer are looped back to the serializer interface. Incoming serial data passes through the input buffer, the CDR, the deserializer, back to the serializer, and finally out through the output buffer.

Only one of the loopbacks can be enabled at the same time.

The following illustration shows the loopback paths.

FIGURE 2-1: SERDES1G LOOPBACK MODES



2.2.3 SYNCHRONOUS ETHERNET

The SERDES1G block can recover the clock from the received data and apply the clock to either a per-SerDes recovered clock output pin (SERDES1G_COMMON_CFG.RECO_SEL_A) or a common recovered clock output pin that may also be driven by another SerDes macros as well (SERDES1G_COMMON_CFG.RECO_SEL_B). Note that only one macro should drive the common recovered clock output pin at the same time.

In addition, it is possible to squelch the recovered clock if the associated PCS cannot detect valid data (SERDES1G_COMMON_CFG.SE_AUTO_SQUELCH_A_ENA and SERDES1G_COMMON_CFG.SE_AUTO_SQUELCH_B_ENA).

For more information about Synchronous Ethernet, see [Section 2.16, Layer-1 Timing](#).

2.2.4 SERDES1G DESERIALIZER CONFIGURATION

The SERDES1G block includes digital control logic that interacts with the analog modules within the block and compensates for the frequency offset between the received data and the internal high-speed reference clock. To gain high jitter performance, the phase regulation is a PI-type regulator, whose proportional (P) and integrative (I) characteristics can be independently configured. The integrative part of the phase regulation loop is configured in `SERDES1G_DES_CFG.DES_PHS_CTRL`. The limits of the integrator are programmable, allowing different settings for the integrative regulation while guaranteeing that the proportional part still is stronger than the integrative part. Integrative regulation compensates frequency modulation from DC up to cut-off frequency. Frequencies above the cut-off frequency are compensated by the proportional part. The time constant of the integrator is controlled independently of the proportional regulation by `SERDES1G_DES_CFG.DES_BW_HYST`. The `DES_BW_HYST` register field is programmable in a range from 3 to 7. The lower the configuration setting, the smaller the time-constant of the integrative regulation. For normal operation, configure `DES_BW_HYST` to 5.

The cut-off-frequency is calculated to:

$$f_{co} = 1/(2 \times \text{PI} \times 128 \times \text{PLL period} \times 32 \times 2^{(\text{DES\_BW\_HYST} + 1 - \text{DES\_BW\_ANA})})$$

$$\text{PLL period} = 1/(\text{data rate})$$

The integrative regulator can compensate a static frequency offset within the programmed limits down to a remaining frequency error of below 4 ppm. In steady state, the integrator toggles between two values around the exact value, and the proportional part of the phase regulation takes care of the remaining phase error.

The loop bandwidth for the proportional part of the phase regulation loop is controlled by configuring `SERDES1G_DES_CFG.DES_BW_ANA`.

The fastest loop bandwidth setting (lowest configuration value) results in a loop bandwidth that is equal to the maximum frequency offset compensation capability. For improved jitter performance, use a setting with sufficient margin to track the expected frequency offset rather than using the maximum frequency offset. For example, if a 100 ppm offset is expected, use a setting that is four times higher than the offset.

The following table provides the limits for the frequency offset compensation. The values are theoretical limits for input signals without jitter, because the actual frequency offset compensation capability is dependent on the toggle rate of the input data and the input jitter. Note that only applicable configuration values are listed.

TABLE 2-11: SERDES1G LOOP BANDWIDTH

DES_BW_ANA	Limits
4	1953 ppm
5	977 ppm
6	488 ppm
7	244 ppm

In the event of an 180° phase jump of the incoming data signal, the sampling stage of the deserializer may become stuck. To prevent this situation, the SERDES1G provides a 180° deadlock protection mechanism (`SERDES1G_DES_CFG.DES_MBTR_CTRL`). If this mechanism is enabled, a small frequency offset is applied to the phase regulation loop. The offset is sufficient to move the sampling point out of the 180° deadlock region, while at the same time, small enough to allow the regulation loop to compensate when the sample point is within the data eye.

2.2.5 SERDES1G SERIALIZER CONFIGURATION

The serializer provides the ability to align the phase of the internal clock and data to a selected source (`SERDES1G_SER_CFG.SER_ENALI`). The phase align logic is used when SERDES1G operates in the facility loopback mode.

2.2.6 SERDES1G INPUT BUFFER CONFIGURATION

The SERDES1G input buffer supports the following configurable options:

- 100BASE-FX mode support
- Signal detection, threshold configurable
- Configurable equalization including corner frequency configuration for the equalization filter
- DC voltage offset compensation
- Configurable common mode voltage (CMV) termination
- Selectable hysteresis, configurable hysteresis levels

When the SerDes interface operates in 100BASE-FX mode, the input buffer of the SERDES1G macro must also be configured for 100BASE-FX (SERDES1G_IB_CFG.IB_FX100_ENA).

The input buffer provides an option to configure the threshold level of the signal detect circuit to adapt to different input amplitudes. The signal detect circuit can be configured by SERDES1G_IB_CFG.IB_ENA_DETLEV and SERDES1G_IB_CFG.IB_DET_LEV.

The SERDES1G block offers options to compensate for channel loss. Degraded signals can be equalized, and the corner frequency of the equalization filter can be adapted to the channel behavior. The equalization settings are configured by SERDES1G_IB_CFG.IB_EQ_GAIN and SERDES1G_IB_CFG.IB_CORNER_FREQ.

The SERDES1G block can compensate for possible DC-offset, which can distort the received input signal, by enabling SERDES1G_IB_CFG.IB_ENA_OFFSET_COMP during normal reception.

The common-mode voltage (CMV) input termination can be set to either an internal reference voltage or to V_{DD_A} . To allow external DC-coupling of the input buffer to an output buffer, set the CMV input termination to the internal reference voltage, with internal DC-coupling disabled. SERDES1G_IB_CFG.IB_ENA_DC_COUPLING controls internal DC-coupling, and SERDES1G_IB_CFG.IB_ENA_CMV_TERM controls CMV input termination. The following modes are defined by CMV input termination and DC-coupling.

- SGMII compliant mode with external AC coupling (IB_ENA_DC_COUPLING = 0, IB_ENA_CMV_TERM = 1)
- Optional-mode with external DC-coupling to another Microchip output buffer, which can operate DC-coupled to the input buffer (IB_ENA_DC_COUPLING = 0, IB_ENA_CMV_TERM = 0)
- 100BASE-FX low frequency mode (IB_ENA_DC_COUPLING = 1, IB_ENA_CMV_TERM = 1)

The SERDES1G macro supports input hysteresis, which is required for some standards (SGMII). The hysteresis function is enabled by SERDES1G_IB_CFG.IB_ENA_HYST, and hysteresis levels are defined by SERDES1G_IB_CFG.IB_HYST_LEV.

Note Hysteresis and DC offset compensation cannot be enabled at the same time. For more information, see IB_ENA_OFFSET_COMP in the SERDES1G_IB_CFG register.

2.2.7 SERDES1G OUTPUT BUFFER CONFIGURATION

The SERDES1G output buffer supports the following configurable options.

- Configurable amplitude settings
- Configurable slew rate control
- 3 dB de-emphasis selectable
- Idle mode

The output amplitude of the output buffer is controlled by SERDES1G_OB_CFG.OB_AMP_CTRL. It can be adjusted in 50 mV steps from 0.4 V to 1.1 V peak-to-peak differential. The output amplitude also depends on the output buffer's supply voltage. For more information about dependencies between the maximum achievable output amplitude and the output buffer's supply voltage, refer to the electrical section for the dependencies between the maximum achievable output amplitude and the output buffer's supply voltage.

Adjust the slew rate adjustment using SERDES1G_OB_CFG.OB_SLP.

The output buffer supports a fixed 3 dB de-emphasis (SERDES1G_SER_CFG.SER_DEEMPH).

The output buffer supports an idle mode (SERDES1G_SER_CFG.SER_IDLE), which results in an differential peak-to-peak output amplitude of less than 30 mV.

2.2.8 SERDES1G CLOCK AND DATA RECOVERY (CDR) IN 100BASE-FX

To enable clock and data recovery when operating the SERDES1G in 100BASE-FX mode, set the following registers.

- SERDES1G_MISC_CFG.DES_100FX_CPMD_ENA = 1
- SERDES1G_IB_CFG.IB_FX100_ENA = 1
- SERDES1G_DES_CFG.DES_CPMD_SEL = 2

2.2.9 ENERGY EFFICIENT ETHERNET

The SERDES1G block supports Energy Efficient Ethernet as defined in IEEE 802.3az. To enable the low power modes, SERDES1G_MISC_CFG.TX_LPI_MODE_ENA and SERDES1G_MISC_CFG.RX_LPI_MODE_ENA must be set. At this point, the attached PCS takes full control over the high-speed output and input buffer activity.

2.2.10 SERDES1G DATA INVERSION

The data streams in the transmit and the receive direction can be inverted using SERDES1G_MISC_CFG.TX_DATA_INV_ENA and SERDES1G_MISC_CFG.RX_DATA_INV_ENA. Effectively this allows for swapping the P and N lines of the high-speed serial link.

2.3 SERDES6G

The SERDES6G is a high-speed SerDes interface that operates at 100 Mbps (100BASE-FX), 1 Gbps (SGMII/SerDes), and 2.5 Gbps (SGMII). The 100BASE-FX mode is supported by oversampling. Port modules 8 and 9 connect to a SERDES6G.

The following table lists the registers associated with SERDES6G.

TABLE 2-12: SERDES6G REGISTERS

Registers	Description	Replication
SERDES6G_COMMON_CFG	Common configuration	Per SerDes
SERDES6G_DES_CFG	Deserializer configuration	Per SerDes
SERDES6G_IB_CFG	Input buffer configuration	Per SerDes
SERDES6G_IB_CFG1	Input buffer configuration	Per SerDes
SERDES6G_SER_CFG	Serializer configuration	Per SerDes
SERDES6G_OB_CFG	Output buffer configuration	Per SerDes
SERDES6G_OB_CFG1	Output buffer configuration	Per SerDes
SERDES6G_PLL_CFG	PLL configuration	Per SerDes
SERDES6G_MISC_CFG	Miscellaneous configuration	Per SerDes

For increased performance in specific application environments, SERDES6G supports the following:

- Baud rate support, configurable from 1 Gbps to 2.5 G, for quarter and half rate modes
- Programmable loop bandwidth and phase regulation for the deserializer
- Configurable input buffer feature such as signal detect/loss of signal (LOS) options
- Configurable output buffer features such as programmable de-emphasis, amplitude drive levels and slew rate control
- Synchronous Ethernet support
- Loopbacks for system test

2.3.1 SERDES6G BASIC CONFIGURATION

The SERDES6G is enabled in SERDES6G_COMMON_CFG.ENA_LANE. By default, the SERDES6G is held reset and must be released from reset before the interface is active. This is done through SERDES6G_COMMON_CFG.SYS_RST and SERDES6G_MISC_CFG.LANE_RST.

2.3.1.1 SERDES6G PLL Frequency Configuration

To operate the SERDES6G block at the correct frequency, configure the internal macro as follows. PLL calibration is enabled through SERDES6G_PLL_CFG.PLL_FSM_ENA.

- Configure SERDES6G_PLL_CFG.PLL_FSM_CTRL_DATA in accordance with data rates listed in the following tables.
- Set SYS_RST = 0 (active) and PLL_FSM_ENA = 0 (inactive).
- Set SYS_RST = 1 (deactive) and PLL_FSM_ENA = 1 (active).

TABLE 2-13: PLL CONFIGURATION

Mode	SERDES6G_PLL_CFG.PLL_FSM_CTRL_DATA
SGMII/SerDes, 1 Gbps data	60
SGMII, 2.5 Gbps data	48

2.3.1.2 SERDES6G Frequency Configuration

The following table lists the range of data rates that are supported by the SERDES6G block.

TABLE 2-14: SERDES6G FREQUENCY CONFIGURATION REGISTERS

Configuration	SGMII/ SerDes 1 Gbps	SGMII 2.5 Gbps
SERDES6G_PLL_CFG.PLL_ROT_FRQ	0	1
SERDES6G_PLL_CFG.PLL_ROT_DIR	1	0
SERDES6G_PLL_CFG.PLL_ENA_ROT	0	1
SERDES6G_COMMON_CFG.QRATE	1	0
SERDES6G_COMMON_CFG.HRATE	0	1

2.3.2 SERDES6G LOOPBACK MODES

The SERDES6G interface supports two different loopback modes for testing and debugging data paths: equipment loopback and facility loopback.


Equipment loopback (SERDES6G_COMMON_CFG.ENA_ELOOP) Data is looped back from serializer output to the deserializer input, and the receive clock is recovered. The equipment loopback includes all transmit and receive functions, except for the input and output buffers. The Tx data can still be observed on the output.

Facility loopback (SERDES6G_COMMON_CFG.ENA_FLOOP) The clock and parallel data output from deserializer are looped back to the serializer interface. Incoming serial data passes through the input buffer, the CDR, the deserializer, back to the serializer, and finally out through the output buffer.

Only one of the loopbacks can be enabled at the same time.

The following illustration shows the loopback paths for the SERDES6G.

FIGURE 2-2: SERDES6G LOOPBACK MODES



2.3.3 SYNCHRONOUS ETHERNET

The SERDES6G block can recover the clock from the received data and apply the clock to either a per-SerDes recovered clock output pin (SERDES6G_COMMON_CFG.RECO_SEL_A) or a common recovered clock output pin that may be driven by another SerDes macros as well (SERDES6G_COMMON_CFG.RECO_SEL_B). Note that only one macro should drive the common recovered clock output pin at the same time.

In addition, it is possible to squelch the recovered clock if the associated PCS cannot detect valid data (SERDES6G_COMMON_CFG.SE_AUTO_SQUELCH_A_ENA and SERDES6G_COMMON_CFG.SE_AUTO_SQUELCH_B_ENA).

For more information about synchronous Ethernet, see [Section 2.16, Layer-1 Timing](#).

2.3.4 SERDES6G DESERIALIZER CONFIGURATION

The SERDES6G block includes digital control logic that interacts with the analog modules within the block and compensates for the frequency offset between the received data and the internal high-speed reference clock. To gain high jitter performance, the phase regulation is a PI-type regulator, whose proportional (P) and integrative (I) characteristics can be independently configured.

The integrative part of the phase regulation loop is configured in SERDES6G_DES_CFG.DES_PHS_CTRL. The limits of the integrator are programmable, allowing different settings for the integrative regulation while guaranteeing that the proportional part still is stronger than the integrative part. Integrative regulation compensates frequency modulation from DC up to cut-off frequency. Frequencies greater than the cut-off frequency are compensated by the proportional part.

The DES_BW_HYST register field controls the time constant of the integrator independently of the proportional regulator. The range of DES_BW_HYST is programmable as follows:

- Full rate mode = 3 to 7
- Quarter-rate mode = 1 to 7

The lower the configuration setting the smaller the time-constant of the integrative regulation. For normal operation, configure DES_BW_HYST to 5.

The cut-off-frequency is calculated to:

$$f_{co} = 1/(2 \times \pi \times 128 \times \text{PLL period} \times 32 \times 2^{(\text{DES\_BW\_HYST} + 1 - \text{DES\_BW\_ANA})})$$

$$\text{PLL period} = 1/(n \times \text{data rate})$$

where, $n = 1$ (full rate mode) or 4 (quarter-rate mode)

The integrative regulator can compensate a static frequency offset within the programmed limits down to a remaining frequency error of below 4 ppm. In steady state, the integrator toggles between two values around the exact value, and the proportional part of the phase regulation takes care of the remaining phase error. The loop bandwidth for the proportional part of the phase regulation loop is controlled by configuring SERDES6G_DES_CFG.SW.ANA.

The fastest loop bandwidth setting (lowest configuration value) results in a loop bandwidth that is equal to the maximum frequency offset compensation capability. For improved jitter performance, use a setting with sufficient margin to track the expected frequency offset rather than using the maximum frequency offset. For example, if a 100 ppm offset is expected, use a setting that is four times higher than the offset.

The following table provides the limits for the frequency offset compensation. The values are theoretical limits for input signals without jitter, because the actual frequency offset compensation capability is dependent on the toggle rate of the input data and the input jitter. Note that only applicable configuration values are listed. HRATE and QRATE are the configuration settings of SERDES6G_COMMON_CFG.HRATE and SERDES6G_COMMON_CFG.QRATE.

TABLE 2-15: SERDES6G LOOP BANDWIDTH

DES_BW_ANA	Limits when HRATE = 1 QRATE = 0	Limits when HRATE = 0 QRATE = 1
2		1953 ppm
3	1953 ppm	977 ppm
4	977 ppm	488 ppm
5	488 ppm	244 ppm
6	244 ppm	122 ppm
7	122 ppm	61 ppm

In the event of an 180° phase jump of the incoming data signal, the sampling stage of the deserializer may become stuck. To prevent this situation, the SERDES6G provides a 180° deadlock protection mechanism (SERDES6G_DES_CFG.DES_MBTR_CTRL). If this mechanism is enabled, a small frequency offset is applied to the phase regulation loop. The offset is sufficient to move the sampling point out of the 180° deadlock region, while at the same time, small enough to allow the regulation loop to compensate when the sample point is within the data eye.

2.3.5 SERDES6G SERIALIZER CONFIGURATION

The serializer provides the ability to align the phase of the internal clock and data to a selected source (SERDES6G_SER_CFG.SER_ENALI). The phase align logic is used when SERDES6G operates in the facility loopback mode.

2.3.6 SERDES6G INPUT BUFFER CONFIGURATION

The SERDES6G input buffer supports configurable options for:

- Automatic input voltage offset compensation
- Loss of signal detection

The input buffer is typically AC-coupled. As a result, the common-mode termination is switched off (SERDES6G_IB_CFG1.IB_CTERM_ENA). In order to support type-2 loads (DC-coupling at 1.0 V termination voltage) according to the OIF CEI specifications, common-mode termination must be enabled.

The sensitivity of the level detect circuit can be adapted to the input signal's characteristics (amplitude and noise). The threshold value for the level detect circuit is set in SERDES6G_IB_CFG.IB_VBCOM. The default value is suitable for normal operation.

When the SerDes interface operates in 100BASE-FX mode, the input buffer of the SERDES6G macro must also be configured for 100BASE-FX (SERDES6G_IB_CFG.IB_FX100_ENA).

During test or reception of low data rate signals (for example, 100BASE-FX), the DC-offset compensation must be disabled. For all other modes, the DC-offset compensation must be enabled for optimized performance. DC-offset compensation is controlled by SERDES6G_IB_CFG1.IB_ENA_OFFSAC and SERDES6G_IB_CFG1.IB_ENA_OFFSDC.

2.3.7 SERDES6G OUTPUT BUFFER CONFIGURATION

The SERDES6G output buffer supports the following configurable options.

- Amplitude control
- De-emphasis and output polarity inversion
- Slew rate control
- Skew adjustment
- Idle mode

The maximum output amplitude of the output buffer depends on the output buffer's supply voltage. For interface standards requiring higher output amplitudes (backplane application or interface to optical modules, for example), the output buffer can be supplied from a 1.2 V instead of a 1.0 V supply. By default, the output buffer is configured for 1.2 V mode, because enabling the 1.0 V mode when supplied from 1.2 V must be avoided. The supply mode is configured by SERDES6G_OB_CFG.OB_ENA1V_MODE.

The output buffer supports a four-tap pre-emphasis realized by one pre-cursor, the center tap, and two post cursors. The pre-cursor coefficient, C0, is configured by SERDES6G_SER_CFG.OB_PREC. C0 is a 5-bit value, with the most significant bit defining the polarity. The lower 4-bit value is hereby defined as B0. The first post-cursor coefficient, C2, is configured by SERDES6G_OB_CFG.OB_POST0. C2 is a 6-bit value, with the most significant bit defining the polarity. The lower 5-bit value is hereby defined as B2. The second post-cursor coefficient, C3, is configured by SERDES6G_SER_CFG.OB_POST1. C3 is 5-bit value, with the most significant bit defining the polarity. The lower 4-bit value is hereby defined as B3. The center-tap coefficient, C1, is a 6-bit value. Its polarity can be programmed by SERDES6G_OB_CFG.OB_POL, which is hereby defined as p1. For normal operation SERDES6G_OB_CFG.OB_POL must be set to 1. The value of the 6 bits forming C1 is calculated by the following equation.

$$\text{Equation 1: } C1: (64 - (B0 + B2 + B3)) \times p1$$

The output amplitude is programmed by SERDES6G_OB_CFG1.OB_LEV, which is a 6-bit value. This value is internally increased by 64 and defines the amplitude coefficient K. The range of K is therefore 64 to 127. The differential peak-peak output amplitude is given by $8.75 \text{ mV} \times K$. The maximum peak-peak output amplitude depends on the data stream and can be calculated to:

$$\text{Equation 2: } H(Z) = 4.375 \text{ mVpp} \times K \times (C0 \times z^1 + C1 \times z^0 + C2 \times z^{-1} + C3 \times z^{-2})/64$$

with z^n denoting the current bits of the data pattern defining the amplitude of Z. The output amplitude also depends on the output buffer's supply voltage. For more information about the dependencies between the maximum achievable output amplitude and the output buffer's supply voltage, see [Section 6.0, Electrical Specifications](#).

The configuration bits are summarized in the following table.

TABLE 2-16: DE-EMPHASIS AND AMPLITUDE CONFIGURATION

Configuration	Value	Description
OB_PREC	Signed 5-bit value	Pre-cursor setting C0. Range is –15 to 15.
OB_POST0	Signed 6-bit value	First post-cursor setting C2. Range is –31 to 31.
OB_POST1	Signed 5-bit value	Second post-cursor setting C3. Range is –15 to 15.
OB_LEV	Unsigned 6-bit value	Amplitude coefficient, $K = OB\_LEV + 64$. Range is 0 to 63.
OB_POL	0 1	Non-inverting mode. Inverting mode.

The output buffer provides the following additional options to configure its behavior.

- Idle mode: Enabling idle mode (SERDES6G_OB_CFG.OB_IDLE) results in a remaining voltage of less than 30 mV at the buffers differential outputs.
- Slew Rate: Slew rate can be controlled by two configuration settings. SERDES6G_OB_CFG.OB_SR_H provides coarse adjustments, and SERDES6G_OB_CFG.OB_SR provides fine adjustments.
- Skew control: In 1 Gbps SGMII mode, skew adjustment is controlled by SERDES6G_OB_CFG1.OB_ENA_CAS. Skew control is not applicable to other modes.

2.3.8 SERDES6G CLOCK AND DATA RECOVERY (CDR) IN 100BASE-FX

To enable clock and data recovery when operating SERDES6G in 100BASE-FX mode, set the following register fields:

- SERDES1G_MISC_CFG.DES_100FX_CPMD_ENA = 1
- SERDES1G_IBJ_CFG.IB_FX100_ENA = 1
- SERDES1G_DES_CFG.DES_CPMD_SEL = 2

2.3.9 ENERGY EFFICIENT ETHERNET

The SERDES6G block supports Energy Efficient Ethernet as defined in IEEE Draft P802.3az. To enable the low power modes, SERDES6G_MISC_CFG.TX_LPI_MODE_ENA and SERDES6G_MISC_CFG.RX_LPI_MODE_ENA must be set. At this point, the attached PCS takes full control over the high-speed output and input buffer activity.

2.3.10 SERDES6G DATA INVERSION

The data streams in the transmit and the receive direction can be inverted using SERDES6G_MISC_CFG.TX_DATA_INV_ENA and SERDES6G_MISC_CFG.RX_DATA_INV_ENA. This effectively allows for swapping the P and N lines of the high-speed serial link.

2.3.11 SERDES6G SIGNAL DETECTION ENHANCEMENTS

Signal detect information from the SERDES6G macro is normally directly passed to the attached PCS. It is possible to enable a hysteresis such that the signal detect condition must be active or inactive for a certain time before it is signaled to the attached PCS.

The signal detect assertion time (the time signal detect must be active before the information is passed to a PCS) is programmable in SERDES6G_DIG_CFG.SIGDET_AST. The signal detect de-assertion time (the time signal detect must be inactive before the information is passed to a PCS) is programmable in SERDES6G_DIG_CFG.SIGDET_DST.

2.3.12 HIGH-SPEED I/O CONFIGURATION BUS

The high-speed SerDes macros are configured using the high-speed I/O configuration bus (MCB), which is a serial bus connecting the configuration register set with all the SerDes macros. The HSIO::MCB_SERDES1G_ADDR_CFG register is used for SERDES1G macros and HSIO::MCB_SERDES6G_ADDR_CFG register is used for SERDES6G macros. The configuration busses are used for both writing to and reading from the macros.

The SERDES6G macros are programmed as follows:

- Program the configuration registers for the SERDES6G macro. For more information about configuration options,

see [Section 2.3, SERDES6G](#).

- Transfer the configuration from the configuration registers to one or more SerDes macros by writing the address of the macro (MCB_SERDES6G_ADDR_CFG.SERDES6G_ADDR) and initiating the write access (MCB_SERDES6G_ADDR_CFG.SERDES6G_WR_ONE_SHOT).
- The SerDes macro address is a mask with one bit per macro so that one or more macros can be programmed at the same time.
- The MCB_SERDES6G_ADDR_CFG.SERDES6G_WR_ONE_SHOT are automatically cleared when the writing is done.

The configuration and status information in the SERDES6G macros can be read as follows:

- Transfer the configuration and status from one or more SerDes macros to the configuration registers by writing the address of the macro (MCB_SERDES6G_ADDR_CFG.SERDES6G_ADDR) and initiating the read access (MCB_SERDES6G_ADDR_CFG.SERDES6G_RD_ONE_SHOT).
- The SerDes macro address is a mask with one bit per macro so that configuration and status information from one or more macros can be read at the same time. When reading from more than one macro, the results from each macro are OR'ed together.
- The MCB_SERDES6G_ADDR_CFG.SERDES6G_RD_ONE_SHOT are automatically cleared when the reading is done.

The SERDES1G macros are programmed similarly to the SERDES6G macros, except that MCB_SERDES1G_ADDR_CFG must be used for register access. For more information about configuration options, see [Section 2.2, SERDES1G](#).

2.4 Statistics

The following table lists the registers for the statistics module.

TABLE 2-17: COUNTER REGISTERS

Register	Description	Replication
SYS:STAT:CNT	Data register for reading port and service counters. The current view is specified in SYS::STAT_CFG.STAT_VIEW.	Per counter for specified port and service
SYS::STAT_CFG.STAT_VIEW	Sets the current counter view.	None
SYS::STAT_CFG.STAT_CLEAR_SHOT	Clears counters per counter group.	None
QSYS::STAT_CNT_CFG.TX_GREEN_CNT_MODE QSYS::STAT_CNT_CFG.TX_YELLOW_CNT_MODE	Controls whether to counts bytes or frames for Tx priority counters.	None
QSYS::STAT_CNT_CFG.DROP_GREEN_CNT_MODE QSYS::STAT_CNT_CFG.DROP_YELLOW_CNT_MODE	Controls whether to counts bytes or frames for drop priority counters.	None
ANA::AGENCTRL.GREEN_COUNT_MODE ANA::AGENCTRL.YELLOW_COUNT_MODE ANA::AGENCTRL.RED_COUNT_MODE	Controls whether to counts bytes or frames for Rx priority counters.	None
REW::STAT_CFG	Controls how to index the egress service statistics.	None

VSC7416-01 implements a statistics block containing the following counter groups:

- Receive statistics available per physical ingress port
- Transmit statistics available per physical egress port
- FIFO Drop statistics available per physical ingress port
- Service statistics available per ingress service point (ISDX). The ISDX is the outcome of the ingress service classification in VCAPs IS0 and IS1 and the path group protection.
- Service statistics available per egress service point (ESDX). The ESDX is the outcome of the egress service clas-

sification in VCAP ES0.

Each counter is 32 bits wide, which is large enough to ensure a wrap-around time longer than 13 seconds for port statistics on a 2.5G port and longer than 30 seconds for port statistics on a 1G port.

Each switch core port has 44 Rx counters, 18 FIFO drop counters, and 31 Tx counters. Each ISDX has 10 counters and each ESDX has 4 counters. The following lists the counters.

2.4.1 PORT STATISTICS

The following table defines the per-port available Rx counters and lists the counter's address in the common statistics block.

TABLE 2-18: RECEIVE COUNTERS IN THE STATISTICS BLOCK

Type	Short Name	Counter Address	Description
Rx	c_rx_oct	0x00	Received octets in good and bad frames.
Rx	c_rx_uc	0x01	Number of good broadcasts.
Rx	c_rx_mc	0x02	Number of good multicasts.
Rx	c_rx_bc	0x03	Number of good unicasts.
Rx	c_rx_short	0x04	Number of short frames with valid CRC (<64 bytes).
Rx	c_rx_frag	0x05	Number of short frames with invalid CRC (<64 bytes).
Rx	c_rx_jabber	0x06	Number of long frames with invalid CRC (according to MAXLEN.MAX_LENGTH).
Rx	c_rx_crc	0x07	Number of CRC errors, alignment errors and RX_ER events.
Rx	c_rx_symbol_err	0x08	Number of frames with RX_ER events.
Rx	c_rx_sz_64	0x09	Number of 64-byte frames in good and bad frames.
Rx	c_rx_sz_65_127	0x0A	Number of 65-127-byte frames in good and bad frames.
Rx	c_rx_sz_128_255	0x0B	Number of 128-255-byte frames in good and bad frames.
Rx	c_rx_sz_256_511	0x0C	Number of 256-511-byte frames in good and bad frames.
Rx	c_rx_sz_512_1023	0x0D	Number of 512-1023-byte frames in good and bad frames.
Rx	c_rx_sz_1024_1526	0x0E	Number of 1024-1526-byte frames in good and bad frames.
Rx	c_rx_sz_jumbo	0x0F	Number of 1527-MAXLEN.MAX_LENGTH-byte frames in good and bad frames. Counter is only applicable if MAXLEN.MAX_LENGTH > 1526.
Rx	c_rx_pause	0x10	Number of received pause frames.
Rx	c_rx_control	0x11	Number of MAC control frames received.
Rx	c_rx_long	0x12	Number of long frames with valid CRC (according to MAXLEN.MAX_LENGTH).
Rx	c_rx_cat_drop	0x13	Number of frames dropped due to classifier rules.
Rx	c_rx_red_prio_0	0x14	Number of received frames classified to QoS class 0 and discarded by a policer.
Rx	c_rx_red_prio_1	0x15	Number of received frames classified to QoS class 1 and discarded by a policer.
Rx	c_rx_red_prio_2	0x16	Number of received frames classified to QoS class 2 and discarded by a policer.
Rx	c_rx_red_prio_3	0x17	Number of received frames classified to QoS class 3 and discarded by a policer.

TABLE 2-18: RECEIVE COUNTERS IN THE STATISTICS BLOCK (CONTINUED)

Type	Short Name	Counter Address	Description
Rx	c_rx_red_prio_4	0x18	Number of received frames classified to QoS class 4 and discarded by a policer
Rx	c_rx_red_prio_5	0x19	Number of received frames classified to QoS class 5 and discarded by a policer.
Rx	c_rx_red_prio_6	0x1A	Number of received frames classified to QoS class 6 and discarded by a policer.
Rx	c_rx_red_prio_7	0x1B	Number of received frames classified to QoS class 7 and discarded by a policer.
Rx	c_rx_yellow_prio_0	0x1C	Number of received frames classified to QoS class 0 and marked yellowby a policer
Rx	c_rx_yellow_prio_1	0x1D	Number of received frames classified to QoS class 1 and marked yellowby a policer
Rx	c_rx_yellow_prio_2	0x1E	Number of received frames classified to QoS class 2 and marked yellowby a policer
Rx	c_rx_yellow_prio_3	0x1F	Number of received frames classified to QoS class 3 and marked yellowby a policer
Rx	c_rx_yellow_prio_4	0x20	Number of received frames classified to QoS class 4 and marked yellowby a policer
Rx	c_rx_yellow_prio_5	0x21	Number of received frames classified to QoS class 5 and marked yellowby a policer
Rx	c_rx_yellow_prio_6	0x22	Number of received frames classified to QoS class 6 and marked yellowby a policer
Rx	c_rx_yellow_prio_7	0x23	Number of received frames classified to QoS class 7 and marked yellowby a policer
Rx	c_rx_green_prio_0	0x24	Number of received frames classified to QoS class 0 and marked green by a policer.
Rx	c_rx_green_prio_1	0x25	Number of received frames classified to QoS class 1 and marked green by a policer.
Rx	c_rx_green_prio_2	0x26	Number of received frames classified to QoS class 2 and marked greenby a policer.
Rx	c_rx_green_prio_3	0x27	Number of received frames classified to QoS class 3 and marked green by a policer.
Rx	c_rx_green_prio_4	0x28	Number of received frames classified to QoS class 4 and marked green by a policer.
Rx	c_rx_green_prio_5	0x29	Number of received frames classified to QoS class 5 and marked green by a policer.
Rx	c_rx_green_prio_6	0x2A	Number of received frames classified to QoS class 6 and marked green by a policer.
Rx	c_rx_green_prio_7	0x2B	Number of received frames classified to QoS class 7 and marked green by a policer.

The following table defines the per-port available Tx counters and lists the counter address.

TABLE 2-19: TX COUNTERS IN THE STATISTICS BLOCK

Type	Short Name	Counter Address	Description
Tx	c_tx_oct	0x40	Transmitted octets in good and bad frames.
Tx	c_tx_uc	0x41	Number of good unicasts.
Tx	c_tx_mc	0x42	Number of good multicasts.
Tx	c_tx_bc	0x43	Number of good broadcasts.

TABLE 2-19: TX COUNTERS IN THE STATISTICS BLOCK (CONTINUED)

Type	Short Name	Counter Address	Description
Tx	c_tx_col	0x44	Number of transmitted frames experiencing a collision. An excessive collided frame gives 16 counts.
Tx	c_txdrop	0x45	Number of frames dropped due to excessive collisions or late collisions.
Tx	c_txpause	0x46	Number of transmitted pause frames.
Tx	c_tx_sz_64	0x47	Number of 64-byte frames in good and bad frames.
Tx	c_tx_sz_65_127	0x48	Number of 65-127-byte frames in good and bad frames.
Tx	c_tx_sz_128_255	0x49	Number of 128-255-byte frames in good and bad frames.
Tx	c_tx_sz_256_511	0x4A	Number of 256-511-byte frames in good and bad frames.
Tx	c_tx_sz_512_1023	0x4B	Number of 512-1023-byte frames in good and bad frames.
Tx	c_tx_sz_1024_1526	0x4C	Number of 1024-1526-byte frames in good and bad frames.
Tx	c_tx_sz_jumbo	0x4D	Number of 1527-MAXLEN.MAX_LENGTH-byte frames in good and bad frames.
Tx	c_tx_yellow_prio_0	0x4E	Number of transmitted frames classified to QoS class 0 with DP level 1.
Tx	c_tx_yellow_prio_1	0x4F	Number of transmitted frames classified to QoS class 1 with DP level 1.
Tx	c_tx_yellow_prio_2	0x50	Number of transmitted frames classified to QoS class 2 with DP level 1.
Tx	c_tx_yellow_prio_3	0x51	Number of transmitted frames classified to QoS class 3 with DP level 1.
Tx	c_tx_yellow_prio_4	0x52	Number of transmitted frames classified to QoS class 4 with DP level 1.
Tx	c_tx_yellow_prio_5	0x53	Number of transmitted frames classified to QoS class 5 with DP level 1.
Tx	c_tx_yellow_prio_6	0x54	Number of transmitted frames classified to QoS class 6 with DP level 1.
Tx	c_tx_yellow_prio_7	0x55	Number of transmitted frames classified to QoS class 7 with DP level 1.
Tx	c_tx_green_prio_0	0x56	Number of transmitted frames classified to QoS class 0 with DP level 0.
Tx	c_tx_green_prio_1	0x57	Number of transmitted frames classified to QoS class 1 with DP level 0.
Tx	c_tx_green_prio_2	0x58	Number of transmitted frames classified to QoS class 2 with DP level 0.
Tx	c_tx_green_prio_3	0x59	Number of transmitted frames classified to QoS class 3 with DP level 0.
Tx	c_tx_green_prio_4	0x5A	Number of transmitted frames classified to QoS class 4 with DP level 0.
Tx	c_tx_green_prio_5	0x5B	Number of transmitted frames classified to QoS class 5 with DP level 0.
Tx	c_tx_green_prio_6	0x5C	Number of transmitted frames classified to QoS class 6 with DP level 0.

TABLE 2-19: TX COUNTERS IN THE STATISTICS BLOCK (CONTINUED)

Type	Short Name	Counter Address	Description
Tx	c_tx_green_prio_7	0x5D	Number of transmitted frames classified to QoS class 7 with DP level 0.
Tx	c_tx_aged	0x5E	Number of frames dropped due to frame aging.

The following table defines the per-port available FIFO drop counters and lists the counter address.

TABLE 2-20: FIFO DROP COUNTERS IN THE STATISTICS BLOCK

Type	Short Name	Counter Address	Description
Drop	c_dr_local	0x80	Number of frames discarded due to no destinations.
Drop	c_dr_tail	0x81	Number of frames discarded due to no more memory in the queue system (tail drop).
Drop	c_dr_yellow_prio_0	0x82	Number of FIFO discarded frames classified to QoS class 0 with DP level 1
Drop	c_dr_yellow_prio_1	0x83	Number of FIFO discarded frames classified to QoS class 1 with DP level 1
Drop	c_dr_yellow_prio_2	0x84	Number of FIFO discarded frames classified to QoS class 2 with DP level 1
Drop	c_dr_yellow_prio_3	0x85	Number of FIFO discarded frames classified to QoS class 3 with DP level 1
Drop	c_dr_yellow_prio_4	0x86	Number of FIFO discarded frames classified to QoS class 4 with DP level 1
Drop	c_dr_yellow_prio_5	0x87	Number of FIFO discarded frames classified to QoS class 5 with DP level 1
Drop	c_dr_yellow_prio_6	0x88	Number of FIFO discarded frames classified to QoS class 6 with DP level 1
Drop	c_dr_yellow_prio_7	0x89	Number of FIFO discarded frames classified to QoS class 7 with DP level 1
Drop	c_dr_green_prio_0	0x8A	Number of FIFO discarded frames classified to QoS class 0 with DP level 0.
Drop	c_dr_green_prio_1	0x8B	Number of FIFO discarded frames classified to QoS class 1 with DP level 0.
Drop	c_dr_green_prio_2	0x8C	Number of FIFO discarded frames classified to QoS class 2 with DP level 0.
Drop	c_dr_green_prio_3	0x8D	Number of FIFO discarded frames classified to QoS class 3 with DP level 0.
Drop	c_dr_green_prio_4	0x8E	Number of FIFO discarded frames classified to QoS class 4 with DP level 0.
Drop	c_dr_green_prio_5	0x8F	Number of FIFO discarded frames classified to QoS class 5 with DP level 0.
Drop	c_dr_green_prio_6	0x90	Number of FIFO discarded frames classified to QoS class 6 with DP level 0.
Drop	c_dr_green_prio_7	0x91	Number of FIFO discarded frames classified to QoS class 7 with DP level 0.

2.4.2 SERVICE STATISTICS

Service statistics work on frames per color. The color is given as the DP level after the frames are policed by the service policer, port policer and VCAP policer:

- Green frames: Frames with DP level 0 after the combined policer action.
- Yellow frames: Frames with DP level 1 after the combined policer action.

- Red frames: Frames discarded by the combined policer action.

The following table defines the per-ISDX available ingress service counters and lists the counter address.

TABLE 2-21: INGRESS SERVICE COUNTERS IN THE STATISTICS BLOCK

Type	Short Name	Counter Address	Description
ISDX	c_is_green_oct	0xC0	Number of octets in green frames classified to ISDX that are forwarded by the device to one or more destination ports.
ISDX	c_is_green_pkt	0xC1	Number of green frames classified to ISDX that are forwarded by the device to one or more destination ports.
ISDX	c_is_yellow_oct	0xC2	Number of octets in yellow frames classified to ISDX that are forwarded by the device to one or more destination ports.
ISDX	c_is_yellow_pkt	0xC3	Number of yellow frames classified to ISDX that are forwarded by the device to one or more destination ports.
ISDX	c_is_red_oct	0xC4	Number of octets in received frames classified to ISDX that are discarded by a policer.
ISDX	c_is_red_pkt	0xC5	Number of received frames classified to ISDX that are discarded by a policer.
ISDX	c_is_drop_green_oct	0xC6	Number of octets in green frames classified to ISDX that are not discarded by a policer but discarded by the queue system or the analyzer.
ISDX	c_is_drop_green_pkt	0xC7	Number of green frames classified to ISDX that are not discarded by a policer but discarded by the queue system or the analyzer.
ISDX	c_is_drop_yellow_oct	0xC8	Number of octets in yellow frames classified to ISDX that are not discarded by a policer but discarded by the queue system or the analyzer.
ISDX	c_is_drop_yellow_pkt	0xC9	Number of yellow frames classified to ISDX that are not discarded by a policer but discarded by the queue system or the analyzer.

Only one set of the service counters (frames and octets) is triggered per frame. For instance, if the c_is_green_pkt and c_is_green_oct counters are incrementing, it implies that the frame is forwarded by the device and that it is not discarded by any of the policers, the queue system, or the analyzer.

The service drop counters are triggered if a frame is discarded by the device after the policers. The following blocks can discard a frame:

- The queue system due to congestion.
- The analyzer due to the forwarding decision (destination port set is empty).

The following table defines the per-ESDX available egress service counters and lists the counter address.

TABLE 2-22: EGRESS SERVICE COUNTERS IN THE STATISTICS BLOCK

Type	Short Name	Counter Address	Description
ESDX	c_es_green_oct	0x100	Number of octets in transmitted frames classified to ESDX with DP level 0.
ESDX	c_es_green_pkt	0x101	Number of transmitted frames classified to ESDX with DP level 0.
ESDX	c_es_yellow_oct	0x102	Number of octets in transmitted frames classified to ESDX with DP level 1.
ESDX	c_es_yellow_pkt	0x103	Number of transmitted frames classified to ESDX with DP level 1.

By default, the egress service counters are indexed using the ESDX from the ES0 action and in case, there is no ES0 hit, the frame's ISDX is used. This behavior can be changed (REW::STAT_CFG) with the following options on how to index the egress service counters:

- Use the ESDX from the ES0 action. If there is no ES0 hit, use the frame's ISDX. This is the default behavior.
- Use the ESDX from the ES0 action. If there is no ES0 hit, do not count the frame in the egress service statistics.
- Always use the frame's ISDX even when there is an ES0 hit. This enables per-VLAN counters when the ISDX ingress classification at the same time assigns a unique ISDX per VID.
- Use eight times the frame's source port plus the frame's QoS class. This enables per-ingress port QoS counters.

2.4.3 ACCESSING AND CLEARING COUNTERS

The counters are accessed through SYS:STAT:CNT using the counter address given for each counter in the tables in the preceding sections. Only a subset of the counters for all ports and services are addressable at the same time. The current subset is programmed in STAT_CFG.STAT_VIEW and includes port counters (Rx, Tx, drop), ingress service counters, and egress service counters for the number programmed in STAT_CFG.STAT_VIEW. This implies that STAT_VIEW is interpreted as a port number, an ISDX value, and an ESDX value at the same time.

The following illustration shows the layout of the counter subset that are addressable through SYS:STAT:CNT. To change the view, write a new port number/service index to STAT_CFG.STAT_VIEW.

FIGURE 2-3: COUNTER LAYOUT (SYS:STAT:CNT) PER VIEW

Counter Groups		
	Bit 31	Bit 0
0x000 ⋮ 0x02A	Rx, Port=STAT_CFG.STAT_VIEW	c_rx_oct ⋮ c_rx_green_prio_7
	Unused	
0x040 ⋮ 0x05E	TX, Port=STAT_CFG.STAT_VIEW	c_tx_oct ⋮ c_tx_aged
	Unused	
0x080 ⋮ 0x091	Drop, Port=STAT_CFG.STAT_VIEW	c_dr_local ⋮ c_dr_green_prio_7
	Unused	
0x0C0 ⋮ 0x0C9	Ingress Service, ISDX=STAT_CFG.STAT_VIEW	c_is_green_oct ⋮ c_is_drop_yellow_pkt
	Unused	
0x100 ⋮ 0x103	Egress Service, ESDX=STAT_CFG.STAT_VIEW	c_es_green_oct ⋮ c_es_yellow_pkt

Read Example To read the number of good unicast frames transmitted on port 3 (counter c_tx_uc), set the current view in STAT_CFG.STAT_VIEW to 3 and read SYS:STAT:CNT[0x41]. Note that with the current view, one could also read Rx and drop counters on port 3, as well as ingress service counters for ISDX=3 and egress service counters for ESDX=3.

The counters can be cleared per counter group per view. Writing to register STAT_CFG.STAT_CLEAR_SHOT clears all counters associated with the counter groups specified by STAT_CFG.STAT_CLEAR_SHOT for the view specified in STAT_CFG.STAT_VIEW.

Clear Example To clear all ingress service counters in ISDX=8 and all egress service counters in ESDX=8, set the current view in STAT_CFG.STAT_VIEW to 8 and write 0x18 to STAT_CFG.STAT_CLEAR_SHOT. This clears all the associated service counters but not the port counters on port 8.

It is possible to select whether to count frames or bytes for the following specific counters.

- The Rx priority counters (c_rx_red_prio_*, c_rx_yellow_prio_*, c_rx_green_prio_*, where x is 0 through 7).
- The Tx priority counters (c_tx_yellow_prio_*, c_tx_green_prio_*, where x is 0 through 7).
- The Drop priority counters (c_dr_yellow_prio_*, c_dr_green_prio_*, where x is 0 through 7).

The Rx priority counters are programmed through ANA::AGENCTRL, and the Tx and drop priority counters are programmed through QSYS::STAT_CNT_CFG. When counting bytes, the frame length excluding inter frame gap and preamble is counted.

For testing purposes, all counters are both readable and writable. All counters wrap around to 0 when reaching the maximum.

For more information about how the counters map to relevant MIBs, see [Section 4.2.4, Port Counters](#).

2.5 MPLS-TP Classifier

The following table lists the registers associated with the MPLS-TP classifier.

TABLE 2-23: MPLS-TP CLASSIFIER REGISTERS

Register	Description	Replication
SYS::PORT_MODE.MPLS_ENA	Enables MPLS-TP classifier on port.	Per port
ANA::FRAME_EXTRACT	Configures frame extraction for MPLS label stack exceptions.	Per exception (7)
ANA::RSVD_LABELS_ENA	Configures exception handling of reserved labels.	Per reserved label (16)
ANA::EXCP_VPROFILE_CONF	Configures virtual port profile for MPLS exceptions.	None
ANA::TC_MAP_TBL	Configures mapping of MPLS TC values to internal QoS class and drop precedence level. Eight different mapping tables are available.	8
ANA::OAM_CONF	Configuration of the CPU extraction queue for OAM frames. Configuration of the VCCV2 filter.	None
ANA::SEGMENT_OAM	Configuration of the segment OAM filter.	Per port

The MPLS-TP classifier is enabled per port in (SYS::PORT_MODE.MPLS_ENA). It then receives all incoming MPLS-encapsulated Ethernet frames. A frame is treated as MPLS encapsulated if the frame's EtherType is 0x8847 or 0x8848. MPLS-encapsulated frames can include zero or one VLAN tag in front of the MPLS EtherType. Non-MPLS frames or MPLS-encapsulated frames on ports where the MPLS-TP classifier is not enabled, are not forwarded to the MPLS-TP classifier but are instead forwarded directly to the basic classification.

The MPLS-TP classifier has the option to determine a number of frame properties affecting the forwarding of the frame through the switch. The total set of properties are listed in the following table.

TABLE 2-24: MPLS-TP CLASSIFICATION PROPERTIES

Property	Description
ISDX	Ingress Service index. The ISDX is input to the IS1 lookup and used in service-based forwarding, statistics, and rewrite operations. The MPLS classifier can only use the first 512 ISDX value.
Virtual port profile	The virtual port profile selects the configuration used in the basic classifier (registers in the ANA:PORT group). By default, the basic classifier uses the physical ingress port number; however, the MPLS-TP classifier has the option to point to any of the 64 available profiles.
VID	Port-based VID. This overrides the configuration of the port-based VID (ANA::VLAN_CFG.VLAN_VID). The basic classifier can change the VID if VLAN aware. In addition, the IS1 lookups may change the VID.
PCP	Port-based PCP. This overrides the configuration of the port-based PCP (ANA::VLAN_CFG.VLAN_PCP). The basic classifier can change the PCP if VLAN aware. In addition, the IS1 lookups may change the PCP.
DEI	Port-based DEI. This overrides the configuration of the port-based DEI (ANA::VLAN_CFG.VLAN_DEI). The basic classifier can change the DEI if VLAN aware. In addition, the IS1 lookups can change the DEI.


TABLE 2-24: MPLS-TP CLASSIFICATION PROPERTIES (CONTINUED)

Property	Description
Tag type	Port-based tag type. This overrides the configuration of port-based tag type (ANA::VLAN_CFG.VLAN_TAG_TYPE). The basic classifier can change the tag type if VLAN aware.
B_PM	B-domain port mask. This mask is unconditionally applied to the forwarding of the frame in the MPLS forwarding step in the analyzer. Ports set in this mask cannot be removed again.
CPU_Q	CPU extraction to use if the CPU port is set in B_PM. This overrides any outcome from the basic classifier.
Classified TTL	Classified time-to-live value from MPLS label stack. The TTL is forwarded with the frame to the rewriter where it is used in the egress MPLS encapsulation of the frame. This cannot be changed in the basic classifier or in IS1.
Classified TC	Classified TC value from MPLS label stack. The TC is forwarded with the frame to the rewriter where it is used in the egress MPLS encapsulation of the frame. This cannot be changed in the basic classifier or in IS1.
Classified pop count	Classified pop count from MPLS label stack. The pop count is forwarded with the frame to the rewriter where it is used in the egress MPLS encapsulation of the frame. This cannot be changed in the basic classifier or in IS1.
Classified S-bit	Classified S-bit from MPLS label stack. The S-bit is forwarded with the frame to the rewriter where it is used in the egress MPLS encapsulation of the frame. This cannot be changed in the basic classifier or in IS1.
Drop precedence	Default drop precedence level on the virtual port. This overrides the configuration of the default DP level in ANA:VPORT:QOS_CFG.DP_DEFAULT_VAL. The basic classifier can change the drop precedence level if QoS classification is enabled. In addition, IS1 can change the DP level.
QoS class	Default QoS class on the virtual port. This overrides the configuration of the default QoS class in ANA:VPORT:QOS_CFG.QOS_DEFAULT_VAL. The basic classifier can change the QoS class if QoS classification is enabled. In addition, IS1 can change the QoS class.
CW present	Control word present after label with S-bit. This instructs the basic classifier to skip a control word when extracting data from the MPLS frame for further classification.

Note Not all frames are classified to the full set of MPLS-TP classification properties. For more information, see [Section 2.5.3, MPLS-TP Classification Result](#).

The following illustration shows an overview of the MPLS-TP classifier. It consists primarily of the two VCAP IS0 look-ups, the MPLS Link Layer (MLL) lookup and the MPLS Label Stack (MLBS) lookup, followed by a processing to classify to the final results.

FIGURE 2-4: MPLS-TP CLASSIFIER OVERVIEW



2.5.1 MPLS LINK LAYER (MLL) LOOKUP

The MLL lookup is the first lookup in the MPLS-TP classifier and it uses the following key to match against the MLL entries in IS0.

TABLE 2-25: IS0 MLL KEY

Field Name	Bit	Width	Description
IGR_PORT	0	4	The port number where the frame was received (0-11).
TTYTYPE	4	2	MPLS link layer VLAN tag type. 0: Untagged. 1: C-tagged. 2: S-tagged. 3: Reserved.
BVID	6	12	Frame's MPLS link layer VID if available, see TTYTYPE.
M_DMACH	18	48	Frame's MPLS link layer destination MAC address.
M_SMACH	66	48	Frame's MPLS link layer source MAC address.
E_TYPE	114	2	MPLS link layer EtherType type. 0: Reserved. 1: Downstream assigned label (0x8847). 2: Upstream assigned label (0x8848). 3: Reserved.

The MLL lookup can either fail or result in a hit. Failing means that no entries were matched. In this case, the MPLS-TP classifier terminates the processing of the MPLS frames and forwards the frame to the basic classifier. The basic classifier continues the classification process of the frame from the start of the frame as if the MPLS-TP classifier had been disabled.

If the MLL lookup matches an entry in IS0, it returns an MLL action vector with the following fields.

TABLE 2-26: IS0 MLL ACTION FIELDS

Action Field	Bit	Width	Description
LL_IDX	0	7	MPLS Link Layer index. Used as input to the MLBS key. Only applicable if MPLS_FWD = 1.
MPLS_FWD	7	1	Forward based on MPLS label stack. This enables the MLBS lookup.
B_PM	8	12	B-domain port mask.
CPU_Q	20	3	CPU extraction queue when frame is forwarded to CPU (bit 11 set in B_PM). Only applicable if MPLS_FWD = 0.
ISDX	23	8	MPLS ISDX value.
RESERVED	31	2	Must be set to 0.
VPROFILE_IDX	33	6	Virtual port profile index, which points out the configuration set used in the basic classifier and analyzer. See Section 2.6.1, Virtual Port Profiles . Only applicable if MPLS_FWD = 0.
SCONFIG_ENA	39	1	For VLAN-tagged frames, use service specific configuration instead of virtual port profile configuration. Only applicable if MPLS_FWD = 0.
CL_VID	40	12	For untagged frames or if SCONFIG_ENA = 0, use this value instead of ANA:VPORT:VLAN_CFG.VLAN_VID. Only applicable if MPLS_FWD = 0.
QOS_DEFAULT_VAL	52	3	For VLAN-tagged frames, use this value instead of ANA:VPORT:QOS_CFG.QOS_DEFAULT_VAL. Only applicable if SCONFIG_ENA = 1 and MPLS_FWD = 0.

TABLE 2-26: ISO MLL ACTION FIELDS (CONTINUED)

Action Field	Bit	Width	Description
DP_DEFAULT_VAL	55	1	For VLAN-tagged frames, use this value instead of ANA:VPORT:QOS_CFG.DP_DEFAULT_VAL. Only applicable if SCONFIG_ENA = 1 and MPLS_FWD = 0.
OAM_ISDX	56	8	ISDX for segment OAM frames. Only applicable if MPLS_FWD = 1.
RESERVED	64	2	Must be set to 0.
ISDX_ADD_REPL	66	1	If set, OAM_ISDX replaces the ISDX for segment OAM frames. Otherwise, OAM_ISDX is added to ISDX for segment OAM frames. Only applicable if MPLS_FWD = 1
HIT_STICKY	67	1	If set, a frame has matched against the associated entry.

The key action from the MLL match is whether MPLS-based forwarding is enabled or not (MLL_action.MPLS_FWD).

2.5.1.1 MPLS Forwarding Disabled (MPLS_FWD = 0)

If MLL_action.MPLS_FWD is cleared, the MPLS-TP classifier terminates the processing of the MPLS frame in the MPLS-TP classifier. The basic classifier continues the processing of the frame from the start of the frame. The following information is returned to the basic classifier from the MPLS-TP classifier.

- Use MLL_action.VPROFILE_IDX as the virtual port profile index. The virtual port profile is used when looking up the configuration in the basic classifier.
- Use MLL_action.ISDX as the frame's classified ISDX.
- Use MLL_action.B_PM as the B-domain forwarding mask.
- Use MLL_action.CPU_Q as the CPU extraction queue if B_PM has the CPU port set.

In addition, the MPLS-TP classifier can overwrite specific configurations in the basic classifier. For VLAN-tagged frames where MLL action SCONFIG_ENA is set, the following classifier configuration is overwritten:

- Use the frame's MPLS link layer VID as the port-based VID (ANA:VPORT:VLAN_CFG.VLAN_VID)
- Use the frame's MPLS link layer PCP as the port-based PCP (ANA:VPORT:VLAN_CFG.VLAN_PCP)
- Use the frame's MPLS link layer DEI as the port-based DEI (ANA:VPORT:VLAN_CFG.VLAN_DEI)
- Use the frame's tag type (C-tag or S-tag) as the port-based tag type (ANA:VPORT:VLAN_CFG.VLAN_TAG_TYPE)
- Use MLL_action.QOS_DEFAULT_VAL as the default QoS class on the port (ANA:VPORT:QOS_CFG.QOS_DEFAULT_VAL)
- Use MLL_action.DP_DEFAULT_VAL as the default drop precedence level on the port (ANA:VPORT:QOS_CFG.DP_DEFAULT_VAL)

For untagged tagged frames or frames for which MLL_action.SCONFIG_ENA is cleared, the following classifier configuration is overwritten:

- Use the frame's MPLS link layer VID as the port-based VID (ANA:VPORT:VLAN_CFG.VLAN_VID).

2.5.1.2 MPLS Forwarding Enabled (MPLS_FWD = 1)

If MLL_action.MPLS_FWD is set, the MPLS-TP classifier continues the processing of the MPLS encapsulation. The following information from the MLL action vector is used:

- MLL_IDX is used as part of the MLBS key. Its purpose is to tie the two lookups together.
- B_PM is always OR'ed with the B_PM action from the MLBS lookup.
- ISDX, OAM_ISDX and ISDX_ADD_REPL are used to determine the ISDX for segment OAM frames, see [Section 2.5.3.4, Segment Reserved Label Frame Classification](#).

Other fields from the MLL action vector are disregarded.

2.5.2 MPLS LABEL STACK (MLBS) LOOKUP

When MPLS forwarding is enabled by the MLL lookup (MPLS_FWD = 1), the MPLS-TP classifier classifies the frame based on the MPLS label stack. Prior to the MLBS lookup, the frame's MPLS label stack is preprocessed.

2.5.2.1 MPLS Label Stack Preprocessing

The MPLS label stack is preprocessed by traversing the label stack, starting at the top of the stack (outer most label) and constructing the MLBS ISO key. The normal flow is that the preprocessing constructs the key with up to three labels in the key but stopping when a label with the S-bit set is encountered. The label with the S-bit set is included in the MLBS key. A maximum of three labels are included in the key.

If a reserved label (label value < 16) is encountered during the first three labels then the following exceptions apply to the normal flow:

- Specific handling of VCCV2 labels can be enabled in OAM_CONF.VCCV2_ENA. The VCCV2 label is defined by OAM_CONF.VCCV2_LABEL. If enabled and the reserved label is a VCCV2 label then the next label following the VCCV2 label in the MPLS label stack is used in the key. The VCCV2 label itself is not part of the key.
- If specific VCCV2 label handling is not enabled or if the reserved label is not a VCCV2 label, then the processing of the label stack stops, and no more labels are added to the key. The reserved label itself is not part of the key. If the reserved label is the outer-most label in the MPLS label stack, the MLBS lookup is cancelled due to lack of labels in the key.

When the MLBS lookup is cancelled due the outer-most label being a reserved label, which is not a VCCV2 reserved label, the frame is copied to the CPU, either as a segment OAM frame or a segment reserved label frame. The following table lists the conditions for classifying a frame as a segment OAM frame or a segment reserved label frame. Note that all conditions must apply.

TABLE 2-27: MPLS SEGMENT FRAME TYPES

Frame Type	Conditions
Segment OAM	First label in label stack is a reserved label, Label not handled as a VCCV2 label, Segment OAM enabled on port (SEGMENT_OAM.SEGMENT_OAM_ENA), Reserved label matches SEGMENT_OAM.SEGMENT_OAM_LBL, Optional check: Reserved label must be bottom of stack (SEGMENT_OAM.SEGMENT_OAM_BOS), An associated channel header (ACH) must be found after the bottom of stack with the first nibble being 1.
Segment reserved label	First label in label stack is a reserved label, Label not handled as a VCCV2 label, Segment OAM disabled or reserved label does not match SEGMENT_OAM.SEGMENT_OAM_LBL, Handling of the reserved label enabled in ANA::RSVD_LBL_ENA.

For more information about the classification results for these frames, see [Section 2.5.3, MPLS-TP Classification Result](#).

2.5.2.2 MPLS Label Stack Lookup

The ISO MLBS key is listed in the following table.

TABLE 2-28: ISO MLBS KEY

Field Name	Bit	Width	Description
LL_IDX	0	7	Link Layer index from the action of the MLL lookup.
LBL_0_VAL	7	20	MPLS label 0 value (outer-most label).
LBL_0_TC	27	3	MPLS label 0 TC bits (outer-most label).
LBL_1_VAL	30	20	MPLS label 1 value.
LBL_1_TC	50	3	MPLS label 1 TC bits.
LBL_2_VAL	53	20	MPLS label 2 value.
LBL_2_TC	73	3	MPLS label 2 TC bits.

All three labels in the key are not necessarily assigned when, for instance, the bottom of the stack is encountered as the first or second label. Unassigned labels in the key can take any value and should be don't cared in the MLBS entries in ISO.

The MLBS lookup can either fail or result in a hit. A fail means that no entries were matched. In this case, the frame is classified as an MLBS-miss frame and is classified accordingly. For more information, see [Section 2.5.3.6, MLBS Miss Frame Classification](#).

If the MLBS lookup matches an entry in IS0, it returns an MLBS action vector with the fields listed in the following table.

TABLE 2-29: IS0 MLBS ACTION FIELDS

Action Field	Bit	Width	Description
ISDX	0	8	MPLS ISDX value.
RESERVED	8	2	Must be set to 0.
B_PM	10	12	B-domain port mask.
CPU_Q	22	3	CPU extraction queue when frame is forwarded to CPU (bit 11 set in B_PM).
OAM_ENA	25	3	This channel is enabled for OAM. 0: No OAM 1: VCCV1 OAM. 2: VCCV2 OAM. 3: VCCV3 OAM. 4: GAL MEP LSP OAM. 5: GAL MIP LSP OAM. 6-7: Reserved.
RESERVED	28	1	Must be set to 0.
RSVD_LBL_VAL	29	4	Use when OAM_ENA = 4 or 5. Configures the RESERVED LABEL used as OAM label for this MEP or MIP. Only the lower 4 bit of the label value is configured.
RSVD_LBL_BOS	33	1	Use with OAM_ENA = 4. Configures if the RESERVED LABEL used as OAM label for this MEP must be bottom of stack.
CW_ENA	34	1	If set, the first nibble following the MPLS label stack bottom must be set to 0. Otherwise error "Expected CW" is generated. If set, the basic classifier skips the control word when extracting data from the frame for further classification.
TC_LABEL	35	2	Selects which label provides the frame's classified TC.
TTL_LABEL	37	2	Selects which label provides the frame's classified TTL.
SWAP_LABEL	39	2	Indicates which label in the label stack is used for forwarding. LER entry: must point to PW label. LSR entry: must point to swap label. GAL MEP OAM entry: Must point to GAL label.
TERMINATE_PW	41	1	If set, this is a terminating PW and the frame is forwarded to the C-domain.
POP_CNT	42	3	Indicates the number of bytes to pop in the rewriter. 0: The frame does not contain a Customer Ethernet layer. This can be used for IP encapsulated PTP frames. 1: Frame is MPLS encapsulated. Do not pop any bytes. 2: Frame is MPLS encapsulated. Pop 14 bytes. 3: Frame is MPLS encapsulated. Pop 18 bytes. 4: Frame is MPLS encapsulated. Pop 22 bytes. 5: Frame is MPLS encapsulated. Pop 26 bytes. 6: Frame is MPLS encapsulated. Pop 30 bytes. 7: Frame is MPLS encapsulated. Pop 34 bytes.
E_LSP	45	1	0: LSP is L-LSP. 1: LSP is E-LSP.

TABLE 2-29: IS0 MLBS ACTION FIELDS (CONTINUED)

Action Field	Bit	Width	Description
TC_MAP	46	3	Selects which TC mapping table to use when mapping TC to drop precedence and QoS class.
L_LSP_COS	49	3	For L-LSP (E_LSP = 0), use L_LSB_QOS as QoS class.
INC_ISDX	52	1	If set, the frame's classified TC value is added to the ISDX value for E-LSP (E_LSP = 1).
OAM_ISDX	53	8	ISDX for OAM frames.
RESERVED	61	2	Must be set to 0.
ISDX_ADD_REPL	63	1	If set, OAM_ISDX replaces the ISDX for OAM frames. Otherwise, OAM_ISDX is added to ISDX for OAM frames.
SWAP_IS_BOS	64	1	If set, the TTL of the swap label is set to 0 and the frame's classified S-bit is set to 0.
VPROFILE_IDX	65	6	Virtual port profile index, which points out the configuration set used in the basic classifier and analyzer.
SCONFIG_ENA	71	1	Use service specific configuration instead of virtual port profile configuration for the following. VLAN_DEI, VLAN_PCP, VLAN_TAG_TYPE. QOS_DEFAULT_VAL, DP_DEFAULT_VAL.
CL_VID	72	12	Use this value instead of ANA:VPORT:VLAN_CFG.VLAN_VID.
VLAN_TAG_TYPE	84	1	Use this value instead of ANA:VPORT:VLAN_CFG.VLAN_TAG_TYPE. Only applicable if SCONFIG_ENA = 1.
VLAN_DEI	85	1	Use this value instead of ANA:VPORT:VLAN_CFG.VLAN_DEI. Only applicable if SCONFIG_ENA = 1.
VLAN_PCP	86	3	Use this value instead of ANA:VPORT:VLAN_CFG.VLAN_PCP. Only applicable if SCONFIG_ENA = 1.
HIT_STICKY	88	1	If set, a frame has matched against the associated entry.

The action vector from the MLBS lookup is post-processed before the final result of the MPLS-TP classification is returned to the basic classifier. Based on frame properties and the MLBS action vector, each frame is classified to one of the frame types listed in the following table.

TABLE 2-30: MLBS FRAME TYPES

Frame Type	Conditions
MLBS miss	Lookup in MLBS fails.
VCCV2 RAL OAM	MLBS_action.OAM_ENA = VCCV2 OAM, Reserved label in frame at level specified by MLBS_action.SWAP_LABEL, Reserved label is a VCCV2 label
GAL MEP LSP OAM	MLBS_action.OAM_ENA = GAL MEP LSP OAM, Reserved label in frame at level specified by MLBS_action.SWAP_LABEL, Reserved label matches MLBS_action.RSVD_LBL_VAL, MLBS_action.RSVD_LBL_BOS = 0 or reserved label is at the bottom of the stack.
Reserved label	Reserved label in frame at level specified by MLBS_action.SWAP_LABEL, Handling of the reserved label enabled in ANA::RSVD_LBL_ENA.
VCCV1 OAM	MLBS_action.TERMINATE_PW = 1, MLBS_action.OAM_ENA = VCCV1 OAM, Associated channel header (ACH) is found after the bottom of stack with the first nibble = 1.

TABLE 2-30: MLBS FRAME TYPES (CONTINUED)

Frame Type	Conditions
VCCV3 OAM	MLBS_action.TERMINATE_PW = 1, MLBS_action.OAM_ENA = VCCV3 OAM, TTL of the swap label = 1.
GAL MIP LSP OAM	MLBS_action.OAM_ENA = GAL MIP LSP OAM, Reserved label matches MLBS_action.RSVD_LBL_VAL, TTL of label pointed to by MLBS_action.TTL_LABEL = 1, Reserved label is at the bottom of the stack, Number of labels in stack = MLBS_action.SWAP_LABEL + 1.
Data	All other frames.

Note that all conditions must apply. By default, a frame is a data frame unless otherwise classified or an error is encountered. For more information about specific error conditions, see [Section 2.5.4, MLBS Error Frame Handling](#).

2.5.3 MPLS-TP CLASSIFICATION RESULT

For each frame type defined in [Table 2-27](#) (segment OAM and segment reserved label frames) and in [Table 2-30](#) (other frame types), the following sections list the specific results of the MPLS-TP classification.

2.5.3.1 Data Frame Classification

The following table lists the classification results for data frames.

TABLE 2-31: DATA FRAME CLASSIFICATION RESULTS

Parameter	Settings
ISDX	MLBS_action.ISDX. When MLBS_action.E_LSP = 1 and MLBS_action.INC_ISDX = 1, the classified TC value is added to the ISDX.
Virtual port profile	MLBS_action.VPROFILE_IDX.
VID	Use MLBS_action.CL_VID as the port-based VID (ANA:VPORT:VLAN_CFG.VLAN_VID).
PCP	When MLBS_action.SCONFIG_ENA = 1, use MLBS_action.VLAN_PCP as the port-based PCP (ANA:VPORT:VLAN_CFG.VLAN_PCP).
DEI	When MLBS_action.SCONFIG_ENA = 1, use MLBS_action.VLAN_DEI as the port-based DEI (ANA:VPORT:VLAN_CFG.VLAN_PCP).
Tag type	When MLBS_action.SCONFIG_ENA = 1, use MLBS_action.VLAN_TAG_TYPE as the port-based tag type (ANA:VPORT:VLAN_CFG.VLAN_TAG_TYPE).
B_PM	MLL_action.B_PM OR MLBS_action.B_PM.
CPU_Q	Use MLBS_action.CPU_Q as CPU extraction queue if MLBS_action.B_PM has the CPU port set, otherwise use MLL_action.CPU_Q if MLL_action.B_PM has the CPU port set.
Classified TTL	Use TTL of label pointed to by MLBS_action.TTL_LABEL decremented with 1. The TTL is forwarded with the frame to the rewriter where it is used in the egress MPLS encapsulation of the frame.
Classified TC	Use TC of label pointed to by MLBS_action.TC_LABEL decremented with 1. The TC is forwarded with the frame to the rewriter where it is used in the egress MPLS encapsulation of the frame.
Classified pop count	MLBS_action.POP_CNT. The pop count is forwarded with the frame to the rewriter where it is used in the egress MPLS encapsulation of the frame.
Classified S-bit	Use S-bit of label pointed to by MLBS_action.SWAP_LABEL. When MLBS_action.SWAP_IS_BOS = 1, classified S-bit is set to 0. The S-bit is forwarded with the frame to the rewriter where it is used in the egress MPLS encapsulation of the frame.

TABLE 2-31: DATA FRAME CLASSIFICATION RESULTS (CONTINUED)

Parameter	Settings
Drop precedence	The default drop precedence level on the port given by ANA:VPORT:QOS_CFG.DP_DEFAULT_VAL is overruled. The MLBS_action.TC_MAP selects which set of mapping tables to use for mapping the classified TC to the default drop precedence level. The mapping tables are given by ANA::TC_MAP_TBL.
QoS class	When MLBS_action.E_LSP = 1, the default QoS class on the port given by ANA:VPORT:QOS_CFG.QOS_DEFAULT_VAL is overruled. The MLBS_action.TC_MAP selects which set of mapping tables to use for mapping the classified TC to the default QoS class. The mapping tables are given by ANA::TC_MAP_TBL. When MLBS_action.E_LSP = 0, use MLBS_action.L_LSP_COS.
CW present	MLBS_action.CW_ENA. This instructs the basic classifier to skip a 4-byte control word after the bottom of stack.

2.5.3.2 Segment OAM Frame Classification

Handling of segment OAM frames is different from the handling of normal OAM frames because there is no MLBS lookup. The following table lists the classification result for segment OAM frames.

TABLE 2-32: SEGMENT OAM FRAME CLASSIFICATION RESULTS

Parameter	Settings
ISDX	When MLL_action.ISDX_ADD_REPL = 1, use MLL_action.OAM_ISDX. Otherwise, use MLL_action.OAM_ISDX + MLL_action.ISDX.
Virtual port profile	When ANA::EXCP_VPROFILE_CONF.OAM_SRC is set, use the physical ingress port number. Otherwise, use ANA::EXCP_VPROFILE_CONF.OAM_VPROFILE.
B_PM	MLL_action.B_PM. Force setting of CPU port.
CPU_Q	ANA::SEGMENT_OAM.SEGMENT_OAM_QUEUE.

2.5.3.3 OAM Frame Classification

The following table lists the classification results for OAM frames including the frame types VCCV1 OAM, VCCV2 RAL OAM, VCCV3 OAM, GAL MEP LSP OAM, and GAL MIP LSP OAM.

TABLE 2-33: OAM FRAME CLASSIFICATION RESULTS

Parameter	Settings
ISDX	When MLBS_action.ISDX_ADD_REPL = 1, use MLBS_action.OAM_ISDX. Otherwise, use MLBS_action.OAM_ISDX + MLBS_action.ISDX.
Virtual port profile	When ANA::EXCP_VPROFILE_CONF.OAM_SRC is set, use MLBS_action.VPROFILE_IDX. Otherwise, use ANA::EXCP_VPROFILE_CONF.OAM_VPROFILE.
B_PM	MLL_action.B_PM. Force setting of CPU port.
CPU_Q	ANA::OAM_CONF.CPU_OAM_QUEUE.

2.5.3.4 Segment Reserved Label Frame Classification

Handling of segment reserved label frames is different from the handling of normal reserved label frames because there is no MLBS lookup. The following table lists the classification result for segment reserved label frames.

TABLE 2-34: SEGMENT RESERVED LABEL FRAME CLASSIFICATION RESULTS

Parameter	Settings
ISDX	MLL_action.ISDX.

TABLE 2-34: SEGMENT RESERVED LABEL FRAME CLASSIFICATION RESULTS (CONTINUED)

Parameter	Settings
Virtual port profile	When ANA::EXCP_VPROFILE_CONF.RSVD_LBL_SRC is set, use MLBS_action.VPROFILE_IDX. Otherwise, use ANA::EXCP_VPROFILE_CONF.RSVD_LBL_VPROFILE.
B_PM	MLL_action.B_PM. Force setting of CPU port.
CPU_Q	Use ANA::RSVD_LABELS_ENA.RSVD_CPU_QUEUE for specific value of reserved label.

2.5.3.5 Reserved Label Frame Classification

The following table lists the classification results for reserved label frames.

TABLE 2-35: RESERVED LABEL FRAME CLASSIFICATION RESULTS

Parameter	Settings
ISDX	MLBS_action.ISDX.
Virtual port profile	When ANA::EXCP_VPROFILE_CONF.RSVD_LBL_SRC is set, use the physical ingress port number. Otherwise, use ANA::EXCP_VPROFILE_CONF.RSVD_LBL_VPROFILE.
B_PM	MLL_action.B_PM. Force setting of CPU port.
CPU_Q	Use ANA::RSVD_LABELS_ENA.RSVD_CPU_QUEUE for specific value of reserved label.

2.5.3.6 MLBS Miss Frame Classification

The following table lists the classification results for frames that fail to match any entries in the MLBS lookup.

TABLE 2-36: MLBS MISS FRAME CLASSIFICATION RESULTS

Parameter	Settings
ISDX	0
Virtual port profile	When ANA::EXCP_VPROFILE_CONF.MLBS_MISS_SRC is set, use MLL_action.VPROFILE_IDX. Otherwise, use MLBS_action.VPROFILE:IDX.
B_PM	MLL_action.B_PM. If ANA::FRAME_EXTRACT.EXTRACT_ENA[6] is set, force setting of CPU port.
CPU_Q	If ANA::FRAME_EXTRACT.EXTRACT_ENA[6] is set, use FRAME_EXTRACT.EXTRACT_QUEUE[6].

2.5.4 MLBS ERROR FRAME HANDLING

During the MPLS-TP analysis, different error conditions may occur. The following table lists the errors that are handled specifically in the MPLS-TP classifier.

TABLE 2-37: MLBS ERROR TYPES

MLBS Error Type	Conditions
TTL Expiry	During GAL MIP LSP OAM frame type classification: MLBS_action.OAM_ENA = GAL MIP LSP OAM, Reserved label matches MLBS_action.RSVD_LBL_VAL, TTL of label pointed to by MLBS_action.TTL_LABEL = 1, Reserved label is not at the bottom of the stack or number of labels in stack is not equal to MLBS_action.SWAP_LABEL + 1. During TTL check for data frames: TTL of label pointed to by MLBS_action.TTL_LABEL is 0 or 1. During TTL check for OAM frames (excluding GAL MIP LSP OAM frames): If ANA::OAM_CONF.OAM_TTL_CHECK is set, TTL of label pointed to by MLBS_action.TTL_LABEL is 0 or 1.
Expected CW	During control word check for data frames terminating a pseudo-wire: MLBS_action.TERMINATE_PW = 1, MLBS_action.CW_ENA = 1, First nibble in control word is different from 0
Expected ACh	During OAM ACh check for OAM frames (excluding VCCV1 OAM frames): An associated channel header (ACH) is not found after the bottom of stack with the first nibble being 1.
Illegal segment OAM	During segment OAM frame type classification: The optional check that reserved label must be bottom of stack fails. During segment reserved label frame type classification: Handling of the reserved label is disabled in ANA::RSVD_LBL_ENA.
Unexpected S-bit clear	During S-bit check for data frames terminating a pseudo-wire: MLBS_action.TERMINATE_PW = 1, S-bit of label pointed to by MLBS_action.SWAP_LABEL is 0.
Unkown reserved label	During classification to OAM frame types: Reserved label at level specified by MLBS_action.SWAP_LABEL, Not a VCCV2 OAM frame, Not a GAL MEP LSP OAM frame, Handling of the reserved label is disabled in ANA::RSVD_LBL_ENA.

The following table lists the classification results for frames that are classified as errors during the MPLS-TP processing.

TABLE 2-38: MPLS ERROR FRAME CLASSIFICATION RESULTS

Parameter	Settings
ISDX	0
Virtual port profile	When ANA::EXCP_VPROFILE_CONF.MLBS_ERR_SRC is set, use MLL_action.VPROFILE_IDX. Otherwise, use ANA::EXCP_VPROFILE_CONF.MLBS_ERR_VPROFILE.
B_PM	Cleared with forced setting of CPU port.
CPU_Q	FRAME_EXTRACT.EXTRACT_QUEUE, per error condition.

2.6 Basic Classifier

The switch core includes a common basic classifier, which determines a number of properties affecting the forwarding of each frame through the switch. These properties are:

- Frame acceptance filtering. Drop illegal frame types.
- QoS classification. Assign one of eight QoS classes to the frame.

- Drop precedence (DP) classification. Assign one of two drop precedence levels to the frame.
- DSCP classification. Assign one of 64 DSCP values to the frame.
- VLAN classification. Extract tag information from the frame or use the port VLAN.
- Link aggregation code generation. Generate the link aggregation code.
- CPU forwarding determination. Determine CPU Forwarding and CPU extraction queue number

The outcome of the classifier is the basic classification result, which can be overruled by more intelligent frame processing in the VCAP IS1. For more information, see [Section 2.7, VCAP](#).

2.6.1 VIRTUAL PORT PROFILES

The classifier takes as input frame data as well as a virtual port profile (VPORT), which points out a configuration setup to use. VSC7416-01 supports 64 virtual port profiles. The configuration setup consists of the register settings in ANA:PORT and ANA:COMMON. The registers in ANA:PORT are replicated per virtual port profile and the virtual port profile thereby offers a better control of the classifier as well as VCAP and analyzer.

By default, the VPORT is equal to the ingress physical port number. However, the MPLS-TP classifier can point out a different virtual port profile based on the MPLS IS0 actions. The MPLS-TP classifier can for instance use this to select different classifier configuration setups per pseudo-wire on the same physical ingress port.

2.6.2 GENERAL DATA EXTRACTION SETUP

This section provides information about the overall settings for data extraction that provides the data input to the classifier, VCAP, analyzer, and rewriter.

The following table lists the registers associated with general data extraction.

TABLE 2-39: GENERAL DATA EXTRACTION REGISTERS

Register	Description	Replication
SYS::PORT_MODE.L3_PARSE_CFG	Enables the use of Layer 3 and 4 protocol information for classification and frame processing.	Per port
SYS::VLAN_ETYPE_CFG	Ethernet Type for S-tags in addition to default value 0x88A8.	None
ANA:PORT.VLAN_CFG.VLAN_INNER_TAG_ENA	Enables using inner VLAN tag for basic classification if available in incoming frame.	Per virtual port profile
ANA:PORT:S1_VLAN_INNER_TAG_ENA	Enables using inner VLAN tag for IS1 key generation if available in incoming frame.	Per virtual port profile per IS1 lookup

For frames that are classified by the MPLS-TP classifier, data is extracted according to the following rules:

- If the MLL lookup failed (no entries were matched), data is extracted from the start of the frame if the MPLS-TP classifier had been disabled.
- If the MLL lookup hit an entry with MPLS-TP forwarding disabled (MLL_action.MPLS_FWD cleared), data is extracted from the start of the frame if the MPLS-TP classifier had been disabled.
- If the MLL lookup hit an entry with MPLS-TP forwarding enabled (MLL_action.MPLS_FWD set), the start of the frame is set to the first byte after the label with the S-bit set. A maximum of 3 labels can be skipped. If the MPLS-TP classifier signals that a control word is present (MLBS_action.CW_ENA), the first nibble of the control word is checked:

If the first nibble is 0, then the control word is also skipped. Data used for further classification in the basic classifier, IS1, IS2, and the analyzer, is extracted from this new starting point expecting a normal Ethernet frame starting with a destination MAC address. This is the case for LER operation when a pseudo-wire has been terminated. For LSR operation further classification using the basic classifier, IS1, and IS2 is not required.

If the first nibble of the control word is 4 or 6, it is really not a control word but rather an IP/MPLS frame where 4 is the version field of an IPv4 header and 6 is the version field of an IPv6 header. In this case, data is extracted with a dummy MAC header in front of the IP header (DMAC=0, SMAC=0, EtherType=0). The frame is treated as an IPv4 or an IPv6 frame through the basic classifier, IS1, and IS2. The purpose of this is to enable IEEE1588 actions on IP/MPLS frames.

In the device, it is programmable which VLAN tags are recognized. The use of Layer-3 and Layer-4 information for classification and forwarding can also be controlled.

The device recognizes three different VLAN tags:

- Customer tags (C-TAGs), which use TPID 0x8100.
- Service tags (S-TAGs), which use TPID 0x88A8 (IEEE 802.1ad).
- Service tags (S-TAGs), which use a custom TPID programmed in SYS::VLAN_ETYPE_CFG.

The device can parse and use information from up to two VLAN tags of any of the kinds described above.

By default, the outer VLAN tag is extracted and used for both the basic classification and the VCAP IS1 key generation. However, for both the basic classification and the VCAP IS1, there is an option to use the inner VLAN tag instead for frames with at least two VLAN tags. For basic classification, this is controlled in VLAN_CFG.VLAN_INNER_TAG_ENA and affects both QoS, DP, and VLAN classification as well as the frame acceptance filter. For IS1, this is controlled per lookup in S1_VLAN_INNER_TAG_ENA. Note that several keys in IS1 always contain both the inner VLAN tag and the outer VLAN tag.

Various blocks in the device uses Layer-3 and Layer-4 information for classification and forwarding. Layer-3 and Layer-4 information can be extracted from a frame with up to two VLAN tags. Frames with more than two VLAN tags are considered non-IP frames.

The actual use of Layer-3 and Layer-4 information for classification, forwarding, and rewriting is enabled in SYS::PORT_MODE.L3_PARSE_CFG. The following blocks are affected by this functionality:

- Basic classification: QoS, DP, and DSCP classification, link aggregation code generation, CPU forwarding
- VCAP: TCAM keys (IS1, IS2) using Layer 3 and Layer4 information
- Analyzer: Flooding and forwarding of IP multicast frames
- Rewriter: Rewriting of IP information

2.6.3 FRAME ACCEPTANCE FILTERING

The following table lists the registers associated with frame acceptance filtering.

TABLE 2-40: FRAME ACCEPTANCE FILTERING REGISTERS

Register	Description	Replication
DEV::PORT_MISC	Configures forwarding of special frames.	Per port
ANA:PORT:DROP_CFG	Configures discarding of illegal frame types.	Per virtual port profile

Based on the configurations in the DROP_CFG and PORT_MISC registers, the classifier instructs the queue system to drop or forward certain frames types, such as:


- Frames with a multicast source MAC address
- Frames with a null source or null destination MAC address (address = 0x000000000000)
- Frames with errors signaled by the MAC (for example, an FCS error)
- MAC control frames
- Pause frames after flow control processing in the MAC.
- Untagged frames (excluding frames with reserved destination MAC addresses from the BPDU, GARP, and Link trace/CCM address ranges).
- Priority S-tagged frames
- Priority C-tagged frames
- VLAN S-tagged frames
- VLAN C-tagged frames

By default, MAC control frames, pause frames, and frames with errors are dropped by the classifier.

The VLAN acceptance filter decides whether a frame's VLAN tagging is allowed on the port. By default, the outer VLAN tag is used as input to the filter, however, there is an option to use the inner VLAN tag instead for double tagged frames (VLAN_CFG.VLAN_INNER_TAG_ENA).

The following illustration shows the flowchart for the VLAN acceptance filter.

FIGURE 2-5: VLAN ACCEPTANCE FILTER



If the frame is accepted by the VLAN acceptance filter, it can still be discarded in other places of the switch, such as the following:

- Policers, due to traffic exceeding a peak information rate.
- IS2 Security TCAM, due to permit/deny rules.
- Analyzer, due to forwarding decisions such as VLAN ingress filtering.
- Queue system, due to lack of resources, frame aging, or excessive collisions.

2.6.4 QOS, DP, AND DSCP CLASSIFICATION

This section provides information about the functions in the QoS, DP, and DSCP classification. The three tasks are described as one, because the tasks have a significant amount of functionality in common.

The following table lists the registers associated with QoS, DP, and DSCP classification.

TABLE 2-41: QOS, DP, AND DSCP CLASSIFICATION REGISTERS

Register	Description	Replication
ANA.PORT.QOS_CFG	Configuration of the overall classification flow for QoS, DP, and DSCP.	Per virtual port profile
ANA:PORT:QOS_P-CP_DEI_MAP_CFG	Mapping from (DEI, PCP) to (DP, QoS).	Per virtual port profile per DEI per PCP
ANA::DSCP_CFG	DSCP configuration per DSCP value.	Per DSCP
ANA::DSCP_REWR_CFG	DSCP rewrite values per DP level and QoS class.	Per DP and per QoS

The basic classification provides the user with control of the QoS, DP, and DSCP classification algorithm. The result of the basic classification are the following frame properties, which follow the frame through the switch:


- The frame's QoS class. This class is encoded in a 3-bit field, where 7 is the highest priority QoS class and 0 is the lowest priority QoS class. The QoS class is used by the queue system when enqueueing frames and when evaluating resource consumptions, for policing, statistics, and rewriter actions.
- The frame's DP level. This level is encoded in a 1-bit field, where frames with DP = 1 have the highest probability of being dropped and frames with DP = 0 have the lowest probability. The DP level is used by the dual leaky bucket policers for measuring committed and peak information rates, for restricting memory consumptions in the queue system, for collecting statistics, and for rewriting priority information in the rewriter. The DP level is incremented by the policers if a frame is exceeding a programmed committed information rate.
- The frame's DSCP. This value is encoded in a 6-bit field. The DSCP value is forwarded with the frame to the rewriter where it is translated and rewritten into the frame. The DSCP value is only applicable to IPv4 and IPv6 frames.

The classifier looks for the following fields in the incoming frame to determine the QoS, DP, and DSCP classification:

- Port default QoS class and DP level. The default DSCP value is the frame's DSCP value. For non-IP frames, the DSCP is 0 and it is not used elsewhere in the switch.
- Priority Code Point (PCP) when the frame is VLAN tagged or priority tagged. There is an option to use the inner tag for double tagged frames (VLAN_CFG.VLAN_INNER_TAG_ENA). Both S-tagged and C-tagged frames are considered.
- Drop Eligible Indicator (DEI) when the frame is VLAN tagged or priority tagged. There is an option to use the inner tag for double tagged frames (VLAN_CFG.VLAN_INNER_TAG_ENA). Both S-tagged and C-tagged frames are considered.
- DSCP (all 6 bits, both for IPv4 and IPv6 packets). The classifier can look for the DSCP value behind up to two VLAN tags.


The following illustration shows the flow chart of basic QoS and DP classification.

FIGURE 2-6: QOS AND DP BASIC CLASSIFICATION FLOW CHART



The following illustration shows the flow chart for basic DSCP classification.

FIGURE 2-7: BASIC DSCP CLASSIFICATION FLOW CHART



The translation part of the DSCP classification is common for both QoS, DP, and DSCP classification.

The basic classified QoS, DP, and DSCP can be overwritten by more intelligent decisions made in the VCAP IS1.

2.6.5 VLAN CLASSIFICATION

The following table lists the registers associated with VLAN classification.

TABLE 2-42: VLAN CONFIGURATION REGISTERS

Register	Description	Replication
ANA:PORT:VLAN_CFG	Configures the virtual port's processing of VLAN information in VLAN-tagged and priority-tagged frames. Configures the port-based VLAN.	Per virtual port profile

The VLAN classification determines a tag header for all frames. The tag header includes the following information:

- Priority Code Point (PCP)
- Drop Eligible Indicator (DEI)
- VLAN Identifier (VID)
- Tag Protocol Identifier (TPID) type (TAG_TYPE). This field informs whether tag used for classification was a C-tag or an S-tag.

The tag header determined by the classifier is carried with the frame through the switch and is used in various places such as the analyzer for forwarding and the rewriter for egress tagging operations.

The device recognizes three kinds of tags based on the TPID, which is the EtherType in front of the tag:

- Customer tags (C-TAGs), which use TPID 0x8100.
- Service tags (S-TAGs), which use TPID 0x88A8 (IEEE 802.1ad).
- Service tags (S-TAGs), which use a custom TPID programmed in SYS::VLAN_ETYPE_CFG.


For customer tags and service tags, both VLAN tags (tags with nonzero VID) and priority tags (tags with VID = 0) are processed.

The tag header is either retrieved from a tag in the incoming frame or from a default port-based tag header. The port-based tag header is configured in ANA:PORT:VLAN_CFG.

For double tagged frames, there is an option to use the inner tag instead of the outer tag (VLAN_CFG.VLAN_INNER_TAG_ENA).

In addition to the tag header, the virtual port decides the number of VLAN tags to pop at egress (VLAN_POP_CNT). If the configured number of tags to pop is greater than the actual number of tags in the frame, the number is reduced to the number of actual tags in the frame.

The following illustration shows the flow chart for basic VLAN classification.

FIGURE 2-8: BASIC VLAN CLASSIFICATION FLOW CHART

The basic classifier can overwrite the basic classified PCP value with the frame's basic classified QoS class coming from the basic QoS classification. This option is enabled in ANA:PORT:PTP_DLY1_CFG.USE_QOS_AS_PCP_ENA. The basic classified tag header can be overwritten by more intelligent decisions made in the VCAP IS1.

2.6.6 LINK AGGREGATION CODE GENERATION

This section provides information about the functions in link aggregation code generation.

The following table lists the registers associated with aggregation code generation.

TABLE 2-43: AGGREGATION CODE GENERATION REGISTERS

Register	Description	Replication
ANA::AGGR_CFG	Configures use of Layer-2 through Layer-4 flow information for link aggregation code generation.	Common

The classifier generates a link aggregation code, which is used in the analyzer when selecting to which port in a link aggregation group a frame is forwarded.

The following contributions to the link aggregation code is configured in the AGGR_CFG register:

- Destination MAC address—use the lower 12 bits of the DMAC.
- Source MAC address—use the lower 12 bits of the SMAC.
- IPv6 flow label—use the 20 bits of the flow label.
- IPv4 source and destination IP addresses—use the lower 8 bits of the SIP and DIP.
- TCP/UDP source and destination port for IPv4 and IPv6 frames—use the lower 8 bits of the SPORT and DPORT.
- ISDX—use the 8 bits of the ISDX. The ISDX used is the resulting ISDX after applying the actions from VCAP IS1 and path protection.
- Random aggregation code—use a pseudo-random number instead of the frame information.

Each of the enabled contributions are XOR'ed together, yielding a 4-bit aggregation code ranging from 0 to 15. For more information about how the aggregation code is used, see [Section 4.3.8, Link Aggregation](#).

2.6.7 CPU FORWARDING DETERMINATION

The following table lists the registers associated with CPU forwarding in the basic classifier.

TABLE 2-44: CPU FORWARDING DETERMINATION

Register	Description	Replication
CPU_FWD_CFG	Enables CPU forwarding for various frame types.	Per virtual port profile
CPU_FWD_BPDU_CFG	Enables CPU forwarding per BPDU address.	Per virtual port profile
CPU_FWD_GARP_CFG	Enables CPU forwarding per GARP address.	Per virtual port profile
CPU_FWD_CCM_CFG	Enables CPU forwarding per CCM/Link trace address	Per virtual port profile
CPUQ_CFG and CPUQ_CFG2	CPU extraction queues for various frame types	None
CPUQ_8021_CFG	CPU extraction queues for BPDU, GARP, and CCM addresses.	None
VRAP_CFG	VLAN configuration of VRAP filter.	None
VRAP_HDR_DATA	Data match against VRAP header.	None
VRAP_HDR_MASK	Mask used to don't care bits in the VRAP header.	None

The classifier has support for determining whether certain frames must be forwarded to the CPU extraction queues. Other parts of the device can also determine CPU forwarding, for example, the MPLS-TP classifier, the analyzer or the VCAP IS2. All events leading to CPU forwarding are OR'ed together, and the final CPU extraction queue mask, which is available to the user, contains the sum of all events leading to CPU extraction. For more information, see [Section 4.7, CPU Extraction and Injection](#).

Upon CPU forwarding by the classifier, the frame type determines whether the frame is redirected or copied to the CPU. Any frame type or event causing a redirection to the CPU cause all front ports to be removed from the forwarding decision - only the CPU receives the frame. When copying a frame to the CPU, the normal forwarding of the frame is unaffected.

The following table lists the standard frame types, with respect to CPU forwarding, that are recognized by the classifier.

TABLE 2-45: FRAME TYPE DEFINITIONS FOR CPU FORWARDING

Frame	Condition	Copy/Redirect
BPDUs frames. Reserved Addresses (IEEE 802.1D 7.12.6)	DMAC = 0x0180C2000000 to 0x0180C20000F (BPDUs and various Slow protocols supporting spanning tree, link aggregation, port authentication)	Redirect/Copy/ Discard
Reserved ALL-BRIDGE address	DMAC = 0x0180C2000010	Redirect/Copy/ Discard
GARP Application Addresses (IEEE 802.1D 12.5)	DMAC = 0x0180C2000020 to 0x0180C200002F	Redirect/Copy/ Discard
CCM/Link Trace Addresses (IEEE P802.1ag)	DMAC = 0x0180C2000030 to 0x0180C200003F	Redirect/Copy/ Discard
IGMP	DMAC = 0x01005E000000 to 0x01005E7FFFFFFF EtherType = IPv4 IP Protocol = IGMP	Redirect
MLD	DMAC = 0x333300000000 to 0x3333FFFFFFFF EtherType = IPv6 IPv6 Next Header = 0 Hop-by-hop options header with the first option being a Router Alert option with the MLD message (Option Type = 5, Opt Data Len = 2, Option Data = 0).	Redirect
IPv4 Multicast Ctrl	DMAC = 0x01005E000000 to 0x01005E7FFFFFFF EtherType = IPv4 IP protocol is not IGMP IPv4 DIP inside 224.0.0.x	Copy
Source port	All frames received on enabled ingress port	Copy
All other frames		

In addition, the classifier can recognize Versatile Register Access Protocol (VRAP) frames and redirect such frames to the CPU. This is a proprietary frame format, which is used for reading and writing switch configuration registers through ethernet frames, see [Section 2.15, VRAP Engine](#).

To determine if a frame is a VRAP frame, the VRAP filter in the classifier performs three checks:

- VLAN check. The filter can be either VLAN unaware or VLAN aware (ANA::VRAP_CFG.VRAP_VLAN_AWARE_ENA). If VLAN unaware, VRAP frames must be untagged. If VLAN aware, VRAP frames must be VLAN tagged and the frame's VID must match a configured value (ANA::VRAP_CFG.VRAP_VID). Double VLAN tagged frames always fail this check.
- EtherType and EPID check. The EtherType must be 0x8880 and the EPID must be 0x0004 (bytes 0 and 1 after the EtherType).
- VRAP header check. The VRAP header (bytes 0, 1, 2, and 3 after the EPID) must match a 32-bit configured value (ANA::VRAP_HDR_DATA) where any bits can be don't cared by a mask (ANA::VRAP_HDR_MASK).

If all three checks are fulfilled, frames are redirected to CPU extraction queue ANA::CPUQ_CFG2.CPUQ_VRAP.

The VRAP filter is enabled in ANA:PORT:CPU_FWD_CFG.

2.7 VCAP


The VCAP is a content-aware packet processor for wire-speed packet inspection for rich implementation of, for example, advanced VLAN and QoS classifications and manipulations, IP source guarding, and security features for wireline and wireless applications.

The following describes four VCAPs implemented in the VSC7416-01 device: IS0, IS1, IS2, and ES0. IS1 and IS2 are generic ingress VCAPs working on the incoming frames while ES0 is an egress VCAP working on all outgoing frames. In addition, the MPLS-TP classifier contains a VCAP, IS0, which is solely used for MPLS-TP classification. For more information about the MPLS-TP classifier, see [Section 2.5, MPLS-TP Classifier](#).

When a VCAP is enabled, each frame is examined to determine the frame type (for example IPv4 TCP frame) so that the frame information is extracted according to the frame type. Together with port-specific configuration and classification results from the basic classification, the extracted frame information makes up an entry key, which is passed to a TCAM and matched against entries in the TCAM.

An entry in the TCAM consists of a pattern and a mask, where the mask allows pattern-matching with the use of “don’t cares”. The first matching entry is then used to select an action. The following illustration provides a functional overview of a general TCAM.

FIGURE 2-9: VCAP FUNCTIONAL OVERVIEW



Each frame results in up to eight ingress VCAP lookups and one egress lookup per destination port. The lookups use different keys and the results determine the frame’s ingress classification, security handling, and egress VLAN manipulation. The eight ingress lookups and the associated VCAPs are:

- MPLS Link Layer classification
VCAP: IS0
Key: MLL
Entry: MPLS Link Layer Control Entry
- MPLS Label Stack classification
VCAP: IS0
Key: MLBS
Entry: MPLS Label Stack Control Entry
- Advanced ingress classification, first lookup
VCAP: IS1
Key: IS1
Entry: IS1 Control Entry
- Advanced ingress classification, second lookup
VCAP: IS1
Key: IS1
Entry: IS1 Control Entry
- Advanced ingress classification, third lookup
VCAP: IS1
Key: IS1
Entry: IS1 Control Entry
- IP source guarding check
VCAP: IS2

Key: SMAC_SIP4 (IPv4 frames) or SMAC_SIP6 (IPv6 frames)

Entry: SMAC_SIP4 Control Entry or SMAC_SIP6 Control Entry

- Security enforcement, first lookup

VCAP: IS2

Key: MAC_ETYPE, MAC_LLC, MAC_SNAP, ARP, IP4_OTHER, IP4_TCP_UDP, IP6_STD, IP6_OTHER, IP6_TCP_UDP, OAM, CUSTOM, depending on frame type

Entry: Access Control Entry

- Security enforcement, second lookup

VCAP: IS2

Key: MAC_ETYPE, MAC_LLC, MAC_SNAP, ARP, IP4_OTHER, IP4_TCP_UDP, IP6_STD, IP6_OTHER, IP6_TCP_UDP, OAM, CUSTOM, depending on frame type

Entry: Access Control Entry

The egress lookup per destination port and associated VCAP is:

- Egress tagging and frame manipulations

VCAP: ES0

Key: ES0

Entry: Egress Control Entry

The MPLS lookups are only carried out for MPLS-encapsulated frames on MPLS enabled ports (ANA::MPLS_ENA). If the MPLS Link Layer lookup fails (no hit), then the lookups in the following VCAPs are done using the entire MPLS encapsulated frame. If the MPLS Link Layer lookup succeeds, the lookups in the following VCAPs is done based on the customer Ethernet frame with the MPLS encapsulation stripped.

The IP source guarding check is only carried out for IP frames.

CPU injected frames are subject to all the above VCAP lookups in IS0, IS1, IS2, and ES0.

With respect to IS1 and IS2, each frame is classified to one of seven overall VCAP frame types. The frame type determines the information to extract from the frame and also which VCAP entries to match against. The following table lists which frame types are used and which VCAP entries the frame types are matched against in IS1 and IS2. **Note** The lookup in ES0 is independent of the frame type, and all frames match against all entries in the TCAM.

TABLE 2-46: IS1 AND IS2 VCAP FRAME TYPES

Frame Type	Condition	IS1 Entries	IS2 Entries
IPv6 Frame	The Type/Len field is equal to 0x86DD. The IP version is 6. Special IPv6 frames: •IPv6 TCP frame: Next header is TCP (0x6) •IPv6 UDP frame: Next header is UDP (0x11) •IPv6 Other frame: Next header is neither TCP nor UDP	Frame type flags: ETYPE_LEN = 1 IP_SNAP = 1 IP4 = 0 TCP_UDP TCP	IP6_TCP_UD P IP6_OTHER IP6_STD
IPv4 Frame	The Type/Len field is equal to 0x800. The IP version is 4. Special IPv4 frames: •IPv4 TCP frame: IP protocol is TCP (0x6) •IPv4 UDP frame: IP protocol is UDP (0x11) •IPv4 Other frame: IP protocol is neither TCP nor UDP	Frame type flags: ETYPE_LEN = 1 IP_SNAP = 1 IP4 = 1 TCP_UDP TCP	IP4_TCP_UD P IP4_OTHER
(R)ARP Frame	The Type/Len field is equal to 0x0806 (ARP) or 0x8035 (RARP).	Frame type flags: ETYPE_LEN = 1 IP_SNAP = 0	ARP
OAM Frame	The Type/Len field is equal to 0x8902, 0x8809, or 0x88EE. In particular, the following protocol data units use this EtherType: •IEEE 802.1ag (CFM) •IEEE 802.3 (EFM) •ITU-T Y.1731 (ETH-OAM) •MEF-16 (E-LMI)	Frame type flags: ETYPE_LEN = 1 IP_SNAP = 0	OAM

TABLE 2-46: IS1 AND IS2 VCAP FRAME TYPES (CONTINUED)

Frame Type	Condition	IS1 Entries	IS2 Entries
SNAP Frame	The Type/Len field is less than 0x600. The Destination Service Access Point field, DSAP is equal to 0xAA. The Source Service Access Point field, SSAP is equal to 0xAA. The Control field is equal to 0x3.	Frame type flags: ETTYPE_LEN = 0 IP_SNAP = 1	MAC_SNAP
LLC Frame	The Type/Len field is less than 0x600 The LLC header does not indicate a SNAP frame.	Frame type flags: ETTYPE_LEN = 0 IP_SNAP = 0	MAC_LL_C
ETTYPE Frame	The Type/Len field is greater than or equal to 0x600. The Type field does not indicate any of the previously mentioned frame types, that is, ARP, RARP, IPv4, or IPv6.	Frame type flags: ETTYPE_LEN = 1 IP_SNAP = 0	MAC_ETTYPE

In addition, Precision Time Protocol (PTP) frames are handled by IS2. The following encapsulations of PTP frames are supported:

- PTP over Ethernet:
ETTYPE frame with Type/Len = 0x88F7.
Matched against MAC_ETTYPE entries.
- PTP over UDP over IPv4:
IPv4 UDP frame with UDP destination port numbers 319 or 320.
Matched against IP4_TCP_UDP entries.
- PTP over UDP over IPv6:
IPv6 UDP frame with UDP destination port numbers 319 or 320.
Matched against IP6_TCP_UDP entries.

PTP frames may be untagged, single tagged, or double tagged. PTP frames may also MPLS encapsulated where both MPLS-TP and IP/MPLS are supported.

For PTP over Ethernet, the following PTP fields are always extracted:

- TransportSpecific (byte 0)
- MessageType (byte 0)
- VersionPTP (byte 1)
- DomainNumber (byte 4)
- FlagField: flags 1, 2, and 7 (byte 6)

Note Byte 0 is the byte immediately following the EtherType, then byte 1, byte 2, and so on.

For PTP over UDP, the following PTP fields are always extracted:

- MessageType (byte 0)
- VersionPTP (byte 1)
- DomainNumber (byte 4)
- FlagField: flags 1, 2, and 7 (byte 6)

Note Byte 0 is the byte immediately following the UDP header, then byte 1, byte 2, and so on.

2.7.1 PORT CONFIGURATION

This section provides information about special port configurations that control the key generation for the VCAPs.

The following table lists the registers associated with port configuration for VCAP.

TABLE 2-47: PORT MODULE CONFIGURATION OF VCAP

Register	Description	Replication
ANA:PORT:VCAP_CFG	Configuration of the key generation for the VCAPs.	Per virtual port profile

TABLE 2-47: PORT MODULE CONFIGURATION OF VCAP (CONTINUED)

Register	Description	Replication
ANA:PORT:VCAP_S1_KEY_CFG	Configuration of the key generation for the VCAP IS1 per lookup.	Per virtual port profile
ANA:PORT:VCAP_S2_CFG	Configuration of the key generation for the VCAP IS2.	Per virtual port profile
REW:PORT:PORT_CFG	Enables VCAP ES0.	Per port

Each port module affects the key generation for VCAPs IS1 and IS2 through the VCAP_CFG, VCAP_S1_KEY_CFG, and VCAP_S2_CFG registers, and the rewriter affects VCAP ES0 through the REW:PORT:PORT_CFG.ES0_ENA register.

2.7.1.1 VCAP IS0 Port Configuration

There is no specific configuration for IS0: All MPLS-encapsulated frames processed by MPLS-TP classifier are subject to an IS0 lookup.

2.7.1.2 VCAP IS1 Port Configuration

The following port configurations are available for IS1:

- Enable lookups in IS1 (VCAP_CFG.S1_ENA). If disabled, frames received by the port module are not matched against rules in VCAP IS1.
- Use destination information rather than source information (VCAP_CFG.S1_DMAC_DIP_ENA). By default, the two advanced classification lookups in IS1 use the source MAC address and source IP address from the incoming frame when generating the key. Through S1_DMAC_DIP_ENA, the corresponding destination information, destination MAC address, and destination IP address can be used instead. This can be controlled per lookup so that, for example, the first lookup applies source information, and the second applies destination information.
- Use inner VLAN tag rather than outer VLAN tag (VCAP_CFG.S1_VLAN_INNER_TAG_ENA). By default, the two advanced classification lookups in IS1 use the outer VLAN tag from the incoming frame when generating the key. Through S1_VLAN_INNER_TAG_ENA, the inner tag for double tagged frames can be used. This can be controlled per lookup so that, for example, the first lookup applies the outer tag, and the second lookup applies the inner tag. For single tagged frames, the outer VLAN tag is always used.
- Select which IS1 key to use for IPv6 frames (VCAP_S1_CFG.S1_KEY_IP6_CFG). This can be controlled per lookup in IS1. IPv6 frames can use any of the six supported keys in IS1. For more information about the IS1 keys, see [Section 2.7.3, VCAP IS1](#).
- Select which IS1 key to use for IPv4 frames (VCAP_S1_CFG.S1_KEY_IP4_CFG). This can be controlled per lookup in IS1. IPv4 frames can use S1_NORMAL, S1_7TUPLE, S1_5TUPLE_IP4, or S1_DBL_VID keys.
- Select which IS1 key to use for non-IP frames (VCAP_S1_CFG.S1_KEY_OTHER_CFG). This can be controlled per lookup in IS1. Non-IP frames can use S1_NORMAL, S1_7TUPLE, or S1_DBL_VID keys.

2.7.1.3 VCAP IS2 Port Configuration

The following port configurations are available for IS2:

- Enable lookups in IS2 (VCAP_S2_CFG.S2_ENA). If disabled, frames received by the port module are not matched against rules in VCAP IS2.
- Default PAG value (VCAP_CFG.PAG_VAL). This PAG value is the initial value. Actions out of IS1 can change the PAG value before it is used in the key for IS2.

Each port module can control a hierarchy of which entry types in IS2 to use for different frame types. This is controllable per lookup. For instance, it is controllable whether IPv6 TCP frames are matched against IP6_TCP_UDP entries, IP6_STD entries, IP4_TCPUDP entries, or MAC_ETYPE entries. Note that matching against an entry type controls how the key is generated.

With reference to the VCAP_S2_CFG register, the following table lists the hierarchy for different frame types.

TABLE 2-48: HIERARCHY OF IS2 ENTRY TYPES

Frame Type	Description
IPv6 TCP and UDP frames	Configuration: S2_IP6_CFG. If S2_IP6_CFG is set to 0, IPv6 TCP and UDP frames are matched against IP6_TCP_UDP entries. If S2_IP6_CFG is set to 1, IPv6 TCP and UDP frames are matched against IP6_STD entries. If S2_IP6_CFG is set to 2, IPv6 TCP and UDP frames are matched against IP4_TCP_UDP entries. If S2_IP6_CFG is set to 3, IPv6 TCP and UDP frames are matched against MAC_ETYPE entries. Note , S2_IP6_CFG also controls the keys generation for IPv6 Other frames.
IPv6 Other frames (non-TCP and non-UDP)	Configuration: S2_IP6_CFG. If S2_IP6_CFG is set to 0, IPv6 Other frames are matched against IP6_OTHER entries. If S2_IP6_CFG is set to 1, IPv6 Other frames are matched against IP6_STD entries. If S2_IP6_CFG is set to 2, IPv6 Other frames are matched against IP4_OTHER entries. If S2_IP6_CFG is set to 3, IPv6 Other frames are matched against MAC_ETYPE entries. Note S2_IP6_CFG also controls the keys generation for IPv6 TCP and UDP frames.
IPv4 TCP and UDP frames	Configuration: S2_IP_TCPUDP_DIS If S2_IP_TCPUDP_DIS is cleared, IPv4 TCP and UDP frames are matched against IP4_TCPUDP entries. If S2_IP_TCPUDP_DIS is set, IPv4 TCP and UDP frames are matched against MAC_ETYPE entries.
IPv4 Other frames (non-TCP and non-UDP)	Configuration: S2_IP_OTHER_DIS If S2_IP_OTHER_DIS is cleared, IPv4 Other frames are matched against IP4_OTHER entries. If S2_IP_OTHER_DIS is set, IPv4 Other frames are matched against MAC_ETYPE entries.
ARP frames	Configuration: S2_ARP_DIS If S2_ARP_DIS is cleared, ARP frames are matched against ARP entries. If S2_ARP_DIS is set, ARP frames are matched against MAC_ETYPE entries.
OAM frames	Configuration: S2_OAM_DIS If S2_OAM_DIS is cleared, OAM frames are matched against OAM entries. If S2_OAM_DIS is set, OAM frames are matched against MAC_ETYPE entries.
SNAP frames	Configuration: S2_SNAP_DIS If S2_SNAP_DIS is cleared, SNAP frames are matched against SNAP entries. If S2_SNAP_DIS is set, SNAP frames are matched against LCC entries.

2.7.1.4 VCAP ES0 Port Configuration

The rewriter configures VCAP ES0 through REW:PORT:PORT_CFG.ES0_ENA. If ES0 is disabled, frames transmitted on the port are not matched against rules in ES0.

2.7.2 VCAP IS0


This section provides a brief overview of the IS0 keys and associated actions. For more information about IS0, see [Section 2.5, MPLS-TP Classifier](#).

2.7.2.1 IS0 Entry Key Encoding

All MPLS-encapsulated frames processed by the MPLS-TP classifier are subject to up to two IS0 lookups. The first lookup is always done using MPLS Link Layer data. Based on the action from first lookup, the second lookup is optionally done based on the MPLS label stack.

The following illustrations show which entry fields are available for each of the two entries.

FIGURE 2-10: IS0 ENTRY TYPE OVERVIEW



VCAP IS0 can hold the following number of IS0 entries:

MLL: 32 MLL entries.

MLBS: 128 MLBS entries.

The following tables provide information about how the IS0 keys are generated.

TABLE 2-49: IS0 MLL KEY

Field Name	Bit	Width	Description
IGR_PORT	0	4	The port number where the frame was received (0-11).
TTYPE	4	2	MPLS link layer VLAN tag type. 0: Untagged, 1: C-tagged, 2: S-tagged, 3: Reserved
BVID	6	12	Frame's MPLS link layer VID if available, see TTYPE.
M_DMACH	18	48	Frame's MPLS link layer destination MAC address
M_SMACH	66	48	Frame's MPLS link layer source MAC address
E_TYPE	114	2	MPLS link layer EtherType type: 0: Reserved 1: Downstream Assigned Label (0x8847) 2: Upstream Assigned Label (0x8848) 3: Reserved

TABLE 2-50: IS0 MLBS KEY

Field Name	Bit	Width	Description
LL_IDX	0	7	Link Layer index from the action of the first lookup.
LBL_0_VAL	7	20	Top MPLS label value
LBL_0_TC	27	3	Top MPLS label TC bits
LBL_1_VAL	30	20	MPLS label 1 value
LBL_1_TC	50	3	MPLS label 1 TC bits
LBL_2_VAL	53	20	MPLS label 2 value
LBL_2_TC	73	3	MPLS label 2 TC bits

2.7.2.2 ISO Action Encoding

Each of the two lookups in ISO has its own action vector. There is one default action vector per lookup.

The following table lists the available fields for the MLL action vector.

TABLE 2-51: ISO MLL ACTION FIELDS

Action Field	Bit	Width	Description
LL_IDX	0	7	MPLS Link Layer index. Used as part of the MLBS key.
MPLS_FWD	7	1	Forward based on MPLS label stack. This enables the MLBS lookup.
B_PM	8	12	B-domain port mask.
CPU_Q	20	3	CPU extraction queue when frame is forwarded to CPU (bit 12 set in B_PM).
ISDX	23	8	MPLS ISDX value.
RESERVED	31	2	Must be set to 0.
VPROFILE_IDX	33	6	Virtual port profile index, which points out the configuration set used in the basic classifier and analyzer for QOS_DEFAULT_VAL, DP_DEFAULT_VAL.
SCONFIG_ENA	39	1	Use service specific configuration instead of virtual port profile configuration.
CL_VID	40	12	Use this value instead of ANA:VPORT:VLAN_CFG.VLAN_VID.
QOS_DEFAULT_VAL	52	3	Use this value instead of ANA:VPORT:QOS_CFG.QOS_DEFAULT_VAL. Only applicable if SCONFIG_ENA=1.
DP_DEFAULT_VAL	55	1	Use this value instead of ANA:VPORT:QOS_CFG.DP_DEFAULT_VAL. Only applicable if SCONFIG_ENA=1.
OAM_ISDX	53	8	ISDX for OAM frames.
RESERVED	61	2	Must be set to 0.
ISDX_ADD_REPL	66	1	If set, OAM_ISDX replaces the ISDX for OAM frames. Otherwise, OAM_ISDX is added to ISDX for OAM frames.
HIT_STICKY	67	1	If set, a frame has matched against the associated entry.

The following table lists the available fields for the MLBS action vector.

TABLE 2-52: ISO MLBS ACTION FIELDS

Action Field	Bit	Width	Description
ISDX	0	8	MPLS ISDX value. Only used when MPLS forwarding is disabled.
RESERVED	8	2	Must be set to 0.
B_PM	10	12	B-domain port mask.
CPU_Q	22	3	CPU extraction queue when frame is forwarded to CPU (bit 12 set in B_PM).
OAM_ENA	25	3	This channel is enabled for OAM: 0: No OAM 1: VCCV1 OAM 2: VCCV2 OAM 3: VCCV3 OAM 4: GAL MEP LSP OAM 5: GAL MIP LSP OAM 6-7: Reserved.

TABLE 2-52: ISO MLBS ACTION FIELDS (CONTINUED)

Action Field	Bit	Width	Description
RESERVED	28	1	Must be set to 0.
RSVD_LBL_VAL	29	4	Use when OAM_ENA =4 or 5. Configures the RESERVED LABEL used as OAM label for this MEP or MIP. Only the lower 4 bit of the label value is configured.
RSVD_LBL_BOS	33	1	Only for use with OAM_ENA = 4. Configures if the RESERVED LABEL used as OAM label for this MEP must be Bottom of Stack.
CW_ENA	34	1	If set, the first nibble following the MPLS label stack bottom must be set to 0, otherwise error "Expected CW" is generated.
TC_LABEL	35	2	Selects which label provides the frame's classified TC.
TTL_LABEL	37	2	Selects which label provides the frame's classified TTL.
SWAP_LABEL	39	2	Indicates which label in the label stack is used for forwarding. LER entry: must point to PW label. LSR entry: must point to swap label. GAL MEP OAM entry: Must point to GAL label.
TERMINATE_PW	41	1	If set, this is a terminating PW and the frame is forwarded to the C-domain.
POP_CNT	42	3	Indicates the number of bytes to pop in the rewriter. 0: The frame does not contain a Customer Ethernet layer. This can be used for IP encapsulated PTP frames. 1: Frame is MPLS encapsulated. Do not pop any bytes. 2: Frame is MPLS encapsulated. Pop 14 bytes. 3: Frame is MPLS encapsulated. Pop 18 bytes. 4: Frame is MPLS encapsulated. Pop 22 bytes. 5: Frame is MPLS encapsulated. Pop 26 bytes. 6: Frame is MPLS encapsulated. Pop 30 bytes. 7: Frame is MPLS encapsulated. Pop 34 bytes.
E_LSP	45	1	0: LSP is L-LSP. 1: LSP is E-LSP.
TC_MAP	46	3	Selects which TC mapping table to use when mapping TC to drop precedence and QoS class.
L_LSP_COS	49	3	For L-LSP (E_LSP = 0), use L_LSB_QOS as QoS class.
INC_ISDX	52	1	If set, the frame's classified TC value is added to the ISDX value.
OAM_ISDX	53	8	ISDX for OAM frames.
RESERVED	61	2	Must be set to 0.
ISDX_ADD_REPL	63	1	If set, OAM_ISDX replaces the ISDX for OAM frames. Otherwise, OAM_ISDX is added to ISDX for OAM frames.
SWAP_IS_BOS	64	1	If set, the TTL of the swap label is set to 0 and the frame's classified S-bit is set to 0.
VPROFILE_IDX	65	6	Virtual port profile index, which points out the configuration set used in the basic classifier and analyzer.
SCONFIG_ENA	71	1	Use service specific configuration instead of virtual port profile configuration for the following. VLAN_DEI, VLAN_PCP, VLAN_TAG_TYPE. QOS_DEFAULT_VAL, DP_DEFAULT_VAL.
CL_VID	72	12	Use this value instead of ANA:VPORT:VLAN_CFG.VLAN_VID.
VLAN_TAG_TYPE	84	1	Use this value instead of ANA:VPORT:VLAN_CFG.VLAN_TAG_TYPE. Only applicable if SCONFIG_ENA = 1.

TABLE 2-52: IS0 MLBS ACTION FIELDS (CONTINUED)

Action Field	Bit	Width	Description
VLAN_DEI	85	1	Use this value instead of ANA:VPORT:VLAN_CFG.VLAN_DEI. Only applicable if SCONFIG_ENA = 1.
VLAN_PCP	86	3	Use this value instead of ANA:VPORT:VLAN_CFG.VLAN_PCP. Only applicable if SCONFIG_ENA = 1.
HIT_STICKY	88	1	If set, a frame has matched against the associated entry.

2.7.3 VCAP IS1

This section provides information about the IS1 keys and associated actions.

IS1 supports six different keys. The main characteristics of these keys are listed in the following table.

TABLE 2-53: OVERVIEW OF IS1 KEYS


IS1 Key	Size	Frame Types	Key Overview
S1_NORMAL	Half	Applicable to all frame types	Source MAC address, source IP address (32 bits) outer VLAN, DSCP, IP protocol, source and destination TCP/UDP ports.
S1_5TUPLE_IP4	Half	Applicable to IPv4 and IPv6 frames	Inner and outer VLAN, source and destination IP addresses (32 bits), DSCP, IP protocol, source and destination TCP/UDP ports.
S1_NORMAL_IP6	Full	Applicable to IPv6 frames only	Similar to S1_NORMAL but with full IPv6 source IP address: Source MAC address, source IPv6 address (128 bits), inner and outer VLAN, DSCP, IP protocol, source and destination TCP/UDP ports.
S1_7TUPLE	Full	Applicable to all frame types	Source and destination MAC addresses, inner and outer VLAN, source and destination IP addresses (64 bits), DSCP, IP protocol, source and destination TCP/UDP ports.
S1_5TUPLE_IP6	Full	Applicable to IPv6 frames only	Similar to S1_5TUPLE_IP4 but with full IPv6 addresses: Inner and outer VLAN, source and destination IP addresses (128 bits), DSCP, IP protocol, source and destination TCP/UDP ports.
S1_DBL_VID	Quad	Applicable to all frame types	Subset of S1_7TUPLE without addresses: Inner and outer VLAN, DSCP, IP protocol, Destination TCP/UDP port.

2.7.3.1 IS1 Entry Key Encoding

All frame types are subject to the three IS1 lookups. The VCAP IS1 port configuration (VCAP_S1_CFG) determines the key generated for each lookup. Keys that are applicable to multiple frame types (for instance S1_NORMAL) contain frame type flags inside the key that indicate the originating frame type. In addition, certain key fields are overloaded with different frame fields depending on the frame type flag settings.

The following illustration shows which entry fields are available for quad, half and full keys in IS1.

FIGURE 2-11: IS1 ENTRY TYPE OVERVIEW



VSC7416-01

VCAP IS1 can hold the following number of IS1 entries:

- S1_NORMAL: Up to 512 entries.
- S1_5TUPLE_IP4: Up to 512 entries.
- S1_NORMAL_IP6: Up to 256 entries.
- S1_7TUPLE: Up to 256 entries.
- S1_5TUPLE_IP6: Up to 256 entries.
- S1_DBL_VID: Up to 1024 entries.

The following table provides information about how the quad key, S1_DBL_VID is generated.

TABLE 2-54: SPECIFIC FIELDS FOR IS1 QUAD KEY S1_DBL_VID

Field Name	Bit	Width	Description
Lookup Information			
LOOKUP	0	2	0: First lookup 1: Second lookup 2: Third lookup
Interface and Miscellaneous Information			
IGR_PORT_MASK	2	12	Ingress port mask. VCAP generated with one bit set in the mask corresponding to the ingress port.
ISDX	14	9	Service index value from MPLS-TP classifier. Only indexes 0 through 255 are valid.
OAM_Y1731	23	1	Set if frame is a Y.1731 OAM frame with EtherType = 0x8902.
L2_MC	24	1	Set if frame's destination MAC address is a multicast address (bit 40 = 1)
L2_BC	15	1	Set if frame's destination MAC address is the broadcast address (FF-FF-FF-FF-FF-FF)
IP_MC	26	1	Set if frame is IPv4 frame and frame's destination MAC address is an IPv4 multicast address (0x01005E0 /25). Set if frame is IPv6 frame and frame's destination MAC address is an IPv6 multicast address (0x3333/16).
Tagging Information			
VLAN_TAGGED	27	1	Set if frame has one or more Q-tags. Independent of port VLAN awareness.
VLAN_D-BL_TAGGED	28	1	Set if frame has two or more Q-tags. Independent of port VLAN awareness.
TPID	29	1	0: Customer TPID 1: Service TPID (88A8 or programmable). S1_NORMAL: TPID is derived from tag pointed to by VCAP_CFG.S1_VLAN_INNER_TAG_ENA.
VID	30	12	Frame's VID if frame is tagged, otherwise port default. S1_NORMAL: VID is taken from tag pointed to by VCAP_CFG.S1_VLAN_INNER_TAG_ENA.
DEI	42	1	Frame's DEI if frame is tagged, otherwise port default. S1_NORMAL: DEI is taken from tag pointed to by VCAP_CFG.S1_VLAN_INNER_TAG_ENA.
PCP	43	3	Frame's PCP if frame is tagged, otherwise port default. S1_NORMAL: PCP is taken from tag pointed to by VCAP_CFG.S1_VLAN_INNER_TAG_ENA.
INNER_TPID	46	1	TPID for inner tag: 0: Customer TPID. 1: Service TPID (88A8 or programmable).
INNER_VID	47	12	Frame's inner VID if frame is double tagged.

TABLE 2-54: SPECIFIC FIELDS FOR IS1 QUAD KEY S1_DBL_VID (CONTINUED)

Field Name	Bit	Width	Description
INNER_DEI	59	1	Frame's inner DEI if frame is double tagged.
INNER_PCP	60	3	Frame's inner PCP if frame is double tagged.
Layer 2 Information			
ETYPE_LEN	63	1	Frame type flag. Set if frame has EtherType >= 0x600 (Frame is type encoded). Otherwise cleared (Frame is length encoded).
ETYPE	64	16	Overloaded field for different frame types: LLC frame: ETYPE = [DSAP, SSAP] SNAP frame: ETYPE = PID[4:3] IPv4 or IPv6 TCP/UDP frame: ETYPE = DPORT IPv4 or IPv6 Other frame: ETYPE = IP protocol ARP or ETYPE frame: ETYPE = Frame's EtherType. Y.1731 OAM frames (EtherType = 0x8902): ETYPE[15:8] = 0x89. ETYPE[7]: Set if the frames was injected with masquerading. ETYPE[6:0]: Encodes the frame's MEG level, see OAM_MEL_FLAGS Table 2-71 . To indicate the overloading of the ETYPE field, OAM_Y1731 is set to 1.
IP_SNAP	80	1	Frame type flag. Set if frame is IPv4, IPv6, or SNAP frame.
IP4	81	1	Frame type flag. Set if frame is IPv4 frame
Layer 3 Information			
L3_FRAGMENT	82	1	Set if IPv4 frame is fragmented (More Fragments flag = 1 or Fragments Offset > 0). Layer 4 information cannot not be trusted.
L3 - FRAG_OFS_GT0	83	1	Set if IPv4 frame is fragmented and it is not the first fragment (Fragments Offset > 0). Layer 4 information cannot not be trusted.
L3_OPTIONS	84	1	Set if IPv4 frame contains options (IP len > 5). IP options are not skipped nor parsed. Layer 4 information cannot not be trusted.
L3_DSCP	85	6	Frame's DSCP value. The DSCP value may have been translated during basic classification. See Section 2.6.4, QoS, DP, and DSCP Classification .
Layer 4 Information			
TCP_UDP	91	1	Frame type flag. Set if frame is IPv4/IPv6 TCP or UDP frame.
TCP	92	1	Frame type flag. Set if frame is IPv4/IPv6 TCP frame.

The following three tables provide information about how the half keys, S1_NORMAL and S1_5TUPLE_IP4, are generated. The first table lists the common fields between the half keys.

TABLE 2-55: IS1 COMMON KEY FIELDS FOR HALF KEYS

Field Name	Bit	Width	Description
Lookup Information			
IS1_TYPE_HALF	0	1	0: S1_NORMAL entries. 1: S1_5TUPLE_IP4 entries.

TABLE 2-55: IS1 COMMON KEY FIELDS FOR HALF KEYS (CONTINUED)

Field Name	Bit	Width	Description
LOOKUP	1	2	0: First lookup. 1: Second lookup. 2: Third lookup.
Interface and Miscellaneous Information			
IGR_PORT_MASK	3	12	Ingress port mask. VCAP generated with one bit set in the mask corresponding to the ingress port.
ISDX	15	9	Service index value from MPLS-TP classifier. Only indexes 0 through 255 are valid.
OAM_Y1731	24	1	Set if frame is Y.1731 OAM frame with EtherType = 0x8902.
L2_MC	25	1	Set if frame's destination MAC address is a multicast address (bit 40 = 1).
L2_BC	26	1	Set if frame's destination MAC address is the broadcast address (FF-FF-FF-FF-FF-FF)
IP_MC	27	1	Set if frame is IPv4 frame and frame's destination MAC address is an IPv4 multicast address (0x01005E0 /25). Set if frame is IPv6 frame and frame's destination MAC address is an IPv6 multicast address (0x3333/16).
Tagging Information			
VLAN_TAGGED	28	1	Set if frame has one or more Q-tags. Independent of port VLAN awareness.
VLAN_D-BL_TAGGED	29	1	Set if frame has two or more Q-tags. Independent of port VLAN awareness.
TPID	30	1	0: Customer TPID. 1: Service TPID (88A8 or programmable). S1_NORMAL: TPID is derived from tag pointed to by VCAP_CFG.S1_VLAN_INNER_TAG_ENA.
VID	31	12	Frame's VID if frame is tagged, otherwise port default. S1_NORMAL: VID is taken from tag pointed to by VCAP_CFG.S1_VLAN_INNER_TAG_ENA.
DEI	43	1	Frame's DEI if frame is tagged, otherwise port default. S1_NORMAL: DEI is taken from tag pointed to by VCAP_CFG.S1_VLAN_INNER_TAG_ENA.
PCP	44	3	Frame's PCP if frame is tagged, otherwise port default. S1_NORMAL: PCP is taken from tag pointed to by VCAP_CFG.S1_VLAN_INNER_TAG_ENA.

TABLE 2-56: SPECIFIC FIELDS FOR IS1 HALF KEY S1_NORMAL

Field Name	Bit	Width	Description
Layer-2 Information			
L2_SMAC	47	48	Frame's source MAC address. Use destination MAC address if VCAP_CFG.S1_DMAC_DIP_ENA is set for ingress port.
ETYPE_LEN	95	1	Frame type flag. Set if frame has EtherType >= 0x600 (Frame is type encoded). Otherwise cleared (Frame is length encoded).

TABLE 2-56: SPECIFIC FIELDS FOR IS1 HALF KEY S1_NORMAL (CONTINUED)

Field Name	Bit	Width	Description
ETYPE	96	16	Overloaded field for different frame types: LLC frame: ETYPE = [DSAP, SSAP] SNAP frame: ETYPE = PID[4:3] IPv4 or IPv6 TCP/UDP frame: ETYPE = DPORT IPv4 or IPv6 Other frame: ETYPE = IP protocol ARP or ETYPE frame: ETYPE = Frame's EtherType. Y.1731 OAM frames (EtherType = 0x8902): ETYPE[15:8] = 0x89. ETYPE[7]: Set if the frames was injected with masquerading. ETYPE[6:0]: Encodes the frame's MEG level, see OAM_MEL_FLAGS Table 2-71. To indicate the overloading of the ETYPE field, OAM_Y1731 is set to 1.
IP_SNAP	11 2	1	Frame type flag. Set if frame is IPv4, IPv6, or SNAP frame.
IP4	11 3	1	Frame type flag. Set if frame is IPv4 frame
Layer-3 Information			
L3_FRAG- MENT	11 4	1	Set if IPv4 frame is fragmented (More Fragments flag = 1 or Fragments Offset > 0). Layer 4 information cannot not be trusted.
L3_- FRAG_OFS_G TO	11 5	1	Set if IPv4 frame is fragmented and it is not the first fragment (Fragments Offset > 0). Layer 4 information cannot not be trusted.
L3_OPTIONS	11 6	1	Set if IPv4 frame contains options (IP len > 5). IP options are not skipped nor parsed. Layer 4 information cannot not be trusted.
L3_DSCP	11 7	6	Frame's DSCP value. The DSCP value may have been translated during basic classification. See Section 2.6.4, QoS, DP, and DSCP Classification .
L3_IP4_SIP	12 3	32	Overloaded fields for different frame types: LLC frame: L3_IP4_SIP = [CTRL, PAYLOAD[0:2]] SNAP frame: L3_IP4_SIP = [PID[2:0], PAYLOAD[0]] IPv4 or IPv6 frame: L3_IP4_SIP = source IP address, bits [31:0] OAM, ARP or ETYPE frame: L3_IP4_SIP = PAYLOAD[0:3] For IPv4 or IPv6 frames, use destination IP address if VCAP_CFG.S1_DMAC_DIP_ENA is set for ingress port.
Layer-4 Information			
TCP_UDP	15 5	1	Frame type flag. Set if frame is IPv4/IPv6 TCP or UDP frame.
TCP	15 6	1	Frame type flag. Set if frame is IPv4/IPv6 TCP frame.
L4_SPORT	15 7	16	TCP/UDP frame's source port.
L4_RNG	17 3	8	Range mask with one bit per range. A bit is set, if the corresponding range is matched. Range types: SPORT, DPORT, SPORT or DPORT, VID, DSCP Input to range checkers: SPORT/DPORT: From frame VID: From frame if tagged, otherwise port's VID DSCP: Translated DSCP from the basic classification. See Section 2.7.6, Range Checkers .

TABLE 2-57: SPECIFIC FIELDS FOR IS1 HALF KEY S1_5TUPLE_IP4

Field Name	Bit	Width	Description
Tagging Information			
INNER_TPID	47	1	TPID for inner tag. 0: Customer TPID. 1: Service TPID (88A8 or programmable).
INNER_VID	48	12	Frame's inner VID if frame is double tagged.
INNER_DEI	60	1	Frame's inner DEI if frame is double tagged.
INNER_PCP	61	3	Frame's inner PCP if frame is double tagged.
Layer-3 Information			
IP4	64	1	Frame type flag. Set if frame is IPv4 frame.
L3_FRAGMENT	65	1	Set if IPv4 frame is fragmented (More Fragments flag = 1 or Fragments Offset > 0). Layer 4 information cannot not be trusted.
L3_FRAG_OFS_GT0	66	1	Set if IPv4 frame is fragmented and it is not the first fragment (Fragments Offset > 0). Layer 4 information cannot not be trusted.
L3_OPTIONS	67	1	Set if IPv4 frame contains options (IP len > 5). IP options are not skipped nor parsed. Layer 4 information cannot not be trusted.
L3_DSCP	68	6	Frame's DSCP value. The DSCP value may have been translated during basic classification, see Section 2.6.4, QoS, DP, and DSCP Classification .
L3_IP4_DIP	74	32	IPv4: destination IP address. IPv6: destination IP address, bits [31:0].
L3_IP4_SIP	106	32	IPv4: source IP address. IPv6: source IP address, bits [31:0].
L3_PROTO	138	8	IPv4: IP protocol. IPv6: next header.
Layer-4 Information			
TCP_UDP	146	1	Frame type flag. Set if frame is IPv4/IPv6 TCP or UDP frame.
TCP	147	1	Frame type flag. Set if frame is IPv4/IPv6 TCP frame.
L4_RNG	148	8	Range mask with one bit per range. A bit is set, if the corresponding range is matched. Range types: SPORT, DPORT, SPORT or DPORT, VID, DSCP Input to range checkers: SPORT/DPORT: From frame VID: From frame if tagged, otherwise port's VID DSCP: Translated DSCP from the basic classification. See Section 2.7.6, Range Checkers .
IP_PAY-LOAD_S1_5TUPLE	156	32	Payload after IPv4 or IPv6 header. For TCP and UDP frames, this field contains the source and destination port numbers.

The following four tables provide information about how the full keys, S1_NORMAL_IP6, S1_7TUPLE, and S1_5TUPLE_IP6, are generated. The first table lists the common fields between the full keys.

TABLE 2-58: IS1 COMMON KEY FIELDS FOR FULL KEYS

Field Name	Bit	Width	Description
Lookup Information			
IS1_TYPE_FULL	0	2	0: S1_NORMAL_IP6 entries. 1: S1_7TUPLE entries. 2: S1_5TUPLE_IP6 entries.
LOOKUP	2	2	0: First lookup. 1: Second lookup. 2: Third lookup.
Interface and Miscellaneous Information			
IGR_PORT_MASK	4	12	Ingress port mask. VCAP generated with one bit set in the mask corresponding to the ingress port.
ISDX	16	9	Service index value from MPLS-TP classifier. Only indexes 0 through 255 are valid.
OAM_Y1731	25	1	Set if frame is a Y.1731 frame with EtherType = 0x8902.
L2_MC	26	1	Set if frame's destination MAC address is a multicast address (bit 40 = 1).
L2_BC	27	1	Set if frame's destination MAC address is the broadcast address (FF-FF-FF-FF-FF-FF).
IP_MC	28	1	Set if frame is IPv4 frame and frame's destination MAC address is an IPv4 multicast address (0x01005E0 /25). Set if frame is IPv6 frame and frame's destination MAC address is an IPv6 multicast address (0x3333/16).
Tagging Information			
VLAN_TAGGED	29	1	Set if frame has one or more Q-tags. Independent of port VLAN awareness.
VLAN_D-BL_TAGGED	30	1	Set if frame has two or more Q-tags. Independent of port VLAN awareness.
TPID	31	1	0: Customer TPID. 1: Service TPID (88A8 or programmable). S1_NORMAL: TPID is derived from tag pointed to by VCAP_CFG.S1_VLAN_INNER_TAG_ENA.
VID	32	12	Frame's VID if frame is tagged, otherwise port default. S1_NORMAL: VID is taken from tag pointed to by VCAP_CFG.S1_VLAN_INNER_TAG_ENA.
DEI	44	1	Frame's DEI if frame is tagged, otherwise port default. S1_NORMAL: DEI is taken from tag pointed to by VCAP_CFG.S1_VLAN_INNER_TAG_ENA.
PCP	45	3	Frame's PCP if frame is tagged, otherwise port default. S1_NORMAL: PCP is taken from tag pointed to by VCAP_CFG.S1_VLAN_INNER_TAG_ENA.
INNER_TPID	48	1	TPID for inner tag. 0: Customer TPID. 1: Service TPID (88A8 or programmable).
INNER_VID	49	12	Frame's inner VID if frame is double tagged.
INNER_DEI	61	1	Frame's inner DEI if frame is double tagged.
INNER_PCP	62	3	Frame's inner PCP if frame is double tagged.

TABLE 2-59: SPECIFIC FIELDS FOR IS1 FULL KEY S1_NORMAL_IP6

Field Name	Bit	Width	Description
Layer 2 Information			
L2_SMAC	65	48	Frame's source MAC address. Use destination MAC address if VCAP_CFG.S1_D-MAC_DIP_ENA is set for ingress port.
Layer 3 Information			
L3_DSCP	113	6	Frame's DSCP value. The DSCP value may have been translated during basic classification, see Section 2.6.4, QoS, DP, and DSCP Classification .
L3_IP6_SIP	119	128	Source IP address. Use destination IP address if VCAP_CFG.S1_D-MAC_DIP_ENA is set for ingress port.
L3_PROTO	247	8	IPv6 next header.
Layer 4 Information			
TCP_UDP	255	1	Frame type flag. Set if frame is IPv4/IPv6 TCP or UDP frame.
TCP	256	1	Frame type flag. Set if frame is IPv4/IPv6 TCP frame.
L4_RNG	257	8	Range mask with one bit per range. A bit is set, if the corresponding range is matched. Range types: SPORT, DPORT, SPORT or DPORT, VID, DSCP. Input to range checkers: SPORT/DPORT: From frame VID: From frame if tagged, otherwise port's VID DSCP: Translated DSCP from the basic classification. See Section 2.7.6, Range Checkers .
IP_PAY-LOAD_S1_IP6	265	112	Payload after IPv6 header. For TCP and UDP frames, this field contains the source and destination port numbers.

TABLE 2-60: SPECIFIC FIELDS FOR IS1 FULL KEY S1_7TUPLE

Field Name	Bit	Width	Description
Layer-2 Information			
L2_DMAC	65	48	Frame's destination MAC address.
L2_SMAC	113	48	Frame's source MAC address.
ETYPE_LEN	161	1	Frame type flag. Set if frame has EtherType >= 0x600 (Frame is type encoded). Otherwise cleared (Frame is length encoded).

TABLE 2-60: SPECIFIC FIELDS FOR IS1 FULL KEY S1_7TUPLE (CONTINUED)

Field Name	Bit	Width	Description
ETYPE	162	16	Overloaded field for different frame types: LLC frame: ETYPE = [DSAP, SSAP] SNAP frame: ETYPE = PID[4:3] IPv4 or IPv6 TCP/UDP frame: ETYPE = DPORT IPv4 or IPv6 Other frame: ETYPE = IP protocol ARP or ETYPE frame: ETYPE = Frame's EtherType. Y.1731 OAM frames (EtherType = 0x8902): ETYPE[15:8] = 0x89. ETYPE[7]: Set if the frames was injected with masquerading. ETYPE[6:0]: Encodes the frame's MEG level, see OAM_MEL_FLAGS Table 2-71 . To indicate the overloading of the ETYPE field, OAM_Y1731 is set to 1.
IP_SNAP	178	1	Frame type flag. Set if frame is IPv4, IPv6, or SNAP frame.
IP4	179	1	Frame type flag. Set if frame is IPv4 frame
Layer-3 Information			
L3_FRAGMENT	180	1	Set if IPv4 frame is fragmented (More Fragments flag = 1 or Fragments Offset > 0). Layer 4 information cannot not be trusted.
L3_FRAG_OFS_GT0	181	1	Set if IPv4 frame is fragmented and it is not the first fragment (Fragments Offset > 0). Layer 4 information cannot not be trusted.
L3_OPTIONS	182	1	Set if IPv4 frame contains options (IP len > 5). IP options are not skipped nor parsed. Layer 4 information cannot not be trusted.
L3_DSCP	183	6	Frame's DSCP value. The DSCP value may have been translated during basic classification, see Section 2.6.4, QoS, DP, and DSCP Classification .
L3_IP6_DIP_MSB	189	16	IPv6 frame: L3_IP6_DIP_MSB = destination IP address, bits [127:112]
L3_IP6_DIP	205	64	Overloaded fields for different frame types: LLC frame: L3_IP6_DIP = PAYLOAD[7:14] SNAP frame: L3_IP6_DIP = PAYLOAD[5:12] IPv4 frame: L3_IP6_DIP = destination IP address, bits [31:0] IPv6 frame: L3_IP6_DIP = destination IP address, bits [63:0] OAM, ARP or ETYPE frame: L3_IP6_DIP = PAYLOAD[8:15]
L3_IP6_SIP_MSB	269	16	IPv6 frame: L3_IP6_SIP_MSB = source IP address, bits [127:112]
L3_IP6_SIP	285	64	Overloaded fields for different frame types: LLC frame: L3_IP6_SIP = [CTRL, PAYLOAD[0:6]] SNAP frame: L3_IP6_SIP = [PID[2:0], PAYLOAD[0:4]] IPv4 frame: L3_IP6_SIP = source IP address, bits [31:0] IPv6 frame: L3_IP6_SIP = source IP address, bits [63:0] OAM, ARP or ETYPE frame: L3_IP6_SIP = PAYLOAD[0:7]
Layer-4 Information			
TCP_UDP	349	1	Frame type flag. Set if frame is IPv4/IPv6 TCP or UDP frame.
TCP	350	1	Frame type flag. Set if frame is IPv4/IPv6 TCP frame.

TABLE 2-60: SPECIFIC FIELDS FOR IS1 FULL KEY S1_7TUPLE (CONTINUED)

Field Name	Bit	Width	Description
L4_SPORT	351	16	TCP/UDP frame's source port.
L4_RNG	367	8	Range mask with one bit per range. A bit is set, if the corresponding range is matched. Range types: SPORT, DPORT, SPORT or DPORT, VID, DSCP Input to range checkers: SPORT/DPORT: From frame VID: From frame if tagged, otherwise port's VID DSCP: Translated DSCP from the basic classification. See Section 2.7.6, Range Checkers .

TABLE 2-61: SPECIFIC FIELDS FOR IS1 FULL KEY S1_5TUPLE_IP6

Field Name	Bit	Width	Description
Layer-3 Information			
L3_DSCP	65	6	Frame's DSCP value. The DSCP value may have been translated during basic classification. See Section 2.6.4, QoS, DP, and DSCP Classification .
L3_IP6_DIP	71	128	IPv6 destination IP address
L3_IP6_SIP	199	128	IPv6 source IP address
L3_PROTO	327	8	IPv6 next header.
Layer-4 Information			
TCP_UDP	335	1	Frame type flag. Set if frame is IPv6 TCP or UDP frame.
TCP	336	1	Frame type flag. Set if frame is IPv6 TCP frame.
L4_RNG	337	8	Range mask with one bit per range. A bit is set, if the corresponding range is matched. Range types: SPORT, DPORT, SPORT or DPORT, VID, DSCP Input to range checkers: SPORT/DPORT: From frame VID: From frame if tagged, otherwise port's VID DSCP: Translated DSCP from the basic classification. See Section 2.7.6, Range Checkers .
IP_PAY-LOAD_S1_5TUPLE	345	32	Payload after IPv6 header. For TCP and UDP frames, this field contains the source and destination port numbers.

Fields not applicable to a certain frame type (for example, L3_OPTIONS for an IPv6 frame) must be set to don't care for entries the frame type can match.

If L3_FRAGMENT or L3_OPTIONS are set to 1 or set to don't care, Layer 4 information cannot be trusted and should be set to don't-care for such entries.

2.7.3.2 IS1 Action Encoding

The VCAP generates an action vector from each of the three IS1 lookups. All matches return the same action vector, independently of the match entry type. The action vectors are combined into one action vector, which is applied to the classification of the frame.

There are no default action vectors for the IS1.

The following table lists the available fields for the IS1 action vector.

TABLE 2-62: IS1 ACTION FIELDS

Action Field	Bit	Width	Description
DSCP_ENA	0	1	If set, use DSCP_VAL as classified DSCP value. Otherwise, DSCP value from basic classification is used.
DSCP_VAL	1	6	See DSCP_ENA.
QOS_ENA	7	1	If set, use QOS_VAL as classified QoS class. Otherwise, QoS class from basic classification is used.
QOS_VAL	8	3	See QOS_ENA.
DP_ENA	11	1	If set, use DP_VAL as classified drop precedence level. Otherwise, drop precedence level from basic classification is used.
DP_VAL	12	1	See DP_ENA.
PAG_OVERRIDE_MASK	13	6	Bits set in this mask will override PAG_VAL from port profile. New PAG = (PAG (input) AND ~PAG_OVERRIDE_MASK) OR (PAG_VAL AND PAG_OVERRIDE_MASK)
OAM_DETECT_SEL	19	2	Configures the detection of OAM frames: 0: Do not detect OAM; frame is treated as data even if EtherType is 0x8902. 1: Detect as OAM if frame is untagged and EtherType is 0x8902. 2: Detect as OAM if frame is single tagged and EtherType is 0x8902. 3: Detect as OAM if frame is double tagged and EtherType is 0x8902. See also OAM_DETECT_ENA.
PAG_VAL	21	6	See PAG_OVERRIDE_MASK.
OAM_DETECT_ENA	27	2	If set to 3, the detection of OAM frame is controlled by OAM_DETECT_SEL. If set to 0, the setting in ANA::VCAP_CFG.OAM_DETECT_CFG is used. Other values of OAM_DETECT_ENA are not supported.
ISDX_REPLACE_ENA	29	1	This action only applies to the first lookup in IS1. Controls the classified ISDX: ISDX_REPLACE_ENA=0: Add ISDX_ADD_VAL to MPLS classified ISDX and use result as new classified ISDX. ISDX_REPLACE_ENA = 1: Replace MPLS classified ISDX with ISDX_VAL value and use as new classified ISDX.
ISDX_ADD_VAL	30	10	This action only applies to the first lookup in IS1. Only indexes 0 through 255 are valid. See ISDX_REPLACE_ENA.
VID_REPLACE_ENA	40	1	Controls the classified VID: VID_REPLACE_ENA=0: Add VID_ADD_VAL to basic classified VID and use result as new classified VID. VID_REPLACE_ENA = 1: Replace basic classified VID with VID_VAL value and use as new classified VID.
VID_ADD_VAL	41	12	See VID_REPLACE_ENA.

TABLE 2-62: IS1 ACTION FIELDS (CONTINUED)

Action Field	Bit	Width	Description
FID_SEL	53	2	Controls the Filter Identifier (FID) used when looking up the MAC table. 0: Disabled: FID = classified VID prepended 0. 1: Use FID_VAL for SMAC lookup in MAC table. 2: Use FID_VAL for DMAC lookup in MAC table. 3: Use FID_VAL for DMAC and SMAC lookup in MAC table.
FID_VAL	55	13	See FID_SEL.
PCP_DEI_ENA	68	1	If set, use PCP_VAL and DEI_VAL as classified PCP and DEI values. Otherwise, PCP and DEI from basic classification are used.
PCP_VAL	69	3	See PCP_DEI_ENA.
DEI_VAL	72	1	See PCP_DEI_ENA.
VLAN_POP_CNT_ENA	73	1	If set, use VLAN_POP_CNT as the number of VLAN tags to pop from the incoming frame. This number is used by the Rewriter. Otherwise, VLAN_POP_CNT from ANA:PORT:VLAN_CFG.VLAN_POP_CNT is used.
VLAN_POP_CNT	74	2	See VLAN_POP_CNT_ENA.
CUSTOM_ACE_TYPE_ENA	76	4	Enables use of custom keys in IS2. Bits 3:2 control second lookup in IS2 while bits 1:0 control first lookup. Encoding per lookup: 0: Disabled. 1: Extract 40 bytes after position corresponding to the location of the IPv4 header and use as key. 2: Extract 40 bytes after SMAC and use as key.
HIT_STICKY	80	1	If set, a frame has matched against the associated entry.

Each lookup returns an action vector if there is a match. The potentially three IS1 action vectors are applied in three steps. First, the action vector from the first lookup is applied, then the action vector from the second lookup is applied to the result from the first action vector, and finally, the action vector from the third lookup is applied to the result from the second action vector. This implies that if two or more lookups return an action of DP_ENA = 1; for example, the DP_VAL from the last lookup is used.

The CUSTOM_ACE_TYPE_ENA action in IS1 enables the use of custom keys in IS2. Two different custom keys are supported:

- CUSTOM_IP: 40 bytes are extracted from the frame starting from byte-position 34. This position corresponds to the first byte after the IPv4 header if the frame is IPv4. The 40 bytes are matched against the CUSTOM entries in IS2.
- CUSTOM_SMAC: 40 bytes are extracted from the frame starting from byte-position 12. This position corresponds to the first byte after the source MAC address. The 40 bytes are matched against the CUSTOM entries in IS2.

Note Up to two VLAN tags (two for double VLAN tagged frames and one for single VLAN tagged frames) are removed from the frame before extracting the data.


2.7.4 VCAP IS2

This section provides information about the IS2 keys, the SMAC_SIP4 and SMAC_SIP6 keys, and associated actions.

2.7.4.1 IS2 Entry Key Encoding

All frame types are subject to the two IS2 lookups. The frame type determines the key entry type. For more information about VCAP frame types, see [Table 2-46](#). The following illustrations show which entry fields are available for each frame type (indicated by the field IS2_TYPE_HALF and IS2_TYPE_FULL) for half and full entries.

FIGURE 2-12: IS2 HALF ENTRY TYPE OVERVIEW





- DS60001705A-page 92

- MAC_SNAP: Up to 512 entries.
- ARP: Up to 512 entries.
- IP4_TCP_UDP: Up to 512 entries.
- IP4_OTHER: Up to 512 entries.
- IP6_STD: Up to 512 entries.
- IP6_TCP_UDP: Up to 256 entries.
- IP6_OTHER: Up to 256 entries.
- OAM: Up to 512 entries.

The following tables provide information about how the IS2 half keys are generated. The first table lists the common fields between the half keys.

TABLE 2-63: IS2 COMMON KEY FIELDS FOR HALF KEYS

Field Name	Bit	Width	Description
Lookup Information			
S2_TYPE_HALF	0	4	0: MAC ETYPE entries 1: MAC LLC entries 2: MAC SNAP entries 3: ARP entries 4: IPv4 TCP/UDP entries 5: IPv4 OTHER entries 6: IPv6 STD entries 7: OAM entries (8: SMAC_SIP6 entries)
FIRST	4	1	Set for first lookup and cleared for second lookup.
Interface Information			
PAG	5	6	Policy association group. Action from VCAP IS1.
RESERVED	11	2	Must be set to don't care.
IGR_PORT_MASK	13	12	Ingress port mask. VCAP generated with one bit set in the mask corresponding to the ingress port.
Tagging and IP Source Guarding Information			
SERVICE_FRM	26	1	Service classified frame. Set if ISDX > 0.
HOST_MATCH	26	1	The action from the SMAC_SIP4 or SMAC_SIP6 look-ups. Used for IP source guarding.
L2_MC	27	1	Set if frame's destination MAC address is a multicast address (bit 40 = 1).
L2_BC	28	1	Set if frame's destination MAC address is the broadcast address (FF-FF-FF-FF-FF-FF).
VLAN_TAGGED	29	1	Set if frame has one or more Q-tags. Independent of port VLAN awareness.
VID	30	12	Classified VID which is the result of the VLAN classification in basic classification and IS1. For frames classified to a service (ISDX > 0), the VID contains the ISDX value.
DEI	42	1	Classified DEI which is the final result of the VLAN classification in basic classification and IS1.
PCP	43	3	Classified PCP which is the final result of the VLAN classification in basic classification and IS1.

TABLE 2-64: IS2 MAC_ETYPE KEY

Field Name	Bit	Width	Description
Layer-2 Information			

TABLE 2-64: IS2 MAC_ETYPE KEY (CONTINUED)

Field Name	Bit	Width	Description
L2_DMAC	46	48	Frame's destination MAC address.
L2_SMAC	94	48	Frame's source MAC address.
ETYPE	142	16	Frame's EtherType. This is the EtherType after up to two VLAN tags.
L2_PAYLOAD0	158	16	Payload bytes 0-1 after the frame's EtherType.
L2_PAYLOAD1	174	8	Payload byte 4 after the frame's EtherType. This is specifically for PTP frames.
L2_PAYLOAD2	182	3	Bits 7, 2, and 1 from payload byte 6 after the frame's EtherType. This is specifically for PTP frames.

TABLE 2-65: IS2 MAC_LLC KEY

Field Name	Bit	Width	Description
Layer-2 Information			
L2_DMAC	46	48	Frame's destination MAC address
L2_SMAC	94	48	Frame's source MAC address
L2_LLC	142	40	LLC header and data after up to two VLAN tags and the type/length field.

TABLE 2-66: IS2 MAC_SNAP KEY

Field Name	Bit	Width	Description
Layer-2 Information			
L2_DMAC	46	48	Frame's destination MAC address
L2_SMAC	94	48	Frame's source MAC address
L2_SNAP	142	40	SNAP header after LLC header (AA-AA-03).

TABLE 2-67: IS2 ARP KEY

Field Name	Bit	Width	Description
Layer-2 Information			
L2_SMAC	46	48	Frame's source MAC address
Layer-3 Information			
ARP_ADDR_SPACE_OK	94	1	Set if hardware address is Ethernet.
ARP_PROTO_SPACE_OK	95	1	Set if protocol address space is IP.
ARP_LEN_OK	96	1	Set if hardware address length = 6 (Ethernet) and IP address length = 4 (IP).
ARP_TARGET_MATCH	97	1	Target hardware address = SMAC (RARP).
ARP_SENDER_MATCH	98	1	Sender hardware address = SMAC (ARP).
ARP_OPCODE_UNKNOWN	99	1	Set if ARP opcode is none of the below are mentioned.
ARP_OPCODE	100	2	0: ARP request 1: ARP reply. 2: RARP request. 3: RARP reply.
L3_IP4_DIP	102	32	Target IPv4 address.

TABLE 2-67: IS2 ARP KEY (CONTINUED)

Field Name	Bit	Width	Description
L3_IP4_SIP	134	32	Sender IPv4 address.
DIP_EQ_SIP	166	1	Set if sender IP address is equal to target IP address.

TABLE 2-68: IS2 IP4_TCP_UDP KEY

Field Name	Bit	Width	Description
Layer-3 and Layer-4 Information			
IP4	46	1	Set if frame is IPv4 frame. IPv6 frames can also use IP4_TCP_UDP entries when IP6_STD entries are disabled.
L3_FRAGMENT	47	1	Set if IP frame is fragmented (More Fragments flag = 1 or Fragments Offset > 0).
L3_FRAG_OFS_GT0	48	1	Set if IP frame is fragmented and it is not the first fragment (Fragments Offset > 0). Such frames do not carry Layer-4 information all Layer-4 information fields in the key are automatically set to don't-care when generating the key.
L3_OPTIONS	49	1	Set if IP frame contains options (IP len > 5). IP options are not skipped nor parsed which implies that Layer-4 information cannot be used. All Layer-4 information fields in the key are automatically set to don't-care when generating the key.
L3_TTL_GT0	50	1	Set if IP TTL is greater than 0.
L3_TOS	51	8	IP TOS field.
L3_IP4_DIP	59	32	IPv4 frames: Destination IPv4 address. IPv6 frames: Source IPv6 address, bits 63:32.
L3_IP4_SIP	91	32	IPv4 frames: Source IPv4 address. IPv6 frames: Source IPv6 address, bit 31:0.
DIP_EQ_SIP	123	1	Set if source IP address is equal to destination IP address. Full addresses are checked, also for IPv6.
TCP	124	1	Set if IP Proto = 6 (TCP).
L4_DPORT	125	16	TCP/UDP destination port.
L4_SPORT	141	16	TCP/UDP source port.
L4_RNG	157	8	Range mask with one bit per range. A bit is set, if the corresponding range is matched. Range types: SPORT, DPORT, SPORT or DPORT, VID, DSCP. Input to range checkers: SPORT, DPORT: From frame VID, DSCP: Classified result from IS1. See Section 2.7.6, Range Checkers .
SPORT_EQ_DPORT	165	1	Set if UDP or TCP source port equals UDP or TCP destination port.
SEQUENCE_EQ0	166	1	TCP: Set if TCP sequence number is 0. PTP over UDP: messageType bit 0.
L4_FIN	167	1	TCP: TCP flag FIN. PTP over UDP: messageType bit 1.
L4_SYN	168	1	TCP: TCP flag SYN. PTP over UDP: messageType bit 2.

TABLE 2-68: IS2 IP4_TCP_UDP KEY (CONTINUED)

Field Name	Bit	Width	Description
L4_RST	169	1	TCP: TCP flag RST. PTP over UDP: messageType bit 3.
L4_PSH	170	1	TCP: TCP flag PSH. PTP over UDP: flagField bit 1 (twoStepFlag).
L4_ACK	171	1	TCP: TCP flag ACK. PTP over UDP: flagField bit 2 (unicastFlag).
L4_URG	172	1	TCP: TCP flag URG. PTP over UDP: flagField bit 7 (reserved).
L4_1588_DOM	173	8	PTP over UDP: domainNumber.
L4_1588_VERSION	181	4	PTP over UDP: version.

Frames with IP options (L3_OPTIONS set to 1 in key) or fragmented frames, which are not the initial fragment (L3_FRAG_OFS_GT0 set to 1 in key), do not carry Layer-4 information. The Layer-4 fields in the key (L4_SPORT, L4_DPORT, L4_RNG, SPORT_EQ_DPORT, SEQUENCE_EQ0, L4_FIN, L4_SYN, L4_RST, L4_PSH, L4_ACK, L4_URG, L4_1588_DOM, and L4_1588_VERSION) are automatically set to don't care.

TABLE 2-69: IS2 IP4_OTHER KEY

Field Name	Bit	Width	Description
Layer-3 Information			
IP4	46	1	Set if frame is IPv4 frame. IPv6 frames can also use IP4_OTHER entries when IP6_STD entries are disabled.
L3_FRAGMENT	47	1	Set if IP frame is fragmented (More Fragments flag = 1 or Fragments Offset > 0)
L3_FRAG_OFS_GT0	48	1	Set if IP frame is fragmented and if it is not the first fragment (Fragments Offset > 0).
L3_OPTIONS	49	1	Set if IPv4 frame contains options (IP len > 5). IP options are not skipped nor parsed, which implies that L3_PAYLOAD contains data from the IP options for IPv4 frames with IP options.
L3_TTL_GT0	50	1	Set if IP TTL is greater than 0.
L3_TOS	51	8	IP TOS field.
L3_IP4_DIP	59	32	IPv4 frames: Destination IPv4 address. IPv6 frames: Source IPv6 address, bits 63:32.
L3_IP4_SIP	91	32	IPv4 frames: Source IPv4 address. IPv6 frames: Source IPv6 address, bit 31:0.
DIP_EQ_SIP	123	1	Set if source IP address is equal to destination IP address. Full addresses are checked, also for IPv6.
L3_PROTO	124	8	IPv4: IP protocol. IPv6: next header.
L3_PAYLOAD	132	56	Bytes 0-6 after IP header.

TABLE 2-70: IS2 IP6_STD KEY

Field Name	Bit	Width	Description
Layer-3 Information			
L3_TTL_GT0	46	1	Set if IP HOPLIMIT is greater than 0.
L3_IP6_SIP	47	128	Frame's source IPv6 address
L3_PROTO	175	8	IPv6 next header.

TABLE 2-71: IS2 OAM KEY

Field Name	Bit	Width	Description
Layer-2 Information			
L2_DMACH	46	48	Frame's destination MAC address
L2_SMACH	94	48	Frame's source MAC address
OAM_MEL_FLAGS	142	7	Encoding of MD level/MEG level (MEL). One bit for each level (lowest level encoded as zero) The following keys can be generated: MEL=0: 0x0000000 MEL=1: 0x0000001 MEL=2: 0x0000011 MEL=3: 0x0000111 MEL=4: 0x0001111 MEL=5: 0x0011111 MEL=6: 0x0111111 MEL=7: 0x1111111 Together with the mask, the following kinds of rules may be created: Exact match. Fx. MEL=2: 0x0000011 Below. Fx. MEL<=4: 0x000XXXX Above. Fx. MEL>=5: 0xXX11111 Between. Fx. 3<= MEL<=5: 0x00XX111, where 'X' means 'don't care'.
OAM_VER	149	5	Frame's OAM version.
OAM_OPCODE	154	8	Frame's OAM opcode.
OAM_FLAGS	162	8	Frame's OAM flags.
OAM_MEPID	170	16	CCM frame's OAM MEP ID.
OAM_C-CM_CNTS_EQ0	186	1	Flag indicating whether dual-ended loss measurement counters in CCM frames are used or not. The flag OAM_C-CM_CNTS_EQ0 is set if the counters TxFCf, RxFCb, and TxFCb are set to all zeros.

The following tables provide information about how the IS2 full keys are generated. The first table lists the common fields between the full keys.

TABLE 2-72: IS2 COMMON KEY FIELDS FOR FULL KEYS

Field Name	Bit	Width	Description
Lookup Information			
S2_TYPE_FULL	0	2	0: IP6_TCP_UDP entries. 1: IP6_OTHER entries. 2: Custom entries.
FIRST	2	1	Set for first lookup and cleared for second lookup.
Interface Information			
PAG	3	6	Policy association group. Action from VCAP IS1.
RESERVED	9	2	Must be set to don't care.
IGR_PORT_MASK	11	12	Ingress port mask. VCAP generated with one bit set in the mask corresponding to the ingress port.
Tagging and IP Source Guarding Information			
SERVICE_FRM	23	1	Service classified frame. Set if ISDX > 0.
HOST_MATCH	24	1	The action from the SMACH_SIP4 or SMACH_SIP6 lookups. Used for IP source guarding.

TABLE 2-72: IS2 COMMON KEY FIELDS FOR FULL KEYS (CONTINUED)

Field Name	Bit	Width	Description
L2_MC	25	1	Set if frame's destination MAC address is a multicast address (bit 40 = 1).
L2_BC	26	1	Set if frame's destination MAC address is the broadcast address (FF-FF-FF-FF-FF-FF).
VLAN_TAGGED	27	1	Set if frame has one or more Q-tags. Independent of port VLAN awareness.
VID	28	12	Classified VID which is the result of the VLAN classification in basic classification and IS1. For frames classified to a service (ISDX > 0), the VID contains the ISDX value.
DEI	40	1	Classified DEI, which is the final result of the VLAN classification in basic classification and IS1.
PCP	41	3	Classified PCP, which is the final result of the VLAN classification in basic classification and IS1.

TABLE 2-73: IS2 IP6_TCP_UDP KEY

Field Name	Bit	Width	Description
Layer-3 and Layer-4 Information			
L3_TTL_GT0	44	1	Set if IP HOPLIMIT is greater than 0.
L3_TOS	45	8	IP TOS field.
L3_IP6_DIP	53	128	Destination IPv6 address.
L3_IP6_SIP	181	128	Source IPv6 address.
DIP_EQ_SIP	309	1	Set if source IP address is equal to destination IP address.
TCP	310	1	Set if IP Proto = 6 (TCP).
L4_DPORT	311	16	TCP/UDP destination port.
L4_SPORT	327	16	TCP/UDP source port.
L4_RNG	343	8	Range mask with one bit per range. A bit is set, if the corresponding range is matched. Range types: SPORT, DPORT, SPORT or DPORT, VID, DSCP. Input to range checkers: SPORT, DPORT: From frame VID, DSCP: Classified result from IS1 See Section 2.7.6, Range Checkers .
SPORT_EQ_DPORT	351	1	Set if UDP or TCP source port equals UDP or TCP destination port.
SEQUENCE_EQ0	352	1	TCP: Set if TCP sequence number is 0. PTP over UDP: messageType bit 0.
L4_FIN	353	1	TCP: TCP flag FIN. PTP over UDP: messageType bit 1.
L4_SYN	354	1	TCP: TCP flag SYN. PTP over UDP: messageType bit 2.
L4_RST	355	1	TCP: TCP flag RST. PTP over UDP: messageType bit 3.
L4_PSH	356	1	TCP: TCP flag PSH. PTP over UDP: flagField bit 1 (twoStepFlag).
L4_ACK	357	1	TCP: TCP flag ACK. PTP over UDP: flagField bit 2 (unicastFlag).

TABLE 2-73: IS2 IP6_TCP_UDP KEY (CONTINUED)

Field Name	Bit	Width	Description
L4_URG	358	1	TCP: TCP flag URG. PTP over UDP: flagField bit 7 (reserved).
L4_1588_DOM	359	8	PTP over UDP: domainNumber.
L4_1588_VERSION	367	4	PTP over UDP: version

TABLE 2-74: IS2 IP6_OTHER KEY

Field Name	Bit	Width	Description
Layer-3 Information			
L3_TTL_GT0	44	1	Set if IP HOPLIMIT is greater than 0.
L3_TOS	45	8	IP TOS field.
L3_IP6_DIP	53	128	Destination IPv6 address.
L3_IP6_SIP	181	128	Source IPv6 address.
DIP_EQ_SIP	309	1	Set if source IP address is equal to destination IP address.
L3_PROTO	310	8	IPv6 next header.
L3_PAYLOAD	318	56	Bytes 0-6 after IP header.

TABLE 2-75: IS2 CUSTOM KEY

Field Name	Bit	Width	Description
Custom Information			
CUSTOM_TYPE	44	1	0: CUSTOM_IP - Custom data extracted from byte 34 (corresponds to the first byte after the IPv4 header if the frame is IPv4). 1: CUSTOM_SMAC - Custom data extracted from byte 12 (first byte after SMAC). Note that up to two VLAN tags are removed before extracting the data.
CUSTOM	45	320	Custom extracted data.

IS2 Action Encoding

The VCAP generates an action vector from each of the two IS2 lookups for each frame.

The first IS2 lookup returns a default action vector per ingress port when no entries are matched, and the second IS2 lookup returns a common default action vector when no entries are matched. There are no difference between an action vector from a match and a default action vector.

The following table lists the available fields for the action vector.

TABLE 2-76: IS2 ACTION FIELDS

Action Field	Bit	Width	Description
HIT_ME_ONCE	0	1	Setting this bit to 1 causes the first frame that hits this action where the HIT_CNT counter is zero to be copied to the CPU extraction queue specified in CPU_QU_NUM. The HIT_CNT counter is then incremented and any frames that hit this action later are not copied to the CPU. To re-enable the HIT_ME_ONCE functionality, the HIT_CNT counter must be cleared.

TABLE 2-76: IS2 ACTION FIELDS (CONTINUED)

Action Field	Bit	Width	Description
CPU_COPY_ENA	1	1	Setting this bit to 1 causes all frames that hit this action to be copied to the CPU extraction queue specified in CPU_QU_NUM.
CPU_QU_NUM	2	3	Determines the CPU extraction queue that is used when a frame is copied to the CPU due to a HIT_ME_ONCE or CPU_COPY_ENA action.
MASK_MODE	5	2	Controls how PORT_MASK is applied. 0: No action from PORT_MASK. 1: Permit/deny (PORT_MASK AND'ed with destination set). 2: Policy forwarding (DMAC lookup replaced with PORT_MASK). 3: Redirect (Previous forwarding decisions (classifier, IS1, ISDX, SRC, AGGR, VLAN, and DMAC lookup) are replaced with PORT_MASK). The CPU port is not touched by MASK_MODE.
MIRROR_ENA	7	1	Setting this bit to 1 causes frames to be mirrored to the mirror target port (ANA::MIRRPORPORTS).
LRN_DIS	8	1	Setting this bit to 1 disables learning of frames hitting this action.
POLICE_ENA	9	1	Setting this bit to 1 causes frames that hit this action to be policed by the ACL policer specified in POLICE_IDX. Only applies to the first lookup.
POLICE_IDX	10	11	Selects VCAP policer used when policing frames (POLICE_ENA). Only indexes 384 through 447 are valid.
POLICE_VCAP_ONLY	21	1	Disable policing from service, QoS, and port policers. Only the VCAP policer selected by POLICE_IDX is active.
PORT_MASK	22	11	Port mask applied to the forwarding decision based on MASK_MODE.

TABLE 2-76: IS2 ACTION FIELDS (CONTINUED)

Action Field	Bit	Width	Description
REW_OP	33	9	Rewriter operation command. The following functions are supported: No operation: REW_OP[3:0] = 0. No operation. Special Rewrite: REW_OP[3:0] = 8. Swap the MAC addresses and clear bit 40 in the new SMAC when transmitting the frame. One-step PTP: REW_OP[2:0] = 2. The frame's residence time is added to the correction field in the PTP frame. The following sub-commands can be encoded: REW_OP[7]: Enables Add/subtract mode. REW_OP[5]: Set if egress delay must be added to correction field at egress. REW_OP[4:3]: Configures if ingress delay must be subtracted from the frame's Rx timestamp. This is not applicable for ingress ports in backplane mode. Bits 4:3 = 0: No delay adjustments. Bits 4:3 = 1: Subtract delay ANA:PORT:PTP_DLY1_CFG. Bits 4:3 = 2: Subtract delay ANA:PORT:PTP_DLY2_CFG. Two-step PTP: REW_OP[2:0] = 3. The frame's departure timestamp is saved in the timestamp FIFO queue at egress. REW_OP[8:3] holds the PTP timestamp identifier used by the CPU. Origin PTP: REW_OP[2:0] = 9. The time of day at the frame's departure time is written into the origin Timestamp field in the PTP frame. Unspecified bits must be set to 0.
LM_TX-CNT_DIS	42	1	If set, Up-MEP loss measurement counters are not updated with this frame. Only applies to the second lookup.
RESERVED	43	1	Reserved. Must be set to 0.
ISDX_ENA	44	1	If set, the frame's classified ISDX is replaced with POLICE_IDX[7:0]. The new ISDX is used in the ES0 lookup. ISDX_ENA overrules the settings of POLICE_ENA and POLICE_VCAP_ONLY for this lookup.
ACL_ID	45	6	Logical ID for the entry. This ID is extracted together with the frame in the CPU extraction header. Only applicable to actions with CPU_COPY_ENA or HIT_ME_ONCE set.
HIT_CNT	51	32	A statistics counter that is incremented by one each time the given action is hit.

The two action vectors from the first and second lookups are combined into one action vector, which is applied in the analyzer. For more information, see [Section 2.8.4, Forwarding Engine](#)). The actions are combined as follows:

- **HIT_ME_ONCE, CPU_COPY_ENA, CPU_QU_NUM:**
If any of the two action vectors have HIT_ME_ONCE or CPU_COPY_ENA set, CPU_COPY_ENA is forwarded to the analyzer. The settings in the action vector from second lookup takes precedence with respect to the CPU extraction queue number.
- **MIRROR_ENA:**
If any of the two action vectors have MIRROR_ENA set, MIRROR_ENA is forwarded to the analyzer.
- **LRN_DIS:**
If any of the two action vectors have LRN_DIS set, LRN_DIS is forwarded to the analyzer.
- **REW_OP:**
The settings in the action vector from the second lookup takes precedence unless the second lookup returns REW_OP[2:0] = 0.

VSC7416-01

- **ACL_ID:**
If both lookups in IS2 hit an entry with CPU_COPY_ENA set, then the resulting value is the addition of the two ACL_ID values. This allow for unique identification of both rules in IS2 through ACL_ID if each lookup manipulates its own part of the ACL_ID. For instance, first lookup returns ACL_ID = 0:15, and second lookup returns ACL_ID = 0, 16, 32, or 48.
- **POLICE_ENA, POLICE_IDX, POLICE_VCAP_ONLY:**
Only applies to actions from the first lookup.

The following table lists the combinations for MASK_MODE and PORT_MASK when combining actions from the first and second lookups.

TABLE 2-77: MASK_MODE AND PORT_MASK COMBINATIONS

First Lookup	Second Lookup			
	No action	Permit/deny	Policy	Redirect
No action	No action	Permit $P^{(1)} = P_2^{(2)}$	Policy $P = P_2$	Redirect $P = P_2$
Permit/deny	Permit $P = P_1^{(3)}$	Permit $P = P_1 \text{ and } P_2$	Policy $P = P_1 \text{ and } P_2$	Redirect $P = P_2$
Policy	Policy $P = P_1$	Policy $P = P_1 \text{ and } P_2$	Policy $P = P_1 \text{ and } P_2$	Redirect $P = P_2$
Redirect	Redirect $P = P_1$	Redirect $P = P_1 \text{ and } P_2$	Redirect $P = P_1 \text{ and } P_2$	Redirect $P = P_2$

1P: Resulting PORT_MASK to analyzer.

2P2: PORT_MASK from second match.


3P1: PORT_MASK from first match.

Policy forwarding for frames matching an IPv4 and IPv6 multicast entry in the MAC table is not possible. Policy forwarding is handled as a permit/deny action for such frames.

SMAC_SIP4 and SMAC_SIP6 Entry Key Encoding

The following illustration shows which entry fields are available for SMAC_SIP4 and SMAC_SIP6 keys in IS2.

FIGURE 2-14: SMAC_SIP ENTRY TYPE OVERVIEW



VCAP IS2 can hold the following number of SMAC_SIP entries.

- SMAC_SIP6: Up to 512 entries.
- SMAC_SIP4: Up to 512 entries.

All IPv6 frames are subject to a SMAC_SIP6 lookup in IS2. The following table lists the SMAC_SIP6 key.

TABLE 2-78: SMAC_SIP6 KEY

Field Name	Bit	Width	Description
Lookup Information			
S2_TYPE_HALF	0	4	8: SMAC_SIP4.
Interface Information			
IGR_PORT	4	4	The port number where the frame was received (0-11).
Layer-2 Information			
L2_SMAC	8	48	Frame's source MAC address.

TABLE 2-78: SMAC_SIP6 KEY (CONTINUED)

Field Name	Bit	Width	Description
Layer-3 Information			
L3_IP6_SIP	56	128	Frame's source IPv6 address.

All IPv4 frames are subject to a SMAC_SIP4 lookup in IS2. The following table lists the SMAC_SIP4 key.

TABLE 2-79: SMAC_SIP4 KEY

Field Name	Bit	Width	Description
Interface Information			
IGR_PORT	0	4	The port number where the frame was received (0-11).
Layer-2 Information			
L2_SMAC	4	48	Frame's source MAC address.
Layer-3 Information			
L3_IP4_SIP	52	32	Frame's source IPv4 address.

SMAC_SIP4 and SMAC_SIP6 Action Encoding

The VCAP generates an action vector from the SMAC_SIP4 or SMAC_SIP6 lookup if there is a match. There is no default action vector if no match.

The following table lists the available fields for the action vector from both the SMAC_SIP4 and SMAC_SIP6 lookups.

TABLE 2-80: SMAC_SIP4 AND SMAC_SIP6 ACTION FIELDS

Action Field	Bit	Width	Description
CPU_COPY_ENA	1	1	Setting this bit to 1 causes all frames that hit this action to be copied to the CPU extraction queue specified in CPU_QU_NUM.
CPU_QU_NUM	2	3	Determines the CPU extraction queue that is used when a frame is copied to the CPU due to a HIT_ME_ONCE or CPU_COPY_ENA action.
FWD_KILL_ENA	5	1	Setting this bit to 1 denies forwarding of the frame forwarding to any front port. The frame can still be copied to the CPU by other actions.
HOST_MATCH	6	1	Used for IP source guarding. If set, it signals that the host is a valid (for instance a valid combination of source MAC address and source IP address). HOST_MATCH is input to the IS2 keys.
HIT_STICKY	7	1	If set, a frame has matched against the associated entry.

The HOST_MATCH flag is used as input into the IS2 keys. This enables further handling of frames either matching or not matching a host pair in the SMAC_SIP table.

2.7.5 VCAP ES0

This section provides information about the ES0 key and associated actions.

2.7.5.1 ES0 Entry Key Encoding

All frames are subject to one ES0 lookup per destination port. The key in ES0 is generated based on the frame's service classification. If ISDX_ES0_KEY_ENA field in the ISDX table is set for the frame's service, then ES0.VID contains the frame's ISDX instead of the classified VID. Whether the ISDX or VID is used in the key is indicated through ES0.KEY_ISDX so that it is possible to distinguishes VLAN entries and service entries.

VCAP ES0 can hold 1024 ES0 entries.

The following table lists the ES0 key.

TABLE 2-81: ES0 KEY

Field Name	Bit	Width	Description
Interface Information			
EGR_PORT	0	4	The port number where the frame is transmitted (0-11).
IGR_PORT	4	4	The port number where the frame was received (0-11).
SERVICE_FRM	8	1	Service classified frame. Set if ISDX > 0.
KEY_ISDX	9	1	If set, the field VID contains the ISDX. KEY_ISDX is selected by the ISDX table (ANA::ISDXTIDX.ISDX_ES0_KEY_ENA).
L2_MC	10	1	Set if frame's destination MAC address is a multicast address (bit 40 = 1).
L2_BC	11	1	Set if frame's destination MAC address is the broadcast address (FF-FF-FF-FF-FF-FF).
Tagging Information			
VID	12	12	Classified VID that is the result of the VLAN classification in basic classification and IS1. If KEY_ISDX is set, the VID contains the ISDX value.
DP	24	1	Frame's drop precedence (DP) level after policing.
PCP	25	3	Classified PCP that is the final result of the VLAN classification in basic classification and IS1.

2.7.5.2 ES0 Action Encoding

The VCAP generates one action vector from the ES0 lookup. The lookup returns a default action vector per egress port when no entries are matched. There are no difference between an action vector from a match and a default action vector.

The following table lists the available action fields for ES0. For more information about how the actions are applied to the VLAN manipulations, see [Section 2.12.1, VLAN Editing](#).

TABLE 2-82: ES0 ACTION FIELDS

Action Field	Bit	Width	Description
PUSH_OUTER_TAG	0	2	Controls outer tagging. 0: No ES0 tag A: Port tag is allowed if enabled on port. 1: ES0 tag A: Push ES0 tag A. No port tag. 2: Force port tag: Always push port tag. No ES0 tag A. 3: Force untag: Never push port tag or ES0 tag A.
PUSH_INNER_TAG	2	1	Controls inner tagging. 0: Do not push ES0 tag B as inner tag. 1: Push ES0 tag B as inner tag.
TAG_A_TPID_SEL	3	2	Selects TPID for ES0 tag A: 0: 0x8100. 1: 0x88A8. 2: Custom (REW:PORT:PORT_VLAN_CFG.PORT_TPID). 3: If IFH.TAG_TYPE = 0 then 0x8100 else custom.
TAG_A_VID_SEL	5	1	Selects VID for ES0 tag A. 0: Classified VID + VID_A_VAL. 1: VID_A_VAL.
TAG_A_PCP_SEL	6	2	Selects PCP for ES0 tag A. 0: Classified PCP. 1: PCP_A_VAL. 2: DP and QoS mapped to PCP (per port table). 3: QoS class.

TABLE 2-82: ES0 ACTION FIELDS (CONTINUED)

Action Field	Bit	Width	Description
TAG_A_DEI_SEL	8	2	Selects PCP for ES0 tag A. 0: Classified DEI. 1: DEI_A_VAL. 2: DP and QoS mapped to PCP (per port table). 3: DP.
TAG_B_TPID_SEL	10	2	Selects TPID for ES0 tag B. 0: 0x8100. 1: 0x88A8. 2: Custom (REW:PORT:PORT_VLAN_CFG.PORT_TPID). 3: If IFH.TAG_TYPE = 0 then 0x8100 else custom.
TAG_B_VID_SEL	12	1	Selects VID for ES0 tag B. 0: Classified VID + VID_B_VAL. 1: VID_B_VAL.
TAG_B_PCP_SEL	13	2	Selects PCP for ES0 tag B. 0: Classified PCP. 1: PCP_B_VAL. 2: DP and QoS mapped to PCP (per port table). 3: QoS class.
TAG_B_DEI_SEL	15	2	Selects PCP for ES0 tag B. 0: Classified DEI. 1: DEI_B_VAL. 2: DP and QoS mapped to PCP (per port table). 3: DP.
VID_A_VAL	17	12	VID used in ES0 tag A. See TAG_A_VID_SEL.
PCP_A_VAL	29	3	PCP used in ES0 tag A. See TAG_A_PCP_SEL.
DEI_A_VAL	32	1	DEI used in ES0 tag A. See TAG_A_DEI_SEL.
VID_B_VAL	33	12	VID used in ES0 tag B. See TAG_B_VID_SEL.
PCP_B_VAL	45	3	PCP used in ES0 tag B. See TAG_B_PCP_SEL.
DEI_B_VAL	48	1	DEI used in ES0 tag B. See TAG_B_DEI_SEL.
ENCAP_ID	49	10	Index into the MPLS encapsulation table (SYS::ENCAP_CTRL). Setting ENCAP_ID to zero disables pushing of MPLS encapsulation.
ENCAP_LEN	59	3	Configures how much of the MPLS encapsulation entry pointed to by ENCAP_ID is pushed. 0: No MPLS push. 1: Does not apply. 2: Push 14 bytes (B-DMAC, B-SMAC, ETYPE). 3: Push 18 bytes (B-DMAC, B-SMAC, ETYPE, label or B-DMAC, B-SMAC, B-VLAN, ETYPE). 4: Push 22 bytes (extra label or control word). 5: Push 26 bytes (extra label or control word). 6: Push 30 bytes (extra label or control word). 7: Push 34 bytes (Maximum push: B-DMAC, B-SMAC, B-VLAN, ETYPE, 3x labels, control word).
MPOP_CNT	62	3	If MPOP_CNT > 0, IFH.MPLS_POP is overruled with MPOP_CNT, which has the same encoding as IFH.MPLS_POP. See Table 2-142 .
ESDX	65	10	Egress service index. Used to index egress service counter set (REW::STAT_CFG). Only indexes 0 through 255 are valid.

TABLE 2-82: ES0 ACTION FIELDS (CONTINUED)

Action Field	Bit	Width	Description
OAM_MEP_IDX_VLD	75	1	Setting this bit to 1 causes the OAM_MEP_IDX to be used.
OAM_MEP_IDX	76	6	OAM MEP index.
IPT_CFG	82	2	Configures protection configuration for ISDX. 0: No protection. 1: Working path: Discard the frame if the path group pointed to by PPT_IDX uses protected path. 2: Protected path: Discard the frame if the path group pointed to by PPT_IDX uses working path. 3: Reserved.
PPT_IDX	84	7	Path protection group. This index points to the path protection entry in the egress path protection table (REW::PPT) used for this ISDX. Only indexes 0 through 22 are valid.
HIT_STICKY	91	1	If set, a frame has matched the associated entry.

2.7.6 RANGE CHECKERS

The following table lists the registers associated with configuring range checkers.

TABLE 2-83: RANGE CHECKER CONFIGURATION

Register	Description	Replication
ANA::VCAP_RNG_TYPE_CFG	Configuration of the range checker types.	None
ANA::VCAP_RNG_VAL_CFG	Configuration of range start and end points.	None

All IS1 entries, together with the IP4_TCP_UDP and IP6_TCP_UDP entries in IS2, contain eight range checker flags (L4_RNG), which are matched against an 8-bit range key. The range key is generated for each frame based on the extracted frame data and the configuration in ANA::VCAP_RNG_TYPE_CFG and ANA::VCAP_RNG_VAL_CFG. Each of the eight range checkers can be configured to one of the following range types:

- TCP/UDP destination port range
Input to the range is the frame's TCP/UDP destination port number.
- TCP/UDP source port range
Input to the range is the frame's TCP/UDP source port number.
- TCP/UDP source and destination ports range. Range is matched if either source or destination port is within range.
Input to the range are the frame's TCP/UDP source and destination port numbers.
- VID range
IS1: Input to the range is the frame's VID or the port VID if the frame is untagged.
IS2: Input to the range is the classified VID.
- DSCP range
IS1: Input to the range is the translated DSCP value from basic classification.
IS2: Input to the range is the classified DSCP value.

For IS2, the range key is only applicable to TCP/UDP frames. For IS1, the range key is generated for any frame types. Specific range types not applicable to a certain frame type (for example, TCP/UDP port ranges for IPv4 Other frames) must be set to don't care in entries the frame type can match.

Range start points and range end points are configured in ANA::VCAP_RNG_VAL_CFG.

2.7.7 VCAP CONFIGURATION

This section provides information about how the VCAPs (IS0, IS1, IS2, and ES0) are configured.

Each VCAP implements its own set of the registers listed in the following two tables.

Entries in a VCAP are accessed indirectly through an entry and action cache. The cache is accessible using the VCAP configuration registers listed in following table. As shown in the following illustration, an entry in the VCAP consists of a TCAM entry and an associated action and counter entry.

The following table lists the registers associated with VCAP configuration.

TABLE 2-84: VCAP CONFIGURATION REGISTERS


Register	Description	Replication
VCAP_UPDATE_CTRL	General configuration register	None
VCAP_MV_CFG	Move configuration	None
VCAP_ENTRY_DAT	Entry data cache	32
VCAP_MASK_DAT	Entry mask cache	32
VCAP_ACTION_DAT	Action data cache	32
VCAP_CNT_DAT	Counter data cache	32
VCAP_TG_DAT	Type-Group cache	None
VCAP_STICKY	Sticky-bit indications	None

Each VCAP has defined various constants and are accessed using the registers listed in the following table.

TABLE 2-85: VCAP CONSTANTS

Register	Description	Replication
ENTRY_WIDTH	Width of entry field	None
ENTRY_CNT	Number of entries	None
ENTRY_SWCNT	Number of subwords	None
ENTRY_TG_WIDTH	Width of type-group field	None
ACTION_DEF_CNT	Number of default actions	None
ACTION_WIDTH	Width of action field	None
CNT_WIDTH	Width of counter field	None

FIGURE 2-15: VCAP CONFIGURATION OVERVIEW



A TCAM entry consists of entry data, an entry mask, and a type-group value. The type-group value is used internally to differentiate between VCAP lookups of different subword sizes. Each TCAM entry has an associated action entry. In addition, the action RAM has an entry for each of the default actions in the VCAP. The entries in the action RAM consists of action data and a counter value.

For a write access, the TCAM and action entry must be written to the cache and then copied from the cache to the TCAM/RAM. For a read access, the TCAM and action entry must first be retrieved from the TCAM/RAM before being read from the cache. When a read or write operation is initiated, it is possible to individually select if the operation should be applied to the TCAM and/or action RAM. When data is moved between the cache and the TCAM/RAM, it is always the entire entry that is moved. For VCAPs with several subwords per entry, this must be taken into account if only a single subword of a TCAM entry should be updated. To modify a single subword, the entire TCAM entry must be read, then the subword must be modified in the cache, and finally the entry must be written back to the TCAM.

VSC7416-01

The cache can hold only one VCAP entry (TCAM and action entry) at a time. After the TCAM and action entry are written to the cache, the cache must be copied to the TCAM and RAM before new entries can be written to the cache.

The following table lists the different parameters for the four VCAPs available in VSC7416-01. The parameters are needed to format the data to be written to the cache. The parameters can also be read in the registers listed in [Table 2-85](#).

TABLE 2-86: VCAP PARAMETERS

VCAP	Entry Width	Number of Entries	Action Width	Number of Default Actions	Counter Width	Subwords	Type-Group Width
IS0	156	192	179	2 ⁽¹⁾	2x1 (sticky)	2	2
IS1	384	256	320	1	4x1 (sticky)	4	2
IS2	384	256	103	13	4x32	4	3
ES0	29	1024	91	11	1 (sticky)	1	1

¹Default actions start at address 384.

2.7.7.1 Creating a VCAP Entry in the Cache

Before a VCAP entry can be created in the TCAM and RAM, the entry must be created in the cache. The cache is accessed through the following 32-bit registers.

- VCAP_ENTRY_DAT
- VCAP_MASK_DAT
- VCAP_ACTION_DAT
- VCAP_CNT_DAT
- VCAP_TG_DAT

Each of the cache registers is replicated 32 times; however, only the bits used by the VCAP are mapped to physical registers. For example, for VCAP IS1, only the lowest 384 bits of VCAP_ENTRY_DAT and VCAP_MASK_DAT is mapped to physical registers. As mentioned previously, a VCAP entry consists of a TCAM entry and an action entry.

The TCAM entry consists of entry data, mask data, a type value, and a type-group value. The entry data prefixed with the type value is written to VCAP_ENTRY_DATA. The mask data is written to VCAP_MASK_DATA, and the type-group value is written to VCAP_TG_DAT. The type and type-group values are used internally in the VCAP to distinguish between the different entry types. The following table lists the type and type-group value for each of the entry types.

TABLE 2-87: ENTRY, TYPE, AND TYPE-GROUP PARAMETERS


VCAP	Entry Type	Entry Width	Subword s	Type Value [width in ()]	Type-Group Value [width in ()]
IS0	MLL	116	1	Not used (0)	1 (2)
IS0	MLBS	76	2	Not used (0)	2 (2)
IS1	S1_NORMAL	180	2	0 (1)	2 (2)
IS1	S1_5TUPLE_IP4	187	2	1 (1)	2 (2)
IS1	S1_NOR-MAL_IP6	374	1	0 (2)	1 (2)
IS1	S1_7TUPLE	373	1	1 (2)	1 (2)
IS1	S1_5TUPLE_IP6	374	1	2 (2)	1 (2)
IS1	S1_DBL_VID	93	4	Not used (0)	3 (2)
IS2	MAC_ETYPE	181	2	0 (4)	2 (2)
IS2	MAC_LCC	178	2	1 (4)	2 (2)
IS2	MAC_SNAP	178	2	2 (4)	2 (2)
IS2	ARP	163	2	3 (4)	2 (2)

TABLE 2-87: ENTRY, TYPE, AND TYPE-GROUP PARAMETERS (CONTINUED)

VCAP	Entry Type	Entry Width	Subword s	Type Value [width in ()]	Type-Group Value [width in ()]
IS2	IP4_TCP_UCP	181	2	4 (4)	2 (2)
IS2	IP4_OTHER	184	2	5 (4)	2 (2)
IS2	IP6_STD	179	2	6 (4)	2 (2)
IS2	OAM	183	2	7 (4)	2 (2)
IS2	IP6_TCP_UDP	369	1	0 (2)	1 (2)
IS2	IP6_OTHER	372	1	1 (2)	1 (2)
IS2	CUSTOM	363	1	2 (2)	1 (2)
IS2	SMAC_SIP4	84	4	Not used (0)	3 (2)
IS2	SMAC_SIP6	180	2	8 (4)	2 (2)
ES0	VID	28	1	Not used (0)	1 (1)

Note that the type value is not used for all entry types. If the type value is not used for an entry type, write the entry data from bit 0 of VCAP_ENTRY_DAT.


As an example of how a TCAM entry is laid out in the cache register, the following illustration shows a TCAM entry of the IP6_TCP_UDP entry type for the VCAP IS2.

FIGURE 2-16: ENTRY LAYOUT IN REGISTER EXAMPLE

Generally, the type value must never be masked. However, by masking the type bits a lookup in the VCAP is able to match several different entry types. For example, the IS2 entry types MAC_ETYPE and MAC_LLC have the binary type values 0000 and 0001, respectively. By masking bit 0, a lookup is able to match both entry types.

The entry type used in the preceding example only has one subword per entry in the TCAM. Creating a TCAM entry with an entry type that has several subwords per TCAM entry is a little more complicated. As shown in the following example, the ARP entry type of the VCAP IS2 is used. The ARP entry type has two subwords per TCAM entry. From [Table 2-87](#), it can be seen that the ARP entry type has a width of 163 bits per subword. A row in the IS2 TCAM is 384 bits wide. For more information, see [Table 2-86](#). Each subword is assigned to half a TCAM row; that is, subword 0 is assigned to bits 0-187 and subword 1 is assigned to bits 188-375. Because the ARP entry only is 167 bits wide, there are 21 unused bits for each subword, as shown in the following illustration. The layout for VCAP_MASK_DAT is similar to VCAP_ENTRY_DAT. In addition, a type-group value is associated to each subword. The type-group values are laid out back-to-back in VCAP_TG_DAT as shown.

FIGURE 2-17: ENTRY LAYOUT IN REGISTER USING SUBWORDS EXAMPLE



To invalidate an entry in the TCAM (so a lookup never matches the entry), set the type-group for the entry to 0. If there are more subwords in the entry, each subword can be individually invalidated by setting its corresponding type-group value to 0.

The action entry is written to VCAP_ACTION_DAT. Similar to an entry data, an action entry also has a prefixed type value. The following table lists the parameters for the different action types available in VCAPs.

TABLE 2-88: ACTION AND TYPE FIELD PARAMETERS


VCAP	Action Type	Action Width	Subwords	Type Value [width in ()]
IS0	MLL_HIT	67	1	0 (1)
IS0	MLBS_HIT	89	2	1 (1)
IS1	S1	80	4	Not used (0)
IS2	BASE_TYPE	51	2	0 (1)
IS2	SMAC_SIP	6	4	1 (1)
ES0	VID	91	1	Not used (0)

An action that is associated with an entry type with several subwords per entry has an equal number of subwords. For actions with several subwords, the subwords are simply concatenated together.

The following illustration shows the action layout in the VCAP_ACTION_DAT register for an SMAC_SIP action entry. The SMAC_SIP has four subwords per row.


The following illustration shows the action layout in the VCAP_ACTION_DAT register for an MLL_HIT and an SMAC_SIP action entry. The MLL_HIT action entry has one subword per row, and the SMAC_SIP has four subwords per row.

FIGURE 2-18: ACTION LAYOUT IN REGISTER EXAMPLE



The counter value associated to the action is written to VCAP_CNT_DAT. VCAP_CNT_DAT contains a counter value for each subword in the TCAM entry. For action entries, the counter values for each subword are simply concatenated together. The counter layout for the VCAP_CNT_DAT register in the ISO VCAP is shown in the following illustration. The VCAP ISO features a 1-bit sticky counter with two subwords.

FIGURE 2-19: COUNTER LAYOUT IN REGISTER EXAMPLE



2.7.7.2 Copying Entries Between Cache and TCAM/RAM

When an entry and associated action is created in the cache, the data in the cache must be copied to a given address in the TCAM and RAM using the VCAP_UPDATE_CTRL register. Use the following procedure.

1. Set VCAP_UPDATE_CTRL.UPDATE_CMD to copy from cache to TCAM/RAM.
2. Set the address for the entry in VCAP_UPDATE_CTRL.UPDATE_ADDR.
3. Set VCAP_UPDATE_CTRL.UPDATE_SHOT to initiate the copy operation. The bit is cleared by hardware when the operation is finished.

Initiating another operation before the UPDATE_SHOT field is cleared is not allowed. The delay between setting the UPDATE_SHOT field and the clearing of that field depends on the type of operation and the traffic load on the VCAP.

By setting the fields UPDATE_ENTRY_DIS, UPDATE_ACTION_DIS, and/or UPDATE_CNT_DIS in the VCAP_UPDATE_CTRL register the writing of the TCAM, action, and/or the counter entry can be disabled.

Copying a VCAP entry from the TCAM/RAM to the cache is done in a similar fashion by setting VCAP_UPDATE_CTRL.UPDATE_CMD to copy from TCAM/RAM to the cache. Note that due to internal mapping of the entry data and mask data, the values that are read back from the TCAM cannot always match with the values that were originally written to the TCAM. The internal mapping that happens is listed in the following table. There are differences, because a masked 1 is read back as a masked 0, which is functionally the same.

TABLE 2-89: INTERNAL MAPPING OF ENTRY AND MASK

Written Entry	Written Mask	Description	Read Entry	Read Mask
0	0	Match-0	0	0
0	1	Match-Any	0	1
1	0	Match-1	1	0
1	1	Match-Any	0	1

If an entry match is not found during a lookup for a given frame, a default action is selected by the VCAP. Default actions and counter values are copied between the cache and the action RAM similar to a regular VCAP entry. The default actions are stored in the RAM right below the last regular action entry; for example, VCAP IS2 has 256 regular entries, so the first default action in VCAP IS2 is stored at address 256, the second at address 257, and so on. For more information about the number of regular VCAP entries in each VCAP, see [Table 2-86](#). When a default action is copied from the cache to the RAM, VCAP_UPDATE_CTRL.UPDATE_ENTRY_DIS must be set to disable the update of the TCAM. If updating of the TCAM is not disabled, the operation may overwrite entries in the TCAM.

The cache can be cleared by setting VCAP_UPDATE_CTRL.CLEAR_CACHE. This sets all replications of VCAP_ENTRY_DAT, VCAP_MASK_DAT, VCAP_ACTION_DAT, VCAP_CNT_DAT, and VCAP_TG_DAT to zeros. The CLEAR_CACHE field is automatically cleared by hardware when the cache is cleared.

2.7.8 ADVANCED VCAP OPERATIONS

The VCAP supports a number of advanced operations that allow easy moving and removal of entries and actions during frame traffic.

2.7.8.1 Moving a Block of Entries and Actions

A number of entries and actions can be moved up or down by several positions in the TCAM and RAM. This is done using the VCAP_UPDATE_CTRL and VCAP_MV_CFG registers.


A Move operation is performed by:

- Setting VCAP_UPDATE_CTRL.UPDATE_ADDR equal to the address of the entry with the lowest address, among the entries that must be moved.
- Setting VCAP_MV_CFG.MV_SIZE to the number of entries that must be moved; $n + 1$ entries are moved.
- Setting VCAP_MV_CFG.MV_NUM_POS to the number of positions the entries must be moved. The entries are moved n positions up or down.
- Setting UPDATE_ENTRY_DIS, UPDATE_ACTION_DIS, and/or UPDATE_CNT_DIS to only move some parts of the VCAP entry.
- Setting VCAP_UPDATE_CTRL.UPDATE_CMD to move up (decreasing addresses) or move down (increasing addresses).
- Initiating the Move operation by setting VCAP_UPDATE_CTRL.VCAP_UPDATE_SHOT.

A new command must not be setup until after the VCAP_UPDATE_CTRL.VCAP_UPDATE_SHOT field has automatically cleared. Also note that the cache is used by the VCAP while a Move operation is being performed. As a result, any value in cache prior to a Move operation is lost, and a write is not permitted to the cache while a Move operation is performed.

The following illustration shows an example of a Move operation.

FIGURE 2-20: MOVE DOWN OPERATION EXAMPLE



A Move operation can be performed hitlessly during frame traffic, that is, all entries and actions are still available during a Move operation, and all hits are counted by the action hit counters. The TCAM entries at the original positions are invalidated after the Move operation is complete.

During heavy frame traffic, it can take some time for a large move operation to complete, because the moving of individual rows are restarted each time a lookup is performed. If it is not important that the hit counters are accurately updated while the move operation is processed, VCAP_UPDATE_CTRL.MV_TRAFFIC_IGN can be set. This prevents the VCAP from restarting moves and consequently, decreases the time it takes for the move operation to complete. It may, however, lead to inaccurate hit counter values. Note that even if MV_TRAFFIC_IGN is set, the VCAP still processes all lookups correctly.

Default actions can also be moved, however, VCAP_UPDATE_CTRL.UPDATE_ENTRY_DIS must be set.

If a row is moved to a negative address (above address 0), the row is effectively deleted. If a block is partly moved above address 0, the block is also only partially deleted. In other words, the rows that are effectively moved to an address below 0 are not deleted. If one or more rows are deleted during a move operation, the sticky bit `VCAP_STICKY.VCAP_ROW_DELETED_STICKY` is set.

2.7.8.2 Initializing A Block of Entries

A block of entries can be set to the value of the cache in a single operation. For example, it can be used to initialize all TCAM, action, and counter entries to a specific value. The block of entries to initialize can also include the default action and counter entries.

To perform an initialization operation:

- Set `VCAP_UPDATE_CTRL.UPDATE_ADDR` equal to the address of the entry with the lowest address, among the entries that should be written.
- Set `VCAP_MV_CFG.MV_SIZE` to the number of entries that must be included in the initialization operation: $n + 1$ entries are included.
- Set `UPDATE_ENTRY_DIS`, `UPDATE_ACTION_DIS`, and/or `UPDATE_CNT_DIS` to select if the TCAM, action RAM, and/or the counter RAM should be excluded from the initialization operation.
- Set `VCAP_UPDATE_CTRL.UPDATE_CMD` to the initialization operation.
- Start the initialization operation by setting `VCAP_UPDATE_CTRL.VCAP_UPDATE_SHOT`.

A new command must not be set up until after the `VCAP_UPDATE_CTRL.VCAP_UPDATE_SHOT` field is automatically cleared neither must the cache be written to before `VCAP_UPDATE_SHOT` is cleared.

2.8 Analyzer

The analyzer module is responsible for a number of tasks:

- Determining the set of destination ports, also known as the forwarding decision, for frames received by port modules. This includes Layer-2 forwarding, CPU-forwarding, mirroring, and SFlow sampling.
- Keeping track of network stations and their MAC addresses through MAC address learning and aging.
- Holding VLAN membership information (configured by CPU) and applying this to the forwarding decision.

The analyzer consists of the following main blocks.

- MAC table
- VLAN table
- ISDX table
- Forwarding Engine

The MAC, VLAN, and ISDX tables are the main databases used by the forwarding engine. The forwarding engine determines the forwarding decision and initiates learning in the MAC table when appropriate.

The analyzer operates on analyzer requests initiated by the port modules. For each received frame, the port module requests the analyzer to determine the forwarding decision. Initially, the analyzer request is directed to the VCAP. The result from the VCAP (the IS2 action) is forwarded to the analyzer along with the original analyzer request. For more information about VCAP, see [Section 2.7, VCAP](#).

The analyzer request contains the following frame information.

- Destination and source MAC addresses.
- Physical port number where the frame was received (referred to as PPORT).
- Logical port number where the frame was received (referred to as LPORT).
By default, LPORT and PPORT are the same. However, when using link aggregation, multiple physical ports map to the same logical port. The LPORT value is configured in `ANA:PORT:PORT_CFG.PORTID_VAL` in the analyzer.
- Virtual port profile number where the frame was received (referred to as VPORT).
By default, VPORT and PPORT are the same. The VPORT points out the configuration set used in the basic classification and in some cases in the analyzer. In MPLS applications, the MPLS-TP classifier can select different VPORT than the default one. The VPORT value is configured in `IS0_ACTION.VPROFILE_IDX` in the MPLS-TP classifier.
- Frame properties derived by the classifier and VCAP IS1:
 - Classified VID
 - Link aggregation code

Basic CPU forwarding
CPU forwarding for special frame types determined by the classifier
Service index ISDX
MPLS B-domain port mask (only applicable for MPLS applications)

Based on this information, the analyzer determines an analyzer reply, which is returned to the ingress port modules. The analyzer reply contains:

- The forwarding decision (referred to as DEST). This mask contains 11 bits, 1 bit for each front port. DEST does not include the CPU port. The CPU port receives a copy of the frame if the CPU extraction queue mask, CPUQ, has any bits set.
- The CPU extraction queue mask (referred to as CPUQ). This mask contains 8 bits, 1 bit for each CPU extraction queue.

The terms PPORT, LPORT, VPORT, DEST and CPUQ, as previously defined, are used throughout the remainder of this section.


2.8.1 MAC TABLE

This section provides information about the MAC table block in the analyzer. The following table lists the registers associated with MAC table access.

TABLE 2-90: MAC TABLE ACCESS

Register	Description	Replication
MACHDATA	MAC address and VID when accessing the MAC table.	None
MACLDATA	MAC address when accessing the MAC table.	None
MACTINDX	Direct address into the MAC table for direct read and write.	None
MACACCESS	Flags and command when accessing the MAC table.	None
MACTOPTIONS	Flags when accessing the MAC table.	None
AUTOAGE	Age scan period.	None
AGENCTRL	Controls the default values for new entries in MAC table.	None
ENTRYLIM	Controls limits on number of learned entries per port.	Per port
LEARNDISC	Counts the number of MAC table entries not learned due lack of storage in the MAC table.	None

The analyzer contains a MAC table with 8192 entries containing information about stations learned by the device. The table is organized as a hash table with four buckets and 2048 rows. Each row is indexed by an 11-bit hash value, which is calculated based on the station's (MAC, VID) pair, as shown in the following illustration.

FIGURE 2-21: MAC TABLE ORGANIZATION

The following table lists the fields for each entry in the MAC table.

TABLE 2-91: MAC TABLE ENTRY

Field	Bits	Description
VALID	1	Entry is valid.
MAC	48	The MAC address of the station (primary key).
VID	13	VLAN identifier that the station is learned with (primary key).
DEST_IDX	6	Destination mask index pointing to a destination mask in the destination mask table (PGID entries 0 through 63).
B_DOM	1	MAC table entry is applicable in the B-domain and only frames received on a B-port (ANA::PORT_CFG.USE_B_DOM_TBL) can match such entries. B_DOM is only applicable when MPLS is enabled.
ENTRY_TYPE	2	Entry type: 0: Normal entry subject to aging. 1: Normal entry not subject to aging (locked). 2: IPv4 multicast entry not subject to aging. Full port set is encoded in MAC table entry. 3: IPv6 multicast entry not subject to aging. Full port set is encoded in MAC table entry.
AGED_FLAG	1	Entry is aged once by an age scan. See Section 2.8.1.2, Age Scan .
MAC_CPU_COPY	1	Copy frames from or to this station to the CPU.
SRC_KILL	1	Do not forward frames from this station. Note This flag is not used for destination lookups.
IGNORE_VLAN	1	Do not use the VLAN_PORT_MASK from the VLAN table when forwarding frames to this station.

Entries in the MAC table can be added, deleted, or updated in three ways:

- Hardware-based learning of source MAC addresses (that is, inserting new (MAC, VID) pairs in the MAC table).
- Age scans (setting AGED_FLAG and deleting entries).

- CPU commands (for example, for CPU-based learning).

2.8.1.1 Hardware-Based Learning

The analyzer adds an entry to the MAC table when learning is enabled, and the MAC table does not contain an entry for a received frame's (SMAC, VID). The new entry is formatted as follows:

- VALID is set.
- MAC is set to the frame's SMAC.
- VID is set to the frame's VID prepended with 0.
- ENTRY_TYPE is set to 0 (normal entry subject to aging).
- DEST_IDX is set to the frame's LPORT.
- B_DOM is set to the ingress port's ANA::PORT_CFG.USE_B_DOM_TBL.
- MAC_CPU_COPY is set to AGENCTRL.LEARN_CPU_COPY.
- SRC_KILL is set to AGENCTRL.LEARN_SRC_KILL.
- IGNORE_VLAN is set to AGENCTRL.LEARN_IGNORE_VLAN.
- All other fields are cleared.

When a frame is received from a known station, that is, the MAC table already contains an entry for the received frame's (SMAC, VID), the analyzer can update the entry as follows.

For entries of entry type 0 (unlocked entries):

- The AGED_FLAG is cleared. This implies the station is active, avoiding the deletion of the entry due to aging.
- If the existing entry's DEST_IDX differs from the frame's LPORT, then the entry's DEST_IDX is set to the frame's LPORT. This implies the station has moved to a new port.

For entries of entry type 1 (locked entries):

- The AGED_FLAG is cleared. This implies the station is active.

Entries of entry types 2 and 3 are never updated, because their multicast MAC addresses are never used as source MAC addresses.

For more information about learning, see [Section 2.8.4.7, SMAC Analysis](#).

2.8.1.2 Age Scan

The analyzer scans the MAC table for inactive entries. An age scan is initiated by either a CPU command or automatically performed by the device with a configurable age scan period (AUTOAGE). The age scan checks the flag AGED_FLAG for all entries in the MAC table. If an entry's AGED_FLAG is already set and the entry is of entry type 0, the entry is removed. If the AGED_FLAG is not set, it is set to 1. The flag is cleared when receiving frames from the station identified by the MAC table entry (for more information, see [Section 2.8.1.1, Hardware-Based Learning](#)).

For CPU initiated age scans, it is selectable whether B-domain entries, C-domain entries, or all entries are aged (ANA::ANAGEFIL). For auto-aging, both B-domain and C-domain entries are always subject to aging.

2.8.1.3 CPU Commands

The following table lists the set of commands that a CPU can use to access the MAC table. The MAC table command is written to MACACCESS.MAC_TABLE_CMD. Some commands require the registers MACLDATA, MACHDATA, and MACTINDX to be preloaded before the command is issued. Some commands return information in MACACCESS, MACLDATA, and MACHDATA.

TABLE 2-92: MAC TABLE COMMANDS

Command	Purpose	Use
LEARN	Insert/learn new entry in MAC table. Position given by (MAC, VID)	Configure MAC and VID of the new entry in MACHDATA and MACLDATA. Configure remaining entry fields in MACACCESS. The location in the MAC table is calculated based on (MAC, VID).
FORGET	Delete/unlearn entry given by (MAC, VID).	Configure MAC and VID in MACHDATA and MACLDATA.
AGE	Start age scan.	No preload required. Issue command.

TABLE 2-92: MAC TABLE COMMANDS (CONTINUED)

Command	Purpose	Use
READ	Read entry pointed to by (row, column).	Configure row (0-2047) and column (0-3) of the entry to read in: MACTINDX.INDEX (row). MACTINDX.BUCKET (column). MACACCESS.VALID must be set to 0. When MAC_TABLE_CMD changes to IDLE, MACHDATA, MACLDATA, and MACACCESS contain the information read.
LOOKUP	Lookup entry pointed to by (MAC, VID).	Configure MAC and VID of station to look up in MACHDATA and MACLDATA. MACACCESS.VALID must be 1. Issue a READ command. When MAC_TABLE_CMD changes to IDLE, success of the lookup is indicated by MACACCESS.VALID. If successful, MACACCESS contains the entry information.
WRITE	Write entry, MAC table position given by (row, column).	Configure MAC and VID of the new entry in MACHDATA and MACLDATA. Configure remaining entry fields in MACACCESS. The location in the MAC table is given by row and column in MACTINDX.
INIT	Initialize the table.	No preload required. Issue command.
GET_NEXT	Get the smallest entry in the MAC table numerically larger than the specified (MAC, VID). The VID and MAC are evaluated as a 60-bit number with the VID being most significant.	Configure MAC and VID of the starting point for the search in MACHDATA and MACLDATA. MACACCESS.B_DOM selects whether B-domain or C-domain entries are searched. When MAC_TABLE_CMD changes to IDLE, success of the search is indicated by MACACCESS.VALID. If successful, MACHDATA, MACLDATA, and MACACCESS contain the information read.
IDLE	Indicate that MAC table is ready for new command.	No preload required.

2.8.1.4 Known Multicasts

From a CPU, entries can be added to the MAC table with any content. This makes it possible to add a known multicast address with multiple destination ports:

- Set the MAC and VID in MACHDATA and MACLDATA.
- Set MACACCESS.ENTRY_TYPE = 1, because this is not an entry subject to aging.
- Set MACACCESS.AGED_FLAG to 0.
- Set MACACCESS.DEST_IDX to an unused value.
- Set the destination mask in the destination mask table pointed to by DEST_IDX to the desired ports.

Example All frames in VLAN 12 with MAC address 0x010000112233 are to be forwarded to ports 8, 9, and 12.

This is done by inserting the following entry in the MAC table:

```
VID = 12
MAC = 0x010000112233
ENTRY_TYPE = 1
VALID = 1
AGED_FLAG = 0
DEST_IDX = 40
```

and configuring the destination mask table PGID[40] = 0x1300.

IPv4 and IPv6 multicast entries can be programmed differently without using the destination mask table. This is described in the following subsection.

2.8.1.5 IPv4 Multicast Entries

MAC table entries with the ENTRY_TYPE = 2 settings are interpreted as IPv4 multicast entries.

IPv4 multicasts entries match IPv4 frames, which are classified to the specified VID, and which have DMAC = 0x01005Exxxxxx, where xxxxxx is the lower 24 bits of the MAC address in the entry.

Instead of a lookup in the destination mask table (PGID), the destination set is programmed as part of the entry MAC address. This is shown in the following table.

TABLE 2-93: IPV4 MULTICAST DESTINATION MASK

Destination Ports	Record Bit Field
Ports 10-0	MAC[34-24]

Example All IPv4 multicast frames in VLAN 12 with MAC 01005E112233 are to be forwarded to ports 3, 8, and 9. This is done by inserting the following entry in the MAC table entry:

VALID = 1
VID = 12
MAC = 0x000308112233
ENTRY_TYPE = 2
DEST_IDX = 0

2.8.1.6 IPv6 Multicast Entries

MAC table entries with the ENTRY_TYPE = 3 settings are interpreted as IPv6 multicast entries. IPv6 multicasts entries match IPv6 frames, which are classified to the specified VID, and which have DMAC=0x3333xxxxxxx, where xxxxxxxx is the lower 32 bits of the MAC address in the entry.

Instead of a lookup in the destination mask table (PGID), the destination set is programmed as part of the entry MAC address. This is shown in the following table.

TABLE 2-94: IPV6 MULTICAST DESTINATION MASK

Destination Ports	Record Bit Field
Ports 10 through 0	MAC [42-32]

Example All IPv6 multicast frames in VLAN 12 with MAC 333300112233 are to be forwarded to ports 3, 8, and 9.

This is done by inserting the following entry in the MAC table entry:

VID = 12
MAC = 0x030800112233
ENTRY_TYPE = 3
VALID

1
DEST_IDX = 0

Port, VLAN, and Domain Filter

The following table lists the registers associated with the ageing filter.

TABLE 2-95: VID/PORT/DOMAIN FILTERS

Register	Description	Replication
ANAGEFIL	Port, VLAN, and domain filter for limiting the target for aging and search operations on MAC table.	None

The ANAGEFIL register can be used to only hit specific VLANs, ports, or B-domains/C-domains when doing certain operations. If the filter is enabled, it affects the Manual age scan command (MACACCESS.MAC_TABLE_CMD = AGE).

The GET_NEXT MAC table command. For more information, see [Section 2.8.1.3, CPU Commands](#).

Note: The B-domain/C-domain filter does not apply to GET_NEXT. When initiating GET_NEXT, either a B-domain or a C-domain search is selected through MACACCESS.B_DOM. It is not possible to search both B-domain and C-domain entries at the same time.

When two or more filters are enabled at the same time, for example, port and domain, all conditions must be fulfilled before the operation (aging, GET_NEXT) is carried out.

2.8.1.7 Shared VLAN Learning

The following table lists the location of the Filter Identifier (FID) used for shared VLAN learning.

TABLE 2-96: FID DEFINITION REGISTERS

Register	Description	Replication
IS1_ACTION.FID_SEL	Specifies the use of IS1_ACTION.FID_VAL for the DMAC lookup, the SMAC lookup, or for both lookups.	Per IS1 entry
IS1_ACTION.FID_VAL	FID value used when FID_SEL>0. This FID takes precedence.	Per IS1 entry
ANA::PORT_CFG.USE_B_DOM_TBL	Enables a port as a B-port. B-ports use FID_MAP.FID_B_VAL for VID-to-FID mappings.	Per virtual port profile
ANA::FID_MAP.FID_C_VAL	VID-to-FID mapping table for C-ports. 64 FIDs are programmable. FID_C_VAL=0 effectively disables the mapping implying that FID_C_VAL=VID.	Per VID
ANA::FID_MAP.FID_B_VAL	VID-to-FID mapping table for B-ports. 64 FIDs are programmable. FID_B_VAL=0 effectively disables the mapping implying that FID_B_VAL=VID.	Per VID
AGENCTRL.FID_MASK	Combines multiple VIDs in the MAC table.	None

In the default configuration, the device is set up to do Independent VLAN Learning (IVL), that is, MAC addresses are learned separately on each VLAN. The device also supports Shared VLAN Learning (SVL), where a MAC table entry is shared among a group of VLANs. For shared VLAN learning, a MAC address and a Filter Identifier (FID) define each MAC table entry. A set of VIDs then map to the FID.

The device supports shared VLAN learning in three ways:

- Through the IS1 actions FID_SEL and FID_VAL specifying the FID to use.
- Through the per-VID mapping table FID_MAP.
- Through the AGENCTRL.FID_MASK, which controls a general mapping between FID and VIDs.

The IS1 action FID_SEL selects whether to use the FID_VAL for the DMAC lookup, for the SMAC lookup (learning), or for both lookups. If set for a lookup, the FID_VAL replaces the VID when calculating the hash key into the MAC table and when comparing with the entry's VID. If used during the SMAC lookup, new entries are learned using the FID_VAL. If an IS1 action returns a FID_SEL > 0, it overrules the use of the FID mapping table for the frame. In addition, FID_SEL > 0 overrules the use of the FID_MASK for the MAC table lookups specified in FID_SEL.

The FID mapping table, FID_MAP, maps the frame's classified VID to a FID. The mapping depends on whether the frame is received on a B-port or a C-port. If USE_B_DOM_TBL is set then FID_MAP.FID_B_VAL is used, otherwise FID_MAP.FID_C_VAL is used. If the returned FID is larger than 0, then the FID overrules the use of the FID_MASK for both the DMAC and SMAC lookups in the MAC table. Learning is done using the returned FID.

If neither IS1 nor the FID_MAP table have instructed changes to the FID, then the FID_MASK is applied. The 12-bit FID_MASK masks out the corresponding bits in the VID. The FID used for learning and lookup is therefore calculated as $FID = VID \text{ AND } (\text{NOT } FID\_MASK)$. The FID used in the MAC table is 13 bits so the calculated FID is prepended with 0. Bit 13 in the FID in the MAC table is only selectable through the FID_ENA action out of IS1.

All VIDs mapping to the same FID share the same MAC table entries.

Example Configure all MAC table entries to be shared among all VLANs.

This is done by setting FID_MASK to 11111111111.

Example Split the MAC table into two separate databases: one for even VIDs and one for odd VIDs.

This is done by setting FID_MASK to 11111111110.

2.8.1.8 Learn Limit

The following table lists the registers associated with controlling the number of MAC table entries per port.

TABLE 2-97: LEARN LIMIT DEFINITION REGISTERS

Register	Description	Replication
ENTRYLIM	Configures maximum number of unlocked entries in the MAC table per ingress port.	Per port
PORT_CFG.LIMIT_CPU	If set, learn frames exceeding the limit are copied to the CPU.	Per virtual port profile
PORT_CFG.LIMIT_DROP	If set, learn frames exceeding the limit are discarded.	Per virtual port profile
LEARNDISC	The number of MAC table entries that could not be learned due to a lack of storage space.	None

The ENTRYLIM.ENTRYLIM register specifies the maximum number of unlocked entries in the MAC table that a port is allowed to use. Locked and IPMC entries are not taken into account.

After the limit is reached, both auto-learning and CPU-based learning on unlocked entries are denied. A learn frame causing the limit to be exceeded can be copied to the CPU (PORT_CFG.LIMIT_DROP) and the forwarding to other front ports can be denied (PORT_CFG.LIMIT_DROP).

The ENTRYLIM.ENTRYSTAT register holds the current number of entries in the MAC table. MAC table aging and manual removing of entries through the CPU cause the current number to be reduced. If a MAC table entry moves from one port to another port, this also reduces the current number. If the move causes the new port's limit to be exceeded, the entry is denied and removed from the MAC table.

The LEARNDISC counts all events where a MAC table entry is not created or updated due to a learn limit.

2.8.1.9 MAC Table Domains

The following table lists the registers associated with controlling the B-domain in the MAC table.

TABLE 2-98: B-DOMAIN REGISTERS

Register	Description	Replication
ANA::PORT_CFG.USE_B_DOM_TBL	Enables a port as a B-port. B-ports use MAC table entries with MACACCESS.B_DOM set.	Per virtual port profile
MACACCESS.B_DOM	If set, the MAC table entry is a B-domain entry.	None

The MAC table supports both B-domains and C-domains for MPLS-TP applications. Per port profile it is configurable which MAC table domain to use (USE_B_DOM_TBL).

When learning and moving MAC table entries, the ingress port's domain setting controls whether B- or C-domain MAC table entries are affected.

2.8.2 VLAN TABLE

The following table lists the registers associated with the VLAN Table.

TABLE 2-99: VLAN TABLE ACCESS

Register	Description	Replication
VLANTIDX	VID to access, and VLAN flags.	None
VLANACCESS	VLAN port mask for VID and command for access.	None

The analyzer has a VLAN table that contains information about the members of each of the 4096 VLANs. The following table lists fields for each entry in the VLAN table.

TABLE 2-100: FIELDS IN THE VLAN TABLE

Field	Bits	Description
VLAN_PORT_MASK	11	One bit for each port. Set if port is member of VLAN. The CPU port is always a member of all VLANs.
VLAN_MIRROR	1	Mirror frames received in the VLAN. See Section , Mirroring .

TABLE 2-100: FIELDS IN THE VLAN TABLE (CONTINUED)

Field	Bits	Description
VLAN_SRC_CHK	1	VLAN ingress filtering. If set, frames classified to this VLAN are dropped if PPORT is not member of the VLAN.
VLAN_LEARN_DISABLED	1	Disable learning in the VLAN.
VLAN_PRIV_VLAN	1	Set VLAN to private.

By default, all ports are members of all VLANs. This default can be changed through a CPU command. The following table lists the set of commands that a CPU can issue to access the VLAN table. The VLAN table command is written to VLANACCESS.VLAN_TBL_CMD.

TABLE 2-101: VLAN TABLE COMMANDS

Command	Purpose	Use
INIT	Initialize the table	Issue command. When VLAN_TBL_CMD changes to IDLE, initialization has completed and all ports are member of all VLANs. All flags are cleared.
READ	Read VLAN table entry for specific VID.	Configure the VLAN to read from in VLANTIDX.V_INDEX. When VLAN_TBL_CMD changes to IDLE, VLANACCESS and VLANTIDX contain the information read.
WRITE	Write VLAN table entry for specific VID.	Configure the VLAN to write to in VLANTIDX.V_INDEX. Configure the content of the VLAN record in VLANACCESS.VLAN_PORT_MASK VLANTIDX.VLAN_MIRROR VLANTIDX.VLAN_SRC_CHK VLANTIDX.VLAN_LEARN_DISABLED VLANTIDX.VLAN_PRIV_VLAN
IDLE	Indicate that VLAN table is ready for new command	No preload required.

2.8.3 ISDX TABLE

The following table lists the registers associated with the ISDX Table.

TABLE 2-102: ISDX TABLE ACCESS

Register	Description	Replication
ISDXTIDX	ISDX to access, and ISDX flags.	None
ISDXACCESS	ISDX port mask for ISDX and command for access.	None

The analyzer has an ISDX table that contains information about the service handling of each of the 2561024 service indexes. The following table lists fields for each entry in the ISDX table.

TABLE 2-103: FIELDS IN THE ISDX TABLE

Field	Bits	Description
ISDX_PORT_MASK	11	One bit for each port. Set if frames belonging to the service must be forwarded to the port. This enables service-based forwarding.
ISDX_MSTI	6	Associated MST instance.
ISDX_ES0_KEY_ENA	1	Enables ISDX lookup in ES0. Otherwise, a normal VID lookup is done.
ISDX_SDLBI	8	Service dual leaky bucket policer index.

By default, all ports are destinations in all services. This default can be changed through a CPU command. The following table lists the set of commands that a CPU can issue to access the ISDX table. The ISDX table command is written to ISDXACCESS.ISDX_TBL_CMD.

TABLE 2-104: ISDX TABLE COMMANDS

Command	Purpose	Use
INIT	Initialize the table	Issue command. When ISDX_TBL_CMD changes to IDLE, initialization has completed and all ports are destinations in all services. All other fields are set to default values.
READ	Read ISDX table entry for specific ISDX value	Configure the ISDX to read from in ISDXTIDX.ISDX_INDEX. When ISDX_TBL_CMD changes to IDLE, ISDXACCESS and ISDXTIDX contain the information read.
WRITE	Write ISDX table entry for specific ISDX	Configure the ISDX to write to in ISDXTIDX.ISDX_INDEX. Configure the content of the ISDX record in ISDXACCESS.ISDX_PORT_MASK ISDXTIDX.ISDX_MSTI ISDXTIDX.ISDX_ES0_KEY_ENA ISDXTIDX.ISDX_SDLBI
IDLE	Indicate that ISDX table is ready for new command	No preload required.

2.8.4 FORWARDING ENGINE

The analyzer determines the set of ports to which each frame is forwarded, in several configurable steps. The resulting destination port set can include any number of front ports, excluding the CPU port. The CPU port is handled through the CPU extraction queue mask.

The analyzer request from the port modules is passed through all the processing steps of the forwarding engine. As each step is carried out, the destination port set (DEST) and CPU extraction queue mask (CPUQ) are built up.


In addition to the forwarding decision, the analyzer determines which frames are subject to learning (also known as learn frames). Learn frames trigger insertion of a new entry in the MAC table or update of an existing entry. Learning is presented as part of the forwarding, because in some cases, learning changes the normal forwarding of a frame, such as secure learning.

During the processing, the analyzer determines a local frame property. The learning-disabled flag, LRN_DIS is used in the SMAC Learning step:

- If the learning-disabled flag is set, learning based on (SMAC, VID) is disabled.
- If the learning-disabled flag is cleared, learning is conducted according to the configuration in the SMAC learning step.

The following illustration shows the configuration steps in the analyzer.

FIGURE 2-22: ANALYSIS STEPS



DMAC Analysis

During the DMAC analysis step, the (DMAC, VID) pair is looked up in the MAC table to get the first input to the calculation of the destination port set. For more information about the MAC table, see [Section 2.8.1, MAC Table](#).

The following table lists the registers associated with the DMAC analysis step.

TABLE 2-105: DMAC ANALYSIS REGISTERS

Register	Description	Replication
FLOODING.FLD_UNICAST	Index into the PGID table used for flooding of unicast frames.	None
FLOODING.FLD_BROADCAST	Index into the PGID table used for flooding of broadcast frames.	None
FLOODING.FLD_MULTICAST	Index into the PGID table used for flooding of multicast frames, not flooded by the IPMC flood masks.	None
FLOODING_IPMC.FLD_MC4_CTRL	Index into the PGID table used for flooding of IPv4 multicast control frames.	None
FLOODING_IPMC.FLD_MC4_DATA	Index into the PGID table used for flooding of IPv4 multicast data frames.	None
FLOODING_IPMC.FLD_MC6_CTRL	Index into the PGID table used for flooding of IPv6 multicast control frames.	None
FLOODING_IPMC.FLD_MC6_DATA	Index into the PGID table used for flooding of IPv6 multicast data frames.	None
PGID[63:0]	Destination and flooding masks table.	64
AGENCTRL.IGNORE_DMAC_FLAGS	Controls the use of MAC table flags from (DMAC, VID) entry and flooding flags.	None
CPUQ_CFG	Configuration of CPU extraction queues	None
PORT_CFG.USE_B_DOM_TBL	Configuration of MAC table domain.	Per virtual port profile

The (DMAC, VID) pair is looked up in the MAC table. If USE_B_DOM_TBL is set, only MAC table entries with B_DOM set are searched. Otherwise, MAC table entries with B_DOM cleared are searched. A match is found when the (DMAC, VID) pair matches that of an entry and USE_B_DOM_TBL matches the entry's B_DOM flag.

If a match is found, the entry is returned and DEST is determined based on the MAC table entry. For more information, see [Section 2.8.1, MAC Table](#).

If an entry is found in the MAC table entry of ENTRY_TYPE 0 or 1 and the CPU port is set in the PGID pointed to by the MAC table entry, CPU extraction queue PGID.DST_PGID is added to the CPUQ.

If an entry is not found for the (DMAC, VID) in the MAC table, the frame is flooded. The forwarding decision is set to one of the seven flooding masks defined in ANA::FLOODING or ANA::FLOODING_IPMC, based on one of the flood type definitions listed in the following table.

TABLE 2-106: FORWARDING DECISIONS BASED ON FLOOD TYPE

Frame Type	Condition
IPv4 multicast data	DMAC = 0x01005E000000 to 0x01005E7FFFFFFF EtherType = IPv4 IP protocol is not IGMP IPv4 DIP outside 224.0.0.x
IPv6 multicast data	DMAC = 0x333300000000 to 0x3333FFFFFFFF EtherType = IPv6 IPv6 DIP outside 0xFF02::/16
IPv4 multicast control	DMAC = 0x01005E000000 to 0x01005E7FFFFFFF EtherType = IPv4 IP protocol is not IGMP IPv4 DIP inside 224.0.0.x
IPv6 multicast control	DMAC = 0x333300000000 to 0x3333FFFFFFFF EtherType = IPv6 IPv6 DIP inside 0xFF02::/16

TABLE 2-106: FORWARDING DECISIONS BASED ON FLOOD TYPE (CONTINUED)

Frame Type	Condition
Broadcast	DMAC = 0xFFFFFFFFFFFFFFF non-IPv4-multicast-data non-IPv6-multicast-data non-IPv4-multicast-control non-IPv6-multicast-control
Multicast	Bit 40 in DMAC = 1 non-broadcast non-IPv4-multicast-data non-IPv6-multicast-data non-IPv4-multicast-control non-IPv6-multicast-control
Unicast	Bit 40 in DMAC = 0

In addition, the MAC table flag MAC_CPU_COPY is processed. If MAC_CPU_COPY is set, the CPUQ_CFG.CPUQ_MAC is added to CPUQ.

The processing of this flag can be disabled through AGENCTRL.IGNORE_DMAL_FLAGS.

Next, CPU-forwarding from the MPLS-TP classifier, the basic classifier, and the IS2 SMAC_SIP lookup is processed if:

- The MPLS-TP classifier decided to copy the frame to the CPU (CPU port is set in B_PM), the CPU extraction queue from the MPLS-TP classifier, B_CPUQ, is added to CPUQ.
- The basic classifier decided to copy the frame to the CPU, the corresponding CPU extraction queue is added to CPUQ.
- The basic classifier decided to redirect the frame to the CPU, DEST is cleared and the corresponding CPU extraction queue is added to CPUQ.
- The IS2 SMAC_SIP result decided to copy the frame to the CPU (SMAC_SIP_ACTION.CPU_COPY_ENA), the corresponding CPU extraction queue, SMAC_SIP_ACTION.CPU_QU_NUM is added to CPUQ.
- The IS2 SMAC_SIP result decided to discard the frame (SMAC_SIP_ACTION.FWD_KILL_ENA set), DEST is cleared.

For more information about frame type definitions for CPU forwarding in the basic classifier, see [Table 2-45](#).

2.8.4.1 VLAN Analysis

During the VLAN analysis step, VLAN configuration is taken into account. As a result, ports can be removed from the forwarding decision. For more information about VLAN configuration, see [Section 2.8.2, VLAN Table](#).

The following table lists the registers associated with VLAN analysis.

TABLE 2-107: VLAN ANALYSIS REGISTERS

Register	Description	Replication
VLANMASK	If PPORT is set in this mask, and PPORT is not member of the VLAN to which the frame is classified, DEST is cleared. This is also called VLAN ingress filtering.	None
PORT_CFG.RECV_ENA	If this bit is cleared for PPORT, forwarding from this port to other front ports is disabled, and DEST is cleared.	Per virtual port profile
PGID[106:80]	Source port mask. Port mask per port, which specifies allowed destination ports for frames received on PPORT. By default, a port can forward to all other ports except itself.	Per port
ISOLATED_PORTS	Private VLAN mask. Isolated ports are cleared in this mask.	None
COMMUNITY_PORTS	Private VLAN mask. Community ports are cleared in this mask.	None
ADV-LEARN.VLAN_CHK	If set and VLAN ingress filtering clears DEST, then SMAC learning is disabled.	None

TABLE 2-107: VLAN ANALYSIS REGISTERS (CONTINUED)

Register	Description	Replication
USE_B_DOM_TBL	If set, B-domain processing is enabled. This implies that a B-VLAN mask is applied to the forwarding result. Otherwise, a C-VLAN mask is applied. The C-VLAN mask is dependent on whether the frame is classified to a service.	Per virtual port profile
PORT_B_DOMAIN	B-port mask. If PPORT is set in this mask, the port is a B-port. Otherwise, it is a C-port.	None

The frame's VID is used as an address for lookup in the VLAN table. The returned VLAN mask is processed into a final VLAN mask taking B-ports, C-ports, and services into account:

- If B-domain processing is enabled for the VPORT (USE_B_DOM_TBL is set), C-ports are removed from the VLAN (VLAN_PORT_MASK).
- If B-domain processing is disabled for the VPORT (USE_B_DOM_TBL is cleared), B-ports are removed from the VLAN mask (VLAN_PORT_MASK). If the frame is a service frame, i.e. the frame is classified to an ISDX>0, B-ports are added to the VLAN mask.

The VLAN mask is now applied to the forwarding decision as follows:

- All ports that are not members of the VLAN are removed from DEST, except if the (DMAC, VID) match in the MAC table has VLAN_IGNORE set, or if there is no match in the MAC table and AGENCTRL.FLOOD_IGNORE_VLAN is set.

Note These two exceptions are skipped if AGENCTRL.IGNORE_DMAC_FLAGS is set.

- If the VLAN_PRIV_VLAN flag in the VLAN table is set, the VLAN is private, and isolated and community ports must be treated differently. An isolated port is identified as an ingress port for which PPORT is cleared in the ISOLATED_PORTS register. A community port is identified as an ingress port for which PPORT is cleared in the COMMUNITY_PORTS register. For frames received on an isolated port, all isolated and community ports are removed from the forwarding decision. For frames received on a community port, all isolated ports are removed from the forwarding decision.
- If VLAN ingress filtering is enabled, it is checked whether PPORT is member of the VLAN (VLAN_PORT_MASK). If this is not the case, DEST is cleared.

VLAN ingress filtering is enabled per port in the VLANMASK register or per VLAN in the VLAN_SRC_CHK flag in the VLAN table. If either is set, VLAN ingress filtering is performed.

Next, it is checked whether the ingress port is enabled to forward frames to other front ports and the source mask (PGID[80+PPORT]) is processed as follows:

- If PORT_CFG.RECV_ENA for VPORT is 0, DEST is cleared.
- Any ports, which are cleared in PGID[80+PPORT], are removed from DEST.

Finally, SMAC learning is disabled by setting the LRN_DIS flag when either of the following two conditions is fulfilled as follows:

- VLAN_LEARN_DISABLED is set in the VLAN table for the VLAN.
- A frame is subject to VLAN ingress filtering (frame dropped due to PPORT not being member of VLAN), and ADV_LEARN.VLAN_CHK is set.

2.8.4.2 Aggregation

During the aggregation step, link aggregation is handled. The following table lists the registers associated with aggregation.

TABLE 2-108: ANALYZER AGGREGATION REGISTERS

Register	Description	Replication
PGID[79:64]	Aggregation mask table.	16

The aggregation step ensures that when a frame is destined for an aggregation group, it is forwarded to exactly one of the group's member ports.

For non-aggregated ports, there is a one-to-one correspondence between logical port (LPORT) and physical port (PPORT). The aggregation step does not change the forwarding decision.

For aggregated ports, all physical ports in the aggregation group map to the same logical port, and the entry in the destination mask table for the logical port includes all physical ports, which are members of the aggregation group. As a result, all but one member port must be removed from the destination port set.

The link aggregation code generated in the classifier is used to look up an aggregation mask in the aggregation masks table. Finally, ports that are cleared in the selected aggregation mask are removed from DEST.

For more information about link aggregation, see [Section 4.3.8, Link Aggregation](#).

2.8.4.3 MSTP Analysis

During the MSTP analysis step, MSTP configuration is taken into account. As a result, ports can be removed from the forwarding decision.

The following table lists the registers associated with MSTP analysis.

TABLE 2-109: MSTP ANALYSIS REGISTERS

Register	Description	Replication
MSTI_STATE[63:0]	VLAN-based MSTi forwarding mask table. The table is looked up with frame's FID. If FID>63 then FID=0 is used. If PPORT is not set in this mask, DEST is cleared. Otherwise, ports cleared in this mask are removed from DEST.	64
MSTI_STATE[127:64]	ISDX-based MSTi forwarding mask table. The table is looked up with key ISDX_TIDX.ISDX_MSTI+64. If PPORT is not set in this mask, DEST is cleared. Otherwise, ports cleared in this mask are removed from DEST.	64

Two port masks are retrieved from the MSTP forwarding mask table. The first lookup used the frame's FID as key while the second lookup uses the ISDX_MSTI pointer from the ISDX table. Both MSTP masks are now applied to the forwarding decision as follows:

- All ports that are not set in both MSTP masks are removed from DEST. This check ensures that frames are not forwarded to a port not in the MSTP forwarding state.
- If PPORT is not set in both MSTP masks then DEST is cleared. This check ensures that frames are not forwarded from a port not in the MSTP forwarding state.

2.8.4.4 Path Protection

This step handles path protection. The following table lists the registers associated with path protection.

TABLE 2-110: ANALYZER PATH PROTECTION REGISTERS

Register	Description	Replication
ANA::IPT	ISDX protection table, one entry per service.	256
ANA::PPT	Path protection status, one bit per path protection group.	23

The path protection step allows E-Line services to be path-protected. Each service is member of a path protection group and a path protection group is configured to either use the working path or the protected path (PPT.PATH_STAT). When a path fails, all services mapping to the same path protection group are protected by only reconfiguring the path group.

During this step, the ISDX protection table is looked up using the frame's classified ISDX. The ISDX protection table returns the protection configuration (IPT.IPT_CFG), which configures the protection scheme used for the ISDX:

- No protection. The service is not protected and the frame forwarding is not changed.
- Upstream protection. If the associated path protection group is using the protect path, then the frame's classified ISDX is replaced by the protected ISDX (IPT.ISDX_P). This affects looking up the ISDX table, the VCAP IS2, and the VCAP ES0. In addition, it affects the selection of the ingress service statistics. If the associated path protection group is using the working path then the ISDX is untouched.
- Downstream protection, working. This implies that the classified ISDX belongs to the working path. If the associated path protection group is using the protect path then the frame is discarded (DEST is cleared).
- Downstream protection, protect. This implies that the classified ISDX belongs to the protect path. If the associated path protection group is using the working path then the frame is discarded (DEST is cleared).

2.8.4.5 Service Forwarding

This step handles service forwarding. The following table lists the registers associated with service handling.

TABLE 2-111: ANALYZER SERVICE REGISTERS

Register	Description	Replication
ISDXTIDX.ISDX_FORCE_ENA	If set, forwarding of the frame is based on the service port mask only.	256
ISDXACCESS.ISDX- _PORT_MASK	Service port mask.	256

The service step ensures that a frame can be service-forwarded based on the service port mask.

If forced service forwarding is enabled (ISDX_FORCE_ENA is set), the forwarding decisions based on the destination MAC address and the VLAN are replaced with the service port mask (ISDXACCESS.ISDXC_PORT_MASK). Source port filtering and aggregation filtering are still applied to the forwarding result.

For normal service forwarding (ISDX_FORCE_ENA is cleared), all ports that are cleared in ISDX_PORT_MASK are removed from DEST. This constrains the flooding domain of the service within the VLAN.

2.8.4.6 VCAP Action Handling

VCAP IS2 actions are processed during the VCAP IS2 action handling step. The following table lists the processing of the VCAP actions. The order of processing is from top to bottom.

TABLE 2-112: VCAP IS2 ACTION PROCESSING

IS2 Action Field	Description
CPU_COPY_ENA CPU_QU_NUM	If CPU_COPY_ENA is set, the CPU_QU_NUM bit is set in CPUQ.
HIT_ME_ONCE CPU_QU_NUM	If HIT_ME_ONCE is set and the HIT_CNT counter is zero, the CPU_QU_NUM bit is set in CPUQ.
LRN_DIS	If set, learning is disabled (LRN_DIS flag is set).
POLICE_ENA POLICE_IDX	If POLICE_ENA is set (only applies to first lookup), the POLICE_IDX instructs which policer to use for this frame. See Section 2.9, Policers .
POLICE_V- CAP_ONLY	If POLICE_VCAP_ONLY is set (only applies to first lookup), the only active policer for this frame is the VCAP policer. Other policers (service, QoS, port) are disabled. See Section 2.9, Policers .
MASK_MODE PORT_MASK	The following actions are defined for MASK_MODE. 0: No action. 1: Permit. Ports cleared in PORT_MASK are removed from DEST. 2: Policy. DEST from the DMAC analysis step is replaced with PORT_MASK 3: Redirect. DEST as the outcome of the DMAC, VLAN, service, and aggregation analysis steps is replaced with PORT_MASK.
MIRROR_ENA	If MIRROR_ENA is set, mirroring is enabled. This is used in the mirroring step. See Section , Mirroring .

2.8.4.7 SMAC Analysis

During the SMAC analysis step, the MAC table is searched for a match against the (SMAC, VID), and the MAC table is updated due to learning. Either the B-MAC table or the C-MAC table is searched. The learning part is skipped if the LRN_DIS flag was set by any of the previous steps.

The following table lists the registers associated with SMAC learning.

TABLE 2-113: SMAC LEARNING REGISTERS

Register	Description	Replication
PORT_CFG.LEARN_ENA	If set for VPORT, learning is skipped (that is, LEARNAUTO, LEARNCPU, LEARN-DROP, LIMIT_CPU, LIMIT_DROP, LOCKED_PORTMOVE_CPU, and LOCKED_PORTMOVE_DROP are ignored).	Per virtual port profile
PORT_CFG.LEARNAUTO	If set for VPORT, hardware-based learning is performed.	Per virtual port profile
PORT_CFG.LEARNCPU	If set for VPORT, learn frames are copied to the CPU.	Per virtual port profile
PORT_CFG.LEARNDROP	If set for VPORT, the CPU drops or forwards learn frames.	Per virtual port profile
PORT_CFG.LIMIT_CPU	If set for VPORT, learn frames for which PPORT exceeds the port's limit are copied to the CPU.	Per virtual port profile
PORT_CFG.LIMIT_DROP	If set for VPORT, learn frames for which PPORT exceeds the port's limit are discarded.	Per virtual port profile
PORT_CFG.LOCKED_PORTMOVE_CPU	If set for VPORT, frames triggering a port move of a locked entry are copied to the CPU.	Per virtual port profile
PORT_CFG.LOCKED_PORTMOVE_DROP	If set for VPORT, frames triggering a port move of a locked entry are discarded.	Per virtual port profile
PORT_CFG.USE_B_DOM_TBL	If set for VPORT, the B-MAC table is searched. Otherwise, the C-MAC table is searched.	Per virtual port profile
AGENCTRL.IGNORE_SMAC_FLAGS	Controls the use of the MAC table flags from (SMAC, VID) entry.	None

The B-MAC table is searched if the port has PORT_CFG.USE_B_DOM_TBL set. Otherwise the C-MAC table is searched.

Three different type of learn frames are identified:

- **Normal learn frames.** Frames for which an entry for the (SMAC, VID) is not found in the MAC table or the (SMAC, VID) entry in the MAC table is unlocked and has a DEST_IDX different from LPORT. In addition, the learn limit for the LPORT must not be exceeded (ENTRYLIM).
- **Learn frames exceeding the learn limit.** Same condition as for normal learn frames except that the learn limit for the LPORT is exceeded (ENTRYLIM)
- **Learn frames triggering a port move of a locked MAC table entry.** Frames for which the (SMAC, VID) entry in the MAC table is locked and has a DEST_IDX different from LPORT.

For all learn frames, the following must apply before learning related processing is applied.

- Learning is enabled by PORT_CFG.LEARN_ENA.
- The LRN_DIS flag from previous processing steps must be cleared, which implies the following:
 - Learning is not disabled due to VLAN ingress filtering
 - Learning is not disabled due to VCAP IS2 action
 - Learning is enabled for the VLAN (VLAN_LEARN_DISABLED is cleared in the VLAN table)

In addition, learning must not be disabled due to the ingress policer having policed the frame. For more information, see [Section 2.9, Policers](#).

If learning is enabled, learn frames are processed according to the setting of the following configuration parameters.

Normal learn frames

- Automatic learning. If PORT_CFG.LEARNAUTO is set for VPORT, the (SMAC, VID) entry is automatically added

to the MAC table in the domain being searched.

- Drop learn frames. If `PORT_CFG.LEARNDROP` is set for VPORT, DEST is cleared for learn frames. Therefore, learn frames are not forwarded on any ports. This is used for secure learning, where the CPU must verify a station before forwarding is allowed.
- Copy learn frames to the CPU. If `PORT_CFG.LEARNCPU` is set for VPORT, the CPU port is added to DEST for learn frames and `CPUQ_CFG.CPUQ_LRN` is set in CPUQ. This is used for CPU based learning.

Learn frames exceeding the learn limit

- Drop learn frames. If `PORT_CFG.LIMIT_DROP` is set for VPORT, DEST is cleared for learn frames. As a result, learn frames are not forwarded on any ports.
- Copy learn frames to the CPU – If `PORT_CFG.LIMIT_CPU` is set for VPORT, the CPU port is added to DEST and `CPUQ_CFG.CPUQ_LRN` is set in CPUQ for learn frames.

Learn frames triggering a port move of a locked MAC table entry

- Drop learn frames. If `PORT_CFG.LOCKED_PORTMOVE_DROP` is set for VPORT, DEST is cleared for learn frames. Therefore, learn frames are not forwarded on any ports.
- Copy learn frames to the CPU. If `PORT_CFG.LOCKED_PORTMOVE_CPU` is set for VPORT, the CPU port is added to DEST and `CPUQ_CFG.CPUQ_LOCKED_PORTMOVE` is added to CPUQ.

Finally, if a match is found in the MAC table for the (SMAC, VID), adjustments can be made to the forwarding decision.

- If the (SMAC, VID) match in the MAC table has `SRC_KILL` set, DEST is cleared.
- If the (SMAC, VID) match in the MAC table has `MAC_CPU_COPY` set, `CPUQ_CFG.CPUQ_MAC_COPY` is added to CPUQ.

The processing of the MAC table flags from the (SMAC, VID) match can be disabled through `AGENCY_CTRL.IGNORE_SMAC_FLAGS`.

2.8.4.8 Storm Policers

The storm policers are activated during the storm policers step. The following table lists the registers associated with storm policers.

TABLE 2-114: STORM POLICER REGISTERS

Register	Description	Replication
<code>STORMLIMIT_CFG</code>	Enables policing of various frame types.	4
<code>STORMLIMIT_BURST</code>	Configures maximum allowed rates of the different frame types.	None

The analyzer contains four storm policers that can limit the maximum allowed forwarding frame rate for various frame types. The storm policers are common to all ports and, as a result, measure the sum of traffic forwarded by the switch. A frame can activate several policers, and the frame is discarded if any of the activated policers exceed a configured rate.

Each policer can be configured to a frame rate ranging from 1 frame per second to 1 million frames per second.

The following table lists the available policers.

TABLE 2-115: STORM POLICERS

Policer	Description
Broadcast	Flooded frames with DMAC = 0xFFFFFFFFFFFF.
Multicast	Flooded frames with DMAC bit 40 set, except broadcasts.
Unicast	Flooded frames with DMAC bit 40 cleared.
Learn	Learn frames copied or redirected to the CPU due to learning (<code>LOCKED_PORTMOVE_CPU</code> , <code>LIMIT_CPU</code> , <code>LEARNCPU</code>).

For each of the policers, a maximum rate is configured in `STORMLIMIT_CFG` and `STORMLIMIT_BURST`:

- `STORM_UNIT` chooses between a base unit of 1 frame per second or 1 kiloframes per second.
- `STORM_RATE` sets the rate to 1, 2, 4, 8, ..., 1024 times the base unit (`STORM_UNIT`).
- `STORM_BURST` configures the maximum number of frames in a burst.
- `STORM_MODE` specifies how the policer affects the forwarding decision. The options are:
When policing, clear CPUQ.

When policing, clear DEST.

When policing, clear DEST and CPUQ.

Frames where the DMAC lookup returned a PGID with the CPU port set are always forwarded to the CPU even when the frame is policed by the storm policers. For more information, see [Section , DMAC Analysis](#).

2.8.4.9 sFlow Sampling

This process step handles sFlow sampling. The following table lists the registers associated with sFlow sampling.

TABLE 2-116: SFLOW SAMPLING REGISTERS

Register	Description	Replication
SFLOW_CFG	Configures sFlow samplers (type and rates).	Per port
CPUQ_CFG.CPUQ_SFLOW	CPU extraction queue for sFlow sampled frames.	None

sFlow is a standard for monitoring high-speed switch networks through statistical sampling of incoming and outgoing frames. Each port in the device can be setup as an sFlow agent monitoring the particular link and generating sFlow data. If a frame is sFlow sampled, it is copied to the sFlow CPU extraction queue (CPUQ_SFLOW).

An sFlow agent is configured through SFLOW_CFG with the following options:

- SF_RATE specifies the probability that the sampler copies a frame to the CPU. Each frame being candidate for the sampler has the same probability of being sampled. The rate is set in steps of 1/4096.
- SF_SAMPLE_RX enables incoming frames on the port as candidates for the sampler.
- SF_SAMPLE_TX enables outgoing frames on the port as candidates for the sampler.

The Rx and Tx can be enabled independently. If both are enabled, all incoming and outgoing traffic on the port is subject to the statistical sampling given by the rate in SF_RATE.

Mirroring

This processing step handles mirroring. The following table lists the registers associated with mirroring.

TABLE 2-117: MIRRORING REGISTERS

Register	Description	Replication
ADVLEARN.LEARN_MIRROR	For learn frames, ports in this mask (mirror ports) are added to DEST.	None
AGENCTRL.MIRROR_CPU	Mirror all frames forwarded to the CPU port module.	None
PORT_CFG.SRC_MIRROR_ENA	Mirror all frames received on an ingress port (ingress port mirroring).	Per port
EMIRRORMASK	Mirror frames that are to be transmitted on any ports set in this mask (egress port mirroring).	None
VLANTIDX.VLAN_MIRROR	Mirror all frames classified to a specific VID.	Per VLAN
IS2_ACTION.MIRROR_ENA	Mirror when an IS2 action is hit.	Per VCAP IS2 entry
MIRRORPORTS	When mirroring a frame, ports in this mask are added to DEST.	None
AGENC-TRL.CPU_CPU_KILL_ENA	Clear the CPU port if source port is the CPU port and the CPU port is set in DEST.	None

Frames subject to mirroring are identified based on the following mirror probes:

- Learn mirroring if ADVLEARN.LEARN_MIRROR is set and frame is a learn frame.
- CPU mirroring if AGENCTRL.MIRROR_CPU is set and the CPU port is set in DEST.
- Ingress mirroring if PORT_CFG.SRC_MIRROR_ENA is set.
- Egress mirroring if any port set in EMIRRORMASK is also set in DEST.
- VLAN mirroring if VLAN_MIRROR set in the VLAN table entry.
- VCAP mirroring if an action is hit that requires mirroring.

The following adjustment is made to the forwarding decision for frames subject to mirroring:

- Ports set in MIRRORPORTS are added to DEST.

If the CPU port is set in the MIRRORPORTS, CPU extraction queue CPUQ_CFG.CPUQ_MIRROR is added to the CPUQ.

For learn frames with learning enabled, all ports in ADVLEARN.LEARN_MIRROR are added to DEST. For more information, see [Section 2.8.4.7, SMAC Analysis](#).

For more information about mirroring, see [Section 4.3.10, Mirroring](#).

Finally, if AGENCTRL.CPU_CPU_KILL_ENA is set, the CPU port is removed if the ingress port is the CPU port itself. This is similar to source port filtering done for front ports and prevents the CPU from sending frames back to itself.

2.8.4.10 MPLS Handling

This step handles MPLS forwarding. The purpose is to add MPLS selected B-ports to the destination set. This is done by adding all ports set in the B_PM port mask from the MPLS-TP classifier to DEST.

2.8.5 ANALYZER MONITORING

Miscellaneous events in the analyzer can be monitored, which can provide an understanding of the events during the processing steps. The following table lists the registers associated with analyzer monitoring.

TABLE 2-118: ANALYZER MONITORING

Register	Description	Replication
ANMOVED	ANMOVED[n] is set when a known station has moved to port n.	None
ANEVENTS	Sticky bit register for various events.	None
LEARNDISC	The number of learn events that failed due to a lack of storage space in the MAC table.	None

Port moves, defined as a known station moving to a new port, are registered in the ANMOVED register. A port move occurs when an existing MAC table entry for (MAC, VID) is updated with new port information (DEST_IDX). Such an event is registered in ANMOVED by setting the bit corresponding to the new port.

Continuously occurring port moves may indicate a loop in the network or a faulty link aggregation configuration.

A list of 24 events, such as frame flooding or policer drop, can be monitored in ANEVENTS.

The LEARNDISC counter registers every time an entry in the MAC table cannot be made or if an entry is removed due to lack of storage.

2.9 Policers

The device has 392 policers that can be allocated to services, ingress ports, QoS classes per port, and VCAP IS2 entries. The policers limit the bandwidth of received frames by discarding frames exceeding configurable rates. All policers are MEF compliant dual leaky bucket policers that are capable of handling committed and excess peak information rates.

Each frame can hit up to three policers: One port policer, one VCAP policer and either a service policer or a QoS policer. Frames can either be policed by a service policer or a QoS policer. The order in which the policers are applied to a frame is programmable.

In addition to the policers, the VSC7416-01 device also supports a number of storm policers and an egress scheduler with shaping capabilities. For more information, see [Section 2.8.4.8, Storm Policers](#) and [Section 2.11, Scheduler and Shapers](#).

The following table lists the registers associated with policer control.

TABLE 2-119: POLICER CONTROL REGISTERS

Register	Description	Replication
ANA:PORT:POL_CFG	Enables use of port and QoS policers.	Per virtual port profile
ANA::ISDXTSDLIDX.ISDX_BI	Service policer index.	Per ISDX
ANA:POL:POL_PIR_CFG	Configures the policer's peak information rate.	448
ANA:POL:POL_CIR_CFG	Configures the policer's committed information rate.	448
ANA:POL:POL_MODE_CFG	Configures the policer's mode of operation.	448

TABLE 2-119: POLICER CONTROL REGISTERS (CONTINUED)

Register	Description	Replication
ANA:POL:POL_PIR_STATE	Current state of the peak information rate bucket.	448
ANA:POL:POL_CIR_STATE	Current state of the committed information rate bucket.	448
ANA:PORT:POL_FLOWC	Flow control settings.	Per port
ANA::POL_HYST	Hysteresis settings.	None


2.9.1 POLICER ALLOCATION

The different policer types are assigned a policer from the pool of policers the following ways:

- Port policers. Frames received on physical port 'p' use policer 'p'. Each of physical ports can be assigned to its own policer.
- Service policers. Frames classified to ISDX 's' use policer $128 + \text{ISDXTID}[s].\text{ISDX\_SDLBI}$. Each of the 256 ingress service points can be assigned a policer. Note that $\text{ISDX\_SDLBI}=0$ disables service policing and allows QoS policing to be used instead. Two different ingress service points can point to the same service policer.
- QoS policers. Frames classified to QoS class 'q' on physical port 'p' use policer $32 + 8 \times 'p' + 'q'$. Each of the eight per-port QoS classes per port can be assigned to its own policer.
- VCAP IS2 policers. Policers 384 through 447 are VCAP policers. The action $\text{IS2\_ACTION.POLICE\_IDX}$ points to the policer that is used.

The policer pool layout is illustrated in the following drawing.

FIGURE 2-23: POLICER POOL LAYOUT



By default, none of the policers from the pool are allocated.

Port policers are allocated through ANA:PORT:POL_CFG.PORT_POL_ENA per virtual port profile and QoS policers are allocated through ANA:PORT:POL_CFG.QUEUE_POL_ENA per QoS class per virtual port profile. The virtual port profile provides the option to enable for instance port policing for some virtual port profiles and disable for others. The actual port policer used is always given by the physical port where the frame was received.

Service policers are allocated per ISDX by setting the service policer index in ANA::ISDXTIDX.ISDX_SDLBI to a non-zero value. Service policing takes precedence over QoS class policing. This implies that a frame is policed by the service policer when the frame's associated service policer index is non-zero and only policed by the QoS policer when the frame is not service policed and the frame's QoS policer is enabled. It is not possible to police by both a service policer and a QoS policer. Non-service frames (ISDX=0) can also point to a service-policer if ISDX_SDLBI is non-zero for the ISDX=0. This policer then polices all non-service traffic across all ports.

Finally, VCAP IS2 policers are allocated by creating IS2 rules with POLICE_ENA and POLICE_IDX actions. IS2 policers actions are only valid in the first lookup in IS2. The VCAP can point to any unused policer in addition to the dedicated VCAP policers.

Any frame received by the MAC and forwarded to the classifier is applicable to policing. Frames with errors, pause frames, or MAC control frames are not forwarded by the MAC and, as a result, they are not accounted for in the policers. That is, they are not policed and are not adding to the rate measured by the policers.

In addition, the following special frame types can bypass the policers:

- If ANA:PORT:POL_CFG.POL_CPU_REDIR_8021 is set, frames being redirected to the CPU due to the classifier detecting the frames as being BPDU, ALLBRIDGE, GARP, or CCM/Link trace frames are not policed.
- If ANA:PORT:POL_CFG.POL_CPU_REDIR_IP is set, frames being redirected to the CPU due to the classifier detecting the frames as being IGMP or MLD frames are not policed.

These frames are still considered part of the rates being measured so the frames add to the relevant policer buckets but they are never discarded due to policing.

The VCAP IS2 has the option to disable the port policing, the QoS class policing, and the service policing. This is done with the action POLICE_VCAP_ONLY. If POLICE_VCAP_ONLY is set, only a VCAP assigned policer can police the frame. The other policers are inactive meaning the frame does not add to the policers' buckets and the frame is never discarded due to policing by the policers.

The order in which the policers are executed is controlled through ANA:PORT:POL_CFG.POL_ORDER. The order can take the following main modes:

- **Serial** The policers are checked one after another. If a policer is closed, the frame is discarded and the subsequent policer buckets are not updated with the frame. The serial order is programmable.
- **Parallel with independent bucket updates** The three policers are working in parallel independently of each other. Each frame is added to a policer bucket if the policer is open, otherwise the frame is discarded. A frame may be added to one policer although another policer is closed.
- **Parallel with dependent bucket updates** The three policers are working in parallel but dependent on each other with respect to bucket updates. A frame is only added to the policer buckets if all three policers are open.

2.9.2 POLICER BURST AND RATE CONFIGURATION

Each of the 392 policers are MEF-compliant dual leaky bucket policers. This implies that each policer supports the following configurations:

- Committed Information Rate (CIR). Specified in POL_CIR_CFG.CIR_RATE in steps of 33.3 kbps. Maximum rate is 1.09 Gbps. If higher bandwidths are required, the policer for 2.5G ports must be disabled.
- Committed Burst Size (CBS). Specified in POL_CIR_CFG.CIR_BURST in steps of 4 kilobytes. Maximum is 252 kilobytes.
- Excess Information Rate (EIR). Specified in POL_PIR_CFG.PIR_RATE in steps of 33.3 kbps. Maximum rate is 1.09 Gbps. If higher bandwidths are required, the policer for 2.5G ports must be disabled.
- Excess Burst Size (EBS). Specified in POL_PIR_CFG.PIR_BURST in steps of 4 kilobytes. Maximum is 252 kilobytes.
- Coupling flag. If POL_MODE_CFG.DLB_COUPLED is set, frames classified as yellow (DP level = 1) are allowed to use of the committed information rate when not fully used by frames classified as green (DP level = 0). If cleared, the rate of frames classified as yellow are bounded by EIR.
- Color mode. Color-blind or color-aware. A policer always obey the frame color assigned by the classifier. To achieve color-blindness, the classifier must be set up to classify all incoming frames to DP level = 0.

The following parameters can also be configured per policer:

- The leaky bucket calculation can be configured to include or exclude preamble and inter-frame gap through configuration of POL_MODE_CFG.IPG_SIZE.
- Each policer can be configured to measure frame rates instead of bit rates (POL_MODE_CFG.FRM_MODE). The rate unit can be configured to 100 frames per second or 1 frame per second.
- Each policer can operate as a single leaky bucket by disabling POL_MODE_CFG.CIR_ENA. When operating as a single leaky bucket, the POL_PIR_CFG register controls the rate and burst of the policer.

By default, a policer discards frames while the policer is closed. A discarded frame is neither forwarded to any ports (including the CPU) nor is it learned.

Each port policer, however, has the option to run in flow control where the policer instructs the MAC to issue flow control pause frames instead of discarding frames. This is enabled in ANA:PORT:POL_FLOWC. Common for all port policers, POL_HYST.POL_FC_HYST specifies a hysteresis, which controls when the policer can re-open after having closed.

To improve fairness between small and large frames being policed by the same policer, POL_HYST.POL_DROP_HYST specifies a hysteresis that controls when the policer can re-open after being closed. By setting it to a value larger than the maximum transmission unit, it guarantees that when the policer opens again, all frames have the same chance of being accepted. This setting only applies to policers working in drop mode.

The current fill level of the dual leaky buckets can be read in POL_PIR_STATE and POL_CIR_STATE. The unit is 0.5 bits.


2.10 Shared Queue System

The device includes a shared queue system with one ingress queue per ingress port and 2616 shared egress queues. The queue system has 1024 kilobytes of buffer.

Frames are linked into the ingress queues after frame analysis. Each egress port module selected by the frame analysis receives a link to the frame and stores the link in the appropriate egress queue selected by mapping various frame properties to a queue number. The transfer from ingress to egress is extremely efficient with a transfer time of 12.8 ns per frame copy (equivalent to a transfer rate from ingress to egress of about 50 Gbps for 64-byte frames and about 1 terabit per second (Tbps) for 1518-byte frames). Each egress port module has a scheduler that selects between the egress queues when transmitting frames.

The following illustration shows the shared queue system.

FIGURE 2-24: QUEUE SYSTEM OVERVIEW



Resource depletion can prevent one or more of the frame copies from the ingress queue to the egress queues. If a frame copy cannot be made due to lack of resources, the ingress port's flow control mode determines the behavior as follows:

- Ingress port is in drop mode: The frame copy is discarded.
- Ingress port is in flow control mode: The frame is held back in the ingress queue and the frame copy is made when the congestion clears.

For more information about special configurations of the shared queue system with respect to flow control, see [Section 2.10.7, Ingress Pause Request Generation](#).

2.10.1 BUFFER MANAGEMENT

A number of watermarks control how much data can be pending in the egress queues before the resources are depleted. There are no watermarks for the ingress queues, except for flow control, because the ingress queues are empty most of the time due to the fast transfer rates from ingress to egress. For more information, see [Section 2.10.7, Ingress Pause Request Generation](#). When the watermarks are configured properly, congested traffic does not influence the forwarding of non-congested traffic.

The memory is split into two main areas:

- A reserved memory area. The reserved memory area is subdivided into areas per port per QoS class per direction (ingress/egress).
- A shared memory area, which is shared by all traffic.

For setting up the reserved areas, egress watermarks exist per port and per QoS class for both ingress and egress. The following table lists the reservation watermarks.

TABLE 2-120: RESERVATION WATERMARKS

Register	Description	Replication
BUF_Q_RSRV_E	Configures the reserved amount of egress buffer per QoS class per egress port.	Per QoS class per egress port
BUF_P_RSRV_E	Configures the reserved amount of egress buffer shared among the egress port's allocated egress queues.	Per egress port
BUF_Q_RSRV_I	Configures the reserved amount of egress buffer per ingress port per QoS class across all egress ports.	Per ingress port per QoS class
BUF_P_RSRV_I	Configures the reserved amount of egress buffer per ingress port shared among the eight QoS classes.	Per ingress port

All the watermarks, including the ingress watermarks, are compared against the memory consumptions in the egress queues. For example, the ingress watermarks in BUF_Q_RSRV_I compare against the total consumption of frames across all egress queues received on the specific ingress port and classified to the specific QoS class. The ingress watermarks in BUF_P_RSRV_I compare against the total consumption of all frames across all egress queues received on the specific ingress port.

The reserved areas are guaranteed minimum areas. A frame cannot be discarded or held back in the ingress queues if the frame's reserved areas are not yet used.

The shared memory area is the area left when all the reservations are taken out. The shared memory area is shared between all ports, however, it is possible to configure a set of watermarks per QoS class and per drop precedence level (green/yellow) to stop some traffic flows before others. The following table lists the sharing watermarks.

TABLE 2-121: SHARING WATERMARKS

Register	Description	Replication
BUF_PRI-O_SHR_E	Configures how much of the shared memory area that egress frames with the given QoS class are allowed to use.	Per QoS class
BUF_COL_SHR_E	Configures how much of the shared memory area that egress frames with the given drop precedence level are allowed to use.	Per drop precedence level
BUF_PRI-O_SHR_I	Configures how much of the shared memory area that ingress frames with the given QoS class are allowed to use.	Per QoS class
BUF_COL_SHR_I	Configures how much of the shared memory area that ingress frames with the given drop precedence level are allowed to use.	Per drop precedence level

The sharing watermarks are maximum areas in the shared memory that a given traffic flow can use. They do not guarantee anything.

When a frame is enqueued into the egress queue system, the frame first consumes from the queue's reserved memory area, then from the port's reserved memory area. When all the frame's reserved memory areas are full, it consumes from the shared memory area.

The following provides some simple examples on how to configure the watermarks and how that influences the resource management.

- Setting BUF_Q_RSRV_E (egress port = 7, QoS class = 4) to 2 kilobytes guarantees that traffic destined for port 7 classified to QoS class 4 have room for 2 kilobytes of frame data before frames can get discarded.
- Setting BUF_Q_RSRV_I (ingress port = 7, QoS class = 4) to 2 kilobytes guarantees that traffic received on port 7 classified to QoS class 4 have room for 2 kilobytes of frame data before frames can get discarded.
- Setting BUF_P_RSRV_I (ingress port 7) to 10 kilobytes guarantees that traffic received on port 7 have room for 10 kilobytes of data before frames can get discarded.
- The three above reservations reserve in total 14 kilobytes of memory (2 + 2 + 10 kilobytes) for port 7. If the same reservations are made for all ports, there are $1024 - 11 \times 14 = 870$ kilobytes left for sharing. If the sharing water-

marks are all set to 870 kilobytes, all traffic groups can consume memory from the shared memory area without restrictions.


If, instead, setting BUF_PRIO_SHR_E(QoS class = 7) to 800 kilobytes and the other watermarks BUF_PRIO_SHR_E(QoS class = 0:6) to 700 kilobytes guarantees that traffic classified to QoS class 7 has 100 kilobytes extra buffer. The buffer is shared between all ports.

- The BUF_PRIO_SHR_I, BUF_PRIO_SHR_E, REF_PRIO_SHR_I, and REF_PRIO_SHR_E watermarks can be used for guaranteeing shared resources for each individual QoS class. This is done by setting QSYS::RES_QOS_MODE so that the watermarks operate on the current consumption of each QoS class instead of the total use of the shared resources.

2.10.2 FRAME REFERENCE MANAGEMENT

Each frame in an egress queue consumes a frame reference, which is a pointer element that points to the frame's data in the memory and to the pointer element belonging to the next frame in the queue. The following illustrations shows how the frame references are used for creating the queue structure.

FIGURE 2-25: FRAME REFERENCE



The shared queue system holds a table of 11000 frame references. The consumption of frame references is controlled through a set of watermarks. The set of watermarks is the exact same as for the buffer control. The frame reference watermarks are prefixed REF_. Instead of controlling the amount of consumed memory, they control the number of frame references. Both reservation and sharing watermarks are available. For more information, see [Table 2-120](#) and [Table 2-121](#).

When a frame is enqueued into the shared queue system, the frame consumes first from the queue's reserved frame reference area, then from the port's reserved frame reference area. When all the frame's reserved frame reference areas are full, it consumes from the shared frame reference area.

2.10.3 RESOURCE DEPLETION CONDITION

A frame copy is made from an ingress port to an egress port when both a memory check and a frame reference check succeed. The memory check succeeds when at least one of the following conditions is met.

- Ingress memory is available: BUF_Q_RSRV_I or BUF_P_RSRV_I are not exceeded.
- Egress memory is available: BUF_Q_RSRV_E or BUF_P_RSRV_E are not exceeded.
- Shared memory is available: None of BUF_PRIO_SHR_E, BUF_COL_SHR_E, BUF_PRIO_SHR_I, or BUF_COL_SHR_I are exceeded.

The frame reference check succeeds when at least one of the following conditions is met.

- Ingress frame references are available: REF_Q_RSRV_I or REF_P_RSRV_I are not exceeded.
- Egress frame references are available: REF_Q_RSRV_E or REF_P_RSRV_E are not exceeded.
- Shared frame references are available: None of REF_PRIO_SHR_E, REF_COL_SHR_E, REF_PRIO_SHR_I, or REF_COL_SHR_I are exceeded.

2.10.4 CONFIGURATION EXAMPLE

This section provides an example of how the watermarks can be configured for a QoS-aware switch with no color handling and the effects of the settings.


TABLE 2-122: WATERMARK CONFIGURATION EXAMPLE

Watermark	Value	Comment
BUF_Q_RSRV_I	500 bytes	Guarantees that a port is capable of receiving at least one frame in all QoS classes. Note It is not necessary to assign a full MTU, because the watermarks are checked before the frame is added to the memory consumption.
BUF_P_RSRV_I	0	No additional guarantees for the ingress port.
BUF_Q_RSRV_E	200 bytes	Guarantees that all QoS classes are capable of sending a non-congested stream of traffic through the switch.
BUF_P_RSRV_E	40 kilobytes	Guarantees that all egress ports have 40 kilobytes of buffer, independently of other traffic in the switch. This is the most demanding reservation in this setup, reserving 440 kilobytes of the total 1024 kilobytes.
BUF_COL_SHR_E BUF_COL_SHR_I	Maximum	Effectively disables frame coloring as watermark is never reached.
BUF_PRIO_SHR_E BUF_PRIO_SHR_I	82 kilobytes to 103 kilobytes	The different QoS classes are cut-off with 3 kilobytes distance (82, 85, 88, 91, 94, 97, 100, and 103 kilobytes). This gives frames with higher QoS classes a larger part of the shared buffer area. Effectively, this means that the burst capacity is 92 kilobytes for frames belonging to QoS class 0 and up to 113 kilobytes for frame belonging to QoS class 7.
REF_Q_RSRV_E REF_Q_RSRV_I	4	For both ingress and egress, this guarantees that four frames can be pending from and to each port.
REF_P_RSRV_E REF_P_RSRV_I	20	For both ingress and egress, this guarantees that an extra 20 frames can be pending, shared between all QoS classes within the port.
REF_COL_SHR_E REF_COL_SHR_I	Maximum	Effectively disables frame coloring as watermark is never reached.
REF_PRIO_SHR_E REF_PRIO_SHR_I	2350 - 2700	The different QoS classes are cut-off with a distance of 50 frame references (2350, 2400, 2450, 2500, 2550, 2600, 2650, and 2700). This gives frames with higher QoS classes a larger part of the shared reference area.

2.10.5 WATERMARK PROGRAMMING AND CONSUMPTION MONITORING

The watermarks previously described are all found in the SYS::RES_CFG register. The register is replicated 1024 times. The following illustration the organization.

FIGURE 2-26: WATERMARK LAYOUT



The illustration shows the watermarks available for the BUF_XXX_I group of watermarks. For the other groups of watermarks (BUF_XXX_I, REF_XXX_I, BUF_XXX_E, and REF_XXX_E), the exact same set of watermarks are available.

For monitoring the consumption of resources, SYS::RES_STAT provides information about current use and the maximum use since the last read of the register. The information is available for each of the watermarks listed and the layout of the RES_STAT register follows the layout of the watermarks. SYS::MMGT.FREECNT holds the amount of free memory in the shared queue system and SYS::EQ_CTRL.FP_FREE_CNT holds the number of free frame references in the shared queue system.

2.10.6 ADVANCED RESOURCE MANAGEMENT

A number of additional handles into the resource management system are available for special use of the device. They are described in the following table.

TABLE 2-123: RESOURCE MANAGEMENT

Resource Management	Description
Forced drop of egress and ingress frames	QSYS::EGR_DROP_MODE QSYS::SWITCH_PORT_MODE.INGRESS_DROP_MODE. If either an ingress port or an egress port in a frame transfer are configured for drop mode, congestion results in frame discards. Otherwise frames are held back in the ingress queues with potential head-of-line blocking effects. Normally all egress ports are set to non-drop-mode while the ingress drop mode reflects whether or not the port is configured for flow control.
Prevent ingress port from using of the shared resources.	QSYS::IGR_NO_SHARING. For frames received on ports set in this mask, the shared watermarks are considered exceeded. This prevents the port from using more resources than allowed by the reservation watermarks.
Prevent egress port from using of the shared resources.	QSYS::EGR_NO_SHARING. For frames switched to ports set in this mask the shared watermarks are considered exceeded. This prevents the port from using more resources than allowed by the reservation watermarks.

TABLE 2-123: RESOURCE MANAGEMENT (CONTINUED)

Resource Management	Description
Weighted Random Early Detection (WRED)	QSYS::RED_PROFILE. It is possible to discard frames with increasing probability as the consumption of shared resources per QoS class per drop precedence level increases. QSYS::RED_PROFILE configures a low and a high watermark per QoS class per drop precedence level. The probability of discarding a frame increases linearly from 0% when the consumption is at the low watermark to 100% when the consumption exceeds the high watermark.
Prevent dequeuing	QSYS::PORT_MODE.DEQUEUE_DIS. Each egress port can disable dequeuing of frames from the egress queues.

2.10.7 INGRESS PAUSE REQUEST GENERATION

During resource depletion, the shared queue system either discards frames when the ingress port operates in drop mode, or holds back frames when the ingress port operates in flow control mode. The following describes special configuration for the flow control mode.

The shared queue system is enabled for holding back frames during resource depletion in SYS:PORT:PAUSE_CFG.PAUSE_ENA. In addition, this enables the generation of pause requests to the port module based on memory consumptions. The MAC uses the pause request to generate pause frames or create back pressure collisions to halt the link partner. This is done according to the MAC configuration. For more information about MAC configuration, see [Section 2.1.2, MAC](#).

The shared queue system generates the pause request based on the ingress port's memory consumption and also based on the total memory consumption in the shared queue system. This enables a larger burst capacity for a port operating in flow control while not jeopardizing the non-dropping flow control.

Generating the pause request partially depends on a memory consumption flag, TOT_PAUSE, which is set and cleared under the following conditions:

- The TOT_PAUSE flag is set when the total consumed memory in the shared queue system exceeds the SYS:PORT:PAUSE_TOT_CFG.PAUSE_TOT_START watermark.
- The TOT_PAUSE flag is cleared when the total consumed memory in the shared queue system is below the SYS:PORT:PAUSE_TOT_CFG.PAUSE_TOT_STOP watermark.

The pause request is asserted when both of the following conditions are met.

- The TOT_PAUSE flag is set.
- The ingress port memory consumption exceeds the SYS:PORT:PAUSE_CFG.PAUSE_START watermark.

The pause request is deasserted the following condition is met:

- The ingress port's consumption is below the SYS:PORT:PAUSE_CFG.PAUSE_STOP watermark.

2.10.8 TAIL DROPPING

The shared queue system implements a tail dropping mechanism where incoming frames are discarded if the port's memory consumption and the total memory consumption exceed certain watermarks. Tail dropping implies that the frame is discarded unconditionally. All ports in the device are subject to tail dropping. It is independent of whether the port is in flow control mode or in drop mode.

Tail dropping can be effective under special conditions. For example, tail dropping can prevent an ingress port from consuming all the shared memory when pause frames are lost or when the link partner is not responding to pause frames.

The shared queue system initiates tail dropping by discarding the incoming frame if the following two conditions are met at any point while writing the frame data to the memory.

- If the Ingress port memory consumption exceeds the SYS:PORT:ATOP_CFG.ATOP watermark
- If the total consumed memory in the shared queue system exceeds the SYS:PORT:ATOP_TOT_CFG.ATOP_TOT watermark

2.10.9 TEST UTILITIES

This section describes some of test utilities that are built into the shared queue system.

Each egress port can enable a frame repeater (SYS::REPEATER), which means that the head-of-line frames in the egress queues are transmitted but not dequeued after transmission. As a result, the scheduler sees the same frames again and again while the repeater function is active.

The SYS:PORT:PORT_MODE.DEQUEUE_DIS disables both transmission and dequeuing from the egress queues when set.

2.10.10 ENERGY EFFICIENT ETHERNET


This section provides information about the functions of Energy Efficient Ethernet in the shared queue system. The following table lists the registers associated with Energy Efficient Ethernet.

TABLE 2-124: ENERGY EFFICIENT ETHERNET CONTROL REGISTERS

Register	Description	Replication
DEV::EEE_CFG	Enables configuration of Energy Efficient Ethernet. Status bit indicating that egress port is in the LPI state.	Per port
QSYS::EEE_CFG	Configures fast queues.	Per port
QSYS::EEE_THRES	Configures bytes and frame thresholds.	None

The shared queue system supports Energy Efficient Ethernet (EEE) as defined by IEEE Draft P802.3az by initiating the Low Power Idle (LPI) mode during periods of low link use. EEE is controlled per port by an egress queue state machine that monitors the queue fillings and ensures correct wake-up and sleep timing. The egress queue state machine is responsible for informing the PCS in the port module of changes in EEE states (active, sleep, low power idle, and wake up).

FIGURE 2-27: LOW POWER IDLE OPERATION



Energy Efficient Ethernet is enabled per port through DEV::EEE_CFG.EEE_ENA.

By default, the egress port is transmitting enqueued data. This is the active state. If none of the port's egress queues have enqueued data for the time specified in DEV::EEE_CFG.EEE_TIMER_HOLDOFF, the egress port instructs the PCS to enter the EEE sleep state.

When data is enqueued in any of the port's egress queues, a timer (DEV::EEE_CFG.EEE_TIMER_AGE) is started. If one of the following conditions is met, the port enters the wake up state.

- A queue specified as high priority (QSYS:PORT:EEE_CFG.EEE_FAST_QUEUES) has any data to transmit.
- The total number of frames in the port's egress queues exceeds QSYS::EEE_THRES.EEE_HIGH_FRAMES.
- The total number of bytes in the port's egress queues exceeds QSYS::EEE_THRES.EEE_HIGH_FRAMES.
- The time specified in DEV::EEE_CFG.EEE_TIMER_AGE has passed. PCS is instructed to wake up.

To ensure that PCS, PHY, and link partner are resynchronized after waking up; the egress port holds back transmission of data until the time specified in DEV::EEE_CFG.EEE_TIMER_WAKEUP has passed. After this time interval, the port resumes transmission of data.

The status bit DEV::EEE_CFG.PORT_LPI is set while the egress port holds back data due to LPI (from the sleep state to the wake up state, both included).

2.11 Scheduler and Shapers

The following table lists the registers associated with the scheduler and egress shaper control.

TABLE 2-125: SCHEDULER AND EGRESS SHAPER CONTROL REGISTERS

Register	Description	Replication
QSYS::QMAP	Configuration of the mapping of frames to egress queues.	Per port
QSYS::ISDX_SGRP	Configuration of Service group index used in service-based egress scheduling.	Per ISDX
QSYS::CIR_CFG	Configuration of committed information rate and burst.	692
QSYS::EIR_CFG	Configuration of excess information rate and burst.	256
QSYS::SE_CFG	Configuration of the scheduler mode.	256
QSYS::SE_DWRR_CFG	Configuration of DWRR scheduler's costs.	256
QSYS::SE_CONNECT	Configuration of the hierarchial scheduler.	256
QSYS::SE_DLB_SENSE	Configuration of DLB shaper.	256
QSYS::CIR_STATE	Status of the committed information rate bucket.	692
QSYS::EIR_STATE	Status of the excess information rate bucket.	256


The egress scheduler consists of a 256 scheduler elements (SE). Each scheduler elements has 12 inputs and 1 output. It contains a scheduler, which can be strict or mixed with a round robin based scheduling algorithm. The output port of a scheduler elements contains a dual leaky bucket shaper.

The scheduler elements can be connected together through configuration to achieve hierarchical Quality-of-Service (H-QoS) with flexible queueing and multiple levels of scheduling. Examples of H-QoS schedulers are:

- **Two-level priority-based fairness.** First level (towards the egress port) schedules between QoS classes while second level (towards the egress queues) schedules between the ingress ports. This is the default configuration of the device.
- **Two-level port-based fairness.** First level (towards the egress port) schedules between ingress ports while second level (towards the egress queues) schedules between the QoS classes.
- **Three-level service-group-based fairness.** First and second level (towards the egress port) schedules between up to 96 service groups while third level (towards the egress queues) schedules between the QoS classes. A service group can for instance be an EVC or an LSP.

The following illustration provides an overview of the egress scheduling system.

FIGURE 2-28: EGRESS SCHEDULER OVERVIEW



The queue system contains 2616 egress queues. The egress queues connect to the inputs on the first 218 scheduler elements in the following way: Inputs 0 through 11 on scheduler element n connect to queues $12 \times n$ through $12 \times n + 11$, where $n = 0, \dots, 217$. This is fixed and cannot be reconfigured.

The outputs of the 218 fixed allocated scheduler elements can connect either directly to an egress port or to an input of one of the remaining 38 scheduler elements. The optional use of 38 remaining scheduler elements is fully configurable. They can for instance all be allocated to one or two ports to achieve advanced three level scheduling.


Each frame is stored in an egress queue for each of the frame's destination ports. The mapping to one of the 2616 egress queues is based on variety of frame properties such as source port, destination port, QoS class, and a service group index. The service group index is a number between 0 and 127 based on the frame's ISDX.

The egress scheduler contains 436 single leaky bucket shapers. The shapers are placed at input 6 and 7 for each of the 218 fixed allocated scheduler elements. These shapers are primarily applicable when the queue mapping is based on the frame's QoS class. In such a scenario, these shapers allow the user to configure a maximum rate for QoS class 6 and 7.

2.11.1 SCHEDULER ELEMENT

The following illustration shows the scheduler element.

FIGURE 2-29: SCHEDULER ELEMENT



Each of the 256 scheduler elements has 12 inputs and 1 output. The inputs can either be connected to frame queues or to other scheduler elements. The output connects to either another scheduler element or an egress port. The scheduler element contains a scheduler and a dual leaky bucket shaper that control which input is scheduled to the output port.

By default, the scheduler operates in strict priority between all 12 inputs. The inputs are searched in the following prioritized order: Input 11 has highest priority followed by 10, 9, 8, 7, 6, 5, 4, 3, 2, 1, and 0.

In addition, the scheduler can operate either in a byte-based deficit weighted round robin (DWRR) mode or a frame-based round robin (RR) mode. This is an overall configuration for the scheduler element and cannot be selected per input. In DWRR mode, the participating inputs are given a weight and the scheduler selects frames from these inputs according to the weights. The DWRR is byte-based and takes the lengths of the frames into account. In RR mode, the participating inputs are selected one after another. The RR is frame-based and does not take the length of the frames into account.

The scheduler supports a mixed mode where some inputs operate in strict priority and others operate in either DWRR or RR. Any number of inputs can be assigned to either group but strict priority inputs must always be selected from highest numbered inputs.

Each scheduler element has an associated dual leaky-bucket shaper at the output. The shaper limits the overall transmission bandwidth from the scheduler element. Frames are only scheduled if the shaper is open. The dual leaky bucket shaper operates a single leaky bucket shaper but has the option to push forward with a excess information rate when the queue system is filling up. The dual leaky bucket shaper can also operate as a normal single leaky bucket shaper.

Each scheduling element determines whether it has a frame ready for scheduling based on status information of the scheduling element's inputs. For each input, the scheduling element knows if the input has a frame ready for transmission and if the frame is ready due to a work conservation mode. The overall scheduling algorithm within a scheduling element is as follows:

1. If the output shaper is closed, frames are only scheduled from the element if the element is enabled for work conservation. Otherwise, frames are held back until the shaper opens.
2. If the output shaper is open or in work conservation mode, the scheduling element schedules between inputs that are not work conserving by following the rules for the scheduler configuration: strict inputs are scheduled first then round robin based inputs are scheduled according to either the DWRR algorithm or the RR algorithm.
3. If no frames are scheduled during step 2, a second round of scheduling is performed. Inputs that are work conserving become candidates for the second round of scheduling.

The hierarchy of scheduling elements is traversed from the element connected to the egress port through to the element that connects to an egress queue by recursively deciding which input should be scheduled.

2.11.1.1 Configuration of the Egress Scheduler Hierarchy

The 256 scheduler elements are connected to each other through the configuration in QSYS::SE_CONNECT.

The 218 fixed allocated scheduler elements can specify the following options:


- **Output.** Connect the SE's output to an egress port (SE_TERMINAL, SE_OUTP_IDX) or connect the SE's output to an input on another SE from the pool of shared SEs (SE_OUTP_IDX, SE_OUTP_CON).
- **Number of inputs.** Specify how many of the inputs on the SE are active (SE_INP_CNT). The inputs used are always a consecutive range starting from 0 through SE_INP_CNT.

The 38 scheduler elements in the pool of shared scheduler elements can specify the following options:

- **Output.** Connect the SE's output to an egress port (SE_TERMINAL, SE_OUTP_IDX) or connect the SE's output to an input on another SE from the pool of shared SEs (SE_OUTP_IDX, SE_OUTP_CON).
- **Input.** Specify the SE, which output connects to input 0 on this SE (SE_INP_IDX). Note that the SEs that connect to this SE's inputs must be allocated as a consecutive range starting from SE_INP_IDX to SE_INP_IDX+SE_INP_CNT-1.
- **Number of inputs.** Specify how many of the inputs on the SE are active (SE_INP_CNT). The inputs used are always a consecutive range starting from 0 through SE_INP_CNT.

The following illustration is an example of a H-QoS configuration.

FIGURE 2-30: H-QOS CONFIGURATION



2.11.2 QUEUE MAPPING

For each egress port, there is a programmable mapping of frames to egress queues. Since the egress queues are fixed connected to the first 218 scheduler elements, the queue mapping function (QSYS::QMAP) does not output a queue number but rather a scheduler element number (0 through 217) and an input port number on the scheduler element (0 through 11). The scheduling element number is calculated as follows:

- SE number = QSYS::QMAP.SE_BASE + QSYS::QMAP.SE_IDX_SEL
- SE input port number = QSYS::QMAP.SE_INP_SEL

where QSYS::QMAP.SE_BASE is a constant offset configured per egress port while QSYS::QMAP.SE_IDX_SEL and QSYS::QMAP.SE_INP_SEL are based on frame properties. The following options are available for both SE_IDX_SEL and SE_INP_SEL:

- **Source port.** Use the frame's source port (0 through 11)
- **QoS class.** Use the frame's classified QoS class (0 through 7)
- **Service group number.** Use the frame's service group number (0 through 127). The service group number is a direct mapping from the frame's ISDX to the service group number through the table in QSYS::ISDX_SGRP.
- **Constrained service group number.** Use the frame's service group number modulus 16 (0 through 15). The service group number is a direct mapping from the frame's ISDX to the service group number through the table in QSYS::ISDX_SGRP.

The selector SE_IDX_SEL also have an option to not contributing to SE_BASE. This is used when for instance an egress port has only one scheduler element associated in which case SE_BASE selects the scheduler element.

Both selectors SE_IDX_SEL and SE_INP_SEL are independently programmed per egress port.


The mapping of frames to SE number must be programmed so that the SE number does not exceed 217. Similarly, the mapping of frames to the input port number must be programmed so that the input port number does not exceed 11.

Example: Configuration of a Priority-Fair Egress Scheduler

By default, the egress scheduler is configured the same way for all egress ports. Each port features a two-level priority-fair egress scheduler where the first level of scheduling towards the egress port schedules between QoS classes while second level (towards the egress queues) schedules between the source ports. The queue mapping in this scenario is configured so that for instance egress port 0 uses scheduler elements 0 through 7 with the frame's QoS class as SE selector and the frame's source port as selector for the SE input port number on the scheduler element. The configuration is as follows:

- Egress port n :
 - QSYS:QMAP.SE_BASE = $8 \times n$
 - QSYS:QMAP.SE_IDX_SEL = 5 (use frame's QoS class as SE selector)
 - QSYS:QMAP.SE_INP_SEL = 0 (use frame's source port as input port number).

The following illustration shows the default configuration.

FIGURE 2-31: PRIORITY-FAIR SCHEDULER FOR EGRESS PORT 0

The result of the priority-fair scheduler is that high priority frames take precedence over low priority frames, and for each priority, the ingress ports are serviced with a frame-based round robin algorithm.

2.11.3 EGRESS SHAPERS

Each of the 256 scheduling elements contain an output shaper that can operate as either a single leaky bucket shaper or a dual leaky bucket shaper.

By default, the shaper operates as a single leaky bucket shaper with the following configurations:

- Maximum rate. Specified in CIR_CFG.CIR_RATE in steps of 100160 bps. Maximum is 3.282 Gbps.
- Maximum burst size. Specified in CIR_CFG.CIR_BURST in steps of 4 kilobytes. Maximum is 252 kilobytes.

The shaper can operate byte-based or frame-based (SE_CFG.SE_FRM_MODE). When operating as byte-based, the frame adjustment value HSCH_MISC.FRM_ADJ can be used to program the fixed number of extra bytes to add to each frame transmitted (irrespective of QoS class) in the shaper and DWRR calculations. A value of 20 bytes corresponds to line-rate calculation and accommodates for 12 bytes of inter-frame gap and 8 bytes of preamble. Data-rate based shaping and DWRR calculations are achieved by programming 0 bytes.

By default, while a shaper in a scheduler element is closed, frames from the scheduler element's inputs are not scheduled to the output, even if none of the other scheduler elements have frames to transmit. Each scheduler element can enable a work-conserving mode (SE_CFG.SE_EXC_ENA) in which a second scheduling round is possible. If none of

the scheduler elements with an open shaper have frames for transmission, work-conserving scheduler elements with closed shapers may get a share of the excess bandwidth. The sharing of the excess bandwidth obeys the same configured scheduling rules as for the first round of scheduling.

Each shaper implements two burst modes. By default, a leaky bucket is continuously assigned new credit according to the configured shaper rate. This implies that during idle periods, credit is building up, which allows for a burst of data when there are again data to transmit. This is not convenient in an Audio/Video Bridging (AVB) environment where this behavior enforces a requirement for larger buffers in end-equipment. To circumvent this, each shaper can enable an AVB mode (`SE_CFG.SE_AVB_ENA`) in which credit is only assigned during periods where the scheduler element has data to transmit and is waiting for another scheduler element to finish a transmission. This AVB mode prevents the accumulation of large amount of credits.

The shapers can also operate as a dual leaky bucket designed to comply with MEF standards. By default, the dual leaky bucket shaper forwards frames with the committed information rate and holds back frames in the queue system when CIR is exceeded. However, the dual leaky bucket functionality allows the shaper to push traffic forward with a higher bandwidth (the excess information rate) if the queue system is filling up.

In addition to configuration of the committed information rate in `CIR_CFG`, the excess information rate can be configured as follows:

- Maximum excess rate – Specified in `EIR_CFG.EIR_RATE` in steps of 100160 bps. Maximum is 3.282 Gbps.
- Maximum excess burst size – Specified in `EIR_CFG.EIR_BURST` in steps of 4 kilobytes. Maximum is 252 kilobytes.

The excess information rate is used when certain fill levels in the queue system are exceeded. The following modes for triggering excess information rate are available:

- Priority-based (`SE_DLB_SENSE.SE_DLB_PRIO_ENA=1`) - Excess information rate is used when the egress buffer use for the QoS class defined in `SE_DLB_SENSE.SE_DLB_PRIO` exceeds its threshold `BUF_PRIO_SHR_E` minus 3000 bytes.
- Source-based (`SE_DLB_SENSE.SE_DLB_SPORT_ENA=1`) - Excess information rate is used when the ingress buffer use for the source port defined in `SE_DLB_SENSE.SE_DLB_SPORT` exceeds its threshold `BUF_P_RSRV_I` minus 3000 bytes.
- Destination-based (`SE_DLB_SENSE.SE_DLB_DPORT_ENA=1`) - Excess information rate is used when the egress buffer use for the destination port defined in `SE_DLB_SENSE.SE_DLB_DPORT` exceeds its threshold `BUF_P_RSRV_E` minus 3000 bytes.
- Source/Priority-based (`SE_DLB_SENSE.SE_DLB_PRIO_ENA=1` and `SE_DLB_SENSE.SE_DLB_SPORT_ENA=1`) - Excess information rate is used when the ingress buffer use for the QoS class defined in `SE_DLB_SENSE.SE_DLB_PRIO` for the source port defined in `SE_DLB_SENSE.SE_DLB_SPORT` exceeds its threshold `BUF_Q_RSRV_I` minus 3000 bytes.
- Destination/Priority-based (`SE_DLB_SENSE.SE_DLB_PRIO_ENA=1` and `SE_DLB_SENSE.SE_DLB_DPORT_ENA=1`) - Excess information rate is used when the egress buffer use for the QoS class defined in `SE_DLB_SENSE.SE_DLB_PRIO` for the destination port defined in `SE_DLB_SENSE.SE_DLB_DPORT` exceeds its threshold `BUF_Q_RSRV_E` minus 3000 bytes.

Multiple modes can be enabled at the same time in which case all thresholds must be exceeded before the excess information rate is allowed. Note that if the source/priority-based mode is enabled then priority-based and source-based are also enabled. The same applies to destination/priority-based and destination-based.

2.11.4 DEFICIT WEIGHTED ROUND ROBIN

The DWRR uses a cost-based algorithm compared to a weight-based algorithm. A high cost implies a small share of the bandwidth. DWRR is enabled when `SE_CFG.SE_DWRR_CFG>0` and `SE_CFG.RR_ENA = 0`. The participating inputs are then inputs 0 through `SE_CFG.SE_DWRR_CFG-1`. Anything from 0 to 12 weighted inputs can participate.

Each input is programmed with a cost (`SE_DWRR_CFG.DWRR_COST`). A cost is a number between 1 and 32. The programmable DWRR costs determine the behavior of the DWRR algorithm. The costs result in weights for each input to the scheduler element. The weights are relative to one another, and the resulting share of the egress bandwidth for a particular input is equal to the input's weight divided by the sum of all the inputs' weights. The algorithm is byte-based and takes the frame lengths into account.

Costs are easily converted to weights and vice versa given the following two algorithms. The following algorithms are shown with six participating inputs but can be applied to other configurations as well.

Weights to Costs Given a desired set of weights ($W_0, W_1, W_2, W_3, W_4, W_5$), the costs can be calculated using the following algorithm.

1. Set the cost of the queue with the smallest weight (W_{smallest}) to cost 32.
2. For any other queue Q_n with weight W_n , set the corresponding cost C_n to $C_n = 32 \times W_{\text{smallest}}/W_n$

Costs to Weights Given a set of costs for all queues ($C_0, C_1, C_2, C_3, C_4, C_5$), the resulting weights can be calculated using the following algorithm:

1. Set the weight of the queue with the highest cost (C_{highest}) to 1.
2. For any other queue Q_n with cost C_n , set the corresponding weight W_n to $W_n = C_{\text{highest}}/C_n$

Cost and Weight Conversion Examples

Implement the following bandwidth distributions.

- Input 0: 5% ($W_0 = 5$)
- Input 1: 10% ($W_1 = 10$)
- Input 2: 15% ($W_2 = 15$)
- Input 3: 20% ($W_3 = 20$)
- Input 4: 20% ($W_4 = 20$)
- Input 5: 30% ($W_5 = 30$)

Given the algorithm to get from weights to costs, the following costs are calculated:

- $C_0 = 32$ (Smallest weight)
- $C_1 = 32 \times 5/10 = 16$
- $C_2 = 32 \times 5/15 = 10.67$ (rounded up to 11)
- $C_3 = 32 \times 5/20 = 8$
- $C_4 = 32 \times 5/20 = 8$
- $C_5 = 32 \times 5/30 = 5.33$ (rounded down to 5)

Due to the rounding off, these costs result in the following bandwidth distribution, which is slightly off compared to the desired distribution:

- Input 0: 4.92%
- Input 1: 9.85%
- Input 2: 14.32%
- Input 3: 19.70%
- Input 4: 19.70%
- Input 5: 31.51%

2.11.5 ROUND ROBIN

The round robin (RR) uses a simple round robin algorithm where each participating input is served one after another. RR is enabled when $SE\_CFG.SE\_DWRR\_CFG > 0$ and $SE\_CFG.RR\_ENA = 1$. The participating inputs are then inputs 0 through $SE\_CFG.SE\_DWRR\_CFG - 1$. Anything from 0 to 12 weighted inputs can participate.

The RR algorithm is frame-based and does not take the frame lengths into account.

2.11.6 SHAPING AND DWRR SCHEDULING EXAMPLES

This section provides examples and additional information about the use of the egress shapers and scheduler. The following assumes a scheduler element is connected to the egress queues.

2.11.6.1 Mixing DWRR and Shaping Example

- Output from scheduler element is shaped down to 500 Mbps.
- Queues 7 and 6 are strict and queues 5 through 0 are weighted.
- Queue 7 is shaped to 100 Mbps.
- Queue 6 is shaped to 50 Mbps.
- The following traffic distribution is desired for queue 5 through 0:
Q0: 5%, Q1: 10%, Q2: 15%, Q3: 20%, Q4: 20%, Q5: 30%.

- Each queue receives 125 Mbps of incoming traffic.

The following table lists the DWRR configuration and the resulting egress bandwidth for the various queues.

TABLE 2-126: EXAMPLE OF MIXING DWRR AND SHAPING

Queue	Distribution of Weighted Traffic	Configuration Costs/Weights (Cn/Wn)	Result: Egress Bandwidth
Q0	5%	32 / 1	$1/(1+2+2.9+4+4+6.4) \times (500 - \text{Mbps} - 150 \text{ Mbps}) = 17.2 \text{ Mbps}$
Q1	10%	16 / 2	$2/(1+2+2.9+4+4+6.4) \times (500 - \text{Mbps} - 150 \text{ Mbps}) = 34.5 \text{ Mbps}$
Q2	15%	11 / 2.9	$2.9/(1+2+2.9+4+4+6.4) \times (500 - \text{Mbps} - 50 \text{ Mbps}) = 50.1 \text{ Mbps}$
Q3	20%	8 / 4	$4/(1+2+2.9+4+4+6.4) \times (500 - \text{Mbps} - 150 \text{ Mbps}) = 68.9 \text{ Mbps}$
Q4	20%	8 / 4	$4/(1+2+2.9+4+4+6.4) \times (500 - \text{Mbps} - 150 \text{ Mbps}) = 68.9 \text{ Mbps}$
Q5	30%	5 / 6.4	$6.4/(1+2+2.9+4+4+6.4) \times (500 - \text{Mbps} - 150 \text{ Mbps}) = 110.3 \text{ Mbps}$
Q6			50 = Mbps
Q7			100 = Mbps
Sum:	100%		500 = Mbps

Strict and Work-Conserving Shaping Example

- Output from scheduler element is shaped down to 500 Mbps.
- All queues are strict.
- All queues are shaped to 50 Mbps.
- Queues 6 and 7 are work-conserving (allowed to use excess bandwidth).
- All queues receive 125 Mbps of traffic each.

The following table lists the resulting egress bandwidth for the various queues.

TABLE 2-127: EXAMPLE OF STRICT AND WORK-CONSERVING SHAPING

Queue	Result: Egress Bandwidth
Q0	50 Mbps
Q1	50 Mbps
Q2	50 Mbps
Q3	50 Mbps
Q4	50 Mbps
Q5	50 Mbps
Q6	75 Mbps (Gets the last 25 Mbps of the 100 Mbps in excess not used by queue 7)
Q7	125 Mbps (Gets 75 Mbps of the 100 Mbps in excess limited only by the received rate)
Sum:	500 Mbps

2.12 Rewriter

VSC7416-01 includes a rewriter common for all ports that determines how the egress frame is edited before transmission. The rewriter performs the following editing:

- VLAN editing; tagging of frames and remapping of PCP and DEI.
- DSCP remarking; rewriting the DSCP value in IPv4 and IPv6 frames based on classified DSCP value.
- FCS updating.
- Precision Time Protocol timestamp updating.
- MPLS editing; encapsulation of MPLS link layer and MPLS label stack.
- OAM editing.

Each port module including the CPU port module (CPU port 11 and CPU port 12) has its own set of configuration in the rewriter. Each frame is handled by the rewriter one time per destination port.

Most rewriting functions in the rewriter are independent of each other and can co-exist. For instance, a frame can be MPLS encapsulated while at the same time being DSCP remarked. However, precision time protocol timestamp updating and DSCP remarking are mutually exclusive with the former taking precedence. If an IS2 action returns any PTP actions then DSCP remarking is automatically disabled for the frame.

2.12.1 VLAN EDITING

The following table lists the registers associated with VLAN editing.

TABLE 2-128: VLAN EDITING REGISTERS

Register	Description	Replication
PORT_VLAN_CFG	Port VLAN for egress port. Used for untagged set.	Per port
TAG_CFG	Tagging rules for port tag.	Per port
PORT_CFG.ESO_ENA	Enable lookups in ES0.	Per port
PCP_DEI_QOS_MAP_CFG	Mapping table. Maps DP level and QoS class to new PCP and DEI values.	Per port per QoS per DP

The rewriter performs five steps related to VLAN editing for each frame and destination:

1. VLAN popping - Zero, one, or two VLAN tags are popped from the frame.
2. ES0 lookup - ES0 is looked up for each of the frame's destination ports. The action from ES0 controls the pushing of VLAN tags.
3. VLAN push decision - Deciding the number of new tags to push and which tag source to use for each tag. Tag sources are: Port and ES0 (tag A and tag B).
4. Constructing the VLAN tags - The new VLAN tags are constructed based on the tag sources' configuration.
5. VLAN pushing - the new VLAN tags are pushed.

2.12.1.1 VLAN Popping

The rewriter initially pops the number of VLAN tags specified by the VLAN_POP_CNT parameter received with the frame from the classifier or VCAP IS1. Up to two VLAN tags can be popped. The rewriter itself does not influence the number VLAN tags being popped.

2.12.1.2 ES0 Lookup

For each of the frame's destination ports, VCAP ES0 is looked up using the ES0 key. See [Section 2.7.5, VCAP ES0](#) for more information about ES0. The action from an ES0 hit is used in the following to determine the frame's VLAN editing.


2.12.1.3 VLAN Push Decision

After popping the VLAN tags, the rewriter decides whether to push zero, one, or two new VLAN tags to the outgoing frame according to the port's tagging configuration in register TAG_CFG and the action from a potential VCAP ES0 hit. The up to two tags can originate from either the port (port tag) or ES0 (ES0 tag A and ES0 tag B).

By default, the port can push one tag according to the port's configuration in TAG_CFG. If the ES0 lookup results in an entry being hit, the ES0 action can overrule the port configuration and push two tags by itself (ES0 tag A and ES0 tag B) or it can combine port tagging and ES0 tagging.

The following illustration shows an overview of the available tagging options.

FIGURE 2-32: TAGGING OVERVIEW



(1) Port tag is controlled by REW:PORT:TAG_CFG.TAG_CFG.

The following table lists all combinations of port tagging configuration and ES0 actions that control the number of tags to push and which tag source to use (port, ES0 tag A, or ES0 tag B):

TABLE 2-129: TAGGING COMBINATIONS


ES0_ACTION	TAG_CFG	Tagging action
No ES0 hit	Controls port tag	Port tag is pushed according to TAG_CFG as outer tag. Available options are: Tag all frames with port tag. Tag all frames with port tag, except if classified VID=0 or classified VID=PORT_VLAN.PORT_VID. Tag all frames with port tag, except if classified VID=0 No port tag No inner tag.
PUSH_OUTER_TAG=0 PUSH_INNER_TAG=0	Controls port tag	Port tag is pushed according to TAG_CFG as outer tag. Available options are: Tag all frames with port tag. Tag all frames with port tag, except if classified VID=0 or classified VID=PORT_VLAN.PORT_VID. Tag all frames with port tag, except if classified VID=0 No port tag No inner tag.
PUSH_OUTER_TAG=1 PUSH_INNER_TAG=0	Don't care	ES0 tag A is pushed as outer tag. No inner tag.
PUSH_OUTER_TAG=2 PUSH_INNER_TAG=0	Don't care	Port tag is pushed as outer tag. This overrides port settings in TAG_CFG. No inner tag.
PUSH_OUTER_TAG=3 PUSH_INNER_TAG=0	Don't care	No tags are pushed. This overrides port settings in TAG_CFG.
PUSH_OUTER_TAG=0 PUSH_INNER_TAG=1	Controls port tag	Port tag is pushed according to TAG_CFG as outer tag. The following options are available: Tag all frames with port tag. Tag all frames with port tag, except if classified VID=0 or classified VID=PORT_VLAN.PORT_VID. Tag all frames with port tag, except if classified VID=0 No port tag ES0 tag B is pushed as inner tag. ES0 tag B is effectively the outer tag if the port tag is not pushed.
PUSH_OUTER_TAG=1 PUSH_INNER_TAG=1	Don't care	ES0 tag A is pushed as outer tag. ES0 tag B is pushed as inner tag.

TABLE 2-129: TAGGING COMBINATIONS (CONTINUED)

ES0_ACTION	TAG_CFG	Tagging action
PUSH_OUTER_TAG=2 PUSH_INNER_TAG=1	Don't care	Port tag is pushed as outer tag. This overrides port settings in TAG_CFG. ES0 tag B is pushed as inner tag.
PUSH_OUTER_TAG=3 PUSH_INNER_TAG=1	Don't care	No outer tag is pushed. This overrides port settings in TAG_CFG. ES0 tag B is pushed as inner tag. ES0 tag B is effectively the outer tag, because no outer tag is pushed.

2.12.1.4 Constructing VLAN Tags

When pushing a VLAN tag, the contents of the tag header, including the TPID, is highly programmable. The starting point is the classified tag header coming from the analyzer containing the PCP, DEI, VID and tag type. For each of the fields in the resulting tag, it is programmable how the value is determined. For more information, see [Figure 2-32](#).

FIGURE 2-33: TAG CONSTRUCTION (PORT TAG, ES0 TAG A, ES0 TAG B)

(*1) For port tag, this is classified only.

(*2) Use 0x8100 if classified tag type is 0x8100, otherwise custom.

The port tag, ES0 tag A, and ES0 tag B have individual configurations. For the port tag, the available tag field options are:

Port tag: PCP

- Use the classified PCP.
- Use the egress port's port VLAN (PORT_VLAN.PORT_PCP).
- Map the DP level and QoS class to a new PCP value using the per-port table PCP_DEI_QOS_MAP_CFG.
- Use the QoS class directly.

Port tag: DEI

- Use the classified DEI.
- Use the egress port's port VLAN (PORT_VLAN.PORT_DEI).
- Map the DP level and QoS class to a new DEI value using the per-port table PCP_DEI_QOS_MAP_CFG.
- Use the DP level directly.

Port tag: VID

- Use the classified VID.
- Use the egress port's port VLAN (PORT_VLAN.PORT_VID).

Port tag: TPID

- Use Ethernet type 0x8100 (C-tag).
- Use Ethernet type 0x88A8 (S-tag).
- Use custom Ethernet type programmed in PORT_VLAN.PORT_TPID.
- Use custom Ethernet type programmed in PORT_VLAN.PORT_TPID unless the incoming tag was a C-tag in

which case Ethernet type 0x8100 is used.

Similar options for the ES0 tag A and ES0 tag B are available:

ES0 tag: PCP

- Use the classified PCP.
- Use ES0_ACTION.PCP_A_VAL for ES0 tag A and use ES0_ACTION.PCP_B_VAL for ES0 tag B.
- Map the DP level and QoS class to a new PCP using the per-port table PCP_DEI_QOS_MAP_CFG.
- Use the QoS class directly.

ES0 tag: DEI

- Use the classified DEI.
- Use ES0_ACTION.DEI_A_VAL for ES0 tag A and use ES0_ACTION.DEI_B_VAL for ES0 tag B.
- Map the DP level and QoS class to a new DEI using the per-port table PCP_DEI_QOS_MAP_CFG.
- Use the DP level directly.

ES0 tag: VID

- Use the classified VID incremented with ES0_ACTION.VID_A_VAL for ES0 tag A and use the classified VID incremented with ES0_ACTION.VID_B_VAL for ES0 tag B.
- Use ES0_ACTION.VID_A_VAL for ES0 tag A and use ES0_ACTION.VID_B_VAL for ES0 tag B.

ES0 tag: TPID

- Use Ethernet type 0x8100 (C-tag).
- Use Ethernet type 0x88A8 (S-tag).
- Use custom Ethernet type programmed in PORT_VLAN.PORT_TPID.
- Use custom Ethernet type programmed in PORT_VLAN.PORT_TPID unless the incoming tag was a C-tag in which case Ethernet type 0x8100 is used.

2.12.1.5 VLAN Pushing

In the final VLAN editing step, the VLAN tags derived from the previous steps are pushed to the frame.

2.12.2 DSCP REMARKING

The following table lists the registers associated with DSCP remarking.

TABLE 2-130: DSCP REMARKING REGISTERS

Register	Description	Replication
DSCP_CFG	Selects how the DSCP remarking is done	Per port
DSCP_REMAP_CFG	Mapping table from DSCP to DSCP for DP level = 0.	None
DSCP_REMAP_DP1_CFG	Mapping table from DSCP to DSCP for DP level = 1.	None

The rewriter can remark the DSCP value in IPv4 and IPv6 frames, that is, write a new DSCP value to the DSCP field in the frame.

If a port is enabled for DSCP remarking (DSCP_CFG.DSCP_REWR_CFG), the new DSCP value is derived by using the classified DSCP value from the analyzer (the basic classification or the VCAP IS1) in the ingress port. This DSCP value can be mapped before replacing the existing value in the frame. The following options are available:

- No DSCP remarking - Leave the DSCP value in the frame untouched.
- Update the DSCP value in the frame with the value received from the analyzer
- Update the DSCP value in the frame with the value received from the analyzer remapped through DSCP_REMAP_CFG. This is done independently of the value of the drop precedence level.
- Update the DSCP value in the frame with the value received from the analyzer remapped through DSCP_REMAP_CFG or DSCP_REMAP_DP1_CFG dependent on the drop precedence level. This enables one mapping for green frames and another for yellow frames so that the resulting DSCP value can reflect the color of the frame.

In addition, the IP checksum is updated for IPv4 frames. Note that the IPv6 header does not contain a checksum. As a result, checksum updating does not apply for IPv6 frames.

DSCP remarking is not possible for frames where PTP timestamps are also generated and is automatically disabled.

2.12.3 FCS UPDATING

The following table lists the registers associated with FCS updating.

TABLE 2-131: FCS UPDATING REGISTERS

Register	Description	Replication
PORT_CFG.FCS_UPDATE_NON-CPU_CFG	FCS update configuration for non-CPU injected frames.	Per port
PORT_CFG.FCS_UPDATE_CPU_ENA	FCS update configuration for CPU injected frames.	Per port

The rewriter updates a frame's FCS when required or instructed to do so. Different handling is available for frames injected by the CPU and for all other frames.

For non-CPU injected frames, the following update options are available:

- Never update the FCS.
- Conditional update: Update the FCS if the frame was modified due to PTP timestamping, VLAN tagging, or DSCP remarking.
- Always update the FCS.

In addition, the rewriter can update the FCS for all frames injected from the CPU through the CPU injection queues in the CPU port module:

- Never update the FCS.
- Always update the FCS.

2.12.4 MPLS EDITING

The following table lists the registers associated with MPLS editing.

TABLE 2-132: MPLS EDITING REGISTERS

Register	Description	Replication
REW:PORT:PORT_CFG.ESO_ENA	Enable lookups in ES0.	Per port
SYS::MPLS_QOS_MAP_CFG	Mapping table. Per eight TC selectors, maps DP level and QoS class to new TC level. The TC selector is given by the MPLS encapsulation configuration word.	16
SYS::ENCAP_CTRL	Controls access to the MPLS encapsulation table.	None
SYS::ENCAP_DATA	MPLS encapsulation data in 10 x 32-bit words. It contains 48-bit MPLS encapsulation configuration and 272-bit MPLS encapsulation data (MPLS Link Layer header and MPLS label stack). Up to 1K MPLS encapsulation entries are supported.	10

The rewriter performs five steps related to MPLS editing for each frame and destination:

1. ES0 lookup. ES0 is looked up for each of the frame's destination ports. The action from ES0 controls the MPLS editing.
2. MPLS popping. The existing MPLS link layer header and MPLS label stack is popped
3. MPLS encapsulation table lookup using the ES0 actions. This table returns the new encapsulation as well as configuration that controls MPLS remarking.
4. MPLS pushing. A new MPLS link layer header and MPLS label stack is pushed
5. MPLS remarking. The VLAN tag header in the MPLS link layer and MPLS labels are remarked

For each of the frame's destination ports, VCAP ES0 is lookup using the appropriate key, see [Section 2.7.5, VCAP ES0](#). The action from an ES0 hit is used in the following to determine the frame's MPLS editing. The relevant MPLS ES0 actions are ENCAP_ID, ENCAP_LEN, and MPOP_CNT. If there is no ES0 hit, then the rewriter can only do MPLS popping because MPLS pushing and MPLS remarking require an ES0 hit.

The rewriter initially pops a number of bytes from the MPLS link layer header and MPLS label stack as specified by either the MPLS_POP parameter received with the frame from the MPLS-TP classifier or by the ES0 action MPOP_CNT where ES0 takes precedence if non-zero. Up to 34 bytes (the MPLS link layer header, a VLAN tag, three MPLS labels, and a control word) can be popped from the frame.

The MPLS encapsulation table is programmed with SYS::ENCAP_CTRL and SYS::ENCAP_DATA. The use of the MPLS encapsulation table is triggered by ES0 action ENCAP_ID. The index into the MPLS encapsulation table (SYS::ENCAP_CTRL.ENCAP_ID) must correspond to the ENCAP_ID used in the associated ES0 action. Note that ENCAP_ID = 0 is not available.

The MPLS encapsulation table stores the new MPLS link layer header and the MPLS label stack to push as well as configuration that controls remarking of the certain fields in the pushed data. The overall layout of the MPLS encapsulation table is illustrated in [Figure 2-34](#), . The figure also illustrates how the possible MPLS encapsulations are programmed into the entry.

	idx: 0	idx: 1	idx: 2	idx: 3	idx: 4	idx: 5	idx: 6	idx: 7	idx: 8	idx: 9
ENCAP_DATA	MPLS Remarking Configuration				MPLS Encapsulation Data					
320	288	256	224	192	160	128	96	64	32	0
320	272									0
	DMAC	SMAC	ETYPE							
	DMAC	SMAC	VLAN	ETYPE						
	DMAC	SMAC	ETYPE	LBL0						
	DMAC	SMAC	VLAN	ETYPE	LBL0					
	DMAC	SMAC	ETYPE	LBL0	LBL1					
	DMAC	SMAC	VLAN	ETYPE	LBL0	LBL1				
	DMAC	SMAC	ETYPE	LBL0	CW					
	DMAC	SMAC	VLAN	ETYPE	LBL0	CW				
	DMAC	SMAC	ETYPE	LBL0	LBL1	LBL2				
	DMAC	SMAC	VLAN	ETYPE	LBL0	LBL1	LBL2			
	DMAC	SMAC	ETYPE	LBL0	LBL1	CW				
	DMAC	SMAC	VLAN	ETYPE	LBL0	LBL1	CW			
	DMAC	SMAC	ETYPE	LBL0	LBL1	LBL2	CW			
	DMAC	SMAC	VLAN	ETYPE	LBL0	LBL1	LBL2	CW		

The following table provide information about how the MPLS encapsulation entry is laid out, particularly the MPLS remarking configuration.

TABLE 2-133: MPLS ENCAPSULATION ENTRY

Field Name	Bit	Width	Description
MPLS_ENCAP_DATA	0	272	MPLS encapsulation data. Data is pushed starting from the most significant bits of MPLS_ENCAP_DATA. The data being pushed is inserted in front of the existing frame. The ES0 action ENCAP_LEN defines how many bytes are pushed. Example: Push VLAN-tagged MPLS link layer header with two MPLS labels and a control word: MPLS_ENCAP_DATA[271:224] = DMAC. MPLS_ENCAP_DATA[223:176] = SMAC. MPLS_ENCAP_DATA[175:144] = VLAN. MPLS_ENCAP_DATA[143:128] = ETYPE. MPLS_ENCAP_DATA[127:96] = LABEL0 MPLS_ENCAP_DATA[95:64] = LABEL1. MPLS_ENCAP_DATA[63:32] = CW. MPLS_ENCAP_DATA[31:0] = unused.
TC_SEL_0	272	3	Used when mapping MPLS label 0's TC bits. The frame's DP level and TC_SEL_0 point out which mapping table to use in SYS::MPLS_QOS_MAP_CFG.
TC_MAP_0	275	1	Controls whether to replace the TC bits in MPLS label 0 in MPLS_ENCAP_DATA: If set, the TC bits are replaced with the TC bits resulting from a mapping of the frame's QoS class and DP level, see TC_SEL_0. If cleared and TC_SEL_0[0] is set, the TC bits are replaced with the frame's classified TC bits. MPLS label 0 is the outermost label in the MPLS label stack.
TC_SEL_1	276	3	Used when mapping MPLS label 1's TC bits. The frame's DP level and TC_SEL_1 point out which mapping table to use in SYS::MPLS_QOS_MAP_CFG.
TC_MAP_1	279	1	Controls whether to replace the TC bits in MPLS label 1 in MPLS_ENCAP_DATA: If set, the TC bits are replaced with the TC bits resulting from a mapping of the frame's QoS class and DP level, see TC_SEL_1. If cleared and TC_SEL_1[0] is set, the TC bits are replaced with the frame's classified TC bits.
TC_SEL_2	280	3	Used when mapping MPLS label 2's TC bits. The frame's DP level and TC_SEL_2 point out which mapping table to use in SYS::MPLS_QOS_MAP_CFG.
TC_MAP_2	283	1	Controls whether to replace the TC bits in MPLS label 2 in MPLS_ENCAP_DATA: If set, the TC bits are replaced with the TC bits resulting from a mapping of the frame's QoS class and DP level, see TC_SEL_2. If cleared and TC_SEL_2[0] is set, the TC bits are replaced with the frame's classified TC bits.
RESERVED	284	4	Reserved.
USE_ISBIT_0	288	1	If set, the S-bit in MPLS label 0 in MPLS_ENCAP_DATA is replaced with the frame's classified S-bit from the MPLS-TP classifier.

TABLE 2-133: MPLS ENCAPSULATION ENTRY (CONTINUED)

Field Name	Bit	Width	Description
USE_ISBIT_1	289	1	If set, the S-bit in MPLS label 1 in MPLS_ENCAP_DATA is replaced with the frame's classified S-bit from the MPLS-TP classifier.
USE_ISBIT_2	290	1	If set, the S-bit in MPLS label 2 in MPLS_ENCAP_DATA is replaced with the frame's classified S-bit from the MPLS-TP classifier.
USE_ITTL_0	291	1	If set, the TTL value in MPLS label 0 in MPLS_ENCAP_DATA is replaced with the frame's classified TTL value from the MPLS-TP classifier.
USE_ITTL_1	292	1	If set, the TTL value in MPLS label 1 in MPLS_ENCAP_DATA is replaced with the frame's classified TTL value from the MPLS-TP classifier.
USE_ITTL_2	293	1	If set, the TTL value in MPLS label 2 in MPLS_ENCAP_DATA is replaced with the frame's classified TTL value from the MPLS-TP classifier.
TAGGED	294	1	Set if the MPLS encapsulation data contains a VLAN tag after the SMAC.
USE_IDEI	295	1	If set, the frame's classified DEI replaces the DEI in the VLAN tag in MPLS_ENCAP_DATA. Only valid if TAGGED is set.
USE_IPCP	296	1	If set, the frame's classified PCP replaces the PCP in the VLAN tag in MPLS_ENCAP_DATA. Only valid if TAGGED is set.
USE_DP	297	1	If set, use frame's classified DP as DEI (DEI = DP). This setting overrides the DEI derived from USE_IDEI. Only valid if TAGGED is set.
USE_IVID	298	1	If set, the frame's classified VID replaces the VID in the VLAN tag in MPLS_ENCAP_DATA. Only valid if TAGGED is set.

2.12.4.4 MPLS Pushing

After popping the MPLS data, the rewriter decides whether to push a new MPLS link layer header and a MPLS label stack. This is defined by the ES0 actions ENCAP_LEN and ENCAP_ID. If ENCAP_LEN = 0 or ENCAP_ID = 0, MPLS pushing is effectively disabled.

The ES0 action ENCAP_ID is an index into the MPLS encapsulation table where the contents of the new encapsulation data is stored. The ES0 action ENCAP_LEN defines how many bytes from the MPLS encapsulation table must be pushed. Up to 34 bytes (an MPLS link layer header, a VLAN tag, three MPLS labels, and a control word) can be pushed to the frame. Both the VLAN tag and the control word are optionally, and the number of MPLS labels can be selected.


Given the ENCAP_ID from ES0, the MPLS encapsulation table is looked up and the entry is returned. The entry consists of remarking configuration and the data to push to the frame. The rewriter pushes the number of bytes defined by ENCAP_LEN from the MPLS encapsulation entry.

2.12.4.5 MPLS Remarking


After pushing the new MPLS link layer header and MPLS label stack, the rewriter can remark certain parameters in the pushed data. This is controlled by the MPLS remarking configuration returned by looking up the MPLS encapsulation table. The following illustrates provides an overview of the available remark options.

FIGURE 2-35: MPLS LABELS AND VLAN TAG REMARK OVERVIEW

Remark options for each MPLS label (up to three):



Remark options for MPLS Link Layer VLAN tag:



The remark options are:

- Remark the PCP and DEI bits in the MPLS link layer VLAN tag with the frame's classified PCP and DEI bits.

If USE_IPCP is set (and the MPLS link layer header contains a tag), the PCP received from the classifier together with the frame overwrites the PCP in the pushed VLAN tag.

If USE_IDEI is set (and the MPLS link layer header contains a tag), the DEI received from the classifier together with the frame overwrites the DEI in the pushed VLAN tag.

- Remark the DEI bit in the MPLS link layer VLAN tag with the frame's drop precedence level.

If USE_DP is set (and the MPLS link layer header contains a tag), the DP level received from the classifier together with the frame overwrites the DEI in the pushed VLAN tag. This settings overrules ISE_IDEI settings.

- Remark the VID in the MPLS link layer VLAN tag with the frame's classified VLAN.

If USE_IVID is set (and the MPLS link layer header contains a tag), the VID received from the classifier together with the frame overwrites the VID in the pushed VLAN tag.

- Remark the S-bit in any of the up to three MPLS labels with the frame's classified S-bit from the MPLS-TP classifier. Each label is controlled independently of the others.

If USE_ISBIT_x is set for one of the MPLS labels, the S-bit received from the MPLS-TP classifier with the frame overwrites the S-bit in the pushed MPLS label.

- Remark the S-bit in any of the up to three MPLS labels with the frame's classified S-bit from the MPLS-TP classifier. Each label is controlled independently of the others.

If USE_ISBIT_x is set for one of the MPLS labels, the S-bit received from the MPLS-TP classifier with the frame overwrites the S-bit in the pushed MPLS label.

- Remark the TTL value in any of the up to three MPLS labels with the frame's classified TTL value from the MPLS-TP classifier. Each label is controlled independently of the others.

If USE_ITTL_x is set for one of the MPLS labels, the TTL value received from the MPLS-TP classifier together with the frame overwrites the TTL value in the pushed MPLS label.

- Remark the TC bits in any of the up to three MPLS labels based on a mapping of the frame's QoS class and DP level. Each label is controlled independently of the others.

The mapping of the TC bits uses one of 16 mapping tables (SYS::MPLS_QOS_MAP_CFG). Each mapping table maps the frame's QoS class to a new set of TC bits that are written into the appropriate MPLS label. The mapping table is selected based on the frame's DP level and a TC selector (one of eight) provided by the MPLS remarking

configuration. The selector is available per MPLS label.

For example, to enable rewriting of the TC bits for MPLS label 0, MPLS remarking configuration field TC_MAP_0 must be 1. The field TC_SEL_0 together with the frame's DP level selects which mapping table to use: SYS::MPLS_QOS_MAP_CFG[2 x DP + TC_SEL_0]. The table is looked up with the frame's QoS class and returns the new set of TC bits.

- Remark the TC bits in any of the up to three MPLS labels with the frame's classified TC bits from the MPLS-TP classifier. Each label is controlled independently of the others.

If TC_MAP_x is cleared and bit 0 in TC_SEL_x is set, the TC bits received from the MPLS-TP classifier together with the frame overwrites the TC bits in the pushed MPLS label.

2.12.5 PATH PROTECTION

The rewriter can request that the transmission of a frame is aborted when path protection is in use. The ES0 lookup returns the protection configuration (ES0_ACTION.IPT_CFG), which configures the protection scheme used for the ISDX.

- No protection. The service is not protected and the frame forwarding is not changed.
- Working path. This implies that this ISDX belongs to the working path. If the associated path protection group (ES0_ACTION.PPT_IDX) is using the protected path then the frame transmission is aborted.
- Protected path. This implies that this ISDX belongs to the protected path. If the associated path protection group (ES0_ACTION.PPT_IDX) is using the working path then the frame transmission is aborted.

2.12.6 PTP TIMESTAMPING

The following table lists the registers associated with PTP timestamping.

TABLE 2-134: PTP TIMESTAMPING REGISTERS

Register	Description	Replication
REW:PORT:PTP_CFG	PTP Configuration of egress port	Per port
REW:PORT:PTP_DLY1_CFG	Egress delay configuration	Per port

The rewriter can do various different PTP timestamping in the egress frame:

Residence time. Adds the frame's residence time through the switch to the PTP correction field (byte-offset 8).

Add/subtract. Adds the frame's Tx time to the correction field (byte-offset 8) or subtract the frame's Rx time from the correction field.

Time-of-day. Samples and writes the 80-bit time-of-day into the PTP origin timestamp (byte-offset 34).

Set reserved bytes. Writes the RxTime into the PTP reserved bytes (byte-offset 16).

Clear reserved bytes. Writes zero into the PTP reserved bytes (byte-offset 16).

This can be done for PTP frames over IEEE802.3/Ethernet, PTP frames over UDP over IPv4, and PTP frames over UDP over IPv6. In addition, frames can be MPLS encapsulated and VLAN tagged. The rewriter automatically finds the correct location of the PTP header in the frame.

The rewriter takes the following frame properties from the analyzer as input.

- IS2 rewriter action (IS2_ACTION.REW_OP[2:0]), which can be either one-step, two-step, or origin.
- The frame's Rx timestamp, adjusted for I/O path delays and ingress delays according to ingress actions.
- IS2 rewriter actions for one-step PTP: Add/subtract mode, egress delay adjustment.
- Ingress backplane mode ANA:PORT:PTP_CFG.PTP_BACKPLANE_MODE from the frame's ingress port.

In addition, the following egress port properties are used.

- Egress port delay (REW:PORT:PTP_DLY1_CFG).
- Egress backplane mode (REW:PORT:PTP_CFG.PTP_BACKPLANE_MODE).

The following table shows the resulting rewriter actions for the one-step PTP action, depending on the add/subtract mode and ingress and egress backplane mode settings.

TABLE 2-135: PTP TIMESTAMPING FOR ONE-STEP PTP

PTP Action (from IS2)	Add/Subtract	Ingress Backplane	Egress Backplane	PTP Timestamp Rewriter Actions
One-step	Disabled	Disabled	Disabled	Residence time: Adds frame's residence time to PTP correction field. Egress port delay is added to correction field if specified by IS2.
One-step	Disabled	Disabled	Enabled	Sets reserved bytes: Writes frame's Rx timestamp into the reserved bytes.
One-step	Disabled	Enabled	Disabled	Residence time: Adds frame's residence time to PTP correction field. Egress port delay is added to correction field if specified by IS2. Clear reserved bytes.
One-step	Disabled	Enabled	Enabled	Does not apply.
One-step	Enabled	Disabled	Disabled	Residence time: Adds frame's residence time to PTP correction field. Egress port delay is added to correction field if specified by IS2.
One-step	Enabled	Disabled	Enabled	Add/subtract: Subtracts frame's Rx timestamp from PTP correction field.
One-step	Enabled	Enabled	Disabled	Add/subtract: Adds frame's Tx timestamp to PTP correction field. Egress port delay is added to correction field if specified by IS2. Clear reserved bytes.
One-step	Enabled	Enabled	Enabled	Does not apply.

All actions associated with one-step or origin can be disabled per egress port through REW:PORT:PTP_CFG.PTP_1-STEP_DIS.

The following table shows the resulting rewriter actions for the two-step PTP action, depending on the ingress and egress backplane mode settings.

TABLE 2-136: PTP TIMESTAMPING FOR TWO-STEP PTP

PTP action (from IS2)	Ingress Backplane	Egress Backplane	PTP Timestamp Rewriter Actions
Two-step	Disabled	Disabled	Saves Tx timestamp in PTP timestamp queue.
Two-step	Disabled	Enabled	Sets reserved bytes: Write the frame's Rx timestamp into the reserved bytes.
Two-step	Enabled	Disabled	Saves Tx timestamp in PTP timestamp queue. Clear reserved bytes.
Two-step	Enabled	Enabled	Does not apply.

All actions associated with two-step can be disabled per egress port through REW:PORT:PTP_CFG.PTP_2STEP_DIS.

The following table shows the resulting rewriter actions for the origin PTP action, depending on the ingress and egress backplane mode settings.

TABLE 2-137: PTP TIMESTAMPING FOR ORIGIN PTP

PTP Action (from IS2)	Ingress Backplane	Egress Backplane	PTP Timestamp Rewriter Actions
Origin	Disabled	Disabled	Time-of-day: Write the time-of-day into the PTP origin timestamp field.
Origin	Disabled	Enabled	None.
Origin	Enabled	Disabled	Time-of-day: Write the time-of-day into the PTP origin timestamp field.
Origin	Enabled	Enabled	Does not apply.

The following describes each of the PTP timestamping rewriter actions.

Residence Time. The frame's residence time is calculated as Tx timestamp - Rx timestamp. For information about how the Rx and Tx timestamps are derived, see [Section 2.1.2.8, Rx and Tx Timestamps](#). The residence time is adjusted for an optional egress delay (REW:PORT:PTP_DLY1_CFG) enabled through IS2_ACTION.REW_OP[5] = 1. The resulting residence time is added to the correction field value in the PTP header. The result is written back into the frame.

Add/subtract. Add the frame's Tx time to the correction field or subtract the frame's Rx time from the correction field. This is used in backplane applications. This mode requires a rollover protection mode to handle the situation where the nanosecond counter rolls over during the backplane transfer. The rollover protection mode is configured in SYS::PTP_CFG.PTP_CF_ROLL_MODE and must be the same in both the ingress and the egress backplane unit.

Time-of-Day. The time-of-day consists of a 48-bit seconds part and a 32-bit nanoseconds part. The current time-of-day is derived by sampling the time-of-day seconds counter (SYS::PTP_TOD_LSB, SYS::PTP_TOD_MSB) and the nanoseconds counter (Tx timestamp) at the time of the frame's departure from the switch. Note that the time-of-day value is adjusted for I/O path delay. For more information, see [Section 2.1.2.8, Rx and Tx Timestamps](#).

Set Reserved Bytes. This action write the frame's Rx timestamp into the reserved bytes of the PTP header.

Clear Reserved Bytes. This action clears the reserved byte of the PTP header by setting the four bytes to 0.

In addition, the UDP checksum is handled for IPv4 frames and IPv6 frames after any PTP modifications by the rewriter. For IPv4, the UDP checksum is cleared while for IPv6, the two bytes immediately following the PTP header are updated so that the UDP checksum remains correct.

All internal calculations on nanoseconds timestamps are signed and use 48 bits of precision. When enabling a port for backplane mode, it is possible specify the number of bits to use from the 48-bit counter. This is specified in SYS::PTP_CFG.PTP_STAMP_WID. This means that the timestamp values carried in the reserved bytes or in the correction field in the PTP header can be for instance 30 bits instead of the default 32 bits.

2.12.7 OAM EDITING

The rewriter rewrites OAM PDUs according to the OAM PDU type and the result of the OAM processing in the Versatile OAM Processor (VOP). For more information about the OAM rewrites, see [Section 2.13, Versatile OAM Processor](#).

2.12.8 SPECIAL REWRITER OPERATIONS

VCAP IS2 can trigger a special rewriter operation through IS2_ACTION.REW_OP = 6 (rewriter special), where the frame's MAC addresses are swapped:

- The original destination MAC address becomes the new source MAC address.
- The original source MAC address becomes the new destination MAC address.

In addition, bit 40 is cleared in the new source MAC address to prevent the possibility of transmitting an illegal frame with a multicast source MAC address.

The swapping of MAC addresses is useful when implementing a general hairpinning functionality where an IS2-identified flow is hardware looped back to the source port while swapping the MAC addresses.

2.13 Versatile OAM Processor

VSC7416-01 supports OAM functions such as continuity checking, loopback, frame loss measurements, test signal, frame delay measurements, and frame delay variation measurements as defined in ITU-T Y.1731, MEF E-LMI, ITU-T Y.1564 (Y.156sam), IEEE 802.3ah, and IEEE 802.1ag. Other OAM functions are supported with CPU interaction.

VSC7416-01 implements a Versatile OAM Processor (VOP) that includes 75 Versatile OAM End Points (VOE), each handling the OAM functions of a MEP. Both port, path, and service MEP monitoring levels are supported and a MEP can be either an Up-MEP or a Down-MEP. Incoming OAM frames are processed by the assigned VOE, which include verification of the frame contents against configurations, forwarding based on the MEG level, and processing according to the received OAM type.

There are three different MEP monitoring levels in VSC7416-01:

- Service
- Path
- Port

A given OAM frame belongs either to a service VOE, a path VOE, or a port VOE. A path VOE is in reality just a service VOE dedicated for a path monitoring level.

The service and path VOEs use VOE index 0 through 63, and the port VOEs use the physical port number + 64 as VOE index.

2.13.1 DOWN-MEP OVERVIEW

This section provides an overview of the Down-MEP implementation in VSC7416-01.


The following illustration shows the location of the Down-MEP. The location of the Down-MEP implies that OAM frames transmitted by the Down-MEP must be injected into the Tx data flow of the active port (P2). Similarly, the OAM frames received by the Down-MEP must be extracted from the Rx data flow of the active port.

This implies that any frame loss and frame delay that the frame flow experiences through the queuing system is not reflected in the Down-MEP loss measurements and delay measurements. This is opposed to the Up-MEP measurements. For more information, see [Section 2.13.2, Up-MEP Overview](#).

The illustration also shows the flow of data between two end points (PC_0 and PC_1). All data frames flow between PC_0 and PC_1.

The connection is monitored by an OAM connection between MEP_0 and MEP_1.

FIGURE 2-36: DOWN-MEP LOCATION IN A SWITCH




The VSC7416-01 implementation of a Down-MEP is shown in the following illustration. It shows a bidirectional frame flow being passing through two VSC7416-01 devices monitored by a Down-MEP pair. One MEP is configured as initiator MEP (left side), and the other MEP is configured as remote MEP (right side). The two directions of the bidirectional frame flow have been separated, so that the frame flow from right to left is shown at the top and the frame flow from left to right is shown at the bottom.

The Down-MEP implementation in VSC7416-01 requires that all frames from the flow being monitored to do one VOE lookup in the Tx direction (Tx lookup) and one VOE lookup in the Rx direction (Rx lookup). For more information, see TX_VOE and RX_VOE in [Figure 2-37](#), . For more information about the purpose of each of the two lookups, see [Section 2.13.1.2, Down-MEP Lookups](#).

The following illustration shows that the initiator MEP inserts OAM frames prior to classification in ES0 at DSP_fwd_1. The frame can be assigned a Service TX_VOE by ES0. After the ES0 lookup, the frame causes a Tx lookup in the TX_VOE. If a service TX_VOE is not allocated, the frame does a Tx lookup in the port VOE.

Frames received by the initiator MEP at ASP_ret_1 are assigned ISDX_1, which is used for a lookup in the IPT table. The IPT table can optionally assign a service RX_VOE. If a service RX_VOE is assigned, an Rx lookup is done in the Service VOE. If a service RX_VOE is not allocated, the frame does an Rx lookup in the port VOE.

FIGURE 2-37: DOWN-MEP DETAILS



2.13.1.1 Down-MEP Data Sampling Points

For Delay Measurements (DM), the Tx timestamp must be sampled at the frame's active port (P2) at egress. This is indicated as TX_TS at port P2. The Rx timestamp must be sampled at the active port (P2) at ingress. This is indicated as RX_TS at port P2.

For Loss Measurements (LM), the Tx Data Loss Measurement is implemented in the Down-MEP, and the Tx LM counters are located in the VOE. This is illustrated as the box TX_LM. The Rx Data Loss Measurement is implemented in the Down-MEP, and the Rx LM counters are located in the VOE. This is illustrated as the box RX_LM.

2.13.1.2 Down-MEP Lookups

The purpose of the Rx and the Tx VOE lookups done for a Down-MEP is explained below. At Tx, the following functions are achieved:

- Tx MEL check.
- Service TX_LM counters are updated.
- Path TX_LM counters are updated (optional).
- Port TX_LM counters are updated.
- PDU specific OAM updating of Tx frame.

At Rx, the following functions are achieved:

- DMAC check is performed
- Rx MEL checking
- CCM validation is performed
- PDU specific OAM handling, including loopback function.
- Service RX_LM counters are updated.
- Path RX_LM counters are updated (optional).
- Port RX_LM counters are updated.
- If an OAM PDU is to be copied to the CPU, this is done here.


2.13.2 UP-MEP OVERVIEW

VSC7416-01 supports Up-MEP VOE operation for service VOEs. MAB: Is this true: Path VOEs and port VOEs must always be configured for Down-MEP operation.

This section gives an overview of the Up-MEP implementation in VSC7416-01.

The following illustration shows that an Up-MEP is located at a port that is different from the active port, and the MEP sends OAM frames to the switch core and queue system towards the active port. For Rx, the Up-MEP receives OAM frames from the switch core and queue system. This implies that any frame loss and frame delay the frame flow experiences through the queue system are reflected in the Up-MEP Loss Measurements and Delay Measurements. This is opposed to the Down-MEP measurements. For more information, see [Section 2.13.1, Down-MEP Overview](#).

FIGURE 2-38: UP-MEP LOCATION IN SWITCH




The location of the Up-MEP implies that OAM frames transmitted by the Up-MEP must be injected in to the Rx data flow at the port at which the Up-MEP is located (see port P0 in the illustration). Similarly, the OAM frames received by the Up-MEP must be extracted from the Tx data flow at the port at which the Up-MEP is located.

[Figure 2-39](#), illustrates the implementation of Up-MEP_0 from [Figure 2-38](#), in VSC7416-01. P0 is located in the left side of the drawing and P2 is located in the right side of the illustration. To ease the description of the OAM Tx and OAM Rx functions, the frame flow from P0 to P2 (Tx direction shown at the top of the figure) and the frame flow from P2 to P0 (Rx direction shown at the bottom) have been separated in the illustration.

The Up-MEP implementation in VSC7416-01 requires that all monitored frames from the flow to do two VOE lookups in the Tx direction (TX_IGR_VOE and TX_EGR_VOE) and two VOE lookups in the Rx direction (RX_IGR_VOE and RX_EGR_VOE). For more information about each of the four lookups, see [Section 2.13.2.2, Up-MEP Lookups](#).

The following illustration shows a port called the masquerade port (MASQ port). The masquerade port is a physical port, different from P0 and P2, which is used to insert Tx OAM frames in to the Up-MEP Tx function. For more information, see [Section 2.13.2.3, Masquerade Port](#).

FIGURE 2-39: UP-MEP DETAILS



The loss measurement counters and delay measurement timestamps provided by the VOE are described in the following. The counters and timestamps are hardware accurate, however, the updating of OAM frames must be done in software.

2.13.2.1 Up-MEP Data Sampling Points

For delay measurements (DM), the Tx timestamp (TX_TS) must be sampled at the frame's ingress port (P0). To achieve this, DM OAM frames are injected at the masquerade port, and then TX_TS is sampled when the frame is received on the masquerade port. This is indicated in Figure 2-39, as TX_TS at the masquerade port. The Rx timestamp (RX_TS) must be sampled at the frame's egress port (P0). To achieve this, Rx OAM DMR/DMM frames are forwarded to P0 in the Rx direction for TS_RX sampling but aborted before they reach the front port P0, hence they are not seen externally to the switch. The RX_TS sampling is illustrated in Figure 2-39, at P0 (Data out).

For Loss Measurements (LM), the Tx data loss measurement (TX_LM) must be done after the service policer (S-DLB) in the Tx direction. For this purpose, a set of TX_LM counters are implemented in the analyzer. The TX_LM counters are controlled by the TX_IGR_VOE lookup in the VOE. This is illustrated as a box labeled TX_LM in Figure 2-39, . The Rx data loss measurement (RX_LM) must be done just before the frames are transmitted at P0 in the Rx direction. The RX_LM counters are implemented in the VOE. The RX_LM counters are updated by the RX_EGR_VOE lookup. This is illustrated with label RX_LM pointing to the RX_EGR_VOE lookup in Figure 2-39, .

2.13.2.2 Up-MEP Lookups

For an Up-MEP, each frame in both the Rx and Tx directions is subject to two VOE lookups. The purposes of the in total four VOE lookups are:

TX_IGR_VOE Lookup

This is where the Tx OAM frames are inserted using the masquerade port. The following functions are achieved.

- OAM MEL check (discard rule) is enforced for OAM frames received from P0.
- Loss Measurement information is gathered and sent to the TX_LM data counters (TX_LM in Figure 2-39,) located after the service policers.
- Service level Tx OAM frame counters are updated.

TX_EGR_VOE Lookup

The following functions are achieved:

- If the frame flow monitored by the Up-MEP is part of a path VOE flow, the Tx path VOE counters are updated.
- Port VOE Tx counters are updated.

RX_IGR_VOE Lookup

The following functions are achieved:

- If the frame flow monitored by the Up-MEP is part of a path VOE flow, the Rx path VOE counters are updated.
- Port VOE Rx counters are updated.
- Miscellaneous OAM PDU information is extracted from the frame. This is forwarded to the RX_EGR_VOE lookup for OAM PDU checking. For example, all CCM PDU validation is done in the RX_IGR_VOE lookup, however, the VOE is not updated based on the CCM PDU checking results until the CCM frame reaches the RX_EGR_VOE lookup.
- If any OAM PDU is to be copied to the CPU, it done as part of the RX_IGR_VOE lookup, because frames cannot be copied to the CPU at the RX_EGR_VOE lookup.

RX_EGR_VOE Lookup

The following functions are achieved:

- The OAM PDU is checked for MEL to implement the OAM blocking and OAM processing rules.
- For CCM PDUs, the result of the CCM checking done at the RX_IGR_VOE lookup is evaluated, and interrupt and LOC counters are updated accordingly.
- OAM PDUs are either discarded or forwarded to the loopback port. This depends on the OAM PDU type and whether the Up-MEP is configured for loopback mode or in-service mode. For more information, see [Section 2.13.2.4, Up-MEP Modes](#).
- RX_LM counter is updated.

2.13.2.3 Masquerade Port

VSC7416-01 implements a masquerading functionality that allows frames received at a given port, the masquerade port, to appear to the classifier and analyzer as if it was actually received on another port. This is referred to as masquerading. The masquerade port can be any front port or the CPU port.

By using masqueraded frames, Tx OAM frames can be inserted into the data frame flow right after the ingress data port and are forwarded through the analyzer and the queue system in the same way as data frames belonging to this flow, including passing through the same queues. This allows for exact Up-MEP DM and LM handling.

The masquerading is implemented by including the CPU injection frame header as described in [section Section 2.14.2, Frame Injection](#).

2.13.2.4 Up-MEP Modes

The VSC7416-01 Up-MEP implementation supports two different modes:

- In-service mode
- Loopback mode

The following describes the two Up-MEP modes.


In-Service Mode

The in-service Up-MEP mode is illustrated in [Figure 2-40](#). In this mode the remote MEP performs loopback of bidirectional OAM PDUs in SW by extracting all received OAM PDUs to CPU. User data is supported simultaneously with the Up-MEP OAM frame flow where user data frames flow between PC_0 and PC_1.

At the remote MEP, the CPU must read the required DM or LM counters to update the OAM PDU Response, and insert this response in the Rx direction towards the initiator MEP.

The VOE provides the required LM and DM values to support hardware-accurate updating of the OAM responses.

FIGURE 2-40: UP-MEP IN-SERVICE OPERATION




Loopback Mode

The loopback Up-MEP mode is shown in the following illustration. It is intended for out-of-service testing. The mode requires that the port on which the remote MEP is located (P6 in Figure 2-41,) must be set into loopback, so that all frames being forwarded to P6 are looped back into P6.

In loopback mode, no user data traffic can be forwarded on the frame flow being monitored by the Up-MEPs. To emulate the effects of user data, the initiator MEP can insert test traffic. The test traffic can be either OAM TST frames or data frames.

To emulate unidirectional traffic, the remote MEP can be configured to terminate all the test traffic at the egress port.

FIGURE 2-41: UP-MEP LOOPBACK OPERATION



At the remote MEP, the DMAC and the SMAC are swapped for data frames that pass through the VOE and are forwarded to the loopback port.

2.13.3 VOE SELECTION

This section describes how to select a service, path, or port VOE in VSC7416-01.

2.13.3.1 Assigning a Service VOE

At ingress, the service VOE assignment is used in the following VOE lookups to select which VOE to use:

- Down-MEP Rx lookup (RX_VOE)
- Up-MEP Rx-Ingress lookup (RX_IGR_VOE)
- Up-MEP Tx-Ingress lookup (TX_IGR_VOE)

Per service, the IPT table (ANA:IPT:OAM_MEP_CFG) can assign a service VOE to a service ISDX. The service VOE can optionally be configured to be a path VOE, see section [Section 2.13.3.2, Assigning A Path VOE](#). When a service VOE is not assigned by the IPT table, the lookup defaults to a port VOE lookup, see section [Section 2.13.3.3, Assigning A Port MEP](#).

At egress, the service VOE assignment is used in the following VOE lookups to which VOE to use:

- Down-MEP Tx lookup (TX_VOE)
- Up-MEP Rx-Egress lookup (RX_EGR_VOE)
- Up-MEP Tx-Egress lookup (TX_EGR_VOE)

Per service, the ES0 actions OAM_MEP_IDX_VLD and OAM_MEP_IDX can assign a service VOE. The service VOE can optionally be configured to be a path VOE. For more information, see [Section 2.13.3.2, Assigning A Path VOE](#). If a service VOE is not assigned by ES0, the lookup defaults to a port VOE lookup. For more information, see [Section 2.13.3.3, Assigning A Port MEP](#).

2.13.3.2 Assigning A Path VOE

A service MEP can optionally be configured to be part of a path flow, meaning the traffic and the OAM frames belonging to a service MEP are counted by the path VOE LM counters instead of the service VOE LM counters.

A service VOE is assigned to a path VOE through OAM_MEP::VOE_MAP_CTRL.

Enabling a path VOE takes effect at both ingress and egress.

2.13.3.3 Assigning A Port MEP

If no valid service VOE is assigned to an OAM frame, a port VOE is assigned based on the physical port number. This is true at both ingress and egress.

2.13.4 GENERAL MEP HANDLING

The VOP handles frames with EtherType 0x8902 and a number of VLAN tags in accordance with ANA::VCAP_CFG.OAM_DETECT_CFG as OAM frames. Other frames are treated as data frames.

If a service VOE has been associated with a frame via the IPT table, the service MEP is looked up based on the index received from the IPT table. If no service MEP has been associated, then the port MEP is looked up based on the physical port number offset with 64.

If the MEP is enabled (OAM_MEP:VOE:BASIC_CTRL.VOE_ENA), the following check are done. The received MEL is validated against the configured MEL (OAM_MEP:VOE:MEL_CTRL.MEL_VAL).

- If the received MEL value is larger than the configured value, the OAM frame is treated as a data frame.
- If the received MEL value is smaller than the configured value, an error indication is set (OAM_MEP:VOE:CCM_CFG.OAM_MEL_ERR), and the frame is either discarded or redirected to the CPU. The redirection to the CPU can either be every time the MEL error occurs or with a hit-me-once functionality (OAM_MEP:VOE:MEL_CTRL).

If the received MEL value is equal to the configured value the OAM frame is available for the VOE.

The received DMAC can optionally (OAM_MEP:VOE:BASIC_CTRL.RX_DMACHK_SEL) be compared against a per-VOE unicast MAC address (OAM_MEP:VOE:MEP_UC_MAC_*SB) or a common multicast address OAM_MEP::COMMON_MEP_MC_MAC_*SB. For the multicast address, the least significant bits are not checked.

Frames failing the DMAC address check are discarded and not handled by the MEP.

The following OAM PDU types are handled specifically by VOE:

- Continuity check messages (CCM) with OpCode = 1. Both with and without loss measurement information.
- Loopback reply (LBR) with OpCode = 2.
- Loopback message (LBM) with OpCode = 3.
- Test (TST) with OpCode = 37.

- Loss measurement reply (LMR) with OpCode = 42.
- Loss measurement message (LMM) with OpCode = 43.
- One-way delay measurement (1DM) with OpCode = 45.
- Delay measurement reply (DMR) with OpCode = 46.
- Delay measurement message (DMM) with OpCode = 47.

Each OAM function is enabled through OAM_MEP:VOE:OAM_HW_CTRL.

In addition, six generic OpCodes can be configured for dedicated CPU copy and forwarding (OAM_MEP::OAM_GENERIC_CFG).

2.13.5 ETHERNET CONTINUITY CHECK (ETH-CC) HANDLING

Ethernet Continuity Check (ETH-CC) is a proactive OAM procedure used to ensure that connectivity exists between two MEG End Points (MEPs) within a Managed Entity Group (MEG). CCM works by periodically exchanging CCM frames between the MEPs. A MEP continuously monitors if OAM CCM frames are received from the peer MEP and if no CCM frames have been detected within 3.5 times a configured interval Loss of Continuity (LOC) is declared. ETH-CC also allows detection of unintended connectivity between two MEGs (Mismerge), unintended connectivity within the MEG with an unexpected MEP (Unexpected MEP), and other defect conditions (for example, Unexpected MEG Level, Unexpected Period, and so on).

In VSC7416-01, ETH-CC can be used to detect loss of continuity (LOC) between MEPs at a port level, at a path level, or at a service level.

By default, CCM frames are discarded by the VOE if the MEL is below or equal to the configured MEL (OAM_MEP:VOE:MEL_CTRL.MEL_VAL), unless this is overruled (OAM_MEP:VOE:OAM_FWD_CTRL.CCM_FWD_ENA).

All received CCM frames can optionally be copied to the CPU (OAM_MEP:VOE:OAM_CPU_COPY_CTRL.CCM_CPU_COPY_ENA) to a CPU extraction queue (OAM_MEP::CPU_CFG.CCM_LM_CPU_QU and OAM_MEP::CPU_CFG.CCM_LM_EXT_PORT_ENA).

When hardware-based CCM is enabled (OAM_MEP:VOE:OAM_HW_CTRL.CCM_ENA), OAM CCM frames are used for Loss of Continuity handling. The CCM handling consists of three tasks that operate together:

- Periodical transmission of CCM frames using timed CPU injections
- Reception and validation of received CCM frames
- Scan handling used to detect time since the last CCM reception

These tasks are described in the following sections.

2.13.5.1 CCM Frame Transmission

VSC7416-01 can be configured to periodically transmit frames for each MEP that is enabled for Proactive OAM via CCM by using the frame generation engine. For more information, see [Section 2.14.4, Frame Generation Engine for Periodic Transmissions](#).

The transmitted OAM PDUs must be formatted as defined in ITU-t Y.1731 with the DMAC set to either the MAC address of the peer MEP or with the standardized multicast address with the MEL configured in the lowest three bits of the MAC address. The CCM frames must be encapsulated to ensure that it follows the path that it monitors.

The period specified as part of the common OAM PDU flag field must be configured corresponding to the interval by which the CCM frames are sent with. Transmission period can be one of the following eight values (as specified in ITU-t Y.1731):

TABLE 2-138: CCM PERIODS (Y.1731)

CCM PDU Period Encoding	Period value	Description
0	Invalid value	Invalid value for CCM PDUs
1	3.33 ms	300 frames per second Default transmission period for protection switching application
2	10 ms	100 frames per second

TABLE 2-138: CCM PERIODS (Y.1731) (CONTINUED)

CCM PDU Period Encoding	Period value	Description
3	100 ms	10 frames per second Default transmission period for protection switching application
4	1 s	1 frame per second Default transmission period for fault management application
5	10 s	6 frames per minute
6	1 min	1 frame per minute
7	10 min	6 frames per hour

The RDI bit in transmitted CCM frames is set to the value configured in OAM_MEP:VOE:CCM_CFG.CCM_TX_RDI.

The sequence number in the CCM PDU can optionally be incremented for each transmitted CCM frame (OAM_MEP:VOE:CCM_SEQ_INCR_ENA). When enabled, the transmitted sequence number is set to one higher than the value configured in OAM_MEP:VOE:CCM_TX_SEQ_CFG. At the same time, OAM_MEP:VOE:CCM_TX_SEQ_CFG is automatically incremented after use.

It is possible to transmit CCM frames that carries loss measurement information by either injecting them specifying REW_OP = 0x1C in the frame injection header or by setting OAM_MEP:VOE:CCM_CFG.CCM_LM_INSERT_NXT. For more information, see [Section 2.13.9, Single-Ended Loss Measurement](#).

2.13.5.2 CCM Frame Reception

If OAM CCM handling is enabled for a MEP (OAM_MEP:VOE:OAM_HW_CTRL.CCM_ENA), it must be validated that a received CCM frame is originating from the correct peer. The following defect conditions can be detected:

- If a CCM frame is received with a MEG level lower than the receiving MEP's MEG level, the defect condition "Unexpected MEG level" is detected.
- If a CCM frame is received with the same MEG level but with a MEG ID different than the receiving MEP's MEG ID, the defect condition "mismatch" is detected. The full 48-byte MEG ID can be configured in OAM_MEP:VOE:CCM_MEGID_CFG. The check is enabled in OAM_MEP:VOE:CCM_CFG.CCM_MEGID_CHK_ENA.
- If a CCM frame is received with the same MEG level and a correct MEG ID but with an MEP ID, different than the receiving MEP's MEP ID is received, the defect condition "Unexpected MEP" is detected. The MEG ID is configured in OAM_MEP:VOE:CCM_MEPID_CFG and the check is enabled in OAM_MEP:VOE:CCM_CFG.CCM_MEPID_CHK_ENA.
- If a CCM frame is received with a correct MEG Level, a correct MEG ID, a correct MEP ID, but with a zero period field value, the period is invalid.

The above mentioned defects all results in the receiving CCM frame being discarded and a corresponding sticky bit indication being set (OAM_MEP:VOE:STICKY). The sticky bit can be used to trigger interrupt of the CPU (OAM_MEP:VOE:INTR_ENA). The defect CCM frames can optionally be send to the CPU through configuration OAM_MEP:VOE:BASIC_CTRL.CCM_ERR_CPU_HITME_ONCE and OAM_MEP:VOE:BASIC_CTRL.CCM_ERR_CPU_COPY_ENA.

CCM frames failing at least one of the above mentioned checks are counted as bad CCM frames in OAM_MEP:VOE:CCM_RX_INVLD_FC_CNT. CCM frames passing all the above mentioned checks are counted as valid CCM frames in OAM_MEP:VOE:CCM_RX_VLD_FC_CNT.

Valid CCM frames are used for clearing the CCM miss count (OAM_MEP:VOE:CCM_CFG.CCM_MISS_CNT). The CCM miss count keeps track of the number of half-periods since the last valid CCM frame was received. The following defect conditions can be detected:

- If no valid CCM frame from a peer MEP is received within the interval equal to 3.5 times the receiving MEP's CCM transmission period, loss of continuity with peer MEP is detected. This can be reported via sticky events and can trigger interrupt.
- If a CCM frame is received with a correct MEG Level, a correct MEG ID, a correct MEP ID, but with a period field value different than the receiving MEP's CCM transmission period, "Unexpected Period" is detected. The CCM transmission period is configured in OAM_MEP:VOE:CCM_CFG.CCM_PERIOD.
- If a CCM frame is received with a correct MEG Level, a correct MEG ID, a correct MEP ID, a correct period, but with a priority value different than the receiving MEP's CCM priority, "Unexpected Priority" is detected. The CCM

priority is configured in OAM_MEP:VOE:CCM_CFG.CCM_PRIO.

If RDI is set in the received CCM frame, it is reported via sticky events (OAM_MEP:VOE:STICKY.CCM_RX_RDI_STICKY) and can trigger interrupt.

It is possible to enable sequence checking of received CCM frames (OAM_MEP:VOE:CCM_RX_SEQ_CHK_ENA) so that the received sequence number is compared against a configured value (OAM_MEP:VOE:CCM_RX_SEQ_CFG). The incoming value must be one higher than the configured value. After the check, the configured value is auto incremented. A sticky bit is set for frames failing the sequence check (OAM_MEP:VOE:RX_STICKY.CCM_RX_SEQ_ERR_STICKY).

For more information about handling of CCM loss measurement, see [Section 2.13.8, Ethernet Frame Loss Measurement \(ETH-LM\) Handling](#).

2.13.5.3 CCM Scan and Loss of Continuity Check

Loss of continuity (LOC) is checked on active CCM connections. LOC occurs when no valid CCM frame is received on an active CCM connection for 3.5 CCM periods. For this purpose, the VOP implements seven different LOC timers; one for each of the valid CCM periods listed in [Table 2-138](#). Each LOC timer can be configured to implement a unique CCM timeout period, including, but not limited to, the values listed in [Table 2-138](#). The LOC timers count using a common base tick. It is configurable as an integer number of the VSC7416-01 system clock period, which is 6.4 ns. The base tick is configured in OAM_MEP::CCM_CTRL.CCM_BASE_TICK_CNT. Each of the seven LOC timers can be programmed to implement a given CCM period by specifying a counter as a number of base ticks using the LOC timer specific register OAM_MEP::CCM_PERIOD_CFG[LOC timer].CCM_PERIOD_VAL.

Each VOE with a valid CCM connection is associated with one of the seven LOC timers by configuring a CCM period (OAM_MEP:VOE:CCM_CFG.CCM_PERIOD) where CCM_PERIOD = 1 uses LOC timer 0, CCM_PERIOD = 2 uses LOC timer 1, and so on.

To comply with the Y.1731 requirement that LOC is detected when no valid CCM frame is received for 3.5 CCM periods, the VOE detects when 3.5 CCM periods have passed. The LOC timers measure half-periods and LOC is therefore detected if the LOC timer has measured 7 half-periods without receiving a valid CCM frame. Effectively, this results in LOC after 3.5 CCM periods. The number of half-periods is counted in OAM_MEP:VOE:CCM_CFG.CCM_MISS_CNT.

The CCM miss count is cleared every time a valid CCM frames is received by the VOE, and it is incremented every time the LOC timer associated with the VOE times out. If it reaches a value of 7, LOC is reported by the VOE.

For example, if the CCM period for a given CCM connection is 10 ms, the associated LOC timer must time out after 5 ms. LOC is reported if no valid CCM frame is received by the VOE for 35 ms.

Whenever a VOE enters the LOC state, the sticky bit in OAM_MEP:VOE:STICKY.CCM_LOC_STICKY is asserted. This event can optionally be configured to cause an interrupt using OAM_MEP:VOE:INTR_ENA.CCM_LOC_INTR_ENA.

When a LOC timer times out, a CCM scan is initiated (OAM_MEP::MEP_CTRL.CCM_SCAN_ENA). The CCM scan scans through all the VOEs and increments the CCM_MISS_CNT for those VOEs that are associated with the LOC timer that has timed out. The CCM scan takes between 480 ns and 1400 ns to complete depending on other OAM activity in the VOP.

The CCM scan updates the CCM_MISS_CNT for VOEs using the LOC timer that initiated the scan. Any LOC timer timing out during an ongoing CCM scan triggers a new scan immediately after the ongoing scan is complete.

If a LOC timer times out while a scan is ongoing for the same LOC timer, it indicates that the LOC timer is misconfigured with a too small timer value. Such an event can be detected in the sticky bit OAM_MEP::CCM_SCAN_STICKY.CCM_SCAN_START_DELAYED_STICKY.

The normal timer control can be overruled by forcing a scan (OAM_MEP::CCM_CTRL.CCM_FORCE_HW_SCAN_ENA). Alternatively, it is possible to stop an ongoing scan (OAM_MEP::CCM_CTRL.CCM_FORCE_HW_SCAN_STOP_ENA).

The status bit OAM_MEP::CCM_SCAN_STICKY.CCM_SCAN_ONGOING_STATUS indicates whether a CCM scan is currently in progress.

When a scan starts or completes, it is reported through OAM_MEP::CCM_SCAN_STICKY.CCM_SCAN_COMPLETED_STICKY and OAM_MEP::CCM_SCAN_STICKY.CCM_SCAN_STARTED_STICKY.

2.13.6 ETHERNET LOOPBACK (ETH-LB) HANDLING

Ethernet Loopback (ETH-LB) is an on-demand OAM function that can be used for the following applications:

- To verify bidirectional connectivity of a MEP with a MIP or a peer MEP.

- To perform a bidirectional in-service or out-of-service diagnostics test between a pair of peer MEPs. This includes verifying bandwidth throughput, detecting bit errors, and so on.


There are two ETH-LB types: unicast ETH-LB and multicast ETH-LB. VSC7416-01 only supports unicast loopback between two peers by hardware, although multicast addresses can be used.

To support Ethernet Loopback, LBM and LBR must be enabled for the VOE (OAM_MEP:VOE:OAM_HW_CTRL.LBM_ENA and OAM_MEP:VOE:OAM_HW_CTRL.LBR_ENA).

MAB: Bit out of place? The LBM PDU is inserted by the CPU in the Initiator MEP. The transmitted LBM PDU must be formatted as defined in ITU-t Y.1731 with the DMAC set to either the MAC address of the peer MEP or with a the standardized multicast address with MEL configured in the lowest three bits of the MAC address.

The general Ethernet Loopback function is depicted in the following illustration. The illustration is valid for both Up-MEP and Down-MEP implementations.

FIGURE 2-42: ETHERNET LOOPBACK



VSC7416-01 supports updating the transaction ID in transmitted LBM PDUs (OAM_MEP:VOE:LBM_TX_TRANSID_UPDATE.LBM_TX_TRANSID_UPDATE_ENA, OAM_MEP:VOE:LBM_TX_TRANSID_CFG.LBM_TX_TRANSID). The configured transaction ID is automatically incremented after use.

For incoming LBR PDUs, the received transaction ID is compared against a value one higher than a configured value (OAM_MEP:VOE:LBR_RX_TRANSID_CFG.LBR_RX_TRANSID). The configured value is automatically incremented after use.

The number of received LBR PDUs is counted in OAM_MEP:VOE.LBR_RX_FRM_CNT.

2.13.6.1 Down-MEP Ethernet Loopback

This section provides information about the Down-MEP implementation of Ethernet loopback.

LBM Transmissions

LBM transmissions use standard CPU injections. For more information, see [Section 2.14.2, Frame Injection](#).

LBM Receptions, Loopback, and LBR Transmissions

The LBM is looped back as an LBR with the following frame modifications.

- The LBR frame's DMAC is set to the LBM's SMAC
- The LBR frame's SMAC is updated with the MEP's MAC address (OAM_MEP:VOE:MEP_UC_MAC_MSB and OAM_MEP:VOE:MEP_UC_MAC_LSB).
- The Opcode is changed to LBR.

The LBM frames are looped back to the port they were received on. By default, the frame's classified ISDX is replaced with the value configured in OAM_MEP:VOE:LOOPBACK_CTRL.LB_ISDX. If the service is part of a protected service and the protection is active, the frames are looped to the protect loopback port configured in OAM_MEP:VOE:LOOPBACK_CTRL.LB_PORTNUM_P) and the ISDX is replaced with OAM_MEP:VOE:LOOPBACK_CTRL.LB_ISDX_P.

LBR Receptions

At the initiator MEP, the transaction ID in the LBR is verified, and the LBR frame is counted.

2.13.6.2 Up-MEP Ethernet Loopback

This section provides information about the Up-MEP implementation of Ethernet loopback.

LBM Transmissions

LBM transmissions must use masquerading. For more information, see [Section 2.13.2.3, Masquerade Port](#).

LBM Receptions, Loopback, and LBR Transmissions

LBM Receptions, Loopback, and LBR Transmissions are handled differently, depending on whether it is in In-Service mode or in Loopback mode.

The In-Service mode is configured in OAM_MEP:VOE:MEL_CTRL.UPMEP_LB. In this mode, the CPU must perform the frame updates required for LBM to LBR (swapping MAC addresses and changing the Opcode to LBR). This requires the VOE to be configured to send a copy of the LBM PDU to the CPU (OAM_MEP:VOE:OAM_CPU_COPY_CTRL.LBM_CPU_COPY_ENA). After the frame modification, the LBR frame must be CPU injected into the remote MEP Tx VOE using masquerading.

The loopback mode is configured in OAM_MEP:VOE:MEL_CTRL.UPMEP_LB. In this mode, the VOE performs the frame updates required for LBM to LBR. No CPU intervention is necessary.

LBR Receptions

At the initiator MEP, the transaction ID in the LBR is verified and the LBR frame is counted. Note that for the Up-MEP implementation, only the lower 16 bits of the transaction ID are verified and saved (OAM_MEP:VOE:LBR_RX_TRANSID_CFG.LBR_RX_TRANSID).

2.13.7 ETHERNET TEST SIGNAL (ETH-TEST) HANDLING

The support of Ethernet Test OAM function is implemented using the same configuration registers and counters as the loopback frames. For more information, see [Section 2.13.6, Ethernet Loopback \(ETH-LB\) Handling](#). As a result, for a single MEP, it is not possible to concurrently process LBM/LBR frames and TST frames.

Handling of TST frames is enabled in OAM_MEP:VOE:OAM_FWD_CONTROL.GENERIC_FWD_MASK[6]. If not enabled, TST frames are handled as an unknown PDU type.

The following describes the handling of TST frames.

2.13.7.1 Tx Handling of TST Frames

When Tx handling of TST frames is enabled (OAM_MEP:VOE:OAM_HW_CTRL.LBM_ENA, OAM_MEP:VOE:OAM_CNT_OAM_CTRL.LBM_OAM_CNT_ENA, OAM_MEP:VOE:OAM_CNT_DATA_CTRL.LBM_DATA_CNT_ENA), the sequence number of a TST frame is updated and incremented as described for LBM frames using the same configuration (OAM_MEP:VOE:LBM_TX_TRANSID_UPDATE.LBM_TRANSID_UPDATE_ENA, OAM_MEP:VOE:LBM_TX_TRANSID_CFG.LBM_TX_TRANSID).

2.13.7.2 Rx Handling of TST Frames

The sequence number of a received TST frame is checked against the expected value (OAM_MEP:VOE:LBR_RX_TRANSID_CFG.LBR_RX_TRANSID). The TST PDU can optionally be extracted to the CPU (OAM_MEP:VOE:OAM_CPU_COPY_CTRL.LBR_CPU_COPY_ENA).

Each received TST frame is counted in OAM_MEP:VOE:LBR_RX_FRM_CNT.LBR_RX_FRM_CNT.

2.13.8 ETHERNET FRAME LOSS MEASUREMENT (ETH-LM) HANDLING

The Ethernet Frame Loss Measurement (ETH-LM) is an OAM function, which is used for performance monitoring. ETH-LM collects set of performance counter values that maintain a count of transmitted and received data frames between a pair of MEPs.

ETH-LM is performed by sending and receiving frames with ETH-LM information to a peer MEP. Each MEP performs frame loss measurements. ETH-LM performance counters can be exchanged in two ways:

- Dual-ended ETH-LM, where the LM counters are piggybacked on OAM CCM frames.
- Single-ended ETH-LM, where the LM counters are sent in dedicated LM frames.

2.13.8.1 Performance Counters

A MEP maintains the following two local performance counters for each peer MEP and for each Class of Service (CoS) being monitored for which loss measurements are to be performed.

- TxFCI: counter for in-profile data frames transmitted towards the peer MEP.
- RxFCI: counter for in-profile data frames received from the peer MEP.

The CoS is identified by the frame's classified PCP value. By default, the performance counters do not include the OAM frames transmitted or received by the MEP, but it is possible for each of the supported OAM frame types to configure whether it shall be included in the counters (OAM_MEP:VOE:OAM_CNT_DATA_CTRL). This is only applicable to OAM frames received at the same level as the active VOE. OAM frames with a higher MEL value than the active VOE are always counted as data frames, regardless of the register settings. OAM frames with lower MEL value than the active VOE are discarded and hence are not counted.

If a path MEP is associated with the service MEP, all OAM and data traffic belonging to a service MEP are included in the performance counters for the path MEP. OAM and data frames passing through the service MEP are also included in the port MEP performance counters.

Service and path data traffic are counted in the associated MEP performance counters per CoS.

The location of the counters for a VOE differs whether the VOE is configured for Up-MEP or Down-MEP operation:

For service and path Down-MEP, the LM counters per CoS are given as:

- Rx counters (RX_LM): OAM_MEP:RX_VOE_PM:MEP_RX_FRM_CNT
- Tx counters (TX_LM): OAM_MEP:TX_VOE_PM:MEP_TX_FRM_CNT

For service Up-MEP, the LM counters per CoS are given as:

- Rx counters (RX_LM): OAM_MEP:TX_VOE_PM:MEP_TX_FRM_CNT
- Tx counters (TX_LM): ANA:OAM_UPM_LM_CNT:OAM_UPM_LM_CNT

Port VOEs must always be configured as Down-MEPs and port VOE data traffic is always counted in the associated port performance counters per priority:

- Rx counters (RX_LM): OAM_MEP:PORT_PM:MEP_RX_FRM_CNT
- Tx counters (TX_LM): OAM_MEP:PORT_PM:MEP_TX_FRM_CNT

By default, the receive and transmit performance monitoring counters exist per CoS for each MEP. However, for each VOE and for each CoS, it is possible to select whether a CoS counts in its own counter set or in the CoS value 7 counter set (OAM_MEP::VOE_CNT_CTRL.PRIO_CNT_MASK).

If the service MEP is configured as being part of a path MEP flow, the performance monitoring counters for the path MEP are also updated. Only one level of path counters is supported.

It is configurable if yellow data frames are counted in the LM counters (OAM_MEP::VOE_CNT_CTRL.CNT_YELLOW_ENA).

2.13.8.2 LM Sample Synchronization in Up-MEP

The Up-MEP implementation of the loss measurement requires synchronization between corresponding Rx OAM LM frames and the sample of the LM counters stored in the VOE. This synchronization is based on an LM Rx sequence number, represented in VSC7416-01 by a 5-bit counter. When an LM PDU is received and a counter sample is stored in the VOE, the LM Rx sequence number is incremented and the new value is stored. Further, when the LM PDU is extracted to the CPU, the LM Rx sequence number is sent along in the CPU extraction header in REW_VAL[31:27]. This allows the CPU to establish the relationship between the LM PDU and the counter value stored in the VOE.

The LM Rx sequence number is stored in OAM_MEP:VOE:LOOPBACK_CTRL.LB_ISDX[4:0] and holds the latest value sent to the CPU.

The LM Rx sequence number is shared between the three LM PDUs that require the VOE to save a sample of the Rx performance monitoring counters:

- LMM (In-service mode only)
- LMR
- CCM-LM

The LM Rx sequence numbers are stored in the following registers for the three LM PDUs.

- LMM: OAM_MEP:VOE:CCM_RX_SEQ_CFG.CCM_RX_SEQ[30:26].
- LMR: OAM_MEP:VOE:CCM_RX_SEQ_CFG.CCM_RX_SEQ[25:21]
- CCM-LM: OAM_MEP:VOE:CCM_RX_SEQ_CFG.CCM_RX_SEQ[20:16]

2.13.9 SINGLE-ENDED LOSS MEASUREMENT

For single-ended ETH-LM, the initiator MEP injects LMM frames into the frame flow. The remote MEP modifies the received LMM frames into LMR frames and loopback the LMR frame to the initiator MEP. The initiator MEP extracts the LMR frames to the CPU where the loss measurement calculations are done.

The required MEP operations to support single-ended ETH-LM are as follows:

Initiator MEP (Transmission of LMM)

- Inject LMM frame from CPU in to frame flow at Tx MEP.
- Sample Tx Loss Measurement counter when the LMM frame is transmitted.
- Update outgoing LMM frame with sampled counter (for example, LMM.TxFCf = sampled counter).

Remote MEP (Loopback of LMM/LMR)


- Sample the Rx Loss Measurement counter when LMM is received.
- Set LMM.RxFCf to sampled counter.
- Create LMR by modifying LMM: Update OAM OpCode and swap MAC addresses.
- Sample the Tx Loss Measurement counter when the LMR is transmitted.
- Insert sampled counter into the LMR.TxFCb.
- Transmit the LMR to the initiator MEP.

Initiator MEP (Reception of LMR):

- Sample the Rx Loss Measurement counter when LMR is received.
- Extract LMR and the sampled counter to the CPU for loss measurement calculations.

The following illustration shows the single-ended loss measurement function. It is valid for both Up-MEP and Down-MEP implementations.

FIGURE 2-43: SINGLE-ENDED LOSS MEASUREMENT



The configuration of the loopback function in the remote MEP depends on if the VOE is configured for Down-MEP or Up-MEP operation.

2.13.9.1 Down-MEP Single-Ended Loss Measurement

If hardware handling of LMM frames is enabled (OAM_MEP:VOE:OAM_HW_CTRL.LMM_ENA), LMM frames are either looped back or sent to CPU (OAM_MEP:VOE:LOOPBACK_CTRL.LB_LMM_ENA).

If the LMM frames are looped back as LMR frames, the following frame modifications are done to the LMR frames:

- The DMAC is set to the LMM frame's SMAC
- The SMAC is set to the MEP address (OAM_MEP:VOE:MEP_UC_MAC_MSB, OAM_MEP:VOE:MEP_UC_MAC_LSB)
- The Opcode is set to LMR
- RxFCf is set to the received frame count for this MEP and TxFCb is set to the transmitted frame count for this MEP.

The frames are looped back to the port they were received on and the ISDX is replaced with the value configured in OAM_MEP:VOE:LOOPBACK_CTRL.LB_ISDX, unless they are part of a protected service and the protection is active in which case the frames are looped to a protect loopback port (OAM_MEP:VOE:LOOPBACK_CTRL.LB_PORT_NUM_P) and the ISDX replaced with the value configured in OAM_MEP:VOE:LOOPBACK_CTRL.LB_ISDX_P.


If the LMM frames are sent to CPU, only the RxFCf is set to the received frame counts for this MEP.

If hardware handling of LMR frames is enabled (OAM_MEP:VOE:OAM_HW_CTRL.LMR_ENA), LMR frames are sent to the CPU together with the local frame count RxFCf. RxFCf is written in REW_VAL in the CPU extraction header. The LMR frame can optionally be expanded with RxFCb after TxFCb (OAM_MEP::MEP_CTRL.LMR_UPD_RxFCL_ENA). This is a proprietary mode.

2.13.9.2 Up-MEP Single-Ended Loss Measurement

The loss measurement sample points for a VSC7416-01 Up-MEP single-ended loss measurement operation are shown in the following illustration.

FIGURE 2-44: UP-MEP SINGLE-ENDED LOSS MEASUREMENT SAMPLE POINTS



Note To simplify the illustration, the CPU port is shown twice, even though the same CPU port is used for CPU injection and extraction.

For the Up-MEP, the remote MEP functionality depends on whether the Up-MEP mode is loopback or in-service. The initiator MEP is unaffected of this choice.

The Up-MEP implementation uses the LM Rx sequence number to allow CPU to synchronize between the LM frames copied to the CPU and the LM counter samples available in the VOE. For more information, see [Section 2.13.8.2, LM Sample Synchronization in Up-MEP](#).

Initiator MEP: Injecting LMM Frames from CPU


An LMM frame must be injected into the relevant frame flow from the CPU using masquerading. For more information, see [Section 2.13.2.3, Masquerade Port](#). When the LMM frame is transmitted, the LMM.TxFCf is updated with the TX_LM sample value (ANA:OAM_UPM_LM_CNT:OAM_UPM_LM_CNT_VAL).

Remote MEP: LMM/LMR Loopback Using In-Service Mode

In this mode, loopback must be disabled (OAM_MEP:VOE:MEL_CTRL.UPMEP_LB, OAM_MEP:VOE:LOOPBACK_CTRL.LB_LMM_ENA). The CPU in the remote MEP must do the LMM to LMR update. The VOE samples the RX_LM counter (OAM_MEP:TX_VOE_PM:MEP_TX_FRM_CNT), which is read by the CPU and inserted into the LMM.RxFCf field.

The in-service LMM to LMR loopback mode is shown in the following illustration.

FIGURE 2-45: IN-SERVICE LMM TO LMR OPERATION



The remote MEP must be configured to send a copy of incoming LMM frames to the CPU (OAM_MEP:VOE:OAM_CPU_COPY_CTRL.LMM_CPU_COPY_ENA). When the remote MEP receives the LMM frame, the following is done by the VOE:

- Sample the RX_LM counter and save the value in register OAM_MEP:VOE:CCM_CAP_TX_FCF.CCM_CAP_TX_FCF.
- Increment the LM Rx sequence number and store the value in register OAM_MEP:VOE:CCM_RX_SEQ_CFG.CCM_RX_SEQ[30:26].
- Signal to the CPU that there is a valid LMM RX_LM sample by asserting the valid bit OAM_MEP::UPMEP_LM_CNT_VLD.LMM_RX_CNT_VLD.

If the valid bit is already set when a new LMM frame arrives at the VOE, the following sticky bit OAM_MEP:VOE:UPMEP_LM_CNT_STICKY.LMM_RX_CNT_STICKY is asserted in addition to updating the counter. This allows the CPU to detect that a valid RX_LM sample value has been overwritten before it was read by the CPU.

In the in-service mode, the CPU injects the LMR frame into the Tx VOE using the masquerade port. When transmitting the LMR, the remote MEP updates the LMR.TxFCb with the current value of the TX_LM counter.

Remote MEP: LMM/LMR loopback using loopback mode


In this mode, the Down-MEP loopback must be disabled (OAM_MEP:VOE:LOOPBACK_CTRL.LB_LMM_ENA) while the Up-MEP loopback must be enabled (OAM_MEP:VOE:MEL_CTRL.UPMEP_LB).

The front port where the loopback is performed must be put in loopback mode to loop all outgoing traffic back into the port (DEV::PCS1G_LB_CFG.TBI_HOST_LB_ENA).

The LMM to LMR loopback in loopback mode is shown in the following illustration.



The following illustration depicts the LMR extraction.



ing the valid bit OAM_MEP::UPMEP_LM_CNT_VLD.LMR_RX_CNT_VLD. If the valid bit is already set when a new LMR frame arrives at the VOE, the sticky bit OAM_MEP::VOE:UPMEP_LM_CNT_STICKY.LMR_RX_CNT_STICKY is asserted in addition to updating the counter value. This allows the CPU to detect that a valid RX_LM sample value has been overwritten before it was read by the CPU.

2.13.10 DUAL-ENDED LOSS MEASUREMENT (CCM-LM)

Dual-ended loss measurement is an OAM function, which uses the CCM PDUs to convey the frame counters between the peer MEPs. Hence, dual-ended loss measurement requires an existing CCM flow to be in place. The CCM flow can consist of only CCM-LM frames or CCM frames interleaved with CCM-LM frames. As is the case for CCM frames, the CCM-LM frame flow must be bidirectional.

A CCM-LM frame is a CCM frame used to transport dual-ended loss measurement information. CCM-LM frames are defined as CCM frames for which at least one of the three frame counter fields in the CCM PDU is non-zero.

This section describes the specific dual-ended loss measurement handling of the CCM-LM frames. All the standard handling of CCM frames previously described in [Section 2.13.5, Ethernet Continuity Check \(ETH-CC\) Handling](#), also applies to the CCM-LM frames. This includes RDI handling, sequence number handling for Down-MEPs, and LOC operation.

For dual-ended loss measurement, the initiator MEP injects CCM-LM frames into the frame flow. The CCM-LM frames injected by the initiator MEP contains two fields with loss measurement information received by the CCM-LM frames in the opposite direction:

- CCM-LM.TxFCb: The CCM-LM.TxFCf value from latest received CCM-LM frame.
- CCM-LM.RxFCb: The RX_LM counter value when latest CCM-LM frame was received.

Further, the initiator MEP updates the CCM-LM frame with the TX_LM counter value, which is sampled when the CCM-LM frame is transmitted:


- CCM-LM.TxFCf = TX_LM

The CCM-LM frames are extracted to the CPU at the remote MEP. The remote MEP also samples the RX_LM counter when the CCM-LM frame is received and forwards the RX_LM sample to the CPU for LM calculations.

Upon reception of a CCM-LM frame, the remote MEP stores the following values for use in the CCM-LM flow in the opposite direction:

- CCM-LM.TxFCf
- RX_LM sample (same value as forwarded to the CPU)

The following illustration shows the dual-ended loss measurement function. It illustrates a situation where all the four values required to do the loss measurements are sampled. The CCM-LM function is symmetrical so that the same operation takes place at each end of the CCM-LM connection. The illustration is valid for both Up-MEP and Down-MEP implementations.

FIGURE 2-48: DUAL-ENDED LOSS MEASUREMENT (CCM-LM)

Hardware handling of CCM-LM must be enabled in the VOE (OAM_MEP:VOE:OAM_HW_CTRL.CCM_LM_ENA).

The functionality supported by the VOE for CCM-LM depends on if the VOE is configured for Down-MEP or Up-MEP operation.

2.13.10.1 Down-MEP Dual-Ended Loss Measurement

Initiator MEP: CCM-LM Injections

CCM-LM frames are injected by setting OAM_MEP:VOE:CCM_CFG.CCM_LM_INSERT_NXT. They are handled as normal CCM injected frames but updated with the following loss measurement information:

- CCM-LM.TxFCf updated with TxFCf
- CCM-LM.RxFCb updated with RxFCf captured when the last CCM_LM frame was received.
- CCM-LM.TxFCb updated with TxFCf captured from the last CCM_LM frame received.

It is possible to configure a period for how often CCM frames with loss measurement are transmitted (OAM_MEP:VOE:CCM_CFG.CCM_LM_PERIOD).

Remote MEP: CCM-LM Extractions


If a CCM frame is received with one or more of the fields TxFCb, TxFCf, and RxFCf fields being non-zero, then the frame is handled as a CCM-LM frame and it is sent to the CPU together with the local frame count RxFCf. For more information, see [Section 2.13.8.1, Performance Counters](#). The RxFCf counter is written into the REW_VAL field of the CPU extraction header. The CCM PDU can optionally be updated with RxFCf (OAM_MEP::MEP_CTRL.CCM_UPD_RxFCf-CL_ENA).

The RxFCf value and the received TxFCf value are kept until the next CCM_LM frame is sent to the peer MEP.

2.13.10.2 Up-MEP Dual-Ended Loss Measurement

The following lists some limitations in the Up-MEP implementation of dual-ended loss measurement:

- The TX_LM and RX_LM sample points in the initiator MEP and the remote MEP for the Up-MEP dual-ended loss measurement operation are shown in the following illustration.



The initiator MEP must inject CCM-LM frames using masquerading. When the CCM-LM frame is transmitted, the VOE updates the PDU with CCM-LM.TxFCf = TX LM.

At ingress, the VOE increments the LM Rx sequence number and copies the CCM-LM frame to the CPU. The frame is at the same time forwarded to the VOE at egress. When receiving the CCM-LM frame at egress, the VOE samples the RX_LM counter and stores the value in register OAM_MEP:VOE:CCM_CAP_RX_FCL.CCM_CAP_RX_FCL. In addition, the LM Rx sequence number is stored in register OAM_MEP:VOE:CCM_RX_SEQ_CFG.CCM_RX_SEQ[20:16].

The VOE signals to the CPU that there is a valid CCM-LM RX_LM sample by asserting the valid bit OAM_MEP::UPMEP_P_LM_CNT_VLD.CCM_LM_RX_CNT_VLD. If the valid bit is already set when a new CCM-LM frame arrives at the VOE, the sticky bit OAM_MEP:VOE:UPMEP_LM_CNT_STICKY.CCM_LM_RX_CNT_STICKY is asserted in addition to updating the counter value. This allows the CPU to detect that a valid RX_LM sample value has been overwritten before it was read by the CPU.

Delay measurement is used for the purpose of performance monitoring.

ETH-DM measures frame delay and frame delay variation between two peer MEPs by sending periodic frames with ETH-DM information to the peer MEP and receiving frames with ETH-DM information from the peer MEP. This information is used by the CPU to perform frame delay and frame delay variation measurements.

- One-way delay measurements using 1DM frames (unidirectional OAM flow).
- Two-way delay measurements using DMM and DMR frames (bidirectional OAM flow).

One-way frame delay measurement requires that the clocks at the transmitting MEP and the receiving MEPs are synchronized.

Up-MEP implementation of delay measurements require a synchronization mechanism for the CPU to establish a relationship between OAM PDUs forwarded to the CPU and the timestamps saved in the timestamp queue. The synchronization is based on a 5-bit DM Rx sequence number. The DM Rx sequence number is sampled when a delay measurement OAM PDU is received and the value is sent to the CPU as part of the CPU extraction header. The same value is saved along with the Rx timestamp in the timestamp queue. For more information about the timestamp queue, see [Section 2.17.3, Hardware Timestamping Module](#).

The latest DM Rx sequence number copied to the CPU is also stored in register OAM_MEP:VOE:LOOPBACK_CTRL.LB_ISDX_P[4:0]. The register is incremented upon reception of any of the supported DM PDUs in Up-MEP mode.

The DM Rx sequence number is contained in REW_VAL[31:27] in the CPU extraction header. The DM Rx sequence number is shared for the three different DM PDUs:

- 1DM
- DMM
- DMR

The synchronization mechanism is described below in the description of the Up-MEP functionality. The synchronization is not required for Down-MEP implementations.

2.13.12 ONE-WAY DELAY MEASUREMENTS (1DM)

The initiator MEP injects 1DM OAM PDU frames into the frame flow. The remote MEP extracts the 1DM frames to the CPU and the CPU does the delay measurement calculations. The required MEP operations to support one-way delay measurements are as follows.

At the initiator MEP:


- Inject 1DM frame from CPU into the frame flow at Tx VOE.
- Sample Tx timestamp (TX_TS) when the 1DM frame is transmitted by the VOE.
- Update outgoing 1DM frame with TX_TS (for example, 1DM.TxTSf = TX_TS).

At the remote MEP:

- Sample Rx timestamp (RX_TX) when 1DM frame is received by the VOE.
- Extract 1DM frame and RxTSf to CPU.

The 1DM function is illustrated in [Figure 2-50](#), . It is valid for both Up-MEP and Down-MEP implementations.

FIGURE 2-50: ONE-WAY DELAY MEASUREMENT (1DM)




1DM operation is enabled in the VOE through OAM_MEP:VOE:OAM_HW_CTRL.SDM_ENA.

2.13.12.1 Down-MEP One-Way Delay Measurements

This section describes the 1DM operation for a Down-MEP implementation in VSC7416-01. The TX_TS and RX_TS sample points for the Down-MEP 1DM operation are shown in the following illustration.

FIGURE 2-51: DOWN-MEP 1DM SAMPLE POINTS



Initiator MEP: 1DM injections

1DM frames are injected by the CPU. The 1DM.TxTimeStampf is updated with time the frame is leaving the port.

Remote MEP: 1DM extractions


1DM frames are extracted to the CPU and the RX_TS is sampled by the VOE. The timestamp is handled differently depending on whether the CPU is internal or external:

- If the CPU is connected via a front port, 1DM.RxTSf is updated with the value of RX_TS when transmitted on the CPU front port.
- If the CPU is the internal CPU, RX_TS is available in the CPU extraction header in the REW_VAL field.

2.13.12.2 Up-MEP One-Way Delay Measurements

This section describes the 1DM operation for an Up-MEP implementation in VSC7416-01. The TX_TS and RX_TS sample points for Up-MEP 1DM operation are shown in the following illustration.

FIGURE 2-52: UP-MEP 1DM SAMPLE POINTS



1DM handling works for both in-service and loopback Up-MEP modes.

The Up-MEP 1DM implementation is based on copying the DM Rx sequence number to the CPU along with a copy of the received 1DM frame. The same value of the DM Rx sequence number is inserted into the timestamp queue together with the RX_TS timestamp. The DM Rx sequence number can subsequently be used by the CPU to match the Rx 1DM frame with the timestamp queue entry.

Initiator MEP: 1DM injections

1DM frames must be injected into the relevant frame flow from the CPU using masquerading.

When the OAM 1DM PDU is transmitted, the 1DM.TxTimeStampf is updated with the TX_TS, which is the time the frame was received from the CPU.

Remote MEP: 1DM extractions

Frame extraction to the CPU is done at ingress where a copy of the 1DM frame is copied to the CPU extraction queue. This is done before it is known whether the frame will actually reach the Up-MEP located at the egress port as the frame might be discarded in the queue system. Together with the frame, the current value of the DM Rx sequence number is also sent to the CPU.

When the 1DM frame arrives at egress where the Up-MEP is located, the RX_TS sample is stored in the timestamp queue together with the egress port number, the VOE number, and the sample value of the DM Rx sequence number. The CPU can subsequently read the timestamp queue and use the information for calculating delay measurements. For more information about the timestamp queue, see [Section 2.17.3, Hardware Timestamping Module](#).

2.13.13 TWO-WAY DELAY MEASUREMENT (DMM)

For two-way delay measurements, the initiator MEP injects DMM frames into the frame flow. The DMM frames are looped back at the remote MEP. The loopback operation includes swapping the MAC addresses and modifying the OAM Opcode. Subsequently the remote MEP transmits the DMR back to the initiator MEP. The initiator MEP extracts the received DMR frame to the CPU. The required MEP operations to support two-way delay measurements are as follows.

Transmission at the initiator MEP:

- Inject DMM frame from the CPU into frame flow at the Tx MEP.

- Sample Tx timestamp (TX_TS) when the DMM frame is transmitted.
- Update outgoing the DMM frame with TX_TS (for example, $1DM.TxTSf = TX\_TS$).

Reception at the remote MEP:

- Sample Rx timestamp (RX_TS) when the DMM frame is received. Insert Rx timestamp into DMM.RxTSf.
- Create the DMR frame from the DMM frame.

Transmission at the remote MEP:

- Loop back the DMR frame.
- Sample Tx timestamp (TX_TS) when the DMR frame is transmitted and insert TX_TS into DMR.TxTSb.


Reception at the initiator MEP:

- Sample Rx timestamp (RX_TS) when the DMR frame is received.
- Extract the DMR frame to the CPU.
- Make RX_TS available to the CPU for delay measurement calculations.

Note that according to Y.1731 updating DMM.RxTSf and DMM.TxTSb is optional. Both are supported in VSC7416-01.

The two-way delay measurement function is shown in the following illustration, valid for both Up-MEP and Down-MEP implementations.

FIGURE 2-53: TWO-WAY DELAY MEASUREMENT



The transmitted PDU must be formatted as defined in ITU-t Y.1731 with the DMAC set to either the MAC address of the peer MEP or with a the standardized multicast address with MEL configured in the lowest three bits of the MAC address.

Hardware handling of DMM and DMR is enabled in OAM_MEP:VOE:OAM_HW_CTRL.DMM_ENA and OAM_MEP:VOE:OAM_HW_CTRL.DMR_ENA.

2.13.13.1 Down-MEP Two-Way Delay Measurements

Initiator MEP: DMM injections

It is possible to transmit DMM frames updated with hardware accuracy timestamps.

Remote MEP: DMM reception and DMR transmission

DMM frames are based on configuration either looped or sent to the CPU (OAM_MEP:VOE:LOOPBACK_CTRL.LB_D-MM_ENA).

If the DMM frames are looped back, the DMR frame's DMAC is set to the DMM frame's SMAC, the DMR frame's SMAC is updated with the MEP address (OAM_MEP:VOE:MEP_UC_MAC_MSB, OAM_MEP:VOE:MEP_UC_MAC_LSB), the DMR frame's Opcode is changed to DMR, the RxTimeStampf is updated with time the frame was received at the port, and the TxTimestampb is updated with time the frame leaves the port.

The DMM frames are looped back to the port they were received on and the ISDX is replaced with the value configured in OAM_MEP:VOE:LOOPBACK_CTRL.LB_ISDX, unless part of a protected service and the protection is active in which case the frames are looped to a protect loopback port (OAM_MEP:VOE:LOOPBACK_CTRL.LB_PORTNUM_P) and the ISDX replaced (OAM_MEP:VOE:LOOPBACK_CTRL.LB_ISDX_P).

If the DMM frames are sent to the CPU, only the RxTimeStampf is updated with the time the frames were received at the port.


Initiator MEP: DMR reception

The DMR frames are sent to the CPU and RxTimeStampf is updated with the time the frames were received at the port.

2.13.13.2 Up-MEP Two-Way Delay Measurements

This section describes the two-way delay measurements for an Up-MEP. The TX_TS and RX_TS sample points in the initiator MEP and the remote MEP for VSC7416-01 Up-MEP are shown in the following illustration.

FIGURE 2-54: UP-MEP DMM SAMPLE POINTS



The initiator MEP operation is independent of the configured Up-MEP loopback mode: both in-service mode and loopback mode are supported.

The way the loopback is done in the remote MEP depends on the configured loopback mode. This is explained in further detail in the following sections.

Initiator MEP: DMM injections

DMM frames must be injected into the relevant frame flow from the CPU using masquerading.

When the DMM frames are transmitted, the DMM.TxTimestampf is updated with the TX_TS, which is the time the frames were received from the CPU.


Remote MEP: DMM reception and DMR transmission (in-service mode)

In this mode, the Down-MEP and the Up-MEP loopback must be disabled (OAM_MEP:VOE:MEL_CTRL.UPMEP_LB, OAM_MEP:VOE:LOOPBACK_CTRL.LB_DMM_ENA).

When configured for Up-MEP in-service mode, the remote MEP's CPU must do the DMM to DMR updates. The VOE provides the RX_TS timestamp, which must be read by the CPU and inserted into the DMM.RxTSf field.

The loopback from DMM to DMR in the in-service mode is shown in the following illustration.

FIGURE 2-55: DMM TO DMR LOOPBACK IN IN-SERVICE MODE



In this mode, the remote MEP must be configured to send a copy incoming DMM frames to the CPU (OAM_MEP:VOE:OAM_CPU_COPY_CTRL.DMM_CPU_COPY_ENA). Furthermore, a copy of the DMM frame is sent to the Up-MEP VOE at the egress port.

When the remote MEP receives the DMM, the VOE samples the RX_TS and the DM Rx sequence number, discards the DMM frame, and saves the RX_TS and the DM Rx sequence number in the timestamp queue. Figure 2-55, illustrates that the DM Rx sequence number is sent to the CPU in the CPU extraction header as well as stored in the timestamp queue. The latest value of the DM Rx sequence number forwarded to the CPU is also stored in the register OAM_MEP:VOE:LOOPBACK_CTRL.LB_ISDX_P[4:0]. The DM Rx sequence number automatically increments after being sampled.

When the CPU receives the copy of the DMM frame, it must read the corresponding RX_TS from the timestamp queue. The correct timestamp can be recognized using the following fields from the timestamp queue entry.

- Flag indicating OAM timestamp entry (SYS::PTP_STATUS.PTP_TXSTAMP_OAM)
- DM Rx sequence number (SYS::PTP_STATUS.PTP_MESS_SEQ_ID[4:0])

The CPU performs the required DMM to DMR modifications and sets the field DMR.RxTSf = RX_TS. The CPU injects the DMR frame into the Tx OAM VOE frame flow.


When the remote MEP transmits the DMR towards the initiator MEP, the DMR.TxTSb is set to TX_TS.

Remote MEP: DMM reception and DMR transmission (loopback mode)

In this mode, the Down-MEP loopback must be disabled (OAM_MEP:VOE.LOOPBACK_CTRL.LB_DMM_ENA) while the Up-MEP loopback must be enabled (OAM_MEP:VOE.MEL_CTRL.UPMEP_LB).

The front port where the loopback is performed must be put in loopback mode to loop all outgoing traffic back into the port. This can be done using the register DEV::PCS1G_LB_CFG.TBI_HOST_LB_ENA.

When configured for Up-MEP loopback mode, the remote MEP does the DMM to DMR loopback all in hardware, which is shown in the following illustration.


FIGURE 2-56: DMM TO DMR LOOPBACK IN LOOPBACK MODE

When the remote MEP receives the DMM frame at the egress port, it samples the TX_TS and updates DMM.RxTSf with the value.

The VOE updates the DMR frame and sends the frame out on the loopback port. When the DMR frame is received on the loopback port, the VOE samples the TX_TS and updates the DMR.TxTSb with the value.

Initiator MEP: DMR reception

Received DMR frames are extracted to the CPU (OAM_MEP:VOE:OAM_CPU_COPY_CTRL.DMR_CPU_COPY_ENA). The DMR frame is also forwarded to the egress port where the MEP is located. The RX_TS is sampled and stored in the timestamp queue. This is shown in the following illustration.

FIGURE 2-57: DMR EXTRACTION TO CPU

When receiving the DMR frame at egress, the VOE samples the RX_TS and stores the value in the timestamp queue.

When the CPU receives the copy of the DMR frame, it must read the corresponding RX_TS from the timestamp queue. The correct timestamp can be recognized using the following fields.

- Flag indicating OAM timestamp entry (SYS::PTP_STATUS.PTP_TXSTAMP_OAM)
- DM Rx sequence number (SYS::PTP_STATUS.PTP_MESS_SEQ_ID[4:0])

Based on the DMR fields and the RX_TS sample, the CPU can do the delay measurement calculations.

2.13.14 MIP SUPPORT

It is possible to configure a VOE to operate as a MIP (OAM_MEP:VOE:BASIC_CTRL.VOE_AS_MIP_ENA, OAM_MEP:VOE:BASIC_CTRL.MIP_UNICAST_ONLY_COPY_ENA). MIP is supported for Down-MEP only.

When operating as a MIP, the VOE only reacts to LBM and LTM frames with unicast DMAC addresses. It discards LBR and LTR frames addressed to the VOE. The VOE is transparent to all other OAM OpCodes.

Unicast LTM frames are copied to the CPU (OAM_MEP:VOE:OAM_CPU_COPY_CTRL.LTM_COU_COPY_ENA) while unicast LBM frames are looped back (OAM_MEP:VOE:OAM_HW_CTRL.LBM_ENA).

2.13.15 INTERRUPT HANDLING

Each VOE can be configured to interrupt based on the following events.

- Changed MEG Level is detected
- Changed MEGID received
- Changed MEPID received
- Changed loss of continuity state
- Changed received RDI
- Changed priority received
- Changed period received
- Changed zero period received

Each of these events can be configured to cause interrupt via the corresponding enable bit in OAM_MEP:VOE:INTR_ENA. The combined interrupt status within a VOE is available in OAM_MEP::INTR.

Interrupt out of the VOP is enabled in OAM_MEP:MASTER_INTR_CTRL.OAM_MEP_INTR_ENA. The overall interrupt status for the VOP with all VOEs combined is available as a status signal in OAM_MEP::MASTER_INTR_CTRL.OAM_MEP_INTR.


For each VOE, the interrupt status is the direct results of whether a asserted sticky bit is enabled for interrupt. So when multiple events are causing interrupt, the interrupt status stays asserted until all pending sources are cleared. The interrupts are cleared by writing one to the asserted sticky bits in OAM_MEP:VOE:STICKY.

2.14 CPU Port Module

The CPU port module connects the switch core to the CPU system so that frames can be injected from or extracted to the CPU. It is also possible to use a regular front port as a CPU port. This is known as a Network Processor Interface (NPI).

The following illustration shows how the switch core interfaces to the CPU system through the CPU port module for injection and extraction of frames.

FIGURE 2-58: CPU INJECTION AND EXTRACTION



2.14.1 FRAME EXTRACTION

The following table lists the registers associated with frame extraction.

TABLE 2-139: FRAME EXTRACTION REGISTERS

Register	Description	Replication
QSYS::CPU_GROUP_MAP	Configuration of mapping of extraction queues to CPU ports.	None
SYS::PORT_MODE.INCL_X-TR_HDR	Enables insertion of extraction header. Configures formatting of outgoing frames.	Per CPU port (ports 11 and 12)

In the switch core, CPU extracted frames are forwarded to one of the eight CPU extraction queues. Each of these queues is mapped to one of two CPU ports (port 11 and port 12) through QSYS::CPU_GROUP_MAP. For each CPU port, there is a scheduler working either in strict mode or round robin, which selects between the CPU extraction queues mapped to the same CPU port. In strict mode, higher queue numbers are preferred over smaller queue numbers. In round robin, all queue are serviced one after another.

The two CPU ports contain the same rewriter as regular front ports. The rewriter modifies the frames before sending them to the CPU. In particular, the rewriter inserts an extraction header (SYS::PORT_MODE.INCL_XTR_HDR), which contains relevant side band information about the frame such as the frame's classification result (VLAN tag information, DSCP, QoS class), ingress timestamp, and the reason for sending the frame to the CPU. For more information about the rewriter, see [Section 2.12, Rewriter](#).

The device CPU contains the functionality for reading out the frames. This can be done through the frame DMA or regular register access. The Versatile Register Access Protocol (VRAP) Engine can also be attached to one of the CPU extraction groups. This allows an external device to access internal registers through Ethernet frames.

The following table lists the contents of the CPU extraction header.

TABLE 2-140: CPU EXTRACTION HEADER

Field	Bit	Width	Description
RESERVED	127	1	Reserved.
REW_OP	118	9	Rewriter operation command. REW_OP[2:0] = 3 implies two-step PTP where REW_OP[8:3] contains the PTP timestamp identifier. REW_OP[2:0] = 4 implies OAM. Other settings do not apply.
REW_VAL	86	32	By default, this field contains the Rx time when the frame was received. If REW_OP[2:0] = 4, the following applies. For REW_OP[2:0] = 4 and PDUs LMM (in-service mode only), LMR, and CCM-LM, REW_VAL[31:27] contains the LM Rx sequence number. This is applicable in Up-MEP mode only. For REW_OP[2:0] = 4 and PDUs LMR and CCM-LM, REW_VAL contains the RxFCI sample. This is applicable in Down-MEP mode only. For REW_OP[2:0] = 4 and PDUs 1DM, DMM and DMR, REW_VAL[31:27] contains the DM Rx sequence number. This is applicable in Up-MEP mode only. For REW_OP[2:0] = 4 and PDU 1DM, REW_VAL contains the RX_TS timestamp. This is applicable in Down-MEP mode only.
LLEN	80	6	Frame length in bytes: $60 \times WLEN + LLEN - 80$.
WLEN	72	8	See LLEN.
MPLS_POP	69	3	Indicates number of bytes to pop in the rewriter. 0: Does not apply. 1: Frame is MPLS encapsulated. Do not pop any bytes. 2: Frame is MPLS encapsulated. Pop 14 bytes. 3: Frame is MPLS encapsulated. Pop 18 bytes. 4: Frame is MPLS encapsulated. Pop 22 bytes. 5: Frame is MPLS encapsulated. Pop 26 bytes. 6: Frame is MPLS encapsulated. Pop 30 bytes. 7: Frame is MPLS encapsulated. Pop 34 bytes.
MPLS_TC	66	3	The frame's classified MPLS TC value. Only applicable for MPLS frames.
MPLS_SBIT	65	1	The frame's classified MPLS S-bit value. Only applicable for MPLS frames.
MPLS_TTL	57	8	The frame's classified MPLS TTL value. Only applicable for MPLS frames.
ISDX	47	10	The frame's classified ISDX.
SRC_PORT	43	4	The port number where the frame was received (0-11).
ACL_ID	37	6	If ACL_HIT is set, this value is the combined ACL_ID action of the rules hit in IS2.
RESERVED	36	1	Reserved.

TABLE 2-140: CPU EXTRACTION HEADER (CONTINUED)

Field	Bit	Width	Description
SFLOW_ID	32	4	sFlow sampling ID. 0-11: Frame was sFlow sampled by a Tx sampler on port given by SFLOW_ID. 12: Frame was sFlow sampled by an Rx sampler on port given by SRC_PORT. 13-14: Reserved. 15: Frame was not sFlow sampled.
ACL_HIT	31	1	Set if frame has hit a rule in IS2, which copies the frame to the CPU (IS2 actions CPU_COPY_ENA or HIT_ME_ONCE).
DP	30	1	The frame's drop precedence (DP) level after policing.
LRN_FLAGS	28	2	The source MAC address learning action triggered by the frame. 0: No learning. 1: Learning of a new entry. 2: Updating of an already learned unlocked entry. 3: Updating of an already learned locked entry.
CPUQ	20	8	CPU extraction queue mask (one bit per CPU extraction queue). Each bit set implies the frame was subjected to CPU forwarding to the specific queue.
QOS_CLASS	17	3	The frame's classified QoS class.
TAG_TYPE	16	1	The tag information's associated Tag Protocol Identifier (TPID). The definitions are: 0: C-tag: EtherType = 0x8100. 1: S-tag: EtherType = 0x88A8 or custom value.
PCP	13	3	The frame's classified PCP.
DEI	12	1	The frame's classified DEI.
VID	0	12	The frame's classified VID.

2.14.2 FRAME INJECTION

The following table lists the registers associated with frame injection.

TABLE 2-141: FRAME INJECTION REGISTERS

Register	Description	Replication
SYS::PORT_MODE.INCL_INJ_HDR	Enable parsing of injection header. Configures formatting of incoming frames.	Per CPU port (ports 11 and 12)
QSYS::EQ_PREFER_SRC	Enable preferred arbitration of the CPU port (port 11) over front ports.	CPU port (port 11 only)

The CPU injects frames through the two CPU injection groups that are independent of each other. The injection groups connect to the two CPU ports (port 11 and port 12) in the CPU port module. In CPU port module, each of the two CPU ports have dedicated access to the switch core. Inside the switch core, all CPU injected frames are seen as coming from CPU port (port 11). This implies that both CPU injection groups consume memory resources from the shared queue system for port 11 and that analyzer configuration for port 11 are applied to all frames.

In the switch core, the CPU port can be preferred over other ingress ports when transferring frames to egress queues by enabling precedence of the CPU port (QSYS::EQ_PREFER_SRC).

The first 20 bytes of a frame written to a CPU injection group is an injection header containing relevant side band information about how the frame must be processed by the switch core. The CPU ports must be enabled to expect the CPU injection header (SYS::PORT_MODE.INCL_INJ_HDR).

On a per-frame basis, the CPU controls whether frames injected through the CPU port module are processed by the analyzer. If the frame is processed by the analyzer, it is sent through the processing steps to calculate the destination ports for the frame. If analyzer processing is not selected, the CPU can specify the destination port set and related information to fully control the forwarding of the frame. For more information about the analyzer's processing steps, see [Section 2.8.4, Forwarding Engine](#).

The contents of the CPU injection header is listed in the following table.

TABLE 2-142: CPU INJECTION HEADER

Field	Bit	Width	Description
BYPASS	127	1	When this bit is set, the analyzer processing is skipped for this frame. The destination set is specified in DEST and CPU_QUEUE. Forwarding uses the QOS_CLASS, and the rewriter uses the tag information (POP_CNT, TAG_TYPE, PCP, DEI, VID) for rewriting actions. When this bit is cleared, the analyzer determines the destination set, QoS class, and VLAN classification for the frame through normal frame processing including lookups in the MAC table and VLAN table.
MASQ	126	1	When this bit is set, masquerading is enabled. The classifier, analyzer, and queue system handle the frame as if it was received by the ingress port specified in MASQ_PORT. This field overloads the REW_OP field and should only be used when BYPASS = 0.
MASQ_PORT	122	4	Masquerading port used when MASQ is set. This field overloads the REW_OP field and should only be used when BYPASS = 0.
REW_OP	118	9	Rewriter operation command. Used when BYPASS = 1. The following commands are supported: No operation: REW_OP[3:0] = 0. No operation. Special Rewrite: REW_OP[3:0] = 8. Swap the MAC addresses and clear bit 40 in the new SMAC when transmitting the frame. DSCP remark: REW_OP[2:0] = 1. REW_OP[8:3] contains the frame's classified DSCP to be used at egress for remarking. One-step PTP: REW_OP[2:0] = 2. The frame's residence time is added to the correction field in the PTP frame. The following sub-commands can be encoded: REW_OP[5]: Set if egress delay must be added to residence time. Two-step PTP: REW_OP[2:0] = 3. The frame's departure timestamp is saved in the timestamp FIFO queue at egress. REW_OP[8:3] contains the PTP timestamp identifier used by the timestamp FIFO queue. Identifiers 0 through 3 are pre-allocated to be used by CPU injected frames. OAM: REW_OP[2:0] = 4. Bits 6:3: 0: Not applicable 1: Not applicable 2: CCM 3: CCM with loss measurement 4: LBM 5: LBR 6: LMM 7: LMR 8: DMM 9: DMR 10: SDM Origin PTP: REW_OP[2:0] = 5. The time of day at the frame's departure time is written into the originTimestamp field in the PTP frame. Unspecified bits must be set to 0.

TABLE 2-142: CPU INJECTION HEADER (CONTINUED)

Field	Bit	Width	Description
REW_VAL	86	32	By default, this field contains the frame's receive timestamp. For injected frames, this can be set by the CPU to indicate when the injection started. The rewriter can then calculate a residence time based on REW_VAL and the frame's transmission timestamp. If TFRM_TIMER > 0, then REW_VAL contains the transmission slot for periodic frame transmission (0 through 1023). For OAM LMM frames, REW_VAL contains the RxFCf counter for the LMR frame. This requires OAM to be enabled.
RESERVED	72	14	Reserved.
MPLS_POP	69	3	Popping of the MPLS encapsulation by the rewriter before adding a new MPLS encapsulation. Used when BYPASS = 1. 0: Frame is not MPLS encapsulated 1: Frame is MPLS encapsulated. Do not pop 2: Frame is MPLS encapsulated. Pop 14 bytes (BDA, BSA, ETYPE) 3: Frame is MPLS encapsulated. Pop 18 bytes 4: Frame is MPLS encapsulated. Pop 22 bytes 5: Frame is MPLS encapsulated. Pop 26 bytes 6: Frame is MPLS encapsulated. Pop 30 bytes 7: Frame is MPLS encapsulated. Pop 34 bytes (BDA, BSA, BVLAN, ETYPE, 3 Labels, CW)
DEST	57	12	This is the destination set for the frame. DEST[11] is the CPU. Used when BYPASS = 1.
ISDX	47	10	The frame's classified ISDX. Used when BYPASS = 1.
SRC_PORT	43	4	The port number where the frame was injected (0-12).
RESERVED	41	2	Reserved.
TFRM_TIMER	37	4	Selects timer for periodic transmissions (1 through 8). If TFRM_TIMER=0 then normal injection.
RESERVED	33	4	Reserved.
ISDX-ES0_KEY_ENA	32	1	If set, the ES0 lookup in the rewriter uses the ISDX in the key instead of the VID. Used when BYPASS = 1.
RESERVED	31	1	Reserved.
DP	30	1	The frame's drop precedence (DP) level after policing. Used when BYPASS = 1.
POP_CNT	28	2	Number of VLAN tags that must be popped in the rewriter before adding new tags. Used when BYPASS = 1. 0: No tags must be popped. 1: One tag must be popped. 2: Two tags must be popped. 3: Disable all rewriting of the frame. The rewriter can still update the FCS.
CPUQ	20	8	CPU extraction queue mask (one bit per CPU extraction queue). Each bit set implies the frame must be forwarded by the CPU to the specific queue. Used when BYPASS = 1 and DEST[11] = 1.
QOS_CLASS	17	3	The frame's classified QoS class. Used when BYPASS = 1.

TABLE 2-142: CPU INJECTION HEADER (CONTINUED)

Field	Bit	Width	Description
TAG_TYPE	16	1	The tag information's associated Tag Protocol Identifier (TPID). Used when BYPASS = 1. 0: C-tag: EtherType = 0x8100. 1: S-tag: EtherType = 0x88A8 or custom value.
PCP	13	3	The frame's classified PCP. Used when BYPASS = 1.
DEI	12	1	The frame's classified DEI. Used when BYPASS = 1.
VID	0	12	The frame's classified VID. Used when BYPASS = 1.

2.14.3 NETWORK PROCESSOR INTERFACE (NPI)

The following table lists the registers associated with the network processor interface.


TABLE 2-143: NETWORK PROCESSOR INTERFACE REGISTERS

Register	Description	Replication
QSYS::EXT_CPU_CFG	Configuration of the NPI port number and configuration of which CPU extraction queues are redirected to the NPI.	None
SYS::PORT_MODE.INCL_X-TR_HDR	Enables insertion of extraction header.	Per port
SYS::PORT_MODE.INCL_INJ_HDR	Configuration of NPI ingress mode.	Per port

Any front port can be configured as a network processor interface through which frames can be injected from and extracted to an external CPU. Only one port can be an NPI at the same time. QSYS::EXT_CPU_CFG.EXT_CPU_PORT holds the port number of the NPI.

A dual CPU system is possible where both the internal and the external CPU are active at the same time. Through QSYS::EXT_CPU_CFG.EXT_CPUQ_MSK, it is configurable to which of the eight CPU extraction queues are directed to the internal CPU and which are directed to external CPU. A frame can be extracted to both the internal CPU and the external CPU if the frame is extracted for multiple reasons.

Frames being extracted by the external CPU can have the CPU extraction header inserted in front of the frame (SYS::PORT_MODE.INCL_XTR_HDR). Similarly, frames being injected into the switch core by the external CPU can have the CPU injection header inserted in front of the frame (SYS::PORT_MODE.INCL_INJ_HDR). In addition, there are three different options in terms of inserting a prefix in front of the CPU injection or extraction header. The following illustration shows the different frame formats supported.

FIGURE 2-59: CPU INJECTION AND EXTRACTION PREFIXES

Note that inserting a CPU extraction header in front of a frame disables all other frame modifications done in the rewriter. When injecting frames with an CPU injection header all incoming frames are expected to adhere to the configured prefix mode. The following parsing of the frames takes place:

- No prefix. All incoming frames are parsed as if they have a CPU injection header in front of the frame. Forwarding is based on the instructions in the CPU injection header.
- Short prefix. Incoming frames are checked against the expected format (start of frame must include 0x88800005). If the prefix does not match then an error indication is set (SYS::PORT_MODE.INJ_HDR_ERR). Compliant frames are forwarded based on the instructions in the CPU injection header. Non-compliant frames are forwarded as normal frames where the prefix and CPU injection header equaling the first 20 bytes of the frame are skipped.
- Long prefix. All incoming frames are parsed as if they have a long prefix followed by the CPU injection header in front of the frame. The incoming frames are checked against the expected format (start of frame must include 12 bytes of MAC addresses followed by 0x88800005). If the prefix does not match then an error indication is set (SYS::PORT_MODE.INJ_HDR_ERR). Both compliant and non-compliant frames are forwarded based on the instructions in the CPU injection header.

The external CPU can control forwarding of injected frames by either letting the frame analyze and forward accordingly or directly specifying the destination set. This is controlled through the BYPASS field in the CPU injection header.

2.14.4 FRAME GENERATION ENGINE FOR PERIODIC TRANSMISSIONS

The following table lists the registers associated with the frame generation engine.

TABLE 2-144: FRAME GENERATION ENGINE

Register	Description	Replication
QSYS::TFRM_MISC	Configuration to cancel ongoing periodic transmissions.	None
QSYS::TFRM_PORT_D-LY_ENA	Enable spaced-out transmissions when multiple periodic transmissions are scheduled at the same time.	Per port
QSYS::TFRM_TIMER_CFG_x	Period configuration. These registers configure the transmission period between frames.	8

VSC7416-01 contains a frame generation engine that can periodically transmit frames on a port with a programmable transmission period. This is powerful in OAM applications to generate for example CCM frames and loopback frames. The transmission period can be as low as 200 ns, which implies wire-speed back-to-back transmissions, or as high as several minutes. Up to 1024 periodic transmissions can coexist with up to eight different transmission periods.

To set up a periodic transmission, the CPU must inject a setup frame using the following settings in the injection header:

- REW_VAL set to the selected transmission slot (0 through 1023).
- ACL_ID set to the selected timer (1 through 8). The transmission period is programmed in steps of 198.2 ns in TFRM_TIMER_CFG_x, where x = ACL_ID. Note that injecting a frame with ACL_ID > 0 effectively enables the periodic transmissions. To inject a normal frame, ACL_ID must be set to 0.
- BYPASS set to 1.
- DEST set to the Tx port to which the periodic transmission applies. Only one Tx port per transmission is possible.
- Other fields in the injection header are applicable the same way as for normal injections. Note, that the periodic transmissions are subject to normal rewriter operations before being transmitted on the Tx port, see [Section 2.12, Rewriter](#).

After injected, the setup frame is placed into the selected transmission slot and it is periodically transmitted using the transmission period defined by the selected timer. Every time the transmission period has passed since the last transmission, the frame is scheduled for a new transmission on the associated port. The periodic frame transmission takes precedence over other transmissions on the port and as a consequence the frame is transmitted immediately after a potential ongoing transmission has ended.

If multiple periodic transmissions on a Tx port use the same timer, multiple frames are scheduled for transmission at the same time when the period has passed. By default, these frames are transmitted in a burst, back-to-back. However, it is also possible to space-out these frame over time and thereby allowing the normal data traffic to be interleaved the periodic transmissions. If the spacing is enabled (TFRM_PORT_DLY), timer 8 defines the period between frames in the burst.

Example: Port 5 has defined 10 CCM flows with a 3.3 ms transmission period using timer 1. Spacing is enabled and timer 8 is set to 100 us. When the timer 1 period has passed, 10 frames are up for transmission. The spacing results in the 10 frames being transmitted with 100 us between each frame.

Periodic transmissions are cancelled again by programming the slot number in TFRM_MISC.TIMED_CANCEL_SLOT and setting TFRM_MISC.TIMED_CANCEL_1SHOT. This removes the injected frame from the transmission slot.

Transmission slots 0 through 15 support any transmission periods including back-to-back transmissions while slots 16 through 1023 support transmission periods larger than 30 μ s.

Frames transmitted from the frame generation engine are counted by the Tx counters just like normal frames using the QoS class and ISDX specified in the injection header by the setup frame.

2.15 VRAP Engine

The Versatile Register Access Protocol (VRAP) engine allows external equipment to access registers in VSC7416-01 through any of the Ethernet port's on the device. The VRAP engine interprets incoming VRAP requests, and executes read, write, and read-modify-write commands contained in the frames. The results of the register accesses are reported back to the transmitter through automatically generated VRAP response frames.

VSC7416-01 supports version 1 of VRAP. All VRAP frames, both requests and responses, are standard Ethernet frames. All VRAP protocol fields are big-endian formatted.

The registers listed in the following table control the VRAP engine.

TABLE 2-145: VRAP REGISTERS

Target:Register_group:Register.field	Description	Replication
ANA:CPU_F-WD_CFG:CPU_VRAP_REDIR_ENA	Enable redirection of VRAP frames to CPU.	Per virtual port profile
ANA:CPUQ_CFG2:CPUQ_VRAP	Configure the CPU extraction queue for VRAP frames.	None
QSYS:CPU_GROUP_MAP:CPU_GROUP_MAP	Map VRAP CPU extraction queue to a CPU port. One CPU port must be reserved for VRAP frames.	None

TABLE 2-145: VRAP REGISTERS (CONTINUED)

Target:Register_group:Register.field	Description	Replication
DEVCPU_QS:XTR_GRP_CFG:MODE	Enable VRAP mode for reserved CPU port.	Per CPU port
DEVCPU_QS:INJ_GRP_CFG:MODE	Enable VRAP mode injection for reserved CPU port.	Per CPU port
SYS:PORT_MODE:INCL_INJ_HDR	The injection header is not present for VRAP response frames.	Per port
SYS:PORT_MODE:INCL_XTR_HDR	The extraction header is not present for redirected VRAP frames.	Per port
DEVCPU_GCB:VRAP_ACCESS_STAT	VRAP access status.	None

The VRAP engine processes incoming VRAP frames that are redirected to the VRAP CPU extraction queue by the basic classifier. For more information about the VRAP filter in the classifier, see [Section 2.6.7, CPU Forwarding Determination](#).


The VRAP engine is enabled by allocating one of the two CPU ports as a dedicated VRAP CPU port (DEVCPU_QS:XTR_GRP_CFG:MODE and DEVCPU_QS:INJ_GRP_CFG:MODE). The VRAP CPU extraction queue (ANA:CPUQ_CFG2:CPUQ_VRAP) must be mapped as the only CPU extraction queue to the VRAP CPU port (QSYS:CPU_GROUP_MAP:CPU_GROUP_MAP). In addition, the VRAP CPU port must disable the use of CPU injection and CPU extraction headers (SYS:PORT_MODE:INCL_INJ_HDR and SYS:PORT_MODE:INCL_XTR_HDR).

The complete VRAP functionality can be enabled automatically at chip startup by the use of special chip boot modes, see [Section 3.1, VCore Configurations](#).

The following describes the VRAP frame formats.

2.15.1 VRAP REQUEST FRAME FORMAT

The following illustration shows the format of a VRAP request frame.

FIGURE 2-60: VRAP REQUEST FRAME FORMAT

VRAP request frames can optionally be VLAN tagged with one VLAN tag.

The EtherType = 0x8880 and the Ethernet Protocol Identifier (EPID) = 0x0004 identify the VRAP frames. The subsequent VRAP header is used in both request and response frames.

The VRAP commands included in the request frame are the actual register access commands. The VRAP engine supports the following five VRAP commands:

- **READ:** Returns the 32-bit contents of any register in the device.
- **WRITE:** Writes a 32-bit value to any register in the device.
- **READ-MODIFY-WRITE:** Does read/modify/write of any 32-bit register in the device.
- **IDLE:** Does not access registers but is useful for padding and identification purposes.
- **PAUSE:** Does not access registers but causes the VRAP engine to pause between register access.

Each of the VRAP commands are described in the following sections. Each VRAP request frame can contain multiple VRAP commands. Commands are processed sequentially starting with VRAP command #1, #2, and so on. For more information, see [Figure 2-60](#). There are no restrictions on the order or number of commands in the frame.


2.15.2 VRAP RESPONSE FRAME FORMAT

Having processed a VRAP request frame by executing the register access commands contained in the frame, the VRAP engine generates a VRAP response frame. The response frame uses the swapped DMAC and SMAC addresses in the VRAP request frame and it is transmitted on port where the VRAP request frame was received. If the DMAC in the VRAP

request frame is a multicast DMAC, then bit 40 is cleared before using the DMAC as the SMAC in the VRAP response frame. The results of the VRAP commands are inserted into the response frame in the same order as the VRAP commands were processed.

The following illustration shows the format of a VRAP response frame.

FIGURE 2-61: VRAP RESPONSE FRAME FORMAT




The VRAP response frame follows the VRAP request frame in terms of VLAN tagging: If the VRAP request was tagged, the VRAP response is also tagged using the same VLAN.

Only the READ and IDLE commands require data to be returned to the transmitter of the VRAP request frame. However, even if the VRAP request frame does not contain READ and IDLE commands, a VRAP response frame is generated with enough padding data (zeros) to fulfill the minimum transfer unit size.

2.15.3 VRAP HEADER FORMAT

Both VRAP request and VRAP response frames contain a 32-bit VRAP header. The VRAP header is used to identify the protocol version and to differentiate between VRAP requests and VRAP response frames. VSC7416-01 supports VRAP version 1. The following illustration shows the VRAP header.

FIGURE 2-62: VRAP HEADER FORMAT




Version:

Set to 0x1 in VRAP response frames.
VRAP request frames are ignored if Version <> 1

Flags:

4 bits are used for Flags. Only one flag is defined:



R: 0=Request, 1=Response

Rsv: Set to 0 in VRAP response frames and ignored in VRAP request frames.

Reserved:

Set to 0 in VRAP response frames and ignored in VRAP request frames.


A valid VRAP request frame must use Version = 1 and Flags.R = 1. Frames with an invalid VRAP header are discarded and not processed by the VRAP engine.

2.15.4 VRAP READ COMMAND

The VRAP READ command returns the contents of any 32-bit register inside the device. The 32-bit read-data result is put into the VRAP response frame.

The READ command is 4 bytes wide and consists of one 32-bit address field, which is 32-bit aligned. The 2 least significant bits of the address set to 01. The following illustration shows the request command and the associated response result.

FIGURE 2-63: READ COMMAND




2.15.5 VRAP WRITE COMMAND

The WRITE command writes a 32-bit value to any register inside the device.

The WRITE command is 8 bytes wide and consists of one 32-bit address field, which is 32-bit aligned. The two least significant bits of the address set to 10, followed by one 32-bit write-data field. The following illustration shows the command.

FIGURE 2-64: WRITE COMMAND




2.15.6 VRAP READ-MODIFY-WRITE COMMAND

The READ-MODIFY-WRITE command does read/modify/write-back on any 32-bit register inside the device.

The READ-MODIFY-WRITE command is 12 bytes wide and consists of one 32-bit address field, which is 32-bit aligned. The two least significant bits of the address set to 11 followed by one 32-bit write-data field followed by one 32-bit overwrite-mask field. For bits set in the overwrite mask, the corresponding bits in the write data field are written to the register while bits cleared in the overwrite mask are untouched when writing to the register. The following figure shows the command.

FIGURE 2-65: READ-MODIFY-WRITE COMMAND




2.15.7 VRAP IDLE COMMAND

The IDLE command does not access registers in the device. Instead it just copies itself (the entire command) into the VRAP response. This can be used for padding to fulfill the minimum transmission unit size, or an external CPU can use it to insert a unique code into each VRAP request frame so that it can separate different replies from each other.

The IDLE command is 4 bytes wide and consists of one 32-bit code word with the four least significant bits of the code word set to 0000. The following illustration depicts the request command and the associated response.

FIGURE 2-66: IDLE COMMAND




2.15.8 VRAP PAUSE COMMAND

The PAUSE command does not access registers in the device. Instead it causes the VRAP engine to enter a wait state and remain there until the pause time has expired. This can be used to ensure sufficient delay between VRAP commands when this is needed.

The PAUSE command is 4 bytes wide and consists of one 32-bit code word with the four least significant bits of the code word set to 0100. The wait time is controlled by the 28 most significant bits of the wait command. The time unit is 1 us. The following illustration depicts the PAUSE command.

FIGURE 2-67: PAUSE COMMAND



2.16 Layer-1 Timing

The following table lists the registers associated with Layer-1 timing.

TABLE 2-146: LAYER-1 TIMING CONFIGURATION REGISTERS

Register	Description	Replication
HSIO::SYNC_ETH_CFG0.RECO_CLK_ENA	Enable use of recovered clock output pin.	Per recovered clock output pin
HSIO::SYNC_ETH_CFG1.RECO_CLK_SEL	Select whether to use individual SERDES clock source or common clock source.	Per recovered clock output pin
HSIO::SYNC_ETH_CFG1.RECO_CLK_SEL0	Select which SERDES drives recovered clock output pin 0.	None
HSIO::SYN-C_ETH_CFG2.SEL_RECO_CLK_B	Clock divider for common recovered clock.	None
HSIO::SERDES1G_COMMON_CFG	Recovered clock selection.	Per SERDES1G port
HSIO::SERDES6G_COMMON_CFG	Recovered clock selection.	Per SERDES6G port
HSIO::SERDES1G_MISC_CFG	Clock divider for individual recovered clock.	Per SERDES1G port
HSIO::SERDES6G_MISC_CFG	Clock divider for individual recovered clock.	Per SERDES6G port

VSC7416-01 supports ten recovered clock output pins (RCVRD_CLK[9:0]) that can provide timing sources for external timing circuitry for redundant timing implementations. SERDES1G macros 0 through 7 and SERDES6G macros 0 and 1 can derive a recovered clock from the incoming data streams. If timing is compromised on either of the ten sources, the appropriate clock output can be squelched to assist with fast timing switchover in the clock synchronization circuitry.

The following illustration shows how the recovered clocks from the SERDES macros connect to the output pins RCVRD_CLK[9:0] and the associated configurations.

2.17 Hardware Timestamping

Hardware timestamping provides micronanosecond-accurate frame arrival and departure time stamps, which are used to obtain high precision timing synchronization and timing distribution, as well as significantly better accuracy in performance monitoring measurements than what is obtained from pure software implementations.

All frames are Rx timestamped on arrival with a 48-bit timestamp value using a hardware timer (timestamp) that is implemented in the Media Access Control (MAC) block. The Rx timestamp provides high timestamp accuracy relative to actual arrival time of the first byte of the frame from the PHY device. Within the VCAP IS2, it is decided if the frame and associated Rx timestamp must be redirected or copied to CPU for processing. The frame is forwarded as normal otherwise.

The VCAP IS2 also decides if a Tx timestamp must be triggered for a frame. Given the Rx and Tx timestamps, the frame's residence time inside the switch is calculated. The residence time can be stored in a timestamp queue for the CPU to access (two-step timestamping) or the residence time can be used to update the residence time field inside Precision Time Protocol frames (one-step timestamping).

The Tx timestamp is located at the transmit side of the MAC block as close to the PHY device as possible and provides high accuracy of timestamp relative to when the first byte of the frame is actually transmitted to the PHY.

The device also implements a time of day counter with nanosecond-accuracy. The time of day counter is derived from a one-second timer. The one-second timer generates a pulse per second and is either derived from an adjusted system clock or from external timing equipment.

2.17.1 TIMESTAMP CLASSIFICATION

Frames requiring Rx or Tx timestamping are identified by VCAP IS2. The IS2 action that triggers timestamping is REW_OP[2:0], with the following options:

- One-step timestamping (REW_OP[2:0] = 2), which implies adding the frame's residence time to the correction field.
- Two-step timestamping (REW_OP[2:0] = 3), which implies saving the frame's Tx time in a timestamp queue.
- Origin timestamping (REW_OP[2:0] = 5), which implies writing the time of day into the originTimestamp field.

IS2 can be configured to identify the following frame formats from IEEE 1588-2008:

- Transport of PTP over UDP over IPv4
- Transport of PTP over UDP over IPv6
- Transport of PTP over IEEE 802.3/Ethernet
- Transport of PTP over UDP over IP-MPLS

In addition, frames can be encapsulated in MPLS-TP and Q-in-Q tunnels. For more information about the frame encapsulations and PTP protocol fields supported by the VSC7416-01 device, see [Section 2.7.4, VCAP IS2](#).

2.17.2 MASTER TIMER

The one-second timer generates one synchronization pulse per second, which is used for the time-of-day seconds counter.

The one-second timer can provide a synchronization pulse output or a reference clock output derived from the one-second synchronization pulse. The one-second timer synchronization pulse can also be controlled from an external pin.

The registers listed in the following table control and monitor the one-second timer.


TABLE 2-148: ONE-SECOND TIMER REGISTERS

Target:Register_group:Register.field	Description	Replication
DEVCPU_GCB::PTP_MISC_CFG	GPIO configuration of hardware timer.	1
DEVCPU_GCB::PTP_UPPER_LIMIT_CFG	One-second counter configuration.	1
DEVCPU_GCB::PTP_UPPER_LIMIT_1_-TIME_ADJ_CFG	One-second counter configuration.	1
DEVCPU_GCB::PTP_EXT_SYNC_-TIME_CFG	One-second counter adjustment configuration for external load pulse.	Per 1588 LD/ST GPIO pin

TABLE 2-148: ONE-SECOND TIMER REGISTERS (CONTINUED)

Target:Register_group:Register.field	Description	Replication
DEVCPU_GCB::GEN_EXT_CLK_HIGH_PERIOD_CFG	External clock output configuration.	Per 1588 PPS GPIO pin
DEVCPU_GCB::GEN_EXT_CLK_LOW_PERIOD_CFG	External clock output configuration.	Per 1588 PPS GPIO pin
DEVCPU_GCB::GEN_EXT_CLK_CFG	External clock output configuration.	Per 1588 PPS GPIO pin
DEVCPU_GCB::CLK_ADJ_CFG	One-second counter adjustment configuration.	1
DEVCPU_GCB::PTP_SYNC_INTR_ENA_CFG	Interrupts control	1
DEVCPU_GCB::PTP_CURRENT_TIME_STAT	One-second counter statistics. Current count value.	1
DEVCPU_GCB::EXT_SYNC_CURRENT_TIME_STAT	One-second counter statistics. One-second counter value at the last external synchronization pulse input.	Per 1588 LD/ST GPIO pin
DEVCPU_GCB::PTP_EVT_STAT	One-second timer event statistics.	1

The one-second timer block diagram is shown in the following illustration.

FIGURE 2-69: ONE-SECOND TIMER BLOCK DIAGRAM

DEVCPU_GCB::PTP_MISC_CFG.PTP_ENA enables the one-second timer and must be set for one-second timer synchronization pulse generation.

By default, the one-second timer synchronization pulse is generated internally and with a frequency of one pulse per second (1 PPS) derived from the system clock. Other one-second timer synchronization pulse frequencies are obtained using register DEVCPU_GCB::PTP_UPPER_LIMIT_CFG. Every time a one-second timer synchronization pulse is generated, a sticky bit is set (DEVCPU_GCB::PTP_EVT_STAT.SYNC_STAT) and an interrupt is generated if DEVCPU_GCB::PTP_SYNC_INTR_ENA_CFG.SYNC_STAT_ENA is enabled.

2.17.2.1 One-Second Timer Counter Adjustments

If a one time correction to the one-second timer synchronization pulse is required, the correction time value must be written into register DEVCPU_GCB::PTP_UPPER_LIMIT_1_TIME_ADJ_CFG.PTP_UPPER_LIMIT_1_TIME_ADJ and one shot is enabled in register DEVCPU_GCB::PTP_UPPER_LIMIT_1_TIME_ADJ_CFG.PTP_UPPER_LIMIT_1_TIME_ADJ_SHOT.

The one-second timer can also be controlled by issuing counter corrections to the one-second timer counter. One-second timer counter corrections are enabled in register DEVCPU_GCB::CLK_ADJ_CFG.CLK_ADJ_ENA.

Corrections to the one-second timer counter is controlled by the adjustments counter. The adjustments counter issues ± 1 corrections to the one-second timer counter. The time period between one-second timer corrections is determined by the load value of the adjustments counter. Time periods between corrections ranges from nanoseconds to one second.

The adjustments counter operates as follows:

- When the counter value of the adjustments counter equals the load value (DEVCPU_GCB::CLK_ADJ_CFG.CLK_ADJ), a one-nanosecond correction is generated.
- Up or down corrections are determined by DEVCPU_GCB::CLK_ADJ_CFG.CLK_ADJ_DIR.
- The DEVCPU_GCB::CLK_ADJ_CFG.CLK_ADJ_UPD register controls whether a load value change takes immediate effect or whether it is synchronized to the next one-second timer synchronization pulse.
- When the load value change occurs, a sticky bit is set (DEVCPU_GCB::PTP_EVT_STAT.CLK_ADJ_UPD_STICKY). This sticky does not gate future updates to the load value and is informative only. The adjustment counter is reset by loading all zeros.

2.17.2.2 External Synchronization Pulse Input

A synchronization pulse can be provided as an input to the device by one of the two GPIO pins, 1588 LD0/ST0 or 1588LD1/ST1. This is controlled through register DEVCPU_GCB::PTP_MISC_CFG. When this input is used to control the one-second timer synchronization pulse, the register DEVCPU_GCB::PTP_MISC_CFG.EXT_SYNC_ENA must be set for the corresponding input. If set, the one-second timer counter is reset to the counter value specified in DEVCPU_GCB::PTP_EXT_SYNC_TIME_CFG by the external synchronization pulse.

In addition to loading the one-second timer counter, the GPIO pins, 1588 LD0/ST0 and 1588 LD1/ST1, can be used to store the one-second timer counter value when the external synchronization pulse arrives. This functionality is enabled in the DEVCPU_GCB::PTP_MISC_CFG.EXT_SYNC_CAP_ENA register. The one-second timer counter value is stored in register DEVCPU_GCB::EXT_SYNC_CURRENT_TIME_STAT and DEVCPU_GCB::PTP_EVT_STAT.EXT_SYNC_CURRENT_TIME_STICKY is set. If DEVCPU_GCB::PTP_SYNC_INTR_ENA_CFG.EXT_SYNC_CURRENT_TIME_ENA is set, an interrupt is generated when an external synchronization pulse is received.

2.17.2.3 External Synchronization Pulse Output

The DEVCPU_GCB::PTP_MISC_CFG.EXT_SYNC_OUTP_SEL register controls the one-second timer synchronization pulse that can be provided as an output to the device. The output is provided on either of the GPIO pins, 1588 PPS0 or 1588 PPS1. The output is delayed by one system clock cycle (6.4 ns).

2.17.2.4 Divide by x External Clock

A divide by x version of the one-second timer frequency can be provided as output from the device on the two 1588 PPS pins (1588 PPS0 or 1588 PPS1). External clock frequencies up to 25 MHz are supported. The default clock frequency is 10 kHz. Each GPIO pin can be programmed individually.

The frequency and duty cycle of the external clock is controlled by registers DEVCPU_GCB::GEN_EXT_CLK_HIGH_PERIOD_CFG and DEVCPU_GCB::GEN_EXT_CLK_LOW_PERIOD_CFG. The divide by X counter implements a high period and a low period of the external clock, based on these register values. The clock period of the external clock is calculated as:
$$(\text{GEN\_EXT\_CLK\_HIGH\_PERIOD\_CFG} + \text{GEN\_EXT\_CLK\_LOW\_PERIOD\_CFG}) \times 6.4 \text{ ns.}$$

The duty cycle of the external clock is only 50%/50% if GEN_EXT_CLK_HIGH_PERIOD_CFG and GEN_EXT_CLK_LOW_PERIOD_CFG are configured to the same value.

The DEVCPU_GCB::GEN_EXT_CLK_CFG.GEN_EXT_CLK_SYNC_ENA register controls if the divide by X counter that controls the external clock is synchronized to the one-second timer synchronization pulse.

Divide by X counter adjustments show up directly on the external clock (unfiltered).

The external clock can also be based on the one-second timer (DEVCPU_GCB::GEN_EXT_CLK_HIGH_PERIOD_CFG.MASTER_CNT_CLK_ENA). If enabled, a 31.25 MHz clock is generated, which is locked to the one-second timer. Any adjustments to the one-second timer also affect the external clock.

2.17.3 HARDWARE TIMESTAMPING MODULE

This section explains the functions of the hardware timestamping module. The following table lists the registers associated with the hardware timestamping module.

TABLE 2-149: HARDWARE TIMESTAMPING REGISTERS

Register	Description	Replication
SYS::PTP_TXSTAMP	Timestamp value in timestamp queue.	None
SYS::PTP_NXT	Advancing the timestamp queue.	None
SYS::PTP_STATUS	Timestamp queue status and entry data.	None
ANA::PTP_ID_HIGH	Release of timestamp identifiers, values 32 through 63.	None
ANA::PTP_ID_LOW	Release of timestamp identifiers, values 0 through 31.	None

Each port module contains a hardware timestamping module that measures arrival and departure times based on the master timer. For information about how the Rx and Tx timestamps are derived, see [Section 2.1.2.8, Rx and Tx Timestamps](#).

2.17.3.1 Two-Step Timestamping

Two-step timestamping is performed if the IS2 rewriter action is two-step (IS2_ACTION.REW_OP[2:0] = 3). This action can be applied to any frame, also non-PTP frames, because the frame itself is not modified. The frame's Tx timestamp is stored in a timestamp FIFO queue, which the CPU can access (SYS::PTP_STATUS). The timestamp is common for all egress ports and can contain up to 128 timestamps. Each entry in the timestamp queue contains the following fields:

- SYS::PTP_STATUS.PTP_MESS_VLD: A 1-bit valid bit meaning the entry is ready for reading.
- SYS::PTP_STATUS.PTP_MESS_ID: A 6-bit timestamp identifier. A unique timestamp identifier is assigned to each frame for which one or more Tx timestamps are generated. The timestamp identifier is also available in the CPU extraction header for frames extracted to the CPU. The timestamp identifier overloads the DSCP value in the CPU extraction header. For more information about the CPU extraction header, see [Table 2-140](#). By providing the timestamp identifier in both the timestamp queue and in the extracted frames, the CPU can correlate which timestamps belong to which frames. Note that timestamp identifier value 63 implies that no free identifier could be assigned to the frame. The timestamp entry can therefore not be trusted.
- SYS::PTP_STATUS.PTP_MESS_TXPORT: The port number where the frame is transmitted. When transmitting a frame on multiple ports, there are generated multiple entries in the timestamp queue. Each entry uses the same timestamp identifier but with different Tx port numbers.
- SYS::PTP_STATUS.PTP_TXSTAMP_OAM: For two-step timestamping, this field is always 0. If set, it identifies an entry inserted by the Versatile OAM processor (VOP).
- SYS::PTP_TXSTAMP: The frame's Tx timestamp.

The timestamp queue is a simple FIFO that can be read by the CPU. The timestamp queue provides the following handles for reading:

- Overflow of the queue is signaled through SYS::PTP_STATUS.PTP_OVFL. Overflow implies that one or more timestamps could not be enqueued due to all 128 entries being in use. Timestamps not enqueued are lost.
- The head-of-line entry is read through SYS::PTP_STATUS and SYS::PTP_TXSTAMP.
- Writing to the one-shot register SYS::PTP_NXT, removes the current head-of-line entry and advances the pointer to the next entry in the timestamp queue.

When two-step Tx timestamping is performed for a frame destined for the CPU extraction queues, no entry in the timestamp FIFO queue is made. The frame's Rx timestamp is available through the CPU extraction header.

The timestamp identifiers can take values between 0 to 63. Value 63 implies that all values 0-62 are in use. Values 0 – 3 are pre-assigned to the CPU to be used for injection of frames. The remaining values are assigned by the analyzer to frames requesting timestamping through the VCAP IS2 action. The assigned values must be released again by the CPU by writing to the corresponding bit in ANA::PTP_ID_HIGH (values 32 through 63) or ANA::PTP_ID_LOW (values 0 through 31). The CPU releases a timestamp identifier when it has read the anticipated timestamp entries from the timestamp queue. Note that multicasted frames generate a timestamp entry per egress port using the same timestamp identifier. Each of these entries must be read before the timestamp identifier is released.

Two-step timestamping can be disabled per egress port using REW:PORT:PTP_CFG.PTP_2STEP_DIS. This setting overrides the IS2 action.

2.17.4 TIME OF DAY COUNTER

The time of day counter holds a 48-bit seconds counter. The nanoseconds counter is given by the one-second timer counter (DEVCPU_GCB::PTP_CURRENT_TIME_STAT), and the seconds counter increments based on the one-second synchronization pulse.

The registers listed in the following table are used for monitoring the time-of-day seconds counter.

TABLE 2-150: TIME OF DAY COUNTER REGISTERS

Target:Register_group:Register.file Id	Description	Replication
SYS::PTP_TOD_MSB	Bits 47:32 of the time-of-day seconds counter.	None
SYS::PTP_TOD_LSB	Bits 31:0 of the time-of-day seconds counter.	None

The time-of-day seconds counter is read/writable.

2.18 Clocking and Reset

The following table lists the registers associated with clocking and reset.

TABLE 2-151: CLOCKING AND RESET REGISTERS

Target:Register_group:Register.file Id	Description	Replication
HSIO::PLL5G_CF30	LCPLL configuration continuation	Per PCLL
HSIO::PLL5G_STATUS0	LCPLL status	Per PCLL
DEVCPU_GCB::SOFT_CHIP_RST	Reset of the entire device	None
ICPU_CFG::RESET	CPU reset configuration	None

VSC7416-01 contains two LCPLLs: LCPLL and LCPLL2. LCPLL provides the clocks used by the SerDes, the central part of the switch core, the IEEE 1588 one-second timer, and the VCore-III CPU system. If LCPLL2 is selected by the PLL2_Enable input pin, LCPLL2 provides the clocks used by the central part of the switch core and the IEEE1588 one-second timer while LCPLL provides the clocks for the SerDes and the VCore-III CPU system. By enabling LCPLL2, the IEEE 1588 clock can, in particular, be made independent of a clock recovered by synchronous Ethernet clock.

The LCPLL2 can output a 250 MHz reference clock on the CLKOUT2 pins on the device that can, as an example, be fed to an external PHY device. Enable the clock by setting HSIO::PLL5G_CFG0.CLKOUT2_DIV = 5 and HSIO::PLL5G_CFG3.CLKOUT2_SEL = 4 for LCPLL2.

The reference clock for the LCPLL (RefClk_P and RefClk_N pins) and LCPLL2 (RefClk2_P and RefClk2_N pins) are either differential or single-ended. The frequency can be 25 MHz, 125 MHz, 156.25 MHz, or 250 MHz.

For more information about the reference clock frequency selections, see [Section 7.2.12, System Clock Interface](#).

For more information about reference clock options, see [Section 6.2.1, Reference Clocks](#).

A global software reset is performed with DEVCPU_GCB::SOFT_CHIP_RST.

For more information about the configuration of the CPU frequency and software reset options when using the V-Core-III, see [Section 3.2, Clocking and Reset](#).

For more information about the clock and reset configuration for the Ethernet interfaces in the port modules, see [Section 2.1.2, MAC](#), [Section 2.2, SERDES1G](#), and [Section 2.3, SERDES6G](#). The MAC clock domains are not included in the global reset.

3.1 VCore Configurations

The behavior of the VCore system after reset is determined by four VCore strapping pins that are overlaid on GPIO pins GPIO_22 through GPIO_25. The value of these GPIOs is sampled shortly after releasing reset to the VSC7416-01 device. By using resistors to pull the GPIOs either low or high, software can use these GPIO pins for other functions after reset has been released.

TABLE 3-1: VCore CONFIGURATIONS

VCore_CFG [3:0]	Behavior
0000	Port 0 is enabled for SGMII Aneq advertising 100M FDX NoFC and NPI port is enabled for SERDES NoAneq 1G FDX NoFC. VRAP block is accessible through the NPI port. Automatic boot of VCore CPU is disabled, and SI slave is enabled.
0001	Port 0 and NPI port are enabled for SERDES NoAneq 1G FDX NoFC. VRAP block is accessible through the NPI port. Automatic boot of VCore CPU is disabled, and SI slave is enabled.
0010	Port 1 is enabled for SGMII Aneq advertising 100M FDX NoFC and NPI port is enabled for SERDES NoAneq 1G FDX NoFC. VRAP block is accessible through the NPI port. Automatic boot of VCore CPU is disabled, and SI slave is enabled.
0011	Port 1 and NPI port are enabled for SERDES NoAneq 1G FDX NoFC. VRAP block is accessible through the NPI port. Automatic boot of VCore CPU is disabled, and SI slave is enabled.
0100	Port 8 is enabled for SGMII Aneq advertising 100M FDX NoFC and NPI port is enabled for SERDES NoAneq 1G FDX NoFC. VRAP block is accessible through the NPI port. Automatic boot of VCore CPU is disabled, and SI slave is enabled.
0101	Port 8 and NPI port are enabled for SERDES NoAneq 1G FDX NoFC. VRAP block is accessible through the NPI port. Automatic boot of VCore CPU is disabled, and SI slave is enabled.
0110	Port 9 is enabled for SGMII Aneq advertising 100M FDX NoFC and NPI port is enabled for SERDES NoAneq 1G FDX NoFC. VRAP block is accessible through the NPI port. Automatic boot of VCore CPU is disabled, and SI slave is enabled.
0111	Port 9 and NPI port are enabled for SERDES NoAneq 1G FDX NoFC. VRAP block is accessible through the NPI port. Automatic boot of VCore CPU is disabled, and SI slave is enabled.
1000	NPI port is enabled for SERDES NoAneq 1G FDX NoFC. VRAP block is accessible through the NPI port. Automatic boot of VCore CPU is disabled, and SI slave is enabled.
1001	PCIe 1.x endpoint is enabled. Automatic boot of VCore CPU is disabled, and SI slave is enabled.
1010	MIIM slave is enabled (MIIM slave pins are overlaid on GPIOs.) Automatic boot of VCore CPU is disabled, and SI slave is enabled.
1100	The VCore CPU is enabled (Big Endian mode) and boots from SI (the SI slave is disabled).
1101	The VCore CPU is enabled (Little Endian mode) and boots from SI (the SI slave is disabled).
1111	Automatic boot of VCore CPU is disabled, and SI slave is enabled.

Undefined configurations are reserved and cannot be used. VCore configurations that enable a front port or a PCIe endpoint drive VCore_CFG[3:2] high when the front port or PCIe endpoint is ready to use.

The configuration value determines the reset value for the ICPU_CFG::GENERAL_CTRL register. After startup, the behavior of the VCore system can be modified by changing some of the fields in this register. Common scenarios are:

- After starting the VSC7416-01 device with the VCore CPU disabled, an external CPU can manually boot the

VCore CPU from SI Flash by writing `ICPU_CFG::GENERAL_CTRL.IF_SI_MST_ENA = 1`, `ICPU_CFG::GENERAL_CTRL.BOOT_MODE = 1`, and `ICPU_CFG::GENERAL_CTRL.CPU_DIS = 0`. Endianess is configured by `ICPU_CFG::GENERAL_CTRL.CPU_BE_ENA`. Setting `GENERAL_CTRL.IF_SI_MST_ENA` will disable the SI slave, so the external CPU must be using another interface than SI.

- After starting the VSC7416-01 device with the VCore CPU disabled and loading software into DDR2/3 memory, an external CPU can boot the VCore CPU from DDR2/3 memory by writing `ICPU_CFG::GENERAL_CTRL.BOOT_MODE = 0` and `ICPU_CFG::GENERAL_CTRL.CPU_DIS = 0`. Endianess is configured using `ICPU_CFG::GENERAL_CTRL.CPU_BE_ENA`.
- After automatically booting from SI Flash, the VCore CPU can release the SI interface and enable the SI slave by writing `ICPU_CFG::GENERAL_CTRL.IF_SI_MST_ENA = 0`. This enables SI access from an external CPU to the VSC7416-01 device.

Note A special PCB design is required to make SI work for both Flash and external CPU access.

- The MIIM slave can be manually enabled by writing `ICPU_CFG::GENERAL_CTRL.IF_MIIM_SLV_ENA = 1`. The MIIM slave automatically takes control of the appropriate GPIO pins.

The EJTAG interface of the VCore CPU and the Boundary Scan JTAG controller are both multiplexed onto the JTAG interface of VSC7416-01. When the `JTAG_ICE_nEn` pin is low, the MIPS's EJTAG controller is selected. When the `JTAG_ICE_nEn` pin is high, the Boundary Scan JTAG controller is selected.

3.2 Clocking and Reset

The following table lists the registers associated with reset and the watchdog timer.

TABLE 3-2: CLOCKING AND RESET CONFIGURATION REGISTERS

Register	Description
<code>ICPU_CFG::RESET</code>	VCore reset protection scheme and initiating soft reset of the VCore System and/or CPU.
<code>DEVCPU_GCB::SOFT_CHIP_RST</code>	Initiating chip-level soft reset.
<code>ICPU_CFG::WDT</code>	Watchdog timer configuration and status.

The VCore CPU runs 416.66 MHz, the DDR2/3 controller runs 312.50 MHz, and the rest of the VCore system runs at 208.33 MHz.

The VCore can be soft-reset by setting `RESET.CORE_RST_FORCE`. By default, this resets both the VCore CPU and the VCore system. The VCore system can be excluded from a soft reset by setting `RESET.CORE_RST_CPU_ONLY`; soft-reset using `CORE_RST_FORCE` only then resets the VCore CPU. The frame DMA must be disabled prior to a soft reset of the VCore system. When `CORE_RST_CPU_ONLY` is set, the frame DMA, PCIe endpoint, and DDR2/3 controller are not affected by a soft reset and will continue to operate throughout soft reset of the VCore CPU.

The VCore system comprises all the blocks attached to the VCore Shared Bus (SBA), including the PCIe, DDR2/3, and frame DMA/injection/extraction blocks. The VSC7416-01 device includes all the blocks attached to the Switch Core Register Bus (CSR) including VRAP, SI, and MIIM slaves. For more information about the VCore System blocks, see [Figure 3-1](#).

The VSC7416-01 device can be soft-reset by writing `SOFT_CHIP_RST.SOFT_CHIP_RST`. The VCore system and CPU can be protected from a device soft-reset by writing `RESET.CORE_RST_PROTECT = 1` before initiating soft-reset. In this case, a chip-level soft reset is applied to all other blocks except the VCore system and the CPU. When protecting the VCore system and CPU from a soft reset, the frame DMA must be disabled prior to a chip-level soft reset.

The VCore general purpose registers (`ICPU_CFG::GPR`) and GPIO alternate modes (`DEVCPU_GCB::GPIO_ALT`) are not affected by a soft-reset. These registers are only reset when an external reset is asserted.

3.2.1 WATCHDOG TIMER

The VCore system has a built-in watchdog timer (WDT) with a two-second timeout cycle. The watchdog timer is enabled, disabled, or reset through the WDT register. The watchdog timer is disabled by default.

After the watchdog timer is enabled, it must be regularly reset by software. Otherwise, it times out and causes a VCore soft reset equivalent to setting RESET.CORE_RST_FORCE. Improper use of the WDT.WDT_LOCK causes an immediate timeout-reset as if the watchdog timer had timed out. The WDT.WDT_STATUS field shows if the last VCore CPU reset was caused by WDT timeout or regular reset (possibly soft-reset). The WDT.WDT_STATUS field is updated only during VCore CPU reset.

To enable or to reset the watchdog timer, write the locking sequence, as described in WDT.WDT_LOCK, at the same time as setting the WDT.WDT_ENABLE field.

Because watchdog timeout is equivalent to setting RESET.CORE_RST_FORCE, the RESET.CORE_RST_CPU_ONLY field also applies to watchdog initiated soft reset.

3.3 Shared Bus

The shared bus is a 32-bit address and 32-bit data bus with dedicated master and slave interfaces that interconnect all the blocks in the VCore system. The VCore CPU, PCIe inbound, and VCore SBA access block are masters on the shared bus; only they can start access on the bus.

The shared bus uses byte addresses, and transfers of 8, 16, or 32 bits can be made. To increase performance, bursting of multiple 32-bit words on the shared bus can be performed.

All slaves are mapped into the VCore systems 32-bit address space and can be accessed directly by masters on the shared bus. Two possible mappings of VCore shared bus slaves are boot mode and normal mode.

- Boot mode. Boot mode is active after power-up and reset of the VCore system. In this mode, the SI Flash Boot Master is mirrored into the lowest address region.
- Normal mode. In normal mode, the DDR2/3 controller is mirrored into the lowest address region.

Changing between boot mode and normal mode is done by first writing and then reading ICPU_CFG::GENERAL_CTRL.BOOT_MODE_ENA. A change takes effect during the read.

The VSC7416-01 device supports up to 1 gigabyte (GB) of DDR2/3 memory. By default, only 512 megabytes (MB) is accessible in the memory map. By writing ICPU_CFG::MEMCTRL_CFG.DDR512MBYTE_PLUS = 1, the memory map is changed to accommodate more than 512 MB of memory.

FIGURE 3-2: SHARED BUS MEMORY MAP

Boot Mode (Physical), 512 MB		Normal Mode (Physical), 512 MB	
0x00000000	256 MB Mirror of SI Flash	0x00000000	512 MB Mirror of DDR2/3
0x10000000	256 MB Reserved	0x20000000	512 MB DDR2/3
0x20000000	512 MB DDR2/3	0x40000000	256 MB SI Flash
0x40000000	256 MB SI Flash	0x50000000	512 MB Reserved
0x50000000	512 MB Reserved	0x70000000	256 MB Chip Registers
0x70000000	256 MB Chip Registers	0x80000000	1 GB Reserved
0x80000000	1 GB Reserved	0xC0000000	1 GB PCIe DMA
0xC0000000	1 GB PCIe DMA	0xFFFFFFF	0xFFFFFFF

Boot Mode (Physical), 1 GB ⁽¹⁾		Normal Mode (Physical), 1 GB ⁽¹⁾	
0x00000000	256 MB Mirror of SI Flash	0x00000000	1 GB Mirror of DDR2/3
0x10000000	768 MB Reserved	0x40000000	256 MB SI Flash
0x40000000	256 MB SI Flash	0x50000000	512 MB Reserved
0x50000000	512 MB Reserved	0x70000000	256 MB Chip Registers
0x70000000	256 MB Chip Registers	0x80000000	1 GB DDR2/3
0x80000000	1 GB DDR2/3	0xC0000000	1 GB PCIe DMA
0xC0000000	1 GB PCIe DMA	0xFFFFFFF	0xFFFFFFF

1. To enable support 1 gigabyte (GB) of DDR2/3 memory (and the associated memory map), set ICPU_CFG::MEMCTRL_CFG.DDR_512MBYTE_PLUS.

Note When the VCore system is protected from a soft reset using ICPU_CFG::RESET.CORE_RST_CPU_ONLY, a soft reset will not change shared bus memory mapping. For more information about protecting the VCore system when using a soft reset, see [Section 3.2, Clocking and Reset](#).

If the boot process copies the SI Flash image to DDR2/3, and if the contents of the SI memory and the DDR2/3 memory are the same, software can execute from the mirrored region when swapping from boot mode to normal mode. Otherwise, software must be executing from the fixed SI Flash region when changing from boot mode to normal mode.

The Frame DMA has dedicated access to PCIe Outbound and DDR2/3. This means that access on the SBA to other parts of the VSC7416-01 device, such as register access, does not affect Frame DMA injection/extraction performance.

3.3.1 VCore SHARED BUS ARBITRATION

The following table lists the registers associated with the shared bus arbitration.

TABLE 3-3: SHARED BUS CONFIGURATION REGISTERS

Register	Description
SBA::PL_CPU	Master priorities
SBA::PL_PCIE	Master priorities
SBA::PL_CSR	Master priorities
SBA::WT_EN	Enable of weighted token scheme
SBA::WT_TCL	Weighted token refresh period
SBA::WT_CPU	Token weights for weighted token scheme
SBA::WT_PCIE	Token weights for weighted token scheme
SBA::WT_CSR	Token weights for weighted token scheme

The VCore shared bus arbitrates between masters that want to access the bus. Default is to use a strict prioritized arbitration scheme where the VCore CPU has highest priority. The strict priorities can be changed using registers PL_CPU, PL_PCIE, and PL_CSR.

- *_CPU registers apply to VCore CPU access.
- *_PCIE registers apply to inbound PCIe access.
- *_CSR registers apply to VCore SBA access block access.

It is possible to enable weighted token arbitration scheme (WT_EN). When using this scheme, specific masters can be guaranteed a certain amount of bandwidth on the shared bus. Guaranteed bandwidth that is not used is given to other masters requesting the shared bus.


When weighted token arbitration is enabled, the masters on the shared bus are granted a configurable number of tokens (WT_CPU, WT_PCIE, and WT_CSR) at the start of each refresh period. The length of each refresh period is configurable (WT_TCL). For each clock-cycle that the master uses the shared bus, the token counter for that master is decremented. When all tokens are spent, the master is forced to a low priority. Masters with tokens always take priority over masters with no tokens. The strict prioritized scheme is used to arbitrate between masters with tokens and between masters without tokens.

Example Guarantee That PCIe Can Get 25% Bandwidth. Configure WT_TCL to a refresh period of 2048 clock cycles; the optimal length of the refresh period depends on the scenario, experiment to find the right setting. Guarantee PCIe access in 25% of the refresh period by setting WT_PCIE to 512 ($2048 \times 25\%$). Set WT_CPU and WT_CSR to 0. This gives the VCore CPU and CSR unlimited tokens. Configure PCIe to highest priority by setting PL_PCIE to 15. Finally, enable the weighted token scheme by setting WT_EN to 1. For each refresh period of 2048 clock cycles, PCIe is guaranteed access to the shared bus for 512 clock cycles, because it is the highest priority master. When all the tokens are spent, it is put into the low-priority category.

3.3.2 CHIP REGISTER REGION

Registers in the VCore domain and inside the Switch Core are memory mapped into the Chip registers region of the shared bus memory map. All registers are 32-bit wide and must only be accessed using 32-bit reads and writes. Bursts are supported.

Writes to this region are buffered (there is a one-word write buffer). Multiple back-to-back write access pauses the shared bus until the write buffer is freed up (until the previous writes are done). Reads from this region pause the shared bus until read data is available.

FIGURE 3-3: CHIP REGISTERS DETAILED MEMORY MAP

The registers in the 0x70FFFFFF to 0x71FFFFFF region are physically located inside the VCore System, so read and write access to these registers is fast (done in a few clock cycles). All registers in this region are considered fast registers.

Registers in the 0x71000000 to 0x72000000 region are physically located in other areas of the VSC7416-01 device rather than the VCore system; access to registers in this range takes approximately 0.5 μ s. The DEVCPU_ORG and DEVCPU_QS targets are special; registers inside these two targets are faster; access to these two targets takes approximately 100 ns.

When more than one CPU is accessing registers, the access time may be increased. For more information, see [Section 3.5.1, Register Access and Multimaster Systems](#).

Writes to the Chip Registers region is buffered (there is a one-word write buffer). Multiple back-to-back write access pauses the shared bus until the write buffer is freed up (until the previous write is done). A read access pause the shared bus until read data is available. Doing a write immediately followed by a read requires the write to be done before the read can be started.

3.3.3 SI FLASH REGION

Read access from the SI Flash region initiates Flash formatted read access on the SI pins of the VSC7416-01 device by means of the SI controller. For more information about the SI controller and protocol, see [Section 3.8.1, SI Controller](#). It is not possible to write to the SI Flash region.

3.3.4 DDR2/3 REGION

Read and write access from or to the DDR2/3 region maps to access on the DDR2/3 interface of the VSC7416-01 device by means of the DDR2/3 controller. For more information about the DDR2/3 controller and initialization of DDR2/3 memory, see [Section 3.8.2, DDR2/3 Memory Controller](#).

3.3.5 PCIE REGION

Read and write access from or to the PCIe region maps to outbound read and write access on the PCIe interface of the VSC7416-01 device by means of the PCIe endpoint controller. For more information about the PCIe endpoint controller and how to reach addresses in 32-bit and 64-bit PCIe environments, see [Section 3.6, PCIe Endpoint Controller](#).

3.4 VCore CPU

The VCore CPU system is based on a powerful MIPS24KEc-compatible microprocessor with 16-entry MMU, 32 kilobytes (kB) instruction, and 32 kilobytes data caches.

This section describes how the VCore CPU is integrated into the VCore system. For more information about internal VCore CPU functions, such as bringing up caches, MMU, and so on, see the software board support package (BSP) at <https://www.microchip.com/>.

When the automatic boot in the little-mode or big-endian mode is enabled using the VCore strapping pins, the VCore CPU automatically starts to execute code from the SI Flash at byte address 0.

The following is a typical automatic boot sequence:

1. Speed up the boot interface. For more information, see [Section 3.8.1, SI Controller](#).
2. Initialize the DDR2/3 controller. For more information, see [Section 3.8.2, DDR2/3 Memory Controller](#)
3. Copy code-image from Flash to DDR2/3 memory.
4. Change memory map from boot mode to normal mode For more information, see [Section 3.3, Shared Bus](#).

When automatic boot is disabled, an external CPU can still start the VCore CPU through registers. For more information, see [Section 3.1, VCore Configurations](#).

The boot vector of the VCore CPU is mapped to the start of the KESEG1, which translates to physical address 0x00000000 on the VCore shared bus.


The VCore CPU interrupts are mapped to interrupt inputs 0 and 1, respectively.

3.4.1 LITTLE AND BIG ENDIAN SUPPORT

The VCore system is constructed as a little endian system, and registers descriptions reflect little endian encoding. When big endian mode is enabled, instructions and data are byte-lane swapped just before they enter and when they leave the VCore CPU. This is the standard way of translating between a CPU in big endian mode and a little endian system.


The following illustration shows how the 16-bit value 0x3210 is transferred between the VCore CPU and the VCore shared bus in little endian and big endian modes.

FIGURE 3-4: 16-BIT ACCESS IN LITTLE ENDIAN AND BIG ENDIAN MODES



For 32-bit access, the difference is less obvious. The following illustration shows how the value 0x76543210 is transferred between the VCore CPU and the VCore shared bus in little endian and big endian modes.

FIGURE 3-5: 32-BIT ACCESS IN LITTLE ENDIAN AND BIG ENDIAN MODE



Note The swapping of byte lanes ensures that no matter the endian mode, the VCore CPU is always accessing the appropriate part of the little endian memory system.

In big-endian mode, care must be taken when accessing parts of the memory system that is also used by users other than the VCore CPU. For example, VSC7416-01 device registers are written and read by the VCore CPU, but they are also used by the device (which sees them in little endian mode). The VCore BSP contains examples of code that correctly handles register access for both little and big endian mode.

3.4.2 SOFTWARE DEBUG AND DEVELOPMENT

The VCore CPU has a standard MIPS EJTAG debug interface that can be used for breakpoints, loading of code, and examining memory. When the JTAG_ICE_nEn strapping pin is pulled low, the JTAG interface is attached to the EJTAG controller.

3.5 External CPU Support

An external CPU attaches to the VSC7416-01 device through PCIe, SI, MIIM, or VRAP. Through these interfaces, an external CPU can access (and control) the device. For more information about interfaces and connections to device registers, see [Figure 3-1](#).

Inbound PCIe access is performed on the VCore Shared Bus (SBA) in the same way as ordinary VCore CPU access. By means of the Switch Core Access block it is possible to access the Switch Core Register (CSR) bus. For more information about supported PCIe BAR-regions, see [Section 3.6, PCIe Endpoint Controller](#).

The SI, MIIM, and VRAP interfaces attach directly to the CSR. Through the VCore SBA Access block, it is possible to access the VCore Shared bus (SBA). For more information, see [Section 3.5.4, Access to the VCore Shared Bus](#).

The external CPU can coexist with the internal VCore CPU and hardware-semaphores and interrupts are implemented for inter-CPU communication. For more information, see [Section 3.5.5, Mailbox and Semaphores](#).

3.5.1 REGISTER ACCESS AND MULTIMASTER SYSTEMS

There are three different groups of registers in the VSC7416-01 device:

- Switch Core
- Fast Switch Core
- VCore

The Switch Core registers and Fast Switch Core registers are separated into individual register-targets and attached to the Switch Core Register bus (CSR). The Fast Switch Core registers are placed in the DEVCPU_QS and DEVCPU_ORG register-targets. Access to Fast Switch Core registers is less than 0.1 μ s; other Switch Core registers take no more than 0.5 μ s to access.

The VCore registers are attached directly to the VCore Shared Bus. The access time to VCore registers is negligible (a few clock cycles).

Although multiple masters can access VCore registers and Switch Core registers in parallel without noticeable penalty to the access time, the following exceptions apply.

- When accessing the same Switch Core register target (for example, DEVCPU_GCB), the second master to attempt access has to wait for the first master to finish (round robin arbitration applies.) This does not apply to Fast Switch Core register targets (DEVCPU_QS and DEVCPU_ORG).
- If both the VCore CPU and the PCIe master are performing Switch Core Register bus (CSR) access, both need to be routed through the Switch Core Access block. The second master has to wait for the first master to finish, because shared bus arbitration applies.
- If two or more SI, MIIM, or VRAP masters are performing VCore register access they all need to go through the VCore SBA Access block. Ownership has to be resolved by use of software (for example, by using the build-in semaphores).

The most common multimaster scenario is with an active VCore CPU and an external CPU using either SI or VRAP. In this case, Switch Core register access to targets that are used by both CPUs may see two times the access time (no more than 1 μ s).

3.5.2 SERIAL INTERFACE IN SLAVE MODE

This section provides information about the function of the serial interface (SI) in slave mode.

The following table lists the registers associated with SI slave mode.

TABLE 3-4: SI SLAVE MODE REGISTER

Register	Description
DEVCPU_ORG::IF_CTRL	Configuration of endianness and bit order.
DEVCPU_ORG::IF_CFGSTAT	Configuration of padding.

The serial interface implements a SPI-compatible protocol that allows an external CPU to perform read and write access to register targets in the VSC7416-01 device. Endianness and bit order is configurable, and several options for high frequencies are supported.

The serial interface is available to an external CPU when the VCore CPU does not use the SI for Flash access. For more information, [Section 3.0, VCore System and CPU Interfaces](#).

The following table lists the pins of the Serial interface.

TABLE 3-5: SI SLAVE MODE PINS

Pin Name	I/O	Description
SI_nEn	I	Active low chip select
SI_Clk	I	Clock input
SI_DI	I	Data input (MOSI)
SI_DO	O	Data output (MISO)

SI_DI is sampled on rising edge of SI_Clk. SI_DO is driven on falling edge of SI_Clk. There are no requirements on the logical values of the SI_Clk and SI_DI inputs when SI_nEn is asserted or deasserted, they can be either 0 or 1. SI_DO is only driven during read-access when read-data is shifted out of the VSC7416-01 device.

The external CPU initiates access by asserting chip select and then transmitting one bit read/write indication, one don't care bit, 22 address bits, and 32 bits of write-data (or don't care bits when reading).


With the register address of a specific register (REG_ADDR), the SI address (SI_ADDR) is calculated as:

$$\text{SI\_ADDR} = (\text{REG\_ADDR} \& 0x00FFFFFF) \gg 2$$

Data word endianness is configured through IF_CTRL[0]. The order of the data bits is configured using IF_CTRL[1].

The following illustration shows various configurations for write access. The order of the data bits during writing, as depicted, is also used when VSC7416-01 is transmitting data during read operations.

FIGURE 3-6: WRITE SEQUENCE FOR SI



When reading registers using the SI interface, the device needs to prepare read data after receiving the last address bit. The access time of the register that is read must be satisfied before shifting out the first bit of read data. For information about access time, see [Section 3.5.1, Register Access and Multimaster Systems](#). The external CPU must apply one of the following solutions to satisfy read access time.

- Use **SI_Clk** with a period of minimum twice the access time for the register target. For example, for normal switch core targets (single Master):
 $1/(2 \times 0.5 \mu\text{s}) = 1 \text{ MHz (maximum)}$.
- Pause the **SI_Clk** between shifting of serial address bit 0 and the first data bit with enough time to satisfy the access time for the register target.
- Configure the VSC7416-01 to send out padding bytes before transmitting the read data to satisfy the access time for the register target. For example, 1 dummy byte allows enough read-time for the SI clock to run up to 13 MHz in a single master system, using the following calculation.

The device is configured for inserting padding bytes by writing to **IF_CFGSTAT.IF_CFG**, these bytes are transmitted before the read data. The maximum frequency of the SI clock is calculated as:

$(\text{IF\_CFGSTAT.IF\_CFG} \times 8 - 1.5) / \text{access-time}$. For example, for normal switch core targets (single Master), 1 byte padding give $(1 \times 8 - 1.5) / 0.5 \mu\text{s} = 13 \text{ MHz (maximum)}$. The **SI_DO** output is kept tristated until the actual read data is transmitted.

The following illustrations show options for serial read access. The illustrations show only one mapping of read data, little endian with most significant bit first. Any of the mappings can be configured and apply to read data in the same way as for write data.

FIGURE 3-7: READ SEQUENCE FOR SI_CLK SLOW






FIGURE 3-8: READ SEQUENCE FOR SI_CLK PAUSE**FIGURE 3-9: READ SEQUENCE FOR ONE-BYTE PADDING**

When dummy bytes are enabled (IF_CFGSTAT.IF_CFG), the SI slave logic enables an error check that sends out 0x88888888 and sets IF_CFGSTAT.IF_STAT if the SI master does not provide enough time for register read.

When using SI, the external CPU must configure the IF_CTRL register after power-up, reset, or chip-level soft reset. The IF_CTRL register is constructed so that it can be written no matter the state of the interface. For more information about constructing write-data for this register, see the instructions in IF_CTRL.IF_CTRL.

3.5.3 MIIM INTERFACE IN SLAVE MODE

This section provides the functional aspects of the MIIM slave interface.

The MIIM slave interface allows an external CPU to perform read and write access to the register targets in VSC7416-01 device. Register access is done indirectly, because the address and data fields of the MIIM protocol is less than those used by the register targets. Transfers on the MIIM interface are using the Management Frame Format protocol specified in IEEE 802.3, Clause 22.

The MIIM slave pins on the VSC7416-01 device are overlaid functions on the GPIO interface. MIIM slave mode is enabled by configuring the appropriate VCore_CFG strapping pins. For more information, see [Section 3.1, VCore Configurations](#). When MIIM slave mode is enabled, the appropriate GPIO pins are automatically overtaken. For more information about overlaid functions on the GPIO pins, see [Table 913](#).

The following table lists the pins of the MIIM slave interface.

TABLE 3-6: MIIM SLAVE PINS

Pin Name	I/O	Description
MIIM_SLV_MDC/GPIO_9	I	MIIM slave clock input
MIIM_SLV_MDIO/GPIO_10	I/O	MIIM slave data input/output
MIIM_SLV_ADDR/GPIO_8	I	MIIM slave address select

MIIM_SLV_MDIO is sampled or changed on the rising edge of MIIM_SLV_MDC by the MIIM slave interface.

The MIIM slave can be configured to answer on one of two different PHY addresses using the MIIM_SLV_ADDR pin. Setting the MIIM_SLV_ADDR pin to 0 configures the MIIM slave to use PHY address 0, and setting it to 1 configures the MIIM slave to use PHY address 31.

The MIIM slave has seven 16-bit MIIM registers defined as listed in the following table.

TABLE 3-7: MIIM REGISTERS

Register Address	Register Name	Description
0	ADDR_REG0	Bits 15:0 of the address to read or write. The address field must be formatted as word address.
1	ADDR_REG1	Bits 31:16 of the address to read or write.
2	DATA_REG0	Bits 15:0 of the data to read or write. Returns 0x8888 if a register read error occurred.
3	DATA_REG1	Bits 31:16 of the data to read or write. The read or write operation is initiated after this register is read or written. Returns 0x8888 if read while busy or a register read error occurred.
4	DATA_REG1_INCR	Bits 31:16 of data to read or write. The read or write operation is initiated after this register is read or written. When the operation is complete, the address register is incremented by one. Returns 0x8888 if read while busy or a register read error occurred.
5	DATA_REG1_INERT	Bits 31:16 of data to read or write. Reading or writing to this register does not cause a register access to be initiated. Returns 0x8888 if a register read error occurred.
6	STAT_REG	The status register gives the status of any ongoing operations. Bit 0: Busy. Set while a register read/write operation is in progress. Bit 1: Busy_rd. Busy status during the last read or write operation. Bit 2: Err. Set if a register access error occurred. Others: Reserved.

A 32-bit switch core register read or write transaction over the MIIM interface is done indirectly due to the limited data width of the MIIM frame. First, the address of the register inside the device must be set in the two 16-bit address registers of the MIIM slave using two MIIM write transactions. Afterwards the two 16-bit data registers can be read or written to access the data value of the register inside the device. Thus, it requires up to four MIIM transactions to perform a single read or write operation on a register target.


The address of the register to read/write is set in registers ADDR_REG0 and ADDR_REG1. The data to write to the register pointed to by the address in ADDR_REG0 and addr_reg1 is first written to DATA_REG0 and then to DATA_REG1. When the write transaction to DATA_REG1 is completed, the MIIM slave initiates the switch core register write.

With the register address of a specific register (REG_ADDR), the MIIM address (MIIM_ADDR) is calculated as:

$$\text{MIIM\_ADDR} = (\text{REG\_ADDR} \& 0x00FFFFFF) \gg 2$$

The following illustration shows a single MIIM write transaction on the MIIM interface.


FIGURE 3-10: MIIM SLAVE WRITE SEQUENCE



A read transaction is done in a similar way. First, read the DATA_REG0 and then read the DATA_REG1. As with a write operation. The switch core register read is not initiated before the DATA_REG1 register is read. In other words, the returned read value is from the previous read transaction.

The following illustration shows a single MIIM read transaction on the MIIM interface.

FIGURE 3-11: MIIM SLAVE READ SEQUENCE



3.5.4 ACCESS TO THE VCore SHARED BUS

This section provides information about how to access the VCore shared bus (SBA) from an external CPU attached by means of the VRAP, SI, or MIIM. The following table lists the registers associated with the VCore shared bus access.

TABLE 3-8: VCore SHARED BUS ACCESS REGISTERS

Register	Description
DEVCPU_GCB::VA_CTRL	Status for ongoing accesses
DEVCPU_GCB::VA_ADDR	Configuration of shared bus address
DEVCPU_GCB::VA_DATA	Configuration of shared bus address
DEVCPU_GCB::VA_DATA_INCR	Data register, access increments VA_ADDR
DEVCPU_GCB::VA_DATA_IN- ERT	Data register, access does not start new access

An external CPU perform 32-bit reads and writes to the SBA through the VCore Access (VA) registers. In the VCore system, a dedicated master on the shared bus handles VA access. For information about arbitration between masters on the shared bus, see [Section 3.3.1, VCore Shared Bus Arbitration](#).

The SBA address is configured in VA_ADDR. Writing to VA_DATA starts an SBA write with the 32-bit value that was written to VA_DATA. Reading from VA_DATA returns the current value of the register and starts an SBA read access, when the read-access completes, the result is automatically stored in the VA_DATA register.

The VA_DATA_INCR register behaves similar to VA_DATA, except that after starting an access, the VA_ADDR register is incremented by four (so that it points to the next word address in the SBA domain). Reading from the VA_DATA_INCR register returns the value of VA_DATA, writing to VA_DATA_INCR overwrites the value of VA_DATA.

Note By using VA_DATA_INCR, sequential word addresses can be accessed without having to manually increment the VA_ADDR register between each access.

The VA_DATA_INERT register provides direct access to the VA_DATA value without starting access on the SBA. Reading from the VA_DATA_INERT register returns the value of VA_DATA, writing to VA_DATA_INERT overwrites the value of VA_DATA.

The VCore shared bus is capable of returning error-indication when illegal register regions are accessed. If a VA access results in an error-indication from the SBA, the VA_CTRL.VA_ERR field is set, and the VA_DATA is set to 0x88888888.

Note SBA error indications only occur when non-existing memory regions or illegal registers are accessed. This should not occur during normal operation, so the VA_CTRL.VA_ERR indication is useful during debugging only.

Example Reading from ICPU_CFG::GRP[1] through the VA registers. The ICPU_GPR register is the second register in the SBA VCore Registers region. Set VA_ADDR to 0x70000004, read once from VA_DATA (and discard the read-value). Wait until VA_CTRL.VA_BUSY is cleared, then VA_DATA contains the value of the ICPU_CFG::GRP[1] register. Using VA_DATA_INERT (instead of VA_DATA) to read the data is appropriate because this does not start a new SBA access.

3.5.4.1 Optimized Reading

SBA access is typically much faster than the CPU interface, which is used to access the VA registers. The VA_DATA register (VA_DATA_INCR and VA_DATA_INERT) return 0x88888888 while VA_CTRL.VA_BUSY is set. This means that it is possible to skip checking for busy between read access to SBA.

For example, after initiating a read access from SBA, software can proceed directly to reading from VA_DATA, VA_DATA_INCR, or VA_DATA_INERT

- If the second read is different from 0x88888888; then the second read returned valid read data (the SBA access was done before the second read was performed).
- If the second read is equal to 0x88888888; VA_CTRL must be read.
If VA_CTRL.VA_BUSY_RD is cleared (and VA_CTRL.VA_ERR_RD is also cleared), then 0x88888888 is the actual read data.
If VA_CTRL.VA_BUSY_RD is set, the SBA access was not yet done at the time of the second read. Start over again by repeating the read from VA_DATA.

Optimized reading can be used for single-read and sequential-read access. For sequential reads the VA_ADDR is only incremented on successful (non-busy) reads.

3.5.5 MAILBOX AND SEMAPHORES

This section provides information about the semaphores and mailbox features for CPU-to-CPU communication. The following table lists the registers associated with mailbox and semaphores.

TABLE 3-9: MAILBOX AND SEMAPHORE REGISTERS

Register	Description
DEVCPU_ORG::MAILBOX_SET	Atomic set of bits in the mailbox register.
DEVCPU_ORG::MAILBOX_CLR	Atomic clear of bits in the mailbox register.
DEVCPU_ORG::MAILBOX	Current mailbox state.
DEVCPU_ORG::SEMA_CFG	Configuration of semaphore interrupts.
DEVCPU_ORG::SEMA0	Taking of semaphore 0.
DEVCPU_ORG::SEMA1	Taking of semaphore 1.
DEVCPU_ORG::SEMA0_OWNER	Current owner of semaphore 0.
DEVCPU_ORG::SEMA1_OWNER	Current owner of semaphore 1.

The mailbox is a 32-bit register that can be set and cleared atomically using any CPU interface (including the VCore CPU). The MAILBOX register allows reading (and writing) of the current mailbox value. Atomic clear of individual bits in the mailbox register is done by writing a mask to MAILBOX_CLR. Atomic setting of individual bits in the mailbox register is done by writing a mask to MAILBOX_SET.

The VSC7416-01 device implements two independent semaphores. The semaphores are part of the Switch Core Register Bus (CSR) block and are accessible by means of the fast switch core registers. Semaphore ownership can be taken by interfaces attached to the Switch Core Register Bus (CSR). That is, the VCore, VRAP, SI, and MIIM can be granted ownership. The VCore CPU and PCIe interface share the same semaphore, because they both access the Switch Core Register Bus (CSR) through the Switch Core Access block.

Any CPU attached to an interface can attempt to take a semaphore by reading SEMA0.SEMA0 or SEMA1.SEMA1. After reading one of the SEMA0 or SEMA1 registers, the status for the corresponding semaphore is available in SEMA0_OWNER.SEMA0_OWNER or SEMA1_OWNER.SEMA1_OWNER. For more information about encoding of owners, see the register description.

Note Any interface can release semaphores; it does not have to be the one that has taken the semaphore, this allows implementation of handshaking protocols.

Software interrupt is generated when semaphores are free or taken. Interrupt polarity is configured through SEMA_CFG.SEMA_INTR_POL. Semaphore 0 is hooked up to SW0 interrupt and semaphore 1 is hooked up to SW1 interrupt. For configuration of software-interrupt, see [Section 3.8.11, Interrupt Controller](#).

In addition to interrupting on semaphore state, software interrupt can be manually triggered by writing directly to the ICPU_CFG::INTR_FORCE register.

Software interrupts (SW0 and SW1) can be individually mapped to either the VCore CPU or to external interrupt outputs (to an external CPU).

3.6 PCIe Endpoint Controller

The VSC7416-01 device implements a single-lane PCIe 1.x endpoint that can be hooked up to any PCIe capable system. Via the PCIe interface an external CPU can access (and control) the device. Ethernet frames can be injected and extracted via registers or the device can be configured to DMA Ethernet frames autonomously to/from PCIe memory. The device's DDR2/3 memory is available via the PCIe interface giving external CPU full read and write access to DDR2/3 modules attached to the device. The DDR2/3 controller and memory modules can be initialized either via PCIe or by the VCore CPU.

Both the VCore CPU and Frame DMA can generate PCIe read/write requests to any 32-bit or 64-bit memory region in PCIe memory space. However, it is up to software running on an external CPU to setup or communicate the appropriate PCIe memory mapping information to the VSC7416-01 device.

The defaults for the endpoints capabilities region and the extended capabilities region are listed in the registers list's description of PCIE::CONF_SPACE. The most important parameters are:

- Vendor (and subsystem Vendor) ID: 0x101B, Microchip.
- Device ID: 0xB002, Device Family ID.
- Revision ID: 0x00, Device Family Revision ID.
- Class Code: 0x028000, Ethernet Network Controller.
- Single Function, non-Bridge, INTA and MSI capable Device.

The Device family 0xB002 covers register compatible devices VSC7414-01, VSC7416-01, and VSC7418-01. The software driver must determine actual device ID and revision by reading DEVCPU_GCB::CHIP_ID from the device's memory mapped registers region.

For information about base address registers, see [Section 3.6.3, Base Address Registers, Inbound Requests](#).

The IDs, Class Code, Revision ID, and base address register setups can be customized before enabling the PCIe endpoint. However, it requires a manual bring-up procedure by software running locally on the VCore CPU. For more information, see [Section 3.6.2, Enabling the Endpoint](#).

The endpoint is Power Management capable and implements PCI Bus Power Management Interface Specification Revision 1.2. For more information, see [Section 3.6.6, Power Management](#).

The endpoint is MSI Interrupt capable. Up to four 64-bit messages are supported. Messages can be generated on rising and falling edges of each of the two external VCore interrupt destinations. For more information, see [Section 3.6.4, Outbound Interrupts](#).

For information about all PCI Express capabilities and extended capabilities register defaults, see the PCIE::CONF_SPACE register description.

3.6.1 ACCESSING ENDPOINT REGISTERS

The root complex accesses the PCIe endpoint's configuration registers by PCIe CfgRd/CfgWr requests. The VCore CPU can read configuration registers by means of the PCIE region. For more information, see [Section 3.3.2, Chip Register Region](#). The PCIE region must not be accessed when the PCIe endpoint is disabled.

The PCIe region is used during manual bring-up of PCIe endpoint. By using this region, it is possible to write most of the endpoint's read-only configuration registers, such as Vendor ID. The PCIe endpoint's read-only configuration register values must not be changed after the endpoint is enabled.

The VCore has a few dedicated PCIe registers in the ICPU_CFG:PCIE register group. An external CPU attached through the PCIe interface has to go through BAR0 to reach these registers.

3.6.2 ENABLING THE ENDPOINT

The PCIe endpoint is disabled by default. It can be enabled by one of two ways: either by automatically setting VCore_CFG strapping pins or by manually running software in the VCore CPU.

The recommended approach is using VCore_CFG strapping, because it is fast and does not require special software running on the VCore CPU. The endpoint is ready approximately 50 ms after release of the device's nReset. Until this point, the device ignores any attempts to do link training, and the PCIe output remains idle (tristated).

By using software running on the VCore CPU, it is possible to manually start the PCIe endpoint. Note that PCIe standard specifies a maximum delay from nReset release to working PCIe endpoint so software must enable the endpoint as part of the boot process. For more information, see [Section 3.6.2.1, Manually Starting MAC/PCS and SerDes](#).

The root complex must follow standard PCIe procedures for bringing up the endpoint.

3.6.2.1 Manually Starting MAC/PCS and SerDes

This section provides information about how to manually start up the PCIe endpoint and customize selected configuration space parameters.

The following table lists the registers related to manually bringing up PCIe.

TABLE 3-10: MANUAL PCIe BRING-UP REGISTERS

Register	Description
ICPU_CFG::PCIE_CFG	Disable of automatic link initialization.
HSIO::SERDES6G_COMMON_CFG	SERDES Configuration.
HSIO::SERDES6G_MISC_CFG	SERDES Configuration.
HSIO::SERDES6G_IB_CFG	SERDES Configuration.
HSIO::SERDES6G_IB_CFG1	SERDES Configuration.
HSIO::SERDES6G_OB_CFG	SERDES Configuration.
HSIO::SERDES6G_OB_CFG1	SERDES Configuration.
HSIO::SERDES6G_DES_CFG	SERDES Configuration.
HSIO::SERDES6G_PLL_CFG	SERDES Configuration.
HSIO::SERDES6G_MISC_CFG	SERDES Configuration.
HSIO::MCB_SERDES6G_ADDR_CFG	SERDES Configuration.
PCIE::CONF_SPACE	Customization of device parameters and disable of base address registers.

Disable automatic link training for PCIe endpoint.

1. PCIE_CFG.LTSSM_DIS = 1.

Configure and enable PCIe SERDES for 2G5 mode.

2. SERDES6G_COMMON_CFG = 0x00040090.
3. SERDES6G_MISC_CFG = 0x00000001.
4. SERDES6G_IB_CFG = 0x1003FC00.
5. SERDES6G_IB_CFG1 = 0x00000F1D.
6. SERDES6G_OB_CFG = 0x20000000.
7. SERDES6G_OB_CFG1 = 0x00000030.
8. SERDES6G_DES_CFG = 0x0000C8AA.
9. SERDES6G_PLL_CFG = 0x00003C0A.
10. MCB_SERDES6G_ADDR_CFG = 0x80000004 and wait until bit 31 is cleared.
11. SERDES6G_COMMON_CFG = 0x80040090.
12. MCB_SERDES6G_ADDR_CFG = 0x80000004 and wait until bit 31 is cleared.
13. Wait at least 200 μ s.

14. SERDES6G_MISC_CFG = 0x00000000.
15. SERDES6G_IB_CFG1 = 0x00000F1C.
16. MCB_SERDES6G_ADDR_CFG = 0x80000004 and wait until bit 31 is cleared.

Optionally disable BAR1:

1. PCIE_CFG.PCIE_BAR_WR_ENA = 1
2. CONF_SPACE[5] = 0x00000000
3. PCIE_CFG.PCIE_BAR_WR_ENA = 0

Optionally reduce BAR2 from 2 gigabytes (64-bit) to 128 megabytes (32-bit):

1. PCIE_CFG.PCIE_BAR_WR_ENA = 1
2. CONF_SPACE[6] = 0x07FFFFFF
3. CONF_SPACE[7] = 0x00000000
4. PCIE_CFG.PCIE_BAR_WR_ENA = 0

Optionally change selected PCIe configuration space values:

- Write Vendor ID/Device ID using CONF_SPACE[0].
- Write Class Code/Revision ID using CONF_SPACE[2].
- Write Subsystem ID/Subsystem Vendor ID using CONF_SPACE[11].

Enable automatic link training for PCIe endpoint

5. PCIE_CFG.LTSSM_DIS = 0

The last step enables the endpoint for link training and the root complex will then be able to initialize the PCIe endpoint. After this the CONF_SPACE parameters must not be changed anymore.

3.6.3 BASE ADDRESS REGISTERS, INBOUND REQUESTS

The VSC7416-01 device implements three memory regions. Read and write operations using the PCIe are translated directly to read and write access on the SBA. When manually bringing up the PCIe endpoint, BAR1 can be disabled. For more information, see [Section 3.6.2.1, Manually Starting MAC/PCS and SerDes](#).

TABLE 3-11: BASE ADDRESS REGISTERS

Register	Description
BAR0, 32bit, 32 megabytes	Chip Registers Region. This region maps to the Chip Registers region in the SBA address space. See Section 3.3.2, Chip Register Region . This region only supports 32-bit word-aligned reads and writes. Single and burst accesses are supported.
BAR1, 32bit, 16 megabytes	SI Flash Region. This region maps to the SI Flash region in the SBA address space. See Section 3.3.3, SI Flash Region . This region only supports 32-bit word-aligned reads. Single and burst accesses are supported.
BAR2/3, 64bit, 1 gigabyte	DDR2/3 Region. This region maps to the DDR2/3 region in the SBA address space. See Section 3.3.4, DDR2/3 Region . This region support all access types.

To access the BAR1 region, a Flash must be attached to the SI interface of the VSC7416-01 device. For information about how to set up I/O timing and to program the Flash through BAR0; the Chip Registers Region, see [Section 3.8.1, SI Controller](#).

To access the BAR2/3 region, DDR2/3 modules must be present and initialized. For information about initializing the DDR2/3 controller and modules through BAR0; the Chip Registers, see [Section 3.8.2, DDR2/3 Memory Controller](#). During manual bring up of PCIe, the size of the BAR2/3 region can be reduced to match the size of available DDR2/3 memory. For more information, see [Section 3.6.2.1, Manually Starting MAC/PCS and SerDes](#).

3.6.4 OUTBOUND INTERRUPTS

The VSC7416-01 device supports both Message Signaled Interrupt (MSI) and Legacy PCI Interrupt Delivery. The root complex configures the desired mode using the MSI Enable bit in the PCIe MSI Capability Register Set. For information about the VCore interrupt controller, see [Section 3.8.11, Interrupt Controller](#).

The following table lists the device registers associated PCIe outbound interrupts.

TABLE 3-12: PCIE OUTBOUND INTERRUPT REGISTERS

Register	Description
ICPU_CFG::PCIE_INTR_COMMON_CFG	Interrupt mode and enable.
ICPU_CFG::PCIE_INTR_CFG	MSI parameters.

In legacy mode, one interrupt is supported; select either EXT_DST0 or EXT_DST1 using PCIE_INTR_COMMON_CFG.LEGASY_MODE_INTR_SEL. The PCIe endpoint uses Assert_INTA and Deassert_INTA messages when configured for legacy mode.

In MSI mode, both EXT_DST interrupts can be used, EXT_DST0 is configured through PCIE_INTR_CFG[0] and EXT_DST1 through PCIE_INTR_CFG[1]. Enable message generation on rising and/or falling edges in PCIE_INTR_CFG[n].INTR_RISING_ENA and PCIE_INTR_CFG[n].INTR_FALLING_ENA. Different vectors can be generated for rising and falling edges, configure these through PCIE_INTR_CFG[n].RISING_VECTOR_VAL and PCIE_INTR_CFG[n].FALLING_VECTOR_VAL. Finally each EXT_DST interrupt must be given an appropriate traffic class via PCIE_INTR_CFG[n].TRAFFIC_CLASS.

After the root complex has configured the PCIe endpoint's MSI Capability Register Set and the external CPU has configured how interrupts from the VCore interrupt controller are propagated to PCIe. Then interrupts must be enabled by setting PCIE_INTR_COMMON_CFG.PCIE_INTR_ENA.

3.6.5 OUTBOUND ACCESS

After the PCIe endpoint is initialized, outbound read/write access to PCIe memory space is initiated by reading or writing the SBA's PCIe DMA region.

The following table lists the device registers associated with PCIe outbound access.

TABLE 3-13: OUTBOUND ACCESSES REGISTERS

Register	Description
ICPU_CFG::PCIESLV_SBA	Configuration for SBA outbound requests.
ICPU_CFG::PCIESLV_FDMA	Configuration for FDMA outbound requests.
PCIE::ATU_REGION	Select active region for the ATU_* registers.
PCIE::ATU_CFG1	Configuration of TLP fields.
PCIE::ATU_CFG2	Enable address translation.
PCIE::ATU_TGT_ADDR_LOW	Configuration of outbound PCIe address.
PCIE::ATU_TGT_ADDR_HIGH	Configuration of outbound PCIe address.

The PCIe DMA region is 1 gigabyte (GB); access in this region is mapped to any 1 gigabyte region in 32-bit or 64-bit PCIe memory space by using address translation. Two address translation regions are supported. The recommended approach is to configure the first region for SBA access and the second region for FDMA access.

Address translation works by taking bits [29:0] from the SBA/FDMA address, adding a configurable offset, and then using the resulting address to access into PCIe memory space. Offsets are configurable in steps of 64 kilobytes in the ATU_TGT_ADDR_HIGH and ATU_TGT_ADDR_LOW registers.

The software on the VCore CPU (or other SBA masters) can dynamically reconfigure the window as needed; however, the FDMA does not have that ability, and it must be disabled while updating the 1 gigabyte window that is set up for it.

Note Although the SBA and FDMA both access the PCIe DMA region by addresses 0xC0000000 through 0xFFFFFFFF, the PCIe address translation unit can differentiate between these accesses and apply the appropriate translation region.

3.6.5.1 Outbound SBA Translation Region Configuration

Configure PCIESLV_SBA.SBA_OFFSET = 0 and select address translation region 0 by writing ATU_REGION.ATU_IDX = 0. Set PCIE::ATU_BASE_ADDR_LOW.ATU_BASE_ADDR_LOW = 0x0000 and PCIE::ATU_LIMIT_ADDR.ATU_LIMIT_ADDR = 0x3FFF.

The following table lists the appropriate PCIe header that must be configured before using the SBA's PCIe DMA region. Remaining header fields are automatically handled.

TABLE 3-14: SBA PCIe ACCESS HEADER FIELDS

Header Field	Register::Fields
Attributes	ATU_CFG1.ATU_ATTR
Poisoned Data	PCIESLV_SBA.SBA_EP
TLP Digest Field Present	ATU_CFG1.ATU_TD
Traffic Class	ATU_CFG1.ATU_TC
Type	ATU_CFG1.ATU_TYPE
Byte Enables	PCIESLV_SBA.SBA_BE
Message Code	ATU_CFG2.ATU_MS-G_CODE

Configure the low address of the destination window in PCIe memory space as follows:

- Set ATU_TGT_ADDR_HIGH.ATU_TGT_ADDR_HIGH to bits [63:32] of the destination window. Set to 0 when a 32-bit address must be generated.
- Set ATU_TGT_ADDR_LOW.ATU_TGT_ADDR_LOW to bits [31:16] of the destination window. This field must not be set higher than 0xC000.

Enable address translation by writing ATU_CFG2.ATU_REGION_ENA = 1. SBA access in the PCIe DMA region will then be mapped to PCIe memory space as defined by ATU_TGT_ADDR_LOW and ATU_TGT_ADDR_HIGH.

The header fields and the PCIe address fields can be reconfigured on the fly as needed, however, make sure the SBA region is selected by setting ATU_REGION.ATU_IDX to 0.

3.6.5.2 Configuration of Outbound FDMA Translation Region

Configure PCIESLV_FDMA.FDMA_OFFSET = 1, and select address translation region 1 by writing ATU_REGION.ATU_IDX = 1.

Set PCIE::ATU_BASE_ADDR_LOW.ATU_BASE_ADDR_LOW = 0x4000 and PCIE::ATU_LIMIT_ADDR.ATU_LIMIT_ADDR = 0x7FFF.

FDMA PCIe header must be MRd/MWt type, reorderable, cache-coherent, and without ECRC. Remaining header fields are automatically handled.

TABLE 3-15: FDMA PCIe ACCESS HEADER FIELDS

Header Field	Register::Fields	Suggested Value
Attributes	ATU_CFG1.ATU_ATTR	Set to 2.
TLP Digest Field Present	ATU_CFG1.ATU_TD	Set to 0.
Traffic Class	ATU_CFG1.ATU_TC	Use an appropriate traffic class.
Type	ATU_CFG1.ATU_TYPE	Set to 0.

Configure low address of destination window in PCIe memory space as follows:

- Set ATU_TGT_ADDR_HIGH.ATU_TGT_ADDR_HIGH to bits [63:32] of the destination window. Set to 0 when a 32-bit address must be generated.
- Set ATU_TGT_ADDR_LOW.ATU_TGT_ADDR_LOW to bits [31:16] of the destination window. This field must not be set higher than 0xC000.

Enable address translation by writing ATU_CFG2.ATU_REGION_ENA = 1. The FDMA can be configured to make accesses in the 0xC0000000 – 0xFFFFFFFF address region. These accesses will then be mapped to PCIe memory space as defined by ATU_TGT_ADDR_LOW and ATU_TGT_ADDR_HIGH. The FDMA must be disabled if the address window needs to be updated.

3.6.6 POWER MANAGEMENT

The VSC7416-01 device's PCIe endpoint supports D0, D1, and D3 device power-management states and associated link power-management states. The switch core does not automatically react to changes in the PCIe endpoint's power management states. It is, however, possible to enable a VCore interrupt on device power state changes and then have the VCore CPU software make application-specific changes to the device operation depending on the power management state.

TABLE 3-16: POWER MANAGEMENT REGISTERS

Register	Description
ICPU_CFG::PCIE_STAT	Current power management state.
ICPU_CFG::PCIEPCS_CFG	Configuration of WAKE output and beacon.
ICPU_CFG::PCIE_INTR	PCIe interrupt sticky events.
ICPU_CFG::PCIE_INTR_ENA	Enable of PCIe interrupts.
ICPU_CFG::PCIE_INTR_IDENT	Currently interrupting PCIe sources.
ICPU_CFG::PCIE_AUX_CFG	Configuration of auxiliary power detection.

Because the VSC7416-01 device does not implement a dedicated auxiliary power for the PCIe endpoint, the endpoint is operated from the V_{DD} core power supply. Before the power management driver initializes the VSC7416-01 device, software can “force” auxiliary power detection by writing `PCIE_AUX_CFG = 3`, which causes the endpoint to report that it is capable of emitting Power Management Events (PME) messages in the D3c state.

The current device power management state is available using `PCIE_STAT.PM_STATE`. A change in this field's value sets the `PCIE_INTR.INTR_PM_STATE` sticky bit. To enable this interrupt, set `PCIE_INTR_ENA.INTR_P-M_STATE_ENA`. The current state of the PCIe endpoint interrupt towards the VCore interrupt controller is shown in `PCIE_INTR_IDENT` register (if different from zero, then interrupt is active).

The endpoint can emit PMEs if the `PME_En` bit is set in the PM Capability Register Set and if the endpoint is in power-down mode.

- Outbound request from either SBA or FDMA trigger PME.
- A change in status for an enabled outbound interrupt (either legacy or MSI) triggers PME. This feature can be disabled by setting `ICPU_CFG::PCIE_INTR_COMMON_CFG.WAKEUP_ON_INTR_DIS`.

In the D3 state, the endpoint transmits a beacon. The beacon function can be disabled and instead drive the WAKE output using the overlaid GPIO function. For more information about the overlaid function for this signal, see [Section 7.2.4, General-Purpose Inputs and Outputs](#).

TABLE 3-17: PCIE WAKE PIN

Pin Name	I/O	Description
PCIE_WAKE/GPIO_17	O	PCIe WAKE output

Enable WAKE by setting `PCIEPCS_CFG.BEACON_DIS`. The polarity of the WAKE output is configured in `PCIEPCS_CFG.WAKE_POL`, and the drive scheme is configured in `PCIEPCS_CFG.WAKE_OE`.

3.7 Frame DMA

This section describes the Frame DMA engine (FDMA). By enabling the FDMA, Ethernet frames can be extracted or injected autonomously to or from the device's DDR2/3 memory and/or PCIe memory space. Linked list data structures in memory are used for injecting or extracting Ethernet frames. The FDMA generates interrupts when frame extraction or injection is done and when the linked lists needs updating.

TABLE 3-18: FDMA REGISTERS

Register	Description
DEVCPU_QS::XTR_GRP_CFG	CPU port ownership, extraction direction.
DEVCPU_QS::INJ_GRP_CFG	CPU port ownership, injection direction.
REW::PORT_CFG	Enable IFH for extraction frames and disable FCS recalculation (for injected frames).

TABLE 3-18: FDMA REGISTERS (CONTINUED)

Register	Description
SYS::PORT_MODE	Enable IFH processing for injection frames.
ICPU_CFG::FDMA_CH_CFG	Channel configuration, priorities, and so on.
ICPU_CFG::FDMA_CH_ACTIVATE	Enabling of channels.
ICPU_CFG::FDMA_CH_DISABLE	Disable of channels.
ICPU_CFG::FDMA_CH_STAT	Status for channels.
ICPU_CFG::FDMA_CH_SAFE	Set when safe to update channel linked lists.
ICPU_CFG::FDMA_DCB_LL	Linked list pointer for channels.
ICPU_CFG::FDMA_DCB_LL_PREV	Previous linked list pointer for channels.
ICPU_CFG::FDMA_CH_CNT	Software counters for channels.
ICPU_CFG::FDMA_INTR_LL	NULL pointer event for channels.
ICPU_CFG::FDMA_INTR_LL_ENA	Enable interrupt on NULL pointer event.
ICPU_CFG::FDMA_INTR_FRM	Frame done event for channels.
ICPU_CFG::FDMA_INTR_FRM_ENA	Enable interrupt on frame done event.
ICPU_CFG::FDMA_INTR_SIG	SIG counter incremented event for channels.
ICPU_CFG::FDMA_INTR_SIG_ENA	Enable interrupt on SIG counter event.
ICPU_CFG::MANUAL_INTR	Special manual injection/extraction events.
ICPU_CFG::MANUAL_INTR_ENA	Enable interrupts on special manual injection/extraction events.
ICPU_CFG::FDMA_EVT_ERR	Error event for channels.
ICPU_CFG::FDMA_EVT_ERR_CODE	Error event description.
ICPU_CFG::FDMA_INTR_ENA	Interrupt enable for channels.
ICPU_CFG::FDMA_INTR_IDENT	Currently interrupting channels.
DEVCPU_QS::XTR_FRM_PRUNING	Enable pruning of extraction frames.
ICPU_CFG::FDMA_CH_INJ_TOKEN_CNT	Injection tokens.
ICPU_CFG::FDMA_CH_INJ_TOKEN_TICK_RLD	Periodic addition of injection tokens.
ICPU_CFG::MANUAL_CFG	Configuration of manual injection/extraction.
ICPU_CFG::MANUAL_XTR	Memory region used for manual extraction.
ICPU_CFG::MANUAL_INJ	Memory region used for manual injection.
DEVCPU_QS::INJ_CTRL	Configuration of injection gap.
ICPU_CFG::FDMA_GCFG	Configuration of injection buffer watermark.

The FDMA implements two channels per CPU port: one extraction-channel and one injection-channel. Before a channel can be used, ownership must be mapped to the FDMA. The VSC7416-01 has four FDMA channels:

- FDMA channel 0 corresponds to port 11 extraction direction. Mapping by means of XTR_GRP_CFG[0].MODE.
- FDMA channel 1 corresponds to port 12 extraction direction. Mapping by means of XTR_GRP_CFG[1].MODE.
- FDMA channel 2 corresponds to port 11 injection direction. Mapping by means of INJ_GRP_CFG[0].MODE.
- FDMA channel 3 corresponds to port 12 injection direction. Mapping by means of INJ_GRP_CFG[1].MODE.

The FDMA implements a strict priority scheme. Injection and extraction channels can be assigned individual priorities, which are used when the FDMA has to decide between servicing two or more channels. Channel priority is configured in FDMA_CH_CFG[ch].CH_PRIO. When channels have same priority, the higher channel number takes priority over the lower channel number.

The internal frame header can be added to extracted frames on a CPU port by setting PORT_CFG[n].IFH_INSERT_ENA. The internal frame header is added in front of extracted frames and provides useful switching-information about the extracted frames. For more information about the extraction IFH, see [Section 2.14.1, Frame Extraction](#).

By default, injection frames are analyzed and switched just as any other frame in the system. However, by adding and enabling internal frame header (using `PORT_MODE[n].INCL_INJ_HDR`) for a CPU port, it is possible to control injection parameters directly from the CPU. The internal frame header is added in front of frame data. For more information about the injection IFH, see [Section 2.14.2, Frame Injection](#).

The rewriter recalculates and overwrites the Ethernet FCS for frames that are injected via FDMA. This can be disabled on a destination port basis via `PORT_CFG[n].FCS_UPDATE_CPU_ENA`.

When injection is enabled injection gap0 must be configured to 1 (`INJ_CTRL[0].GAP_SIZE = 1`). If any of the injection channels is injecting to a CPU port without IFH then injection gap1 must also be configured to 1 (`INJ_CTRL[1].GAP_SIZE = 1`). When injection is enabled injection buffer watermark must be configured to 13 (`FDMA_GCFG.INJ_RF_WM 0x0D`).

The FDMA supports a manual mode where the FDMA decision logic is disabled and the FDMA takes and provides data in the order that was requested by an external master (internal or external CPU). The manual mode is a special case of normal FDMA operation where only few of the FDMA features apply. For more information about manual operation, see [Section 3.7.7, Manual Mode](#).

3.7.1 DMA CONTROL BLOCK STRUCTURES

The FDMA processes linked lists of DMA Control Block Structures (DCBs). The DCBs have the same basic structure for both injection or for extraction. A DCB must be placed on a 32-bit word-aligned address in memory. Each DCB must have an associated data block that is placed on a 32-bit word aligned address in memory, the length of the data block must be a complete number of 32-bit words.

An Ethernet frame can be contained inside one data block (if the data block is big enough) or the frame can be spread across multiple data blocks. A data block never contains more than one Ethernet frame. Data blocks that contain start-of-frame have set a special bit in the DCB's status word, likewise for data blocks that contains end-of-frame. The FDMA stores or retrieves Ethernet frame data in network order. This means that the data at byte address (n) of a frame was received just before the data at byte address ($n + 1$).

Frame data inside the DCB's associated data blocks can be placed at any byte-offset and have any byte-length as long as byte-offset and length does not exceed the size of the data block. Byte offset and length is configured using special fields inside the DCB status word. Software can specify offset both when setting up extraction and injection DCBs. Software only specifies length for injection DCBs; the FDMA automatically calculates and updates byte-length during extraction.

Example If a DCB's status word has block-offset 5 and block-length 2, then the DCB's data block contains 2 bytes of frame data placed at second and third bytes inside the second 32-bit word of the DCB's associated data block.

FIGURE 3-12: LAYOUT OF A DCB

word-address n	LLP [31:0]: Linked List Pointer, this field points to the next DCB in the list, set to NULL for end of list. DCBs must be word-aligned, which means LLP must point to a word-aligned address (bit [1:0] = 00).						LLP
$n + 1$	DATAP [31:0]: Data block pointer, this field points to the data block associated with this DCB. DATAP must be !=NULL and must point to a word-aligned address (bit [1:0] = 00).						DATAP
$n + 2$	[31:24]: For SW use. FDMA will not modify.	[23:18]: Must be 0.	17	16	DATAL [15:0]: Number of bytes available in the data block associated with this DCB. DATAL must be a complete number of words (bit [1:0] = 00).		DATAL
$n + 3$	BLOCKO [31:20]: Byte offset from beginning of the data block to the first byte of frame-data.	19	18	17	16	BLOCKL [15:0]: Number of framedata bytes in the data block. BLOCKL does not include offset bytes.	STAT

STAT[16] SOF: Set to 1 if data block contains start of frame.

STAT[17] EOF: Set to 1 if data block contains end of frame.

STAT[18] ABORT: Abort indication.

Extraction: Set to 1 if the frame associated with the data block was aborted.

Injection: If set to 1 when FDMA loads the DCB, it aborts the frame associated with the data block.

STAT[19] PD: Pruned/Done indication.

Extraction: Set to 1 if the frame associated with the data block was pruned.

Injection: The FDMA set this to 1 when done processing the DCB. If set to 1 when FDMA loads the DCB, it is treated as ABORT.

DATAL[16] SIG: If set 1 when FDMA loads the DCB, the CH_CNT_SIG counter is incremented by one.

DATAL[17] TOKEN: Token indication, only used during injection.

If set to 1, the FDMA uses one token (CH_INJ_TOKEN_CNT) when injecting the contents of the DCB. If the token counter is 0 when loading the DCB, then injection is postponed until tokens are made available.

DCBs are linked together via the DCB's LLP field. The last DCB in a chain must have LLP = NULL. Chains consisting of a single DCB are allowed.


3.7.2 ENABLING AND DISABLING FDMA CHANNELS

To enable a channel (ch), write a valid DCB pointer to FDMA_DCB_LLP[ch] and enable the channel by setting FDMA_CH_ACTIVATE.CH_ACTIVATE[ch]. This makes the FDMA load the DCB from memory and start either injection or extraction.

To schedule a channel for disabling, set FDMA_CH_DISABLE.CH_DISABLE[ch]. An active channel does not disable immediately. Instead it waits until the current data block is done, saves the status word, and then disables.

Channels can be in one of three states: DISABLED, UPDATING, or ACTIVE. Channels are DISABLED by default. When the channel is reading a DCB or writing the DCB status word, it is considered to be UPDATING. The status of individual channels is available in FDMA_CH_STAT.CH_STAT[ch].

FIGURE 3-13: CHANNEL STATES



A channel that has `FDMA_DCB_LLIP[ch].DCB_LLIP==NULL` when going from ACTIVE to UPDATING disables itself instead of loading a new DCB. After this it can be re-enabled as previously described. Extraction channels emit an `INTR_LLIP`-event when loading a DCB with `LLP==NULL`. Injection channels emit an `INTR_LLIP`-event when saving status for a DCB that has the `LLP==NULL`.

Note Extraction channel running out of DCBs during extraction is a problem that software must avoid. A hanging extraction channel will potentially be head-of-line blocking other extraction channels.

It is possible to update an active channels LLP pointer and pointers in the DCB chains. Before changing pointers software must schedule the channel for disabling (by writing `FDMA_CH_DISABLE.CH_DISABLE[ch]`) and then wait for the channel to set `FDMA_CH_SAFE.CH_SAFE[ch]`. When done, software must re-activate the channel by setting `FDMA_CH_ACTIVATE.CH_ACTIVATE[ch]`. Setting activate will cancel the deactivate-request, or if the channel has disabled itself in the meantime, it will re-activate the channel.

Note The address of the current DCB is available in `FDMA_DCB_LLIP_PREV[ch]`. This information is useful when modifying pointers for an active channel. The FDMA does not reload the current DCB when re-activated, so if the LLP-field of the current DCB is modified, then software must also modify `FDMA_DCB_LLIP[ch]`.

Setting `FDMA_CH_CFG[ch].DONEEOF_STOP_ENA` disables an FDMA channel and emits LLP-event after saving status for DCBs that contains EOF (after extracting or injecting a complete frame). Setting `FDMA_CH_CFG[ch].DONE_STOP_ENA` disables an FDMA channel and emits LLP-event after saving status for any DCB.

3.7.3 CHANNEL COUNTERS

The FDMA implements three counters per channel: SIG, DCB, and FRM. These counters are accessible through `FDMA_CH_CNT[ch].CH_CNT_SIG`, `FDMA_CH_CNT[ch].CH_CNT_DCB`, and `FDMA_CH_CNT[ch].CH_CNT_FRM`, respectively. For more information about how to safely modify these counters, see register descriptions.

- The SIG (signal) counter is incremented by one each time the FDMA loads a DCB that has the `DATAL.SIG` bit set to 1.
- The FRM (frame) counter is incremented by one each time the FDMA store status word for DCB that has EOF set.

It is a wrapping counter that can be used for software driver debug and development. This counter does not count aborted frames.

- The DCB counter is incremented by one every time the FDMA loads a DCB from memory. It is a wrapping counter that can be used for software driver debug and development.

It is possible to enable channel interrupt whenever the SIG counter is incremented; this makes it possible for software to receive interrupt when the FDMA reaches certain points in a DCB chain. DCB refill for extraction channels can be implemented by using the SIG counter.

3.7.4 FDMA EVENTS AND INTERRUPTS

Each FDMA channel can generate four events: LLP-event, FRM-event, SIG-event, and ERR-event. These events cause a bit to be set in `FDMA_INTR_LLP.INTR_LLP[ch]`, `FDMA_INTR_FRM.INTR_FRM[ch]`, `FDMA_INTR_SIG.INTR_SIG[ch]`, and `FDMA_EVT_ERR.EVT_ERR[ch]`, respectively.

- LLP-event occurs when an extraction channel loads a DCB that has `LLP = NULL` or when an injection channel writes status for a DCB that has `LLP=NULL`. LLP-events are also emitted from channels that have `FDMA_CH_CFG[ch].DONEEOF_STOP_ENA` or `FDMA_CH_CFG[ch].DONE_STOP_ENA` set. For more information, see [Section 3.7.2, Enabling and Disabling FDMA Channels](#).
- FRM-event is indicated when an active channel loads a new DCB and the previous DCB had EOF. The FRM-event is not emitted for channels that are disabled after writing DCB status with EOF.
- SIG-event is indicated whenever the `FDMA_CH_CNT[ch].CH_CNT_SIG` counter is incremented. The SIG (signal) counter is incremented when loading a DCB that has the `DATAL.SIG` bit set.
- ERR-event is an error indication that is set for a channel if it encounters an unexpected problem during normal operation. This indication is implemented to ease software driver debugging and development. A channel that encounters an error will be disabled, depending on the type of error the channel state may be too corrupt for it to be restarted without system reset. When an ERR-event occurs, the `FDMA_EVT_ERR_CODE.EVT_ERR_CODE` shows the exact reason for an ERR-event. For more information about the errors that are detected, see `FDMA_EVT_ERR_CODE.EVT_ERR_CODE`.

Each of the events (LLP, FRM, SIG) can be enabled for channel-interrupt through `FDMA_INTR_LLP_ENA.INTR_LLP_ENA[ch]`, `FDMA_INTR_FRM_ENA.INTR_FRM_ENA[ch]`, and `FDMA_INTR_SIG.INTR_SIG[ch]` respectively. The ERR-event is always enabled for channel-interrupt.


The highest numbered extraction channel supports two additional non-sticky events related to manual extraction: `XTR_SOF_RDY`-event and `XTR_ANY_RDY`-event. An active event causes the following fields to be set: `MANUAL_INTR.INTR_XTR_SOF_RDY` and `MANUAL_INTR.INTR_XTR_ANY_RDY` respectively. For more information, see [Section 3.7.7.1, Manual Extraction](#).

- `XTR_SOF_RDY`-event is active when the next word to be manually extracted contains an SOF indication.
- `XTR_ANY_RDY`-event is active when any word is available for manually extraction.

The above events are enabled in `MANUAL_INTR_ENA.INTR_XTR_SOF_RDY_ENA` and `MANUAL_INTR_ENA.INTR_XTR_ANY_RDY_ENA` respectively.

The highest numbered injection channel supports one additional non-sticky event related to manual injection: `INJ_RDY`-event. This event is active when the injection logic has room for (at least) 16 32-bit words of injection frame data. When this event is active, `MANUAL_INTR.INTR_INJ_RDY` is set. The `INJ_RDY`-event can be enabled for channel interrupt by setting `MANUAL_INTR_ENA.INTR_INJ_RDY_ENA`. For more information, see [Section 3.7.7.2, Manual Injection](#).

FIGURE 3-14: CHANNEL INTERRUPT HIERARCHY



The FDMA_INTR_ENA.INTR_ENA[ch] field enables interrupt from individual channels, FDMA_INTR_IDENT.INTR_IDENT[ch] field shows which channels that are currently interrupting. While INTR_IDENT is non-zero, the FDMA is indicating interrupting towards to the VCore interrupt controller.

3.7.5 FDMA EXTRACTION

During extraction the FDMA extracts Ethernet frame data from the Queuing System and saves this into the data block of the DCB that is currently loaded by the FDMA extraction channel. The FDMA continually process DCBs until it reaches a DCB with LLP = NULL or until it is disabled.

When an extraction channel writes the status word of a DCB it will update SOF/EOF/ABORT/PRUNED-indications and BLOCKL. BLOCKO remains unchanged (write-value is taken from the value that was loaded when the DCB was read).

Aborting of frames from the QS will not occur during normal operation. If the Queuing System is reset during extraction of a particular frame, then ABORT and EOF will be set. Software must discard frames that have ABORT set.

Pruning of frames during extraction is enabled per extraction-port via XTR_FRM_PRUNING[n].PRUNE_SIZE. When enabled Ethernet frames above the configured size will be truncated and PD and EOF will be set.

3.7.6 FDMA INJECTION

During injection the FDMA reads Ethernet frame data from the data block of the DCB that is currently loaded by the FDMA injection channel and injects this into the Queuing System. The FDMA continually process DCBs until it reaches a DCB with LLP = NULL or until it is disabled.

When an injection channel writes the status word of a DCB it will set DONE-indication. All other status word fields remain unchanged (write-value is stored first time the injection channel loads the DCB).

The rate by which the FDMA injects frames can be shaped by using tokens. Each injection channel has an associated token counter (FDMA_CH_INJ_TOKEN_CNT[ich]). A DCB that has the DATAL.TOKEN field set will cause the injection channel to deduct one from the token counter before the data block of the DCB can be injected. If the token counter is already at 0 then the injection channel postpones injection until the channels token counter is set to a value different from 0.

Tokens can be added to the token counter by writing the FDMA_CH_INJ_TOKEN_CNT[ich] register or tokens can be added automatically (with a fixed interval) by using the dedicated token tick counter. Setting FDMA_CH_INJ_TOKEN_TICK_RLD[ich] to a value n (different from 0) will cause one token to be added to that injection channel every $n \times 198.4$ ns.

3.7.7 MANUAL MODE

It is possible to disable the decision making logic of the FDMA extraction path and/or injection paths and thereby give control of the FDMA's extraction and/or injection buffers directly to any master attached to the Shared Bus Architecture. When operating in manual mode DCB structure, Counters and most of the interrupts does not apply.

Manual extraction and injection use hardcoded channel numbers.

- Manual extraction mode uses FDMA channel 1 (port 12 extraction direction.)
- Manual injection mode uses FDMA channel 3 (port 12 injection direction.)


To enable manual extraction, set MANUAL_CFG.XTR_ENA; the FDMA must not be enabled for FDMA extraction when manual extraction mode is active. To enable manual injection set MANUAL_CFG.INJ_ENA; the FDMA must not be enabled for FDMA injection when manual injection mode is active. Extraction and injection paths are independent. For example, FDMA can control extraction at the same time as doing manual injection by means of the CPU.

3.7.7.1 Manual Extraction

Extraction is done by reading one or more blocks of data from the extraction memory area. A block of data is defined by reading one or more data words followed by reading of the extraction status word. The extraction memory area is 16 kilobytes and is implemented as a replicated register region MANUAL_XTR. The highest register in the replication (0xFFF) maps to the extraction status word. The status word is updated by the extraction logic and shows; the number of frame-bytes that has been read, SOF and EOF indications, and PRUNED/ABORT indications.

Note During extraction, a CPU does not know the frame length when extraction is started. This means that the CPU must check for EOF after each block it reads out. When reading a data block from the device, the CPU can burst read from the memory area in such a way that the extraction status word is read as the last word of the burst.

FIGURE 3-15: EXTRACTION STATUS WORD ENCODING



The extraction logic updates the extraction header while the CPU is reading out data for a block. Prior to starting a new data block, the CPU can write the two least significant bits of the BLOCKO field. The BLOCKO value is stored in the extraction logic and takes effect when the new data block is started. Reading the status field always returns the BLOCKO value that applies to the current data block. Unless written (before starting a data block), the BLOCKO is cleared, so by default all blocks have 0 byte offset. The offset can be written multiple times; the last value takes effect.

The CPU can abort frames that it has already started to read out by writing the extraction status field with the ABORT field set. This causes the extraction logic to discard the active frame and remove it from the queuing system.

Reading of block data does not have to be done in one burst. As long as the status word is not read, multiple reads from the extraction region are treated as belonging to the same block.


3.7.7.2 Manual Injection

Injection is done by writing one or more blocks of data to the injection memory area. A block of data is defined by writing an injection status word followed by writing one or more data words. The injection memory area is 16 kilobytes and is implemented as a replicated register region MANUAL_INJ. The first register in this replication maps to the injection status word. The status word for each block defines the following:

- Number of bytes to be injected
- Optional byte-offset before injection data
- SOF and EOF indications
- Optional ABORT indication

Note In general, it makes sense to inject frames as a single large block of data (containing both SOF and EOF). However, because offset/length can be specified individually for each block, injecting frames through several blocks is useful when compensating for offset/word misalignment (for example, when building a frame from different memory regions in the CPU main memory).

FIGURE 3-16: INJECTION STATUS WORD ENCODING



Injection logic updates the PD field of the status word. The status word can be read at any time during injection without affecting current data block transfers. However a CPU is able to calculate how much data that is required to inject a complete block of data (at least BLOCKO + BLOCKL bytes) so reading the status word is for software-development and debug only. Writing status word before finishing the previous data block will cause that frame to be aborted. Frames are also aborted if they are not “started” with SOF or “ended” with EOF indications.

As long as the status word is not written, then multiple writes to the injection region will be treated as data belonging to the same block. This means multiple bursts of data words can be written between each injection status word update.

Software can abort frames by writing to injection status with ABORT field set. No other status-word fields need to be valid. When a frame is aborted, the already injected data is removed from the Queuing System.

3.8 VCore System Peripherals

This section describes the subblocks of the VCore system. They are primarily intended to be used by the VCore CPU. However, an external CPU can also access and control these through the shared bus.

3.8.1 SI CONTROLLER

This section provides information about the SPI boot master that allows booting from a Flash hooked up to the serial interface. For information about using an external CPU to access device registers using the serial interface, see [Section 3.5.2, Serial Interface in Slave Mode](#).

The following table lists the registers associated with the SI controller.

TABLE 3-19: SI CONTROLLER CONFIGURATION REGISTERS

Register	Description
ICPU_CFG::SPI_MST_CFG	Serial interface speed
ICPU_CFG::SW_MODE	Manual control of the serial interface pins

There are four programmable chip selects when the VCore system controls the SI. Through individually mapped memory regions, each chip select can address up to 16 megabytes (MB) of memory. Reading from the memory region for a specific SI chip select generates SI read on that chip select. The VCore CPU can execute code directly from Flash by executing from the SI controller's memory region.

FIGURE 3-17: SI CONTROLLER MEMORY MAP

SI Controller Memory Map	
SI Controller	
+0x01000000	16 MB Chip Select 0, SI_nEn
+0x02000000	16 MB Chip Select 1, SI_nEn1
+0x03000000	16 MB Chip Select 2, SI_nEn2
+0x04000000	16 MB Chip Select 3, SI_nEn3

The SI controller accepts 8-bit, 16-bit, and 32-bit read access with or without bursting. Writing to the SI requires manual control of the SI pins using software. Setting SW_MODE.SW_PIN_CTRL_MODE places all SI pins under software control. Output enable and the value of SI_Clk, SI_DO, SI_nEn[3:0] are controlled through the SW_MODE register. The value of the SI_DI pin is available in SW_MODE.SW_SPI_SDI.

Note The VCore CPU cannot execute code directly from the SI controller's memory region at the same time as manually writing to the serial interface.

The following table lists the serial interface pins.

TABLE 3-20: SERIAL INTERFACE PINS

Pin Name	I/O	Description
SI_nEN SI_nEN1/GPIO_8 SI_nEN2/GPIO_9 SI_nEN3/GPIO_10	O	Active low chip selects. Only one chip select can be active at any time. Chip selects 1 through 3 are overlaid functions on the GPIOs. See Section 7.2.4, General-Purpose Inputs and Outputs .
SI_Clk	O	Clock output.
SI_DO	O	Data output (MOSI).
SI_DI	I	Data input (MISO).

The SI controller does speculative prefetching of data. After reading address n , the SI controller automatically continues reading address $n + 1$, so that the next value is ready if requested by the VCore CPU. This greatly optimizes reading from sequential addresses in the Flash, such as when copying data from Flash into program memory.

FIGURE 3-18: SI READ TIMING IN NORMAL MODE





FIGURE 3-19: SI READ TIMING IN FAST MODE



The default timing of the SI controller operates with most serial interface Flash devices. Use the following process to calculate the optimized SI parameters for a specific SI device:

1. Calculate an appropriate frequency divider value as described in SPI_MST_CFG.CLK_DIV. The SI operates at no more than 25 MHz, and the maximum frequency of the SPI device must not be exceeded. The VCore system frequency in the VSC7416-01 device is 208.33MHz.
2. The SPI device may require a FAST_READ command rather than normal READ when the SI frequency is increased. Setting SPI_MST_CFG.FAST_READ_ENA makes the SI controller use FAST_READ commands.
3. Calculate SPI_MST_CFG.CS_DESELECT_TIME so that it matches how long the SPI device requires chip-select to be deasserted between accesses. This value depends on the SI clock period that results from the SPI_MST_CFG.CLK_DIV setting.

These parameters must be written to SPI_MST_CFG. The CLK_DIV field must either be written last or at the same time as the other parameters. The SPI_MST_CFG register can be configured while also booting up from the SI.

When the VCore CPU boots from the SI interface, the default values of the SPI_MST_CFG register are used until the SI_MST_CFG is reconfigured with optimized parameters. This implies that SI_Clk is operating at approximately 7 MHz, with normal read instructions, and maximum gap between chip select operations to the Flash.

Note The SPI boot master does optimized reading. SI_DI (from the Flash) is sampled just before driving falling edge on SI_CLK (to the Flash). This greatly relaxes the round trip delay requirement for SI_CLK to SI_DI significantly, allowing higher Flash clock frequencies.

3.8.2 DDR2/3 MEMORY CONTROLLER

This section provides information about how to configure and initialize the DDR2/3 memory controller.

The following table lists the registers associated with the memory controller.

TABLE 3-21: DDR2 CONTROLLER REGISTERS

Register	Description
ICPU_CFG::RESET	Reset Control
ICPU_CFG::MEMCTRL_CTRL	Control
ICPU_CFG::MEMCTRL_CFG	Configuration
ICPU_CFG::MEMCTRL_STAT	Status
ICPU_CFG::MEMCTRL_REF_PERIOD	Refresh period configuration
ICPU_CFG::MEMCTRL_ZQCAL	DDR3 ZQ-calibration

TABLE 3-21: DDR2 CONTROLLER REGISTERS (CONTINUED)

Register	Description
ICPU_CFG::MEMCTRL_TIMING0	Timing configuration
ICPU_CFG::MEMCTRL_TIMING1	Timing configuration
ICPU_CFG::MEMCTRL_TIMING2	Timing configuration
ICPU_CFG::MEMCTRL_TIMING3	Timing configuration
ICPU_CFG::MEMCTRL_MR0_VAL	Mode register 0 value
ICPU_CFG::MEMCTRL_MR1_VAL	Mode register 1 value
ICPU_CFG::MEMCTRL_MR2_VAL	Mode register 2 value
ICPU_CFG::MEMCTRL_MR3_VAL	Mode register 3 value
ICPU_CFG::MEMCTRL_TERMRES_CTRL	Termination resistor control
ICPU_CFG::MEMCTRL_DQS_DLY	DQS window configuration
ICPU_CFG::MEMPHY_CFG	SSTL configuration
ICPU_CFG::MEMPHY_ZCAL	SSTL calibration

The memory controller works with JEDEC-compliant DDR2 and DDR3 memory modules. The controller runs 312.50 MHz and supports one or two byte lanes in either single or dual row configurations (one or two chip selects). The controller supports up to 16 address bits as well as 4- or 8-bank modules. The maximum amount of physical memory that can be attached to the controller is 1 gigabyte.

The memory controller supports ECC encoding/decoding by using one of the two lanes for ECC information. Enabling ECC requires that both byte-lanes are populated. When ECC is enabled half the total physical memory is available for software use, the other half is used for ECC information.

The following steps are required to bring up the memory controller:

1. Release the controller from reset by clearing RESET.MEM_RST_FORCE.
2. Configure timing and mode parameters. These depend on the DDR2/DDR3 modules and configuration selected for the VSC7416-01 device. For more information, see [Section 3.8.2.1, DDR2/3 Timing and Mode Configuration Parameters](#).
3. Enable and calibrate the SSTL I/Os. For more information, see [Section 3.8.2.3, Enabling and Calibrating SSTL I/Os](#).
4. Initialize the memory controller and modules. For more information, see [Section 3.8.2.4, Memory Controller and Module Initialization](#).
5. Calibrate the DQS read window. For more information, see [Section 3.8.2.5, DQS Read Window Calibration](#).
6. Optionally configure ECC and 1 gigabyte support. For more information see [Section 3.8.2.6, ECC and 1 Gigabyte Support](#).

Note For selected DDR2/DDR3 modules, the bring-up of the memory controller is already implemented as part of the Board Support Package (BSP). For more information about implementing the bring-up procedure, see *BSP and Tool Chain*, which is available on the Microchip Web site at <https://www.microchip.com/>.

3.8.2.1 DDR2/3 Timing and Mode Configuration Parameters

This section provides information about each of the parameters that must be configured prior to initialization of the memory controller. It provides a quick overview of fields that must be configured and their recommended values. For more information about each field, see the register list.

The memory controller operates at 312.5 MHz, and the corresponding DDR clock period is 3.2 ns (DDR625). When a specific timing value is used in this section (for example, t_{MOD}), it means that the corresponding value from the memory module datasheet is expressed in memory controller clock cycles (divided by 3.2 ns and rounded up and possibly adjusted for minimum clock cycle count).

The memory controller can be configured for either BURST4 or BURST8 operation. A two byte-lane DDR2 configuration can use BURST4. All other configurations (DDR3 and single byte-lane DDR2) must use BURST8 mode.

VSC7416-01

The following table defines general parameters for DDR2 and DDR3 required when configuring the memory controller.

TABLE 3-22: GENERAL DDR2/3 TIMING AND MODE PARAMETERS

Parameter	DDR2	DDR3
RL	CL	CL
WL	RL-1	CWL
MD	t_{MRD}	t_{MOD}
ID	125	t_{XPR}
SD	t_{XSRD}	t_{DLLK}
OW	WL-2	2
OR	RL-3	2
RP	t_{RP} for 4-bank modules $t_{RPA}^{(1)}$ for 8-bank modules	t_{RP}
FAW	1 for 4-bank modules t_{FAW} for 8-bank modules	t_{FAW}
BL	2 for BURST4 mode 4 for BURST8 mode	4
MR0	$((RL < 4) \mid ((t_{WR-1}) < 9) \mid 0x0002)$ for BURST4 mode $((RL < 4) \mid ((t_{WR-1}) < 9) \mid 0x0003)$ for BURST8 mode	$((RL-4) < 4) \mid ((t_{WR-4}) < 9)$
MR1	0x03C2	0x0040
MR2	0x0000	$((WL-5) < 3)$
MR3	0x0000	0x0000

¹Some DDR datasheets do not list t_{RPA} for 8-bank modules; use t_{RP} instead.

The following table lists the memory controller registers that must be configured before initialization. It is recommended to configure all fields listed, even when ODT or dual row configurations are not used.

TABLE 3-23: MEMORY CONTROLLER CONFIGURATION

Register Fields	Description
MEMCTRL_CFG.DDR_MODE	Set to 0 for DDR2 mode. Set to 1 for DDR3 mode.
MEMCTRL_CFG.DDR_WIDTH	Set to 0 when using 1 byte lane. Set to 1 when using 2 byte lanes.
MEMCTRL_CFG.MSB_COL_ADDR	Set to one less than the number of column address bits for the module(s).
MEMCTRL_CFG.MSB_ROW_ADDR	Set to one less than the number of row address bits for the module(s). See field description for example.
MEMCTRL_CFG.BANK_CNT	Set to 0 for 4-bank module(s). Set to 1 for 8-bank module(s).
MEMCTRL_CFG.BURST_LEN	Set to 0 for BURST4 mode. Set to 1 for BURST8 mode.
MEMCTRL_CFG.BURST_SIZE	Set to (MEMCTRL_CFG.DDR_WIDTH & MEMCTRL_CFG.BURST_LEN).
MEMCTRL_CFG.DDR_ECC_ENA	Set to 0 during initialization of the memory controller.
MEMCTRL_CFG.DDR_ECC_COR_ENA	Set to 0 during initialization of the memory controller.
MEMCTRL_CFG.DDR_ECC_ERR_ENA	Set to 0 during initialization of the memory controller.
MEMCTRL_CFG.DDR_512MBYTE_PLUS	Set to 0 during initialization of the memory controller.

TABLE 3-23: MEMORY CONTROLLER CONFIGURATION (CONTINUED)


Register Fields	Description
MEMCTRL_REF_PERIOD.REF_PERIOD	Set to ($t_{REFI}(ns)/3.2ns$) round down to the nearest integer.
MEMCTRL_REF_PERIOD.MAX_PEND_REF	Set to 8.
MEMCTRL_TIMING0.RD_DATA_XFR_DLY	Set to (RL-3).
MEMCTRL_TIMING0.WR_DATA_XFR_DLY	Set to (WL-1).
MEMCTRL_TIMING0.RD_TO_PRECH_DLY	Set to (BL-1).
MEMCTRL_TIMING0.WR_TO_PRECH_DLY	Set to (WL+BL+tWR-1).
MEMCTRL_TIMING0.RAS_TO_PRECH_DLY	Set to ($t_{RAS}(min)$ -1).
MEMCTRL_TIMING0.RD_CS_CHANGE_DLY	Set to BL.
MEMCTRL_TIMING0.WR_CS_CHANGE_DLY	Set to (BL-1).
MEMCTRL_TIMING0.RD_TO_WR_DLY	Set to (RL+BL+1-WL).
MEMCTRL_TIMING1.WR_TO_RD_DLY	Set to (WL+BL+tWTR-1).
MEMCTRL_TIMING1.RAS_TO_CAS_DLY	Set to (t_{RCD} -1).
MEMCTRL_TIMING1.RAS_TO_RAS_DLY	Set to (t_{RRD} -1).
MEMCTRL_TIMING1.PRECH_TO_RAS_DLY	Set to (t_{RP} -1).
MEMCTRL_TIMING1.BANK8_FAW_DLY	Set to (FAW-1).
MEMCTRL_TIMING1.RAS_TO_RAS_SAME_BANK_DLY	Set to (t_{RC} -1).
MEMCTRL_TIMING2.INIT_DLY	Set to (ID-1) during initialization of the memory controller.
MEMCTRL_TIMING2.REF_DLY	Set to (t_{RFC} -1).
MEMCTRL_TIMING2.MDSET_DLY	Set to (MD-1).
MEMCTRL_TIMING2.PRECH_ALL_DLY	Set to (RP-1).
MEMCTRL_TIMING3.WR_TO_RD_CS_CHANGE_DLY	Set to (WL+tWTR-1).
MEMCTRL_TIMING3.LOCAL_ODT_RD_DLY	Set to (RL-1).
MEMCTRL_TIMING3.ODT_WR_DLY	Set to (OW-1).
MEMCTRL_TIMING3.ODT_RD_DLY	Set to (OR-1).
MEMCTRL_MR0_VAL.MR0_VAL	Set to MR0.
MEMCTRL_MR1_VAL.MR1_VAL	Set to MR1.
MEMCTRL_MR2_VAL.MR2_VAL	Set to MR2.
MEMCTRL_MR3_VAL.MR3_VAL	Set to MR3.

The memory controller supports two chip selects and automatically selects modules based on the configured address width and bank count. The resulting pool of memory is one large consecutive region.

3.8.2.2 On-Die Termination

End termination scheme must be configured when writing from the VSC7416-01 device to the memory modules, the memory module furthest away from the device must enable termination on write access. Enable assertion of ODT0 output for all write access (also CS1 writes) by setting `MEMCTRL_TERMRES_CTRL = 0x0000000C`. With the previously specified MR configuration, this will cause termination to be enabled in the CS0 modules for all write operations on the DDR2/3 interface. The following illustration shows the recommended connection of the DDR2/3 modules.

FIGURE 3-20: MEMORY CONTROLLER ODT HOOKUP



If the CS1 modules are placed furthest away from the VSC7416-01 device, setting `MEMCTRL_TERMRES_CTRL = 0x00000030` will enable termination in the CS1 modules for all write access.

3.8.2.3 Enabling and Calibrating SSTL I/Os

Prior to controller initialization, the device's SSTL I/O drivers must be enabled and calibrated to correct drive strength and termination resistor values. It is recommended that the device's I/O drive and termination strengths be 60 Ω /60 Ω .

Complete the following tasks to enable and calibrate the SSTL I/Os:

1. Enable the SSTL mode by clearing `MEMPHY_CFG.PHY_RST` and setting `MEMPHY_CFG.PHY_SSTL_ENA`.
2. Perform calibration with the previously mentioned drive and termination strengths by writing 0xEF to `MEMPHY_ZCAL`.
3. Wait until `MEMPHY_ZCAL.ZCAL_ENA` is cleared (indicates calibration is done).
4. Enable drive of the SSTL I/Os by setting `MEMPHY_CFG.PHY_CHK_OE`, `MEMPHY_CFG.PHY_CL_OE`, and `MEMPHY_CFG.PHY_ODT_OE`.

The SSTL interface is now enabled and calibrated, and the initialization of the memory controller can commence.

3.8.2.4 Memory Controller and Module Initialization

After memory controller parameters are configured, and after the SSTL I/Os are enabled and calibrated, the memory controller and modules are initialized by setting `MEMCTRL_CTRL.INITIALIZE`.

During initialization, the memory controller automatically follows the proper JEDEC defined procedure for initialization and writing of mode registers to the DDR2 or DDR3 memory modules.

The memory controller automatically clears `MEMCTRL_CTRL.INITIALIZE` and sets the `MEMCTRL_STAT.INIT_DONE` field after the controller and the memory module(s) are operational. Software must wait for the `INIT_DONE` indication before continuing to calibrate the read-window.

3.8.2.5 DQS Read Window Calibration

After initialization of the memory controller, writes to the memory are guaranteed to be successful. Reading is not yet possible, because the round trip delay between controller and memory modules are not calibrated.

Calibration of the read-window includes writing a known value to the start of the DDR memory and then continually reading this value while adjusting the DQS window until the correct value is read from the memory. The calibration procedure depends on the number of byte lanes that are used (see MEMCTRL_CFG.DDR_WIDTH).

Note For an ECC system, the calibration procedure for two byte lanes must be used, because the memory controller is initialized without ECC enabled.

Calibration of One Byte Lane

The following procedure applies to systems where MEMCTRL_CFG.DDR_WIDTH is 0:

- Write 0x000000FF to SBA address 0x20000000.
- Set the MEMCTRL_DQS_DLY[0].DQS_DLY field to 0.

Perform the following steps to calibrate the read window. Do not increment the DQS_DLY field beyond its maximum value. If the DQS_DLY maximum value is exceeded, it is an indication something is incorrect, and the memory controller will not be operational.

1. Read byte address 0x20000000. If the content is different from 0xFF, increment MEMCTRL_DQS_DLY[0].DQS_DLY by one, and repeat step 1, else continue to step 2.
2. Read byte address 0x20000000. If the content is different from 0x00, increment MEMCTRL_DQS_DLY[0].DQS_DLY by one, and repeat step 2, else continue to step 3.
3. Decrement MEMCTRL_DQS_DLY[0].DQS_DLY by three.

After the last step, which configures the appropriate DSQ read window, the memory is ready for random read and write operations. For more information about ECC or 1 gigabyte support, see [Section 3.8.2.6, ECC and 1 Gigabyte Support](#).

Calibration of Two Byte Lanes

This procedure applies to systems where MEMCTRL_CFG.DDR_WIDTH is 1:

- Write 0x0000FFFF to SBA address 0x20000000.
- Write 0x00000000 to SBA address 0x20000004
- Set the MEMCTRL_DQS_DLY[0].DQS_DLY field to 0.
- Set the MEMCTRL_DQS_DLY[1].DQS_DLY field to 0.

Perform the following steps to calibrate the read window. Do not increment the DQS_DLY fields beyond their maximum value. If a DQS_DLY maximum value is exceeded, it is an indication something is incorrect, and the memory controller will not be operational.

1. Read byte addresses 0x20000000 and 0x20000001 from memory.

If the contents of byte address 0x20000000 is different from 0xFF, increment MEMCTRL_DQS_DLY[0].DQS_DLY by one.

If the contents of byte address 0x20000001 is different from 0xFF, increment MEMCTRL_DQS_DLY[1].DQS_DLY by one.

If any DQS_DLY was incremented repeat step 1, else continue to step 2.

2. Read byte addresses 0x20000000 and 0x20000001 from memory.

If the contents of byte address 0x20000000 is different from 0x00, increment MEMCTRL_DQS_DLY[0].DQS_DLY by one.

If the contents of byte address 0x20000001 is different from 0x00, increment MEMCTRL_DQS_DLY[1].DQS_DLY by one.

If any DQS_DLY was incremented repeat step 2, else continue to step 3.

3. Decrement MEMCTRL_DQS_DLY[0].DQS_DLY by three, and decrement MEMCTRL_DQS_DLY[1] by three.

The last step configures the appropriate DSQ read windows. After this step the memory is ready for random read and write operations. If ECC or 1 gigabyte is needed see [Section 3.8.2.6, ECC and 1 Gigabyte Support](#).

3.8.2.6 ECC and 1 Gigabyte Support

The memory controller supports ECC, allowing correction of single bit and detection of two bit errors in every byte that is read from memory.

When ECC is enabled, one byte lane is used for data words and the other is used for ECC check-word. After enabling ECC, the amount of memory available to software is half of the physical memory (the lower half of the memory region).

ECC support must be enabled just after read calibration by setting MEMCTRL_CFG.DDR_ECC_ENA = 1 and at the same time setting MEMCTRL_CFG.BURST_SIZE = 0.

When ECC is enabled, the memory controller automatically corrects single-bit errors and detects two-bit errors. It is possible to propagate these events to the SBA as an error. Non-correctable error indications are forwarded when MEMCTRL_CFG.DDR_ECC_ERR_ENA = 1. Correctable “errors” are forwarded when MEMCTRL_CFG.DDR_ECC_COR_ENA = 1. The VCore CPU can generate interrupt on these events so that software can take the appropriate action.

The memory controller supports up to 1 gigabyte of physical memory. To access more than 512 megabytes from the VCore CPU, set MEMCTRL_CFG.DDR_512MBYTE_PLUS = 1 to change the SBA memory map. For more information, see [Section 3.3, Shared Bus](#). When ECC is enabled with 1 gigabyte of physical memory, MEMCTRL_CFG.DDR_512MBYTE_PLUS does not have to be set, because only half of the physical memory is available to software.

3.8.3 TIMERS

This section provides information about the timers. The following table lists the registers associated with timers.

TABLE 3-24: TIMER REGISTERS

Register	Description
ICPU_CFG::TIMER_CTRL	Enable/disable timer
ICPU_CFG::TIMER_VALUE	Current timer value
ICPU_CFG::TIMER_RELOAD_VALUE	Value to load when wrapping
ICPU_CFG::TIMER_TICK_DIV	Common timer-tick divider

There are three decrementing 32-bit timers in the VCore system that run from a common divider. The common divider is driven by a fixed 156.25 MHz clock and can generate timer ticks from 0.1 μ s (10 MHz) to 1 ms (1 kHz), configurable through TIMER_TICK_DIV. The default timer tick is 100 μ s (10 kHz).

Software can access each timer value through the TIMER_VALUE[n] registers. These can be read or written at any time, even when the timers are active.

When a timer is enabled through TIMER_CTRL[n].TIMER_ENA, it decrements from the current value until it reaches zero. An attempt to decrement a TIMER_VALUE[n] of zero generates interrupt and assigns TIMER_VALUE[n] to the contents of TIMER_RELOAD_VALUE[n]. Interrupts generated by the timers are sent to the VCore interrupt controller. From here, interrupts can be forwarded to the VCore CPU or to an external CPU. For more information, see [Section 3.8.11, Interrupt Controller](#).

By setting TIMER_CTRL[n].ONE_SHOT_ENA, the timer disables itself after generating one interrupt. By default, timers will decrement, interrupt, and reload indefinitely (or until disabled by clearing TIMER_CTRL[n].TIMER_ENA).

A timer can be reloaded from TIMER_RELOAD_VALUE[n] at the same time as it is enabled by setting TIMER_CTRL[n].FORCE_RELOAD and TIMER_CTRL[n].TIMER_ENA at the same time.

Example Configure Timer0 So That It Interrupts Every 1 ms. With the default timer tick of 100 μ s ten timer ticks are needed for a timer that wraps every 1 ms. Configure TIMER_RELOAD_VALUE[0] to 0x9, then enable the timer and force a reload by setting TIMER_CTRL[0].TIMER_ENA and TIMER_CTRL[0].FORCE_RELOAD at the same time.

3.8.4 UARTS

This section provides information about the UART (Universal Asynchronous Receiver/Transmitter) controllers. There are two independent UART controller instances in the VCore System: UART and UART2. These instances are identical copies and anywhere in this description the word UART can be replaced by UART2.

The following table lists the registers associated with the UART.


TABLE 3-25: UART REGISTERS

Register	Description
UART::RBR_THR	Receive buffer/transmit buffer/Divisor (low).
UART::IER	Interrupt enable/Divisor (high).

TABLE 3-25: UART REGISTERS (CONTINUED)

Register	Description
UART::IIR_FCR	Interrupt identification/FIFO control.
UART::LCR	Line control.
UART::MCR	Modem control.
UART::LSR	Line status.
UART::MSR	Modem status.
UART::SCR	Scratchpad.
UART::USR	UART status.

The VCore system UART is functionally based on the industry-standard 16550 UART (RS232 protocol). This implementation features a 16-byte receive and a 16-byte transmit FIFO.

FIGURE 3-21: UART TIMING

The number of data-bits, parity, parity-polarity, and stop-bit length are all programmable using LCR.

The UART pins on the devices are overlaid functions on the GPIO interface. Before enabling the UART, the VCore CPU must enable overlaid modes for the appropriate GPIO pins. For more information about enabling the overlaid functionality of the GPIO pins, see [Section 7.2.4, General-Purpose Inputs and Outputs](#).

The following table lists the pins of the UART interfaces.

TABLE 3-26: UART INTERFACE PINS

Pin Name	I/O	Description
UART_RX/GPIO_27	I	UART receive data
UART_TX/GPIO_26	O	UART transmit data
UART2_RX/GPIO_14	I	UART2 receive data
UART2_TX/GPIO_13	O	UART2 transmit data

The baud rate of the UART is derived from the VCore system frequency. The divider value is indirectly set through the RBR_THR and IER registers. The baud rate is equal to the VCore system clock frequency, divided by 16, and multiplied by the value of the baud rate divisor. A divider of zero disables the baud rate generator and no serial communications occur. The default value for the divisor register is zero.

Example Configuring a baud rate of 9600 in a 208.33 MHz VCore System. To generate a baud rate of 9600, the divisor register must be set to 0x54C ($208.33 \text{ MHz} / (16 \times 9600 \text{ Hz})$). Set LCR.DLAB and write 0x4C to RBR_THR and 0x05 to IER (this assumes that the UART is not in use). Finally, clear LCR.DLAB to change the RBR_THR and IER registers back to the normal mode.

By default, the FIFO mode of the UART is disabled. Enabling the 16-byte receive and 16-byte transmit FIFOs (through IIR_FCR) is recommended.

Note Although the UART itself supports RTS and CTS, these signals are not available on the pins of the device.

3.8.4.1 UART Interrupt

The UART can generate interrupt whenever any of the following prioritized events are enabled (through IER):

- Receiver error
- Receiver data available
- Character timeout (in FIFO mode only)
- Transmit FIFO empty or at or below threshold (in programmable THRE interrupt mode)

When an interrupt occurs, the IIR_FCR register can be accessed to determine the source of the interrupt. Note that the IIR_FCR register has different purposes when reading or writing. When reading, the interrupt status is available in bits 0 through 3. For more information about interrupts and how to handle them, see the IIR_FCR register description.

Example Enabling Interrupt When Transmit FIFO is Below One-Quarter Full. To get this type of interrupt, the THRE interrupt must be used. First, configure TX FIFO interrupt level to one-quarter full by setting IIR_FCR.TET to 10; at the same time, ensure that the IIR_FCR.FIFOE field is also set. Set IER.PTIME to enable the THRE interrupt in the UART. In addition, the VCore interrupt controller must be configured for the CPU to be interrupted. For more information, see [Section 3.8.11, Interrupt Controller](#).

3.8.5 TWO-WIRE SERIAL INTERFACE

This section provides information about the functions of the two-wire serial interface controller.

The following table lists the registers associated with the two-wire serial interface.

TABLE 3-27: TWO-WIRE SERIAL INTERFACE REGISTERS

Register	Description
TWI::CFG	General configuration.
TWI::TAR	Target address.
TWI::SAR	Slave address.
TWI::DATA_CMD	Receive/transmit buffer and command.
TWI::SS_SCL_HCNT	Standard speed high time clock divider.
TWI::SS_SCL_LCNT	Standard speed low time clock divider.
TWI::FS_SCL_HCNT	Fast speed high time clock divider.
TWI::FS_SCL_LCNT	Fast speed low time clock divider.
TWI::INTR_STAT	Masked interrupt status.
TWI::INTR_MASK	Interrupt mask register.
TWI::RAW_INTR_STAT	Unmasked interrupt status.
TWI::RX_TL	Receive FIFO threshold for RX_FULL interrupt.
TWI::TX_TL	Transmit FIFO threshold for TX_EMPTY interrupt.
TWI::CLR_*	Individual CLR_* registers are used for clearing specific interrupts. See register descriptions for corresponding interrupts.
TWI::CTRL	Control register.
TWI::STAT	Status register.
TWI::TXFLR	Current transmit FIFO level.
TWI::RXFLR	Current receive FIFO level.
TWI::TX_ABRT_SOURCE	Arbitration sources.
TWI::SDA_SETUP	Data delay clock divider.
TWI::ACK_GEN_CALL	Acknowledge of general call.
TWI::ENABLE_STATUS	General two-wire serial controller status.
ICPU_CFG::TWI_CONFIG	Configuration of SDA hold-delay.
ICPU_CFG::TWI_SPIKE_FILTER_CFG	Configuration of SDA spike filter.

The two-wire serial interface controller is compatible with the industry standard two-wire serial interface protocol. The controller supports standard speed up to 100 kbps and fast speed up to 400 kbps. Multiple bus masters, as well as both 7-bit and 10-bit addressing, are also supported.

By default, the two-wire serial interface controller operates as master only (CFG.MASTER_ENA), however, slave mode can be enabled (CFG.SLAVE_DIS). In slave mode, the controller generates an interrupt when addressed by an external master. For read requests, the controller then halts the two-wire serial bus until the VCore CPU has processed the request and provided a response (reply-data) to the controller. The slave addresses (SAR) of the two-wire serial inter-

face controller must be unique on the two-wire serial interface bus. This must be configured before enabling slave mode. For information about addresses that have a special meaning on the bus, see [Section 3.8.5.1, Two-Wire Serial Interface Addressing](#).

The two-wire serial interface pins on the devices are overlaid functions on the GPIO interface. Before enabling the two-wire serial interface, the VCore CPU must enable overlaid functions for the appropriate GPIO pins. For more information, see [Section 7.2.4, General-Purpose Inputs and Outputs](#).

The following table lists the pins of the two-wire serial interface.

TABLE 3-28: TWO-WIRE SERIAL INTERFACE PINS

Pin Name	I/O	Description
TWI_SCL/GPIO_6	I/O	Two-wire serial interface clock, open-collector output.
TWI_SDA/GPIO_7	I/O	Two-wire serial interface data, open-collector output.
TWI_SCL_Mn/GPIO	I/O	Two-wire serial interface multiplexed clocks (11 instances in total), open-collector outputs. Overlaid functions on GPIO_11 through GPIO_21.


Setting CTRL.ENABLE enables the controller. The controller can be disabled by clearing the CTRL.ENABLE field, there is a chance that disabling is not allowed (at the time when it is attempted); the ENABLE_STATUS register shows if the controller was successful disabled.

Before enabling the controller, the user must decide on either standard or fast mode (CFG.SPEED) and configure clock dividers for generating the correct timing (SS_SCL_HCNT, SS_SCL_LCNT, FS_SCL_HCNT, FS_SCL_LCNT, and SDA_SETUP). The configuration of the divider registers depends on the VCore system clock frequency. The register descriptions explain how to calculate the required values.

Some two-wire serial devices require a hold time on SDA after SCK when transmitting from the two-wire serial interface controller. The device supports a configurable hold delay through the TWI_CONFIG register.

The two-wire serial interface controller has an 8-byte combined receive and transmit FIFO.

FIGURE 3-22: TWO-WIRE SERIAL INTERFACE TIMING FOR 7-BIT ADDRESS ACCESS



During normal operation of the two-wire serial interface controller, the STATUS register shows the activity and FIFO states.

3.8.5.1 Two-Wire Serial Interface Addressing

To configure either 7-bit or 10-bit addressing for master and slave modes, use CFG.MASTER_10BITADDR and CFG.SLAVE_10BITADDR, respectively.

There are a number of reserved two-wire serial interface addresses. The two-wire serial interface controller does not restrict the use of these. However, if they are used out of context, there may be compatibility issues with other two-wire serial devices. The following table lists the two-wire serial interface reserved addresses.

TABLE 3-29: RESERVED TWO-WIRE SERIAL INTERFACE ADDRESSES

Register Address	Description
0000 000	General Call address/START Byte. If the slave is enabled the two-wire serial interface controller places the data in the receive buffer and issues a general call interrupt. The acknowledge response is configurable (through ACK_GEN_CALL).
0000 001	CBUS address. The two-wire serial interface controller ignores this address.
0000 01X	Reserved, do not use.

TABLE 3-29: RESERVED TWO-WIRE SERIAL INTERFACE ADDRESSES (CONTINUED)

Register Address	Description
0000 1XX	Reserved, do not use.
1111 1XX	Reserved, do not use.
1111 0XX	10-bit addressing indication, 7-bit address devices must not use this.

The two-wire serial interface controller can generate both General Call and START Byte. Initiate this through TAR.GC_OR_START_ENA or TAR.GC_OR_START. When operating as master, the target/slave address is configured using the TAR register.

3.8.5.2 Two-Wire Serial Interface Interrupt

The two-wire serial interface controller can generate a multitude of interrupts. All of these are described in the RAW_INTR_STAT register. The RAW_INTR_STAT register contains interrupt fields that are always set when their trigger conditions occur. The INTR_MASK register is used for masking interrupts and allowing interrupts to propagate to the INTR_STAT register. When set in the INTR_STAT register, the two-wire serial interface controller asserts interrupt toward the VCore interrupt controller.

The RAW_INTR_STAT register also specifies what is required to clear the specific interrupts. When the source of the interrupt is removed, reading the appropriate CLR_* register (for example, CLR_RX_OVER) clears the interrupt.

3.8.5.3 Built-in Two-Wire Serial Multiplexer

The VSC7416-01 device has built-in support for connecting to multiple two-wire serial devices that use the same two-wire serial address. This is done by using the multiplexed clock outputs (SCL_Mn) for the two-wire serial devices rather than the SCL. Depending on which device or devices it needs to talk to, software can then enable/disable the various clocks.

From the two-wire serial controllers point of view, it does not know if it is using SCL or SCL_Mn clocks. When using multiplexed mode, software needs to enable/disable individual SCL_Mn connections before initiating the access operation using the two-wire serial controller. Feedback on the clock (from slow two-wire serial devices) is automatically done for the SCL_Mn lines that are enabled.

To enable multiplexed clocks, first configure the TWI_SCL_Mn overlaid mode in the GPIO controller. Individual SCL_Mn clocks are then enabled by writing 1 to the corresponding GPIO output bit (in DEVCPU_GCB::GPIO_OUT)/. To disable the clock, write 0 to the GPIO output bit. Disabled TWI_SCL_Mn clocks are not driven during access and the feedback is disabled.

Note In multiprocessor systems, the DEVCPU_GCB::GPIO_OUT_SET and DEVCPU_GCB::GPIO_OUT_CLR registers can be used to avoid race conditions.

3.8.6 MII MANAGEMENT CONTROLLER

This section provides information about the MII Management controllers. The following table lists the registers associated with the MII Management controllers.

TABLE 3-30: MIIM REGISTERS

Register	Description
DEVCPU_GCB::MII_STATUS	Controller status
DEVCPU_GCB::MII_CMD	Command and write data
DEVCPU_GCB::MII_DATA	Read data
DEVCPU_GCB::MII_CFG	Clock frequency configuration
DEVCPU_GCB::MII_SCAN_0	Auto-scan address range
DEVCPU_GCB::MII_SCAN_1	Auto-scan mask and expects
DEVCPU_GCB::MII_SCAN_LAST_RSLTS	Auto-scan result
DEVCPU_GCB::MII_SCAN_LAST_RSLTS_VLD	Auto-scan result valid
DEVCPU_GCB::MII_SCAN_RSLTS_STICKY	Differences in expected versus read auto-scan

The devices contain two MIIM controllers with equal functionality. Data is transferred on the MIIM interface using the Management Frame Format protocol specified in IEEE 802.3, Clause 22 or the MDIO Manageable Device protocol defined in IEEE 802.3, Clause 45. The Clause 45 protocol differs from the Clause 22 protocol by using indirect register access to increase the address range. The controller supports both Clause 22 and Clause 45.


The MIIM interface pins are overlaid functions on the GPIO interface. Before enabling the MIIM controllers, the overlaid functions for the appropriate GPIO pins must first be enabled. For more information, see [Section 7.2.4, General-Purpose Inputs and Outputs](#).

TABLE 3-31: MIIM MANAGEMENT CONTROLLER PINS

Pin Name	I/O	Description
MDC0/GPIO_11	O	MIIM clock for controller 0
MDIO0/GPIO_12	I/O	MIIM data input/output for controller 0
MDC1/GPIO_9	O	MIIM clock for controller 1
MDIO1/GPIO_10	I/O	MIIM data input/output for controller 1

The MDIO signal is changed or sampled on the falling edge of the MDC clock by the controller. The MDIO pin is tri-stated in-between access and when expecting read data.

FIGURE 3-23: MII MANAGEMENT TIMING



3.8.6.1 Clock Configuration

The frequency of the management interface clock generated by the MIIM controller is derived from the VCore system frequency. The MIIM clock frequency is configurable and is selected with `MII_CFG.MIIM_CFG_PRESCALE`. The calculation of the resulting frequency is explained in the register description for `MII_CFG.MIIM_CFG_PRESCALE`. The maximum frequency of the MIIM clock is 25 MHz.

3.8.6.2 MII Management PHY Access

Reads and writes across the MII management interface are performed through the `MII_CMD` register. Details of the operation, such as the PHY address, the register address of the PHY to be accessed, the operation to perform on the register (for example, read or write), and write data (for write operations), are set in the `MII_CMD` register. When the appropriate fields of `MII_CMD` are set, the operation is initiated by writing 0x1 to `MII_CMD.MIIM_CMD_VLD`. The register is automatically cleared when the MIIM command is initiated. When initiating single MIIM commands, `MII_CMD.MIIM_CMD_SCAN` must be set to 0x0.

When an operation is initiated, the current status of the operation can be read in `MII_STATUS`. The fields `MII_STATUS.MIIM_STAT_PENDING_RD` and `MII_STATUS.MIIM_STAT_PENDING_WR` can be used to poll for completion of the operation. For a read operation, the read data is available in `MII_DATA.MIIM_DATA_RDDATA` after completion of the operation. The value of `MII_DATA.MIIM_DATA_RDDATA` is only valid if `MII_DATA.MIIM_DATA_SUCCESS` indicates no read errors.

The MIIM controller contains a small command FIFO. Additional MIIM commands can be queued as long as `MII_STATUS.MIIM_STAT_OPR_PEND` is cleared. Care must be taken with read operations, because multiple queued read operations will overwrite `MII_DATA.MIIM_DATA_RDDATA`.

Note A typical software implementation will never queue read operations, because the software needs read data before progressing the state of the software. In this case, `MII_STATUS.MIIM_STAT_OPR_PEND` is checked before issuing MIIM read or write commands, for read-operations `MII_STATUS.MIIM_STAT_BUSY` is checked before returning read result.

By default, the MIIM controller operates in Clause 22 mode. To access Clause 45 compatible PHYs, MII_CFG.MIIM_ST_CFG_FIELD and MII_CMD.MIIM_CMD_OPR_FIELD must be set according to Clause 45 mode of operation.

3.8.6.3 PHY Scanning

The MIIM controller can be configured to continuously read certain PHY registers and detect if the read value is different from an expected value. If a difference is detected, a special sticky bit register is set or a CPU interrupt is generated, or both. For example, the controller can be programmed to read the status registers of one or more PHYs and detect if the Link Status changed since the sticky register was last read.

The reading of the PHYs is performed sequentially with the low and high PHY numbers specified in MII_SCAN_0 as range bounds. The accessed address within each of the PHYs is specified in MII_CMD.MIIM_CMD_REGAD. The scanning begins when a 0x1 is written to MII_CMD.MIIM_CMD_SCAN and a read operation is specified in MII_CMD.MIIM_CMD_OPR_FIELD. Setting MII_CMD.MIIM_CMD_SINGLE_SCAN stops the scanning after all PHYs have been scanned one time. The remaining fields of MII_CMD register are not used when scanning is enabled.

The expected value for the PHY register is set in MII_SCAN_1.MIIM_SCAN_EXPECT. The expected value is compared to the read value after applying the mask set in MII_SCAN_1.MIIM_SCAN_MASK. To “don’t care” a bit-position, write a 0 to the mask. If the expected value for a bit position differs from the read value during scanning, and the mask register has a 1 for the corresponding bit, a mismatch for the PHY is registered.

The scan results from the most recent scan can be read in MII_SCAN_LAST_RSLTS. The register contains one bit for each of the possible 32 PHYs. A mismatch during scanning is indicated by a 0. MII_SCAN_LAST_RSLTS_VLD will indicate for each PHY if the read operation performed during the scan was successful. The sticky-bit register MII_SCAN_RSLTS_STICKY has the mismatch bit set for all PHYs that had a mismatch during scanning since the last read of the sticky-bit register. When the register is read, its value is reset to all-ones (no mismatches).

3.8.6.4 MII Management Interrupt

The MII management controllers can generate interrupts during PHY scanning. Each MII management controller has a separate interrupt signal to the interrupt controller. Interrupt is asserted when one or more PHYs have a mismatch during scan. The interrupt is cleared by reading the MII_SCAN_RSLTS_STICKY register, which resets all MII_SCAN_RSLTS_STICKY indications.

3.8.7 GPIO CONTROLLER

This section provides information about the use of GPIO pins. Many of the GPIO pins have overlaid functions. For more information about the alternative functions and how to enable them, see [Section 7.2.4, General-Purpose Inputs and Outputs](#).

The following table lists the registers associated with GPIO.

TABLE 3-32: GPIO REGISTERS

Register	Description
DEVCPU_GCB::GPIO_OUT	Value to drive on GPIO outputs
DEVCPU_GCB::GPIO_OUT_SET	Atomic set of bits in GPIO_OUT
DEVCPU_GCB::GPIO_OUT_CLR	Atomic clear of bits in GPIO_OUT
DEVCPU_GCB::GPIO_IN	Current value on the GPIO pins
DEVCPU_GCB::GPIO_OE	Enable of GPIO output mode (drive GPIOs)
DEVCPU_GCB::GPIO_ALT	Enable of overlaid GPIO functions
DEVCPU_GCB::GPIO_INTR	Interrupt on changed GPIO value
DEVCPU_GCB::GPIO_INTR_ENA	Enable interrupt on changed GPIO value
DEVCPU_GCB::GPIO_INTR_IDENT	Currently interrupting sources

The GPIO pins are individually programmable. GPIOs are inputs by default and can be individually changed to outputs through GPIO_OE. The value of the GPIO pins is reflected in the GPIO_IN register. GPIOs that are in output mode are driven to the value specified in GPIO_OUT.

In a system where multiple different CPU threads (or different CPUs) may work on the GPIOs at the same time, the GPIO_OUT_SET and GPIO_OUT_CLR registers provide a way for each thread to safely control the output value of individual GPIOs, without having to implement locked regions and semaphores.

The GPIO_ALT registers are only reset by external reset to the device. This means that software reset of the DEVCPU_GCB is possible without affecting the mapping of overlaid functions on the GPIOs.

3.8.7.1 GPIO Interrupt

The GPIO controller continually monitors all inputs and set bits in the GPIO_INTR register whenever a GPIO changes its input value. By enabling specific GPIO pins in the GPIO_INTR_ENA register, a change indication from GPIO_INTR is allowed to propagate (as GPIO interrupt) from the GPIO controller to the VCore Interrupt Controller.

The currently interrupting sources can be read from GPIO_INTR_IDENT, this register is the result of a binary AND between the GPIO_INTR and GPIO_INTR_ENA registers.

3.8.8 SERIAL GPIO CONTROLLER

The VSC7416-01 device features a serial GPIO controller (SIO). By using a serial interface, the SIO controller significantly extends the number of available GPIOs with a minimum number of additional pins on the device. The primary purpose of the SIO controller is to connect control signals from SFP modules and to act as an LED controller.

The SIO controller supports up to 128 serial GPIOs (SGPIOs) organized into 32 ports, with four SGPIOs per port.

The following table lists the registers associated with the serial GPIO.

TABLE 3-33: SIO REGISTERS

Register	Description
DEVCPU_GCB::SIO_INPUT_DATA	Input data
DEVCPU_GCB::SIO_INT_POL	Interrupt polarity
DEVCPU_GCB::SIO_PORT_INT_ENA	Interrupt enable
DEVCPU_GCB::SIO_PORT_CONFIG	Output port configuration
DEVCPU_GCB::SIO_PORT_ENABLE	Port enable
DEVCPU_GCB::SIO_CONFIG	General configuration
DEVCPU_GCB::SIO_CLOCK	Clock configuration
DEVCPU_GCB::SIO_INT_REG	Interrupt register

The following table lists the pins of the SIO controller. The pins of the SIO controller are overlaid on GPIO pins. For more information about enabling the overlaid functionality of the GPIO pins, see [Section 7.2.4, General-Purpose Inputs and Outputs](#).

TABLE 3-34: SIO CONTROLLER PINS

Pin Name	I/O	Description
SIO_CLK/GPIO_0	O	SIO clock output, frequency is configurable using SIO_CLOCK.SIO_CLK_FREQ.
SIO_LD/GPIO_1	O	SIO load data, polarity is configurable using SIO_CONFIG.SIO_LD_POLARITY.
SIO_DO/GPIO_2	O	SIO data output.
SIO_DI/GPIO_3	I	SIO data input.

The SIO controller works by shifting SGPIO values out on SIO_DO through a chain of shift registers on the PCB. After shifting a configurable number of SGPIO bits, the SIO controller asserts SIO_LD, which causes the shift registers to apply the values of the shifted bits to outputs. The SIO controller can also read inputs while shifting out SGPIO values on SIO_DO. It also samples the SIO_DI input. The values sampled on SIO_DI are made available to software.

If the SIO controller is only used for outputs, the use of the load signal is optional. If the load signal is omitted, simpler shift registers (without load) can be used, however, the outputs of these registers will toggle during shifting.

When driving LED outputs, it is acceptable that the outputs will toggle when SGPIO values are updated (shifted through the chain). When the shift frequency is fast, the human eye is not able to see the shifting through the LEDs.


The number of shift registers in the chain is configurable. The SIO controller allows enabling of individual ports through SIO_PORT_ENABLE; only enabled ports are shifted out on SI_DO. Ports that are not enabled are skipped during shifting of GPIO values.

Note SIO_PORT_ENABLE allows skipping of ports in the SGPIO output stream that are not in use. The number of GPIOs per (enabled) port is configurable as well, through SIO_CONFIG.SIO_PORT_WIDTH this can be set to 1,2,3, or 4 bits. The number of bits per port is common for all enabled ports, so the number of shift registers on the PCB must be equal to the number of enabled ports times the number of SGPIOs per port.

Enabling of ports and configuration of SGPIOs per port applies to both output mode and input mode. Unlike a regular GPIO port, a single SGPIO position can be used both as output and input. That is, software can control the output of the shift register AND read the input value at the same time. Using SGPIOs as inputs requires load-capable shift registers.

Regular shift registers and load-capable shift-registers can be mixed, which is useful when driving LED indications for integrated PHYs while supporting reading of link status from SFP modules, for example.

FIGURE 3-24: SIO TIMING




The SGPIO values are output in bursts followed by assertion of the SIO_LD signal. Values can be output as a single burst or as continuous bursts separated by a configurable burst gap. The maximum length of a burst is 32×4 data cycles. The burst gap is configurable in steps of approximately 1 ms, from 0 ms to 33 ms through SIO_CONFIG.SIO_BURST_GAP_DIS and SIO_CONFIG.SIO_BURST_GAP.

A single burst is issued by setting SIO_CONFIG.SIO_SINGLE_SHOT. The field is automatically cleared by hardware when the burst is finished. To issue continuous bursts, set SIO_CONFIG.SIO_AUTO_REPEAT. The SIO controller continues to issue bursts until SIO_CONFIG.SIO_AUTO_REPEAT is cleared.

SGPIO output values are configured in SIO_PORT_CONFIG.BIT_SOURCE. The input value is available in SIO_INPUT_DATA.S_IN.

The following illustration shows what happens when the number of SGPIOs per port is configured to 2 (through SIO_CONFIG.SIO_PORT_WIDTH). Disabling of ports (through SIO_PORT_ENABLE) is handled in the same way as disabling the SGPIO ports.

FIGURE 3-25: SIO TIMING WITH SGPIOS DISABLED



The frequency of the SIO_CLK clock output is configured through SIO_CLOCK.SIO_CLK_FREQ. The SIO_LD output is asserted after each burst, this output is asserted for a period of 25 ns to 30 ns. The polarity of SIO_LD is configurable through SIO_CONFIG.SIO_LD_POLARITY.


The SIO_LD output can be used to ensure that outputs are stable when serial data is being shifted through the registers. This can be done by using the SIO_LD output to shift the output values into serial-to-parallel registers after the burst is completed. If serial-to-parallel registers are not used, the outputs will toggle while the burst is being shifted through the chain of shift registers. A universal serial-to-parallel shift register outputs the data on a positive-edge load signal, and a universal parallel-to-serial shift register shifts data when the load pin is high, so one common load signal can be used for both input and output serial ↔ parallel conversion.

The assertion of SIO_LD happens after the burst to ensure that after power up, the single burst will result in well-defined output registers. Consequently, to sample input values one time, two consecutive bursts must be issued. The first burst results in the input values being sampled by the serial-to-parallel registers, and the second burst shifts the input values into the SIO controller.

The port order required in the serial bitstream depends on the physical layout of the shift register chain. Often the input and output port orders must be opposite in the serial streams. The port order of the input and output bitstream is independently configurable in SIO_CONFIG.SIO_REVERSE_INPUT and SIO_CONFIG.SIO_REVERSE_OUTPUT.

The following illustration shows the port order.

FIGURE 3-26: SGPIO OUTPUT ORDER



3.8.8.1 Output Modes

The output mode of each SGPIO can be individually configured in SIO_PORT_CONFIG.BIT_SOURCE. The SIO controller features three output modes:

- Static
- Blink
- Link activity

Static Mode The static mode is used to assign a fixed value to the SGPIO, for example, fixed 0 or fixed 1.

Blink Mode The blink mode makes the SGPIO blink at a fixed rate. The SIO controller features two blink modes that can be set independently. A SGPIO can then be configured to use either blink mode 0 or blink mode 1. The blink outputs are configured in SIO_CONFIG.SIO_BMODE_0 and SIO_CONFIG.SIO_BMODE_1. To synchronize the blink modes between different devices, reset the blink counter using SIO_CONFIG.SIO_BLINK_RESET. The burst toggle mode of blink mode 1 toggles the output with every burst.

TABLE 3-35: BLINK MODES

Register	Description
Blink Mode 0	0: 20 Hz blink frequency 1: 10 Hz blink frequency 2: 5 Hz blink frequency 3: 2.5 Hz blink frequency

TABLE 3-35: BLINK MODES


Register	Description
Blink Mode 1	0: 20 Hz blink frequency 1: 10 Hz blink frequency 2: 5 Hz blink frequency 3: Burst toggle

Link Activity Mode The link activity mode makes the output blink when there is activity on the port module (Rx or Tx). The mapping between SIO port number port module number is 1:1, For example, port 0 is connected to port module 0, port 1 is connected to port module 1, and so on.

The link activity mode uses an envelope signal to gate the selected blinking pattern (blink mode 0 or blink mode 1). When the envelope signal is asserted, the output blinks, and when the envelope pattern is de-asserted, the output is turned off. To ensure that even a single packet makes a visual blink, an activity pulse from the port module is extended to minimum 100 ms. If another packet is sent while the envelope signal is asserted, the activity pulse is extended by another 100 ms. The polarity of the link activity modes can be set in SIO_PORT_CONFIG.BIT_SOURCE.

The following illustration shows the link activity timing.

FIGURE 3-27: LINK ACTIVITY TIMING



3.8.8.2 SIO Interrupt

The SIO controller can generate interrupts based on the value of the input value of the SGPIOs. All interrupts are level sensitive.

Interrupts are enabled using the two registers. Interrupts can be individually enabled for each port in SIO_PORT_INT_ENA.INT_ENA (32 bits) and in SIO_CONFIG.SIO_INT_ENA (4 bits) interrupts are enabled for the four inputs per port. In other words, SIO_CONFIG.SIO_INT_ENA is common for all 32 ports. The polarity of interrupts is configured for each SGPIO in SIO_INT_POL.

The SIO controller has one interrupt output connected to the main interrupt controller, which is asserted when one or more interrupts are active. To determine which SGPIO is causing the interrupt, the CPU must read the sticky bit interrupt register SIO_INT_REG. The register has one bit per SGPIO and can only be cleared by software. A bit is cleared by writing a 1 to the bit position. The interrupt output remains high until all interrupts in SIO_INT_REG are cleared.

3.8.8.3 Loss of Signal Detection

The SIO controller can propagate loss of signal detection inputs directly to the signal detection input of the port modules. This is useful when, for example, SFP modules are connected to the device. The mapping between SIO ports and port modules is the same as for the link activity outputs; port 0 is connected to port module 0, port 1 is connected to port module 1, and so on.

The value of SGPIO bit 0 of each SIO port is forwarded directly to the loss of signal input on the corresponding port module. The port module must enable the loss of signal input locally.

Loss of signal can also be taken directly from overlaid functions on the regular GPIOs. In that case, the input from the SIO controller is ignored. For more information, see [Section 7.2.4, General-Purpose Inputs and Outputs](#).

The polarity of the loss of signal input is configured using SIO_INT_POL, meaning the same polarity must be used for loss of signal detect and interrupt.

3.8.9 FAN CONTROLLER

The VSC7416-01 device includes a fan controller that can be used to control and monitor a system fan. A pulse-width-modulation (PWM) output regulates the fan speed. The fan speed is monitored using a TACHO input. This is especially powerful when combined with the internal temperature sensor.

The following table lists the registers associated with the fan controller.

TABLE 3-36: FAN CONTROLLER REGISTERS

Register	Description
DEVCPU_GCB::FAN_CFG	General configuration
DEVCPU_GCB::FAN_CNT	Fan revolutions counter

The following table lists fan controller pins, which are overlaid on GPIOs. For more information about enabling the overlaid functionality of GPIOs, see [Section 7.2.4, General-Purpose Inputs and Outputs](#).

TABLE 3-37: FAN CONTROLLER PINS

Register	I/O	Description
TACHO/GPIO_4	I	TACHO input for counting revolutions
PWM/GPIO_5	O	PWM fan output

The PWM output can be configured to any of the following frequencies in FAN_CFG.PWM_FREQ: 10 Hz, 20 Hz, 40 Hz, 60 Hz, 80 Hz, 100 Hz, 120 Hz, or 25 kHz.

The low frequencies can be used for driving three-wire fans using a FET/transistor. The 25 kHz frequency can be used for four-wire fans that use the PWM input internally to control the fan. The duty cycle of the PWM output is programmable from 0% to 100%, with 8-bit accuracy. The polarity of the output can be controlled by FAN_CFG.INV_POL, so a duty cycle of 100%, for example, can be either always low or always high.

The PWM output pin can be configured to act as a normal output or as an open-collector output, where the output value of the pin is kept low, but the output enable is toggled. The open-collector output mode is enabled by setting FAN_CFG.PWM_OPEN_COL_ENA.

Note By using open-collector mode, it is possible to externally pull-up to a higher voltage than the maximum GPIO I/O supply. The GPIO pins are 5V-tolerable.

The speed of the fan is measured using a 16-bit counter that counts the rising edges on the TACHO input. A fan usually gives one to four pulses per revolution, depending on the fan type. The counter value is available in the FAN_CNT register. Depending on the value of FAN_CFG.FAN_STAT_CFG, the FAN_CNT register is updated in two different ways:

- If FAN_CFG.FAN_STAT_CFG is set, the FAN_CNT register behaves as a 16-bit wrapping counter that shows the total number of ticks on the TACHO input.
- If FAN_CFG.FAN_STAT_CFG is cleared, the FAN_CNT register is updated one time per second with the number of TACHO ticks received during the last second.

Optionally, the TACHO input is gated by the polarity-corrected PWM output by setting FAN_CFG.GATE_ENA, so that only TACHO pulses received while the polarity corrected PWM output is high are counted. Glitches on the TACHO input can occur right after the PWM output goes high. As a result, the gate signal is delayed by 10 μ s when PWM goes high. There is no delay when PWM goes low, and the length of the delay is not configurable. Software must read the counter value in FAN_CNT and calculate the RPM of the fan.

An example of how to calculate the RPM of the fan is if the fan controller is configured to 100 Hz and a 20% duty cycle, each PWM pulse is high in 2 ms and low in 8 ms. If gating is enabled, the gating of the TACHO input is open in 1.99 ms and closed in 8.01 ms. If the fan is turning with 100 RPMs and gives two TACHO pulses per revolution, it will ideally give 200 pulses per minute. TACHO pulses are only counted in 19.99% of the time, so it will give $200 \times 0.1999 = 39.98$ pulses per minute. If the additional 10 μ s gating time is ignored, the counter value is multiplied by 2.5 to get the RPM value, because there is a 20% duty cycle with two TACHO pulses per revolution. By multiplying with 2.5, the RPM value is calculated to 99.95, which is 0.05% off the correct value (due to the 10 μ s gating time).

3.8.10 MEMORY INTEGRITY MONITOR

Soft errors are a result of natural alpha decay, cosmic radiation, or electrical disturbances in the environment in which the device operates and can happen in all integrated circuits. The chance of soft-errors happening in a memory (RAM) is higher than for flip flop based logic, because the memory structures are physically small (requires less outside force to change values in memory cells than in flip flops). The VSC7416-01 device has built-in protection from soft errors by using error correcting code (ECC) on critical memories. In addition, the VSC7416-01 device allows monitoring and reporting of soft-error events.

The following table lists the registers associated with the memory integrity monitor.

TABLE 3-38: INTEGRITY MONITOR REGISTERS

Register	Description
DEVCPU_GCB::MEMITGR_CTRL	Trigger monitor state changes.
DEVCPU_GCB::MEMITGR_STAT	Current state of the monitor and memory status.
DEVCPU_GCB::MEMITGR_INFO	Shows indication when in DETECT state.
DEVCPU_GCB::MEMITGR_IDX	Shows memory index when in DETECT state.
DEVCPU_GCB::MEMITGR_DIV	Monitor speed.

The memory integrity monitor looks for memory soft-error indications. Correctable (single bit) and non-correctable (multi-bit or parity) indications are detected during memory read and can be read out from the controller. Software can read out indications regularly to test if there are non-correctable errors in the memories to which software has to respond.


The memory integrity monitor operates in three different states: IDLE, LISTEN, and DETECT. After a reset, the monitor starts in the IDLE state.

IDLE The monitor is deactivated and in quiet mode. In IDLE mode, the memories still correct, detect, and store indications locally, but they are not able to report indications to the monitor.

LISTEN In the LISTEN state, the monitor looks for indications in the memories. MEMITGR_STAT.INDICATION is set when indications are detected.

DETECT DETECT state is used when indications are read from the memories. It means that a valid indication is available in MEMITGR_INFO and the corresponding memory index in MEMITGR_IDX.

The current state of the monitor is reported in MEMITGR_STAT.MODE_IDLE, MEMITGR_STAT.MODE_DETECT, and MEMITGR_STAT.MODE_LISTEN. Software initiates transitions between states by setting the one-shot MEMITGR_CTRL.ACTIVATE field. It may take some time to transition from one state to the next. The MEMITGR_CTRL.ACTIVATE field is not cleared before the next state is reached (also the MEMITGR_STAT.MODE_BUSY field is set while transitioning between states).

FIGURE 3-28: MONITOR STATE DIAGRAM

The first time after reset that MEMITGR_CTRL.ACTIVATE is set, the monitor resets the detection logic in all the memories and transitions directly from IDLE to LISTEN state.


Before setting MEMITGR_CTRL.ACTIVATE for the first time, speed up the monitor by setting MEMITGR_DIV.MEM_DIV to the value specified in the registers list. The memories of the PCIe MAC are clock-gated and bypassed when PCIe interface is not enabled; enable monitoring of PCIe MAC memories by setting ICPU_CFG::PCIE_CFG.MEM_RING_CORE_ENA. This field must not be set when PCIe interface is not enabled.

To read out indications, first transition from LISTEN to IDLE, then continue transitioning until the LISTEN state is reached. Every time the monitor ends up in the DETECT state, an indication is available in the MEMITGR_INFO and MEMITGR_IDX registers. Each memory stores one indication. Indications are cleared when they are read by way of the monitor. Each indication contains four flags and one memory address.

- The MEM_ERR flag is set when a non-correctable upset is detected and the corresponding address is available in MEM_ADDR.
- The MEM_ERR_OVF flag is set when a non-correctable upset is detected for which address could not be stored.
- The MEM_COR flag is set when a correctable upset is detected and the corresponding address is available in MEM_ADDR.
- The MEM_COR_OVF flag is set when a correctable upset is detected for which address could not be stored.

Information about non-correctable upsets is prioritized over correctable upsets. Address can only be saved when the monitor is in LISTEN mode. The following flowchart shows how the detection logic sets flags and address.

FIGURE 3-29: MEMORY DETECTION LOGIC



If the MEM_ERR_OVF or MEM_COR_OVF flag is set, it means that there is at least one event for which the address could not be stored.

The following table shows ECC enabled memories in the VSC7416-01 device, their index, and the recommended approach for handling indications. If the controller reports an index that is not mentioned in the list, the recommended approach is to reboot the device.

TABLE 3-39: MEMORIES WITH INTEGRITY SUPPORT

Index	Description
1	VCAP IS1 action RAM. A non-correctable event means that an incorrect action was applied to a frame during the VCAP IS1 lookup. (Events can also be generated as a result of reading actions out of memory and into the VCAP cache.) It is not possible to identify and remove the frame from queuing system. The recommended approach is to correct the VCAP action and resume normal operation. There is a 1:1 relation between the event address and row index in the VCAP. For a non-correctable upset the action or actions at row-index corresponding to the event address must be rewritten (a copy of the actions must be kept in program memory). If address is unknown, rewrite all rows in VCAP IS1. For correctable upset, the action can either be rewritten (as described previously) or it can be read into VCAP cache and written back to action memory. The latter approach will generate another correctable event for this memory. In case of unknown address, rewrite or read/write-back all rows in VCAP IS1.
2	VCAP IS1 sticky RAM. A non-correctable event means that an action was triggered for which the corresponding sticky bit was incorrect. (Events can also be generated as a result of reading sticky bits out of memory and into the VCAP cache.) The sticky bits in IS1 are for software debug only. The VCAP overwrites the sticky bit value whenever an action is triggered; the recommended approach is to resume normal operation.
3	VCAP IS2 action RAM. Same as VCAP IS1 action RAM, but for IS2.

TABLE 3-39: MEMORIES WITH INTEGRITY SUPPORT (CONTINUED)

Index	Description
4	VCAP IS2 counter RAM. A non-correctable event means that an incorrect action may have been applied to a frame during the VCAP IS2 lookup. (Events can also be generated as a result of reading counters out of memory and into the VCAP cache.) The counter value is used for HIT_ME_ONCE enabled actions. It is not possible to identify and remove the Frame from queuing system, the recommended approach is to clear the counter value and resume normal operation. There is a 1:1 relation between the event address and row index in the VCAP. For non-correctable upsets, the counters corresponding to the event address must be cleared. In case of an unknown address, clear all counters in VCAP IS2. For correctable upset, the counter can either be cleared or it can be read into VCAP cache and written back to counter memory. The latter approach will generate another correctable event for this memory.
5	VCAP IS0 action RAM. Same as VCAP IS1 action RAM, but for IS0.
6	VCAP IS0 sticky RAM. Same as VCAP IS1 sticky RAM, but for IS0.
16	VCAP ES0 action RAM. Same as VCAP IS1 action RAM, but for ES0.
17	VCAP ES0 sticky RAM. Same as VCAP IS1 sticky RAM, but for ES0.
38	VCore System extraction group RAM. A non-correctable event means that incorrect frame data was provided to an extraction group. If the event address is in the 0 though 7 range, extraction group 0 is affected, else extraction group 1 is affected. The recommended action depends on the owner of the implicated group (see DEVCPU_QS::XTR_GRP_CFG.MODE). VRAP, an erred register operation may have been performed. The recommended approach is to reboot the device. Manual extraction (DEVCPU_QS::XTR_RD), erred frame data was or is about to be extracted. The recommended approach is to discard extraction frames and resume normal operation. Frame DMA, erred frame data was or is about to be extracted. The recommended approach is to discard extraction frames and resume normal operation. Correctable events should be ignored and normal operation resumed.
39	PCIe Master Response Compose RAM. A non-correctable event means that a PCIe reply (to inbound request) with incorrect packet data was transmitted onto the PCIe interface. It is not possible to identify and remove the PCIe packet. An erred reply to an external CPU attached using PCIe may have sent software into an unexpected state. Recommended approach is to reboot the device. Correctable events should be ignored and normal operation resumed.
40	PCIe Retry Buffer RAM. A non-correctable event means that an outbound PCIe package with incorrect packet data was transmitted or re-transmitted onto the PCIe interface. It is not possible to identify and remove the PCIe packet. An erred outbound PCIe package may have sent external software into unexpected state. Recommended approach is to reboot the device. Correctable events should be ignored and normal operation resumed.
41	PCIe Receive Queue Header RAM. A non-correctable event means that the header for an inbound PCIe package was corrupted, which may send the PCIe MAC into unexpected state. Recommended approach is to reboot the device. Correctable events should be ignored and normal operation resumed.
42	PCIe Receive Queue Data RAM. A non-correctable event means that the data for an inbound PCIe package was corrupted, which may send the device into unexpected state. Recommended approach is to reboot the device. Correctable events should be ignored and normal operation resumed.
23, 26,29, and 39	Spurious ECC indications. These memory indices may report ECC indications during normal device operation. Events should be ignored and normal operation resumed.

TABLE 3-39: MEMORIES WITH INTEGRITY SUPPORT (CONTINUED)

Index	Description
Others	For unlisted indexes, the recommended approach is to reboot the device.

The VCore CPU implements parity protection of all VCore cache structures as specified by MIPS architecture. For information about how to enable parity protection and interrupt on parity upsets, see the MIPS documentation.

Reading from uninitialized memory locations has a high plausibility of triggering non-correctable or correctable indications. This is useful when developing integrity monitor software driver. For example, powering up a system without initializing the VCAPs and reading actions and sticky bits will trigger monitor indications. Note that the contents of memories are not changed by device reset, so power cycle is needed to reset the memories.

3.8.11 INTERRUPT CONTROLLER

This section provides information about the VCore interrupt controller.

The following table lists the registers associated with the interrupt controller.

TABLE 3-40: INTERRUPT CONTROLLER REGISTERS

Register	Description
ICPU_CFG::INTR_RAW	Current value of interrupt inputs
ICPU_CFG::INTR_BYPASS	Force non-sticky function
ICPU_CFG::INTR_TRIGGER	Configure edge or level sensitive events
ICPU_CFG::INTR_FORCE	Force events (for software debug)
ICPU_CFG::INTR_STICKY	Currently logged events
ICPU_CFG::INTR_ENA	Enable of interrupt sources
ICPU_CFG::INTR_ENA_CLR	Atomic clear of bits in INTR_ENA
ICPU_CFG::INTR_ENA_SET	Atomic set of bits in INTR_ENA
ICPU_CFG::INTR_IDENT	Currently enabled and interrupting sources
ICPU_CFG::DST_INTR_MAP	Mapping of interrupt sources to destinations
ICPU_CFG::DST_INTR_IDENT	Currently enabled, mapped, and interrupting sources per destination
ICPU_CFG::DEV_INTR_POL	Polarity of module interrupt inputs
ICPU_CFG::DEV_INTR_RAW	Current value of module interrupts
ICPU_CFG::DEV_INTR_BYPASS	Force non-sticky function for module interrupts
ICPU_CFG::DEV_INTR_TRIGGER	Configure edge or level sensitive events for module interrupts
ICPU_CFG::DEV_INTR_STICKY	Currently logged module interrupt events
ICPU_CFG::DEV_ENA	Enable of module interrupts
ICPU_CFG::DEV_IDENT	Currently interrupting and enabled module interrupts
ICPU_CFG::EXT_INTR_DIR	Direction of external interrupt pins
ICPU_CFG::EXT_INTR_POL	Polarity of external interrupts inputs/outputs
ICPU_CFG::EXT_INTR_DRV	Drive mode for external interrupt outputs

The interrupt controller maps interrupt sources from VCore and switch-core blocks, port modules, and external interrupt inputs to four interrupt destinations. Two interrupt destinations are mapped to the VCore CPU, and two can be transmitted from the VSC7416-01 device using the overlaid functions on GPIOs or using PCIe inband interrupt signaling.

The following table lists the available interrupt sources in VSC7416-01.

TABLE 3-41: INTERRUPT SOURCES

Source Name	Description
DEV	Aggregated port module interrupt, this interrupt is asserted if there is an active and enabled interrupt from any of the VSC7416-01 port modules. See Section 3.8.11.3, Port Module Interrupts . This interrupt has bit index 0 in INTR_* and DST_INTR_* registers.
EXT_SRC0	External interrupt source 0. See Section 3.8.11.4, External Interrupts . This interrupt has bit index 1 in INTR_* and DST_INTR_* registers.
EXT_SRC1	External interrupt source 1. See Section 3.8.11.4, External Interrupts . This interrupt has bit index 2 in INTR_* and DST_INTR_* registers.
TIMER0	Timer 0 interrupt. See Section 3.8.3, Timers . This interrupt has bit index 3 in INTR_* and DST_INTR_* registers.
TIMER1	Timer 1 interrupt. See Section 3.8.3, Timers . This interrupt has bit index 4 in INTR_* and DST_INTR_* registers.
TIMER2	Timer 2 interrupt. See Section 3.8.3, Timers . This interrupt has bit index 5 in INTR_* and DST_INTR_* registers.
UART	UART interrupt. See Section 3.8.4.1, UART Interrupt . This interrupt has bit index 6 in INTR_* and DST_INTR_* registers.
UART2	UART2 interrupt. See Section 3.8.4.1, UART Interrupt . This interrupt has bit index 7 in INTR_* and DST_INTR_* registers.
TWI	Two-wire serial interrupt. See Section 3.8.5.2, Two-Wire Serial Interface Interrupt . This interrupt has bit index 8 in INTR_* and DST_INTR_* registers.
SW0	Software interrupt 0. See Section 3.5.5, Mailbox and Semaphores . This interrupt has bit index 9 in INTR_* and DST_INTR_* registers.
SW1	Software interrupt 1. See Section 3.5.5, Mailbox and Semaphores . This interrupt has bit index 10 in INTR_* and DST_INTR_* registers.
SGPIO	Serial GPIO interrupt. See Section 3.8.8.2, SIO Interrupt . This interrupt has bit index 11 in INTR_* and DST_INTR_* registers.
GPIO	Parallel GPIO interrupt. See Section 3.8.7.1, GPIO Interrupt . This interrupt has bit index 12 in INTR_* and DST_INTR_* registers.
MIIM0	MIIM Controller 0 interrupt. See Section 3.8.6.4, MII Management Interrupt . This interrupt has bit index 13 in INTR_* and DST_INTR_* registers.
MIIM1	MIIM Controller 1 interrupt. See Section 3.8.6.4, MII Management Interrupt . This interrupt has bit index 14 in INTR_* and DST_INTR_* registers.
FDMA	Frame DMA interrupt, see Section 3.7.4, FDMA Events and Interrupts . This interrupt has bit index 15 in INTR_* and DST_INTR_* registers.
OAM_MEP	OAM/MEP interrupt, see <TBA>. This interrupt has bit index 16 in INTR_* and DST_INTR_* registers.
PTP_RDY	Timestamp ready interrupt. See Section 2.17.3, Hardware Timestamping Module . This interrupt has bit index 17 in INTR_* and DST_INTR_* registers.
PTP_SYNC	PTP synchronization interrupt. See Section 2.17.2, Master Timer . This interrupt has bit index 18 in INTR_* and DST_INTR_* registers.
XTR_RDY0	Extraction data ready interrupt 0. See Section 2.14.1, Frame Extraction . This interrupt has bit index 20 in INTR_* and DST_INTR_* registers.
XTR_RDY1	Extraction data ready interrupt 1. See Section 2.14.1, Frame Extraction . This interrupt has bit index 21 in INTR_* and DST_INTR_* registers.
INJ_RDY0	Injection ready interrupt 0. See Section 2.14.2, Frame Injection . This interrupt has bit index 22 in INTR_* and DST_INTR_* registers.

TABLE 3-41: INTERRUPT SOURCES (CONTINUED)

Source Name	Description
INJ_RDY1	Injection ready interrupt 1. See Section 2.14.2, Frame Injection . This interrupt has bit index 23 in INTR_* and DST_INTR_* registers.
PCIE	PCIe interrupt. See Section 3.6.6, Power Management . This interrupt has bit index 24 in INTR_* and DST_INTR_* registers.

The following table lists the available interrupt destinations in VSC7416-01.

TABLE 3-42: INTERRUPT DESTINATIONS


Destination Name	Description
CPU0	Interrupt 0 to VCore CPU. This interrupt has replication index 0 in DST_INTR_* registers.
CPU1	Interrupt 1 to VCore CPU. This interrupt has replication index 1 in DST_INTR_* registers.
EXT_DST0	External interrupt destination 0, see Section 3.8.11.4, External Interrupts . This interrupt has replication index 2 in DST_INTR_* registers.
EXT_DST1	External interrupt destination 1, see Section 3.8.11.4, External Interrupts . This interrupt has replication index 3 in DST_INTR_* registers.

All interrupts, events, and indications inside in the interrupt controller are active high. If an interrupt source supports polarity correction, it is applied before going into the interrupt controller. If an interrupt destination supports polarity correction, it is applied after leaving the interrupt controller.

3.8.11.1 Interrupt Source Configuration

Interrupt sources are handled identically inside the interrupt controller. This section describes interrupt source n , which refers to the bit index of that interrupt source in the INTR_* and DST_INTR_* registers. The following illustration shows the logic associated with a single interrupt source.

FIGURE 3-30: INTERRUPT SOURCE LOGIC



The current value of an interrupt source is available in INTR_RAW.INTR_RAW[n]. INTR_STICKY.INTR_STICKY[n] is set when the interrupt controller detects an interrupt. There are two detection methods:

- When INTR_TRIGGER.INTR_TRIGGER[n] is set to level-activated, the interrupt controller continually sets INTR_STICKY.INTR_STICKY[n] for as long as the interrupt source is active.
- When INTR_TRIGGER.INTR_TRIGGER[n] is set to edge-triggered, the interrupt controller only sets INTR_STICKY.INTR_STICKY[n] when the interrupt source changes value.

Software can clear INTR_STICKY.INTR_STICKY[n] by writing 1 to bit n . However, the interrupt controller will immediately set this bit again if the source input is still active (when INTR_TRIGGER is 0) or if it sees a change on the source input (when INTR_TRIGGER is 1).

The interrupt source is enabled in `INTR_ENA.INTR_ENA[n]`. When `INTR_STICKY.INTR_STICKY[n]` is set and the interrupt is enabled, the interrupt is indicated towards the interrupt destinations. For more information, see [Section 3.8.11.2, Interrupt Destination Configuration](#). An active and enabled interrupt source sets `INTR_IDENT.INTR_IDENT[n]`.


On rare occasions it is desirable to bypass the stickiness of interrupt sources and use `INTR_RAW.INTR_RAW[n]` directly instead of `INTR_STICKY.INTR_STICKY[n]`. Set `INTR_BYPASS.INTR_BYPASS[n]` to enable bypass and ignore `INTR_STICKY` and `INTR_TRIGGER` configurations.

Note The bypass function may be useful for some software interrupt handler architectures. It should only be used for interrupt sources that are guaranteed to be sticky in the source block. For example, the GPIO interrupts that are generated from sticky bits in `DEV_CPU_GCB::GPIO_INTR` may be applicable for the bypass mode.

3.8.11.2 Interrupt Destination Configuration

The four interrupt destinations are handled identically in the interrupt controller. This section describes destination `d`, which refers to the replication index of that interrupt in the `DST_INTR_*` registers. The following illustration shows the logic associated with a single interrupt destination.

FIGURE 3-31: INTERRUPT DESTINATION LOGIC




The interrupt destination can enable individual sources for interrupt by writing a mask to `DST_INTR_MAP[d].DST_INTR_MAP`. When a source is enabled in `DST_INTR_MAP` then an interrupt from this source will be propagated to the interrupt destination.

The currently active and enabled source interrupts for a destination can be seen by reading `DST_INTR_IDENT[d].DST_INTR_IDENT`.

3.8.11.3 Port Module Interrupts

Each port module can generate an interrupt. Because there are too many modules to handle the interrupts in parallel with the other source interrupts in the `INTR_*` registers, the port module interrupts are aggregated in a separate source interrupt hierarchy before being presented to the interrupt controller source logic as the `DEV` source interrupt.

FIGURE 3-32: PORT MODULE INTERRUPT LOGIC



The module interrupt polarity is configurable in `DEV_INTR_POL.DEV_INTR_POL[m]`.

`DEV_INTR_RAW`, `DEV_INTR_TRIGGER`, `DEV_INTR_STICKY`, `DEV_INTR_BYPASS`, `DEV_INTR_ENA`, and `DEV_INTR_IDENT` works in the same way as the `INTR_*` counterparts, see [Section 3.8.11.1, Interrupt Source Configuration](#) for more information.

The final step when handling module interrupts is an aggregation of all individual module interrupts to the `DEV` source interrupt.

3.8.11.4 External Interrupts

The interrupt controller support two external source interrupts and two external destination interrupts. The external interrupts can be mapped to GPIO using overlaid functions. For more information about how to enable these pins, see [Section 7.2.4, General-Purpose Inputs and Outputs](#).

Each external interrupt pins can work either as an interrupt source (external devices may generate interrupt inside the VSC7416-01 device) or as an interrupt destination (the VSC7416-01 device may generate interrupt to external devices). The mode of external interrupt pins is configured in EXT_INTR_DIR and the polarity (active high or low) in EXT_INTR_POL.

Note When an external interrupt destination is mapped to PCIe the external interrupts pin can still be enabled and used as interrupt sources.

TABLE 3-43: EXTERNAL INTERRUPT PINS

Register	I/O	Description
IRQ0/GPIO_28	I/O	External Interrupt 0; source/destination mode is configured by EXT_INTR_DIR.EXT_INTR_DIR[0] and polarity by EXT_INTR_POL.EXT_INTR_POL[0].
IRQ1/GPIO_29	I/O	External Interrupt 1; source/destination mode is configured by EXT_INTR_DIR.EXT_INTR_DIR[1] and polarity by EXT_INTR_POL.EXT_INTR_POL[1].

When an external interrupt pin is configured for destination mode (as an output) it is possible to either drive the output permanently or emulate open-collector output.

- To drive, permanently configure EXT_INTR_DRV[e] = 0.
- To emulate open collector output configure EXT_INTR_DRV[e] = 1 and EXT_INTR_POL[e] = 0. To safely enable open-collector output, the EXT_INTR_DRV and EXT_INTR_POL registers must be configured before enabling the overlaid function in the GPIO controller.

Note Open collector output mode is needed when multiple interrupt sources is hooked up to the same interrupt wire on the PCB. Then the wire will be pulled high with a resistor and each interrupt source can then drive the wire low via open-collector output when they want to signal interrupt.

When an external interrupt pin is configured for destination mode (as an output) then the corresponding interrupt source must not be enabled in the interrupt controller.

4.0 FEATURES

This section provides information about specific features supported by individual blocks in the VSC7416-01 device and how these features are administrated by configurations across the entire device. Examples of various standard features are included such as the support for different spanning tree versions and VLAN operations, as well as more advanced features such as QoS and VCAP.

4.1 Switch Control

This section provides information about the minimum requirements for switch operation.

4.1.1 SWITCH INITIALIZATION

The following initialization sequence is required to ensure proper operation of the switch.

1. Initialize memories:
SYS::RESET_CFG.MEM_ENA = 1.
SYS::RESET_CFG.MEM_INIT = 1.
2. Wait 100 μ s for memories to initialize (SYS::RESET_CFG.MEM_INIT cleared).
3. Enable the switch core:
SYS::RESET_CFG.CORE_ENA = 1.
4. Enable each port module through QSYS:PORT:SWITCH_PORT_MODE.PORT_ENA = 1.

4.2 Port Module Control

This section provides information about the features and configurations for port control, port reset procedures, and port counters.

4.2.1 MAC CONFIGURATION PORT MODE CONTROL

All port modules can be configured independently to the speed and duplex modes listed in the following tables.

TABLE 4-1: MAC CONFIGURATION OF PORT MODES

Configuration	10 Mbps, Half Duplex	10 Mbps, Full Duplex	100 Mbps, Half Duplex	100 Mbps, Full Duplex	1 Gbps, Full Duplex
PORT::CLOCK_CFG.LINK_SPEED	3	3	2	2	1
PORT::MAC_MODE_CFG.FDX_ENA	0	1	0	1	1
PORT::MAC_MODE_CFG.GIGA - MODE_ENA	0	0	0	0	1
SYS:FRONT_PORT_MODE.HDX_MODE	1	0	1	0	0
PORT::MAC_IFG_CFG.TX_IFG	15	15	15	15	5
PORT::MAC_IFG_CFG.RX_IFG1	11		7		
PORT::MAC_IFG_CFG.RX_IFG2	9		9		
PORT::MAC_HDX_CFG.LATE_COL_POS	67		67		
SYS::FRONT_PORT_MODE.HDX_MODE	1	0	1	0	0

4.2.2 SERDES CONFIGURATION PORT MODE CONTROL

Each SerDes port can connect to one of two types of SerDes macros. Ports 8 and 9 connect to a SERDES6G and must be configured according to the following table.

TABLE 4-2: SERDES6G CONFIGURATION

Configuration	SGMII Mode
HSIO::SERDES6G_PLL_CFG.PLL_ROT_FRQ	0
HSIO::SERDES6G_PLL_CFG.PLL_ROT_DIR	1

TABLE 4-2: SERDES6G CONFIGURATION (CONTINUED)

Configuration	SGMII Mode
HSIO::SERDES6G_PLL_CFG.PLL_ENA_ROT	0
HSIO::SERDES6G_COM-MON_CFG.ENA_LANE	1
HSIO::SERDES6G_COMMON_CFG.IF_MODE	1
HSIO::SERDES6G_COMMON_CFG.QRATE	1
HSIO::SERDES6G_COMMON_CFG.HRATE	0

Ports 0 through 7 and 10 connect to a SERDES1G and must be configured according to the following table.

TABLE 4-3: SERDES1G CONFIGURATION

Configuration	SGMII mode
HSIO::SERDES1G_COM-MON_CFG.ENA_LANE	1

4.2.3 PORT RESET PROCEDURE

When changing a switch port's mode of operation or restarting a switch port, the following port reset procedure must be followed:

1. Disable the MAC frame reception in the switch port:
DEV::MAC_ENA_CFG.RX_ENA = 0.
2. Disable traffic being sent to or from the switch port:
QSYS:PORT:SWITCH_PORT_MODE_ENA = 0.
SYS:PORT:FRONT_PORT_MODE.HDX_MODE = 0.
3. Disable shaping to speed up flushing of frames
QSYS::RATE_CFG.SHAPER_RATE = 0,.
4. Flush the queues associated with the port:
REW:PORT:PORT_CFG.FLUSH_ENA = 1.
5. Wait at least the time it takes to receive a frame of maximum length on the port Worst-case delays for 10 kilobyte jumbo frames are:
8 ms on a 10M port
800 μ s on a 100M port
80 μ s on a 1G port.
6. Reset the switch port by setting the following reset bits in CLOCK_CFG:
DEV::CLOCK_CFG.MAC_TX_RST = 1,
DEV::CLOCK_CFG.MAC_RX_RST = 1,
DEV::CLOCK_CFG.PORT_RST = 1,
7. Wait until flushing is complete:
QSYS:PORT:SW_STATUS.EQ_AVAIL must return 0.
8. Clear flushing again:
REW:PORT:PORT_CFG.FLUSH_ENA = 0.
9. Re-enable traffic being sent to or from the switch port:
QSYS:PORT:SWITCH_PORT_MODE.PORT_ENA = 1.
10. Set up the switch port to the new mode of operation. Keep the reset bits in CLOCK_CFG set. For more information about port mode configurations, see [Table 4-1](#).
11. Release the switch port from reset by clearing the reset bits in CLOCK_CFG.

It is not necessary to reset the SerDes macros.

4.2.4 PORT COUNTERS

The statistics collected in each port module provide monitoring of various events. This section describes how industry-standard Management Information Bases (MIBs) can be implemented using the counter set in the VSC7416-01 device. The following MIBs are considered.

- RMON statistics group (RFC 2819)

- IEEE 802.3-2005 Annex 30A counters
- SNMP interfaces group (RFC 2863)
- SNMP Ethernet-like group (RFC 3536)

4.2.4.1 RMON Statistics Group (RFC 2819)

The following table provides the mapping of RMON counters to port counters.

TABLE 4-4: MAPPING OF RMON COUNTERS TO PORT COUNTERS

RMON Counter	RX/TX	Switch Core Implementation
EtherStatsDropEvents	RX	C_RX_CAT_DROP + C_DR_TAIL + sum of C_DR_PRIO_x, where x is 0 through 7.
EtherStatsOctets	RX	C_RX_OCT
EtherStatsPkts	RX	C_RX_SHORT + C_RX_FRAG + C_RX_JAB- BER + C_RX_LONG + C_RX_SZ_64 + C_RX_SZ_65_127 + C_RX_SZ_128_255 + C_RX_SZ_256_511 + C_RX_SZ_512_1023 + C_RX_SZ_1024_1526 + C_RX_SZ_JUMBO
EtherStatsBroadcastPkts	RX	C_RX_BC
EtherStatsMulticastPkts	RX	C_RX_MC
EtherStatsCRCAAlignErrors	RX	C_RX_CRC
EtherStatsUndersizePkts	RX	C_RX_SHORT
EtherStatsOversizePkts	RX	C_RX_LONG
EtherStatsFragments	RX	C_RX_FRAG
EtherStatsJabbers	RX	C_RX_JABBER
EtherStatsPkts64Octets	RX	C_RX_SZ_64
EtherStatsPkts65to127Octets	RX	C_RX_SZ_65_127
EtherStatsPkts128to255Octets	RX	C_RX_SZ_128_255
EtherStatsPkts256to511Octets	RX	C_RX_SZ_256_511
EtherStatsPkts512to1023Octets	RX	C_RX_SZ_512_1023
EtherStatsPkts1024to1518Octets	RX	C_RX_SZ_1024_1526
EtherStatsDropEvents	TX	C_TX_DROP + C_TX_AGE
EtherStatsOctets	TX	C_TX_OCT
EtherStatsPkts	TX	C_TX_SZ_64 + C_TX_SZ_65_127 + C_TX- SZ_128_255 + C_TX_SZ_256_511 + C_TX- SZ_512_1023 + C_TX_SZ_1024_1526 + C_TX_SZ_JUMBO
EtherStatsBroadcastPkts	TX	C_TX_BC
EtherStatsMulticastPkts	TX	C_TX_MC
EtherStatsCollisions	TX	C_TX_COL
EtherStatsPkts64Octets	TX	C_TX_SZ_64
EtherStatsPkts65to127Octets	TX	C_TX_SZ_65_127
EtherStatsPkts128to255Octets	TX	C_TX_SZ_128_255
EtherStatsPkts256to511Octets	TX	C_TX_SZ_256_511
EtherStatsPkts512to1023Octets	TX	C_TX_SZ_512_1023
EtherStatsPkts1024to1518Octets	TX	C_TX_SZ_1024_1526

4.2.4.2 IEEE 802.3-2005 Annex 30A Counters

This section provides the mapping of IEEE 802.3-2005 Annex 30A counters to port counters. Only counter groups with supported counters are listed.

TABLE 4-5: MANDATORY COUNTERS

Counter	Rx/Tx	Switch Core Implementation
aFramesTransmittedOK	TX	C_TX_SZ_64 + C_TX_SZ_65_127 + C_TX- SZ_128_255 + C_TX_SZ_256_511 + C_TX- SZ_512_1023 + C_TX_SZ_1024_1526 + C_TX_SZ_JUMBO
aSingleCollisionFrames	TX	Not available.
aMultipleCollisionFrames	TX	Not available.
aFramesReceivedOK	RX	Sum of C_RX_PRIO_x, where x is 0 through 7.
aFrameCheckSequenceErrors	RX	Not available. C_RX_CRC is the sum of FCS and alignment errors.
aAlignmentErrors	RX	Not available. C_RX_CRC is the sum of FCS and alignment errors.

TABLE 4-6: OPTIONAL COUNTERS

Counter	RX/TX	Switch Core Implementation
aMulticastFramesXmittedOK	TX	C_TX_MC
aBroadcastFramesXmittedOK	TX	C_TX_BC
aMulticastFramesReceivedOK	RX	C_RX_MC
aBroadcastFramesReceivedOK	RX	C_RX_BC
aInRangeLengthErrors	RX	Not available
aOutOfRangeLengthField	RX	Not available
aFrameTooLongErrors	RX	C_RX_LONG

TABLE 4-7: RECOMMENDED MAC CONTROL COUNTERS

Counter	RX/TX	Switch Core Implementation
aMACControlFramesTransmitted	TX	Not available
aMACControlFramesReceived	RX	C_RX_CONTROL
aUnsupportedOpCodesReceived	RX	Not available

TABLE 4-8: PAUSE MAC CONTROL RECOMMENDED COUNTERS

Counter	RX/TX	Switch Core Implementation
aPauseMACControlFramesTransmit- ted	TX	C_TX_PAUSE
aPauseMACControlFramesReceived	RX	C_RX_PAUSE

4.2.4.3 SNMP Interfaces Group (RFC 2863)

The following table provides the mapping of SNMP interfaces group counters to port counters.

TABLE 4-9: MAPPING OF SNMP INTERFACES GROUP COUNTERS TO PORT COUNTERS

Counter	RX/TX	Switch Core Implementation
IfInOctets	RX	C_RX_OCT
IfInUcastPkts	RX	C_RX_UC
IfInNUcastPkts	RX	C_RX_BC + C_RX_MC
IfInBroadcast (RFC 1573)	RX	C_RX_BC
IfInMulticast (RFC 1573)	RX	C_RX_MC
IfInDiscards	RX	C_DR_TAIL + C_RX_CAT_DROP
IfInErrors	RX	C_RX_CRC + C_RX_SHORT + C_RX_FRAG + C_RX_JABBER + C_RX_LONG
IfInUnknownProtos	RX	Always zero.
IfOutOctets	TX	C_TX_OCT
IfOutUcastPkts	TX	C_TX_UC
IfOutNUcastPkts	TX	C_TX_BC + C_TX_MC
IfOutMulticast (RFC 1573)	TX	C_TX_MC
IfOutBroadcast (RFC 1573)	TX	C_TX_BC
IfOutDiscards	TX	Always zero.
IfOutErrors	TX	C_TX_DROP + C_TX_AGE

4.2.4.4 SNMP Ethernet-Like Group (RFC 3536)

The following table provides the mapping of SNMP Ethernet-like group counters to port counters.

TABLE 4-10: MAPPING OF SNMP ETHERNET-LIKE GROUP COUNTERS TO PORT COUNTERS

Counter	RX/TX	Switch Core Implementation
dot3StatsAlignmentErrors	RX	Not available. C_RX_CRC is the sum of FCS and alignment errors.
dot3StatsFCSErrors	RX	Not available. C_RX_CRC is the sum of FCS and alignment errors.
dot3StatsSingleCollisionFrames	TX	Not available.
dot3StatsMultipleCollisionFrames	TX	Not available.
dot3StatsSQETestErrors	RX	Not applicable.
dot3StatsDeferredTransmissions	TX	Not available.
dot3StatsLateCollisions	TX	Not available. C_TX_DROP is the sum of Late collisions and Excessive collisions.
dot3StatsExcessiveCollisions	TX	Not available. C_TX_DROP is the sum of Late collisions and Excessive collisions.
dot3StatsInternalMacTransmitErrors	TX	Not applicable. Always 0.
dot3StatsCarrierSenseErrors	TX	Not available.
dot3StatsFrameTooLongs	RX	C_RX_LONG.
dot3StatsInternalMacReceiveErrors	RX	Not applicable. Always 0.
dot3InPauseFrames	RX	C_RX_PAUSE.
dot3OutPauseFrames	TX	C_TX_PAUSE.

4.3 Layer-2 Switch

This section describes the following Layer-2 switch features.

- Switching
- VLAN and GVRP
- Rapid and Multiple Spanning Tree
- Link aggregation
- Port-based access control
- Mirroring
- SNMP support

4.3.1 BASIC SWITCHING

Basic switching covers forwarding, address learning, and address aging.

4.3.1.1 Forwarding

The device contains a Layer-2 switch and frames are forwarded using Layer-2 information only. Exceptions to this are possible using VCAP capabilities. For example, to provide source-specific IP multicast forwarding.

The switch is designed to comply with the IEEE Bridging standard in Std 802.1D and the IEEE VLAN standard in Std 802.1Q:

- Unicast frames are forwarded to a single destination port that corresponds to the DMAC.
- Multicast frames are forwarded to multiple ports determined by the DMAC multicast group. The CPU configures multicast groups in the MAC table and the port group identifier (PGID) table. A multicast group can span across any set of ports.
- Broadcast frames (DMAC = FF-FF-FF-FF-FF-FF) are, by default, flooded to all ports except the ingress port. Also, in compliance with the standard, a unicast or multicast frame with unknown DMAC is flooded to all ports except the ingress port. It is possible to configure flood masks to restrict the flooding of frames. There are separate flood masks for the following frame types:

Unicast (ANA::FLOODING.FLD_UNICAST)
Layer 2 multicast (ANA::FLOODING.FLD_MULTICAST)
Layer 2 broadcast (ANA::FLOODING.FLD_BROADCAST)
IPv4 multicast data (ANA::FLOODING_IPMC.FLD_MC4_DATA)
IPv4 multicast control (ANA::FLOODING_IPMC.FLD_MC4_CTRL)
IPv6 multicast data (ANA::FLOODING_IPMC.FLD_MC6_DATA)
IPv6 multicast control (ANA::FLOODING_IPMC.FLD_MC6_CTRL)

For frames with a known destination MAC address, the destination mask comes from an entry in the port group identifier table (ANA::PGID). The PGID table contains 92 entries (entry 0 through 91), where entry 0 through 63 are used for destination masks. The remaining PGID entries are used for other parts of the forwarding and are described below.

The following table shows the PGID table organization.

TABLE 4-11: PORT GROUP IDENTIFIER TABLE ORGANIZATION

Entry Type	Number
Unicast entries	0 – 11 (including CPU)
Multicast entries	12 – 63
Aggregation Masks	64 – 79
Source Masks	80 – 91

The unicast entries contains only the port number corresponding to the entry number.

Destination masks for multicast groups must be manually entered through the CPU into the destination masks table. IPv4 and IPv6 multicast entries can also be entered using direct encoding in the MAC table, where the destination masks table is not used. For information about forwarding and configuring destination masks, see [Section 2.1.2, MAC](#).

The aggregation masks ensures that a frame is forwarded to exactly one member of an aggregation group.

For all forwarding decisions, a source mask prevents frames from being sent back to the ingress port. The source mask removes the ingress port from the destination mask.

All ports are enabled for receiving frames by default. This can be disabled by clearing ANA:PORT:PORT_CFG.RECV_ENA.

4.3.1.2 Address Learning

The learning process minimizes the flooding of frames. A frame's source MAC address is learned together with its VID. Each entry in the MAC table is uniquely identified by a (MAC,VID) pair. In the forwarding process, a frame's (DMAC,VID) pair is used as the key for the MAC table lookup.

The learning of unknown SMAC addresses can be either hardware-based or CPU-based. The following list shows the available learn schemes, which can be configured per port:

- **Hardware-based learning** autonomously adds entries to the MAC table without interaction from the CPU. Use the following configuration:
 ANA:PORT:PORT_CFG.LEARN_ENA = 1
 ANA:PORT:PORT_CFG.LEARNCPU = 0
 ANA:PORT:PORT_CFG.LEARNDROP = 0
 ANA:PORT:PORT_CFG.LEARNAUTO = 1
- **CPU-based learning** copies frames with unknown SMACs, or when the SMAC appears on a different port, to the CPU. These frames are forwarded as usual. Use the following configuration.
 ANA:PORT:PORT_CFG.LEARN_ENA = 1
 ANA:PORT:PORT_CFG.LEARNCPU = 1
 ANA:PORT:PORT_CFG.LEARNDROP = 0
 ANA:PORT:PORT_CFG.LEARNAUTO = 0
- **Secure CPU-based learning** is similar to CPU-based learning, except that it allows the CPU to verify the SMAC addresses before both learning and forwarding. Secure CPU-based learning redirects frames with unknown SMACs, or when the SMAC appears on a different port, to the CPU. These frames are not forwarded by hardware. Use the following configuration.
 ANA::PORT_CFG.LEARN_ENA = 1
 ANA::PORT_CFG.LEARNCPU = 1
 ANA::PORT_CFG.LEARNDROP = 1
 ANA::PORT_CFG.LEARNAUTO = 0
- **No learning** where all learn frames are discarded. Frames with known SMAC in the MAC table are forwarded by hardware. Use the following configuration.
 ANA:PORT:PORT_CFG.LEARN_ENA = 1
 ANA:PORT:PORT_CFG.LEARNCPU = 0
 ANA:PORT:PORT_CFG.LEARNDROP = 1
 ANA:PORT:PORT_CFG.LEARNAUTO = 0

Frames forwarded to the CPU for learning can be extracted from the CPU extraction queue configured in ANA:PORT:CPUQ_CFG.CPUQ_LRN.

During CPU-based learning, the rate of frames subject to learning being copied or redirected to the CPU can be controlled with the learn storm policer (ANA::STORMLIMIT_CFG[3]). This policer puts a limit on the number of frames per second that are subject to learning being copied or redirected to the CPU. The learn frames storm policer can help prevent a CPU from being overloaded when performing CPU based learning.

4.3.1.3 MAC Table Address Aging

To keep the MAC table updated, an aging scan is conducted to remove entries that were not recently accessed. This ensures that stations that have moved to a new location are not permanently prevented from receiving frames in their new location. It also frees up MAC table entries occupied by obsolete stations to give room for new stations.

In IEEE 802.1D, the recommended period for aging-out entries in the MAC address table is 300 seconds per entry. The device aging implementation checks for the aging-out of all the entries in the table. The first age scan sets the age bit for every entry in the table. The second age scan removes entries where the age bit has not been cleared since the first age scan. An entry's age bit is cleared when a received frame's (SMAC, VID) matches an entry's (MAC, VID); that is, the station is active and transmits frames. To ensure that 300 seconds is the longest an entry can reside not accessed (and unchanged) in the table, the maximum time between age scans is 150 seconds.

The device can conduct age scans in two ways:

- Automatic age scans
- CPU initiated age scans

When using automatic aging, the time between age scans is set in the ANA::AUTOAGE register in steps of 1 second, in the range from 1 second to 12 days.

When using CPU-initiated aging, the CPU implements the timing between age scans. A scan is initiated by sending an aging command to the MAC address table (ANA::MACACCESS.MAC_TABLE_CMD).

The CPU-controlled age scan process can conveniently be used to flush the entire MAC table by conducting two age scans, one immediately after the other.

Flushing selective MAC table entries is also possible. Incidents that require MAC table flushing are:

- Reconfiguration of Spanning Tree protocol port states, which may cause station moves to occur.
- If there is a link failure notification (identified by a PHY layer device), flush the MAC table on the specific port where the link failed.

To deal with these incidents, the age scan process is configurable to run only for entries learned on a specified port or for a specified VLAN (ANA::ANAGEFIL.VID_VAL). The filters can also be combined to do aging on entries that match both the specific port and the specific VLAN.

Single entries can be flushed from the MAC table by sending the FORGET command to the MAC address table.

4.3.2 STANDARD VLAN OPERATION

This section provides information about configuring and operating the device as a standard VLAN-aware switch. Subsequent sections discuss the switch as a Q-in-Q enabled provider bridge and the use of private VLANs and asymmetric VLANs.

The following table lists the port module registers for standard VLAN operation.

TABLE 4-12: PORT MODULE REGISTERS FOR STANDARD VLAN OPERATION

Register/Register Field	Description	Replication
MAC_TAGS_CFG	Allows tagged frames to be 4 bytes longer than the length configured in MAC_MAXLEN_CFG.	Per port

The following table lists the analyzer configurations and status bits for standard VLAN operation.

TABLE 4-13: ANALYZER REGISTERS FOR STANDARD VLAN OPERATION

Register/Register Field	Description	Replication
DROP_CFG.DROP_UNTAGGED_ENA	Discard untagged frames.	Per port
DROP_CFG.DROP_C_TAGGED_ENA	Discard VLAN tagged frames.	Per port
DROP_CFG.DROP_PRIO_C_TAGGED_ENA	Discard priority tagged frames.	Per port
VLAN_CFG.VLAN_AWARE_ENA	Use incoming VLAN tags in VLAN classification.	Per port
VLAN_CFG.VLAN_POP_CNT	Remove VLAN tags from frames in the rewriter.	Per port
VLAN_CFG.VLAN_DEI VLAN_CFG.VLAN_PCP VLAN_CFG.VLAN_VID	Ingress port VLAN configuration.	Per port
VLANMASK	Per-port VLAN ingress filtering enable.	None
ANEVENTS.VLAN_DISCARD	A sticky bit indicating that a frame was dropped due to lack of VLAN membership of source port.	None
ADVLEARN.VLAN_CHK	Disable learning for frames discarded due to source port VLAN membership check.	None
VLANACCESS	VLAN table command. For indirect access to configuration of the 4096 VLANs.	None
VLANTIDX	VLAN table index. For indirect access to configuration of the 4096 VLANs.	None
AGENCTRL.FID_MASK	Enable shared VLAN learning.	None

TABLE 4-13: ANALYZER REGISTERS FOR STANDARD VLAN OPERATION (CONTINUED)

Register/Register Field	Description	Replication
CPU_FWD_GARP_CFG	Enable capture of frames with reserved GARP DMAC addresses, including GVRP for VLAN registration. Per-address configuration.	Per port
CPUQ_8021_CFG.CPUQ_G-ARP_VAL	CPU queue for captured GARP frames.	Per GARP address

The following table lists the rewriter registers for standard VLAN operation.

TABLE 4-14: REWRITER REGISTERS FOR STANDARD VLAN OPERATION

Register/Register Field	Description	Replication
TAG_CFG	Egress VLAN tagging configuration.	Per port
PORT_VLAN_CFG	Egress port VLAN configuration.	Per port

In a VLAN-aware switch, each port is a member of one or more virtual LANs. Each incoming frame must be assigned a VLAN membership and forwarded according to the assigned VID. The following information draws on the definitions and principles of operations in IEEE 802.1Q. Note that the switch supports more features than mentioned in the following section, which only describes the basic requirements for a VLAN aware switch.

Standard VLAN operation is configured individually per switch port using the following configuration:

- MAC_TAGS_CFG.VLAN_AWR_ENA = 1
MAC_TAGS_CFG.VLAN_LEN_AWR_ENA = 1
- VLAN_CFG.VLAN_AWARE_ENA = 1,
VLAN_CFG.VLAN_POP_CNT = 1.

Each switch port has an Acceptable Frame Type parameter, which is set to Admit Only VLAN tagged frames or Admit All Frames:

- Admit Only VLAN-tagged frames:
DROP_CFG.DROP_UNTAGGED_ENA = 1,
DROP_CFG.DROP_PRIO_C_TAGGED_ENA = 1,
DROP_CFG.DROP_C_TAGGED = 0.
- Admit All Frames:
DROP_CFG.DROP_UNTAGGED_ENA = 0,
DROP_CFG.DROP_PRIO_C_TAGGED_ENA = 0,
DROP_CFG.DROP_C_TAGGED = 0.

Frames that are not discarded are subject to the VLAN classification. Untagged and priority-tagged frames are classified to a Port VLAN Identifier (PVID). The PVID is configured per port in VLAN_CFG.VLAN_VID. Tagged frames are classified to the VID given in the frame's tag. For more information about VLAN classification, see [Section 2.6.5, VLAN Classification](#).

4.3.2.1 Forwarding

Forwarding is always based on the combination of the classified VID and the destination MAC address. By default, all switch ports are members of all VLANs. This can be changed in VLANACCESS and VLANTIDX where port masks per VLAN are set up.

4.3.2.2 Ingress Filtering

VLAN ingress filtering can be enabled per switch port with the register VLANMASK.

The filter checks for all incoming frames to determine if the ingress port is a member of the VLAN to which the frame is classified. If the port is not a member, the frame is discarded. Whenever a frame is discarded due to lack of VLAN membership, the ANEVENTS.VLAN_DISCARD sticky bit is set. To ensure that VLAN ingress filtered frames are not learned, ADVLEARN.VLAN_CHK must be set.

4.3.2.3 GVRP

GARP VLAN Registration Protocol (GVRP) is used to propagate VLAN configurations between bridges. On a GVRP-enabled switch, all GVRP frames must be redirected to the CPU for further processing. The GVRP frames use a reserved GARP MAC address (01-80-C2-00-00-21) and can be redirected to the CPU by setting bit 1 in the analyzer register CPU_FWD_GARP_CFG.

4.3.2.4 Shared VLAN Learning

The device can be configured for either Independent VLAN learning or Shared VLAN learning. Independent VLAN learning is the default.

Shared VLAN learning, where multiple VLANs map to the same filtering database, is enabled through Filter Identifiers (FIDs). Basically, this means that learning is unique for a (MAC, FID) set and that a learned MAC address is learned for all VLANs that map to the FID. Shared VLAN learning is configured in FID_MAP per VLAN. The device supports 64 different FIDs. Any number of VLANs can map to the same FID.

4.3.2.5 Untagging

An untagged set can be configured for each egress port, which defines the VLANs for which frames are transmitted untagged. The untagged set can consist of zero, one, or all VLANs. For all VLANs not in the untagged set, frames are transmitted tagged. The available configurations in the rewriter are:

- The untagged set is empty:
REW:PORT:TAG_CFG.TAG_CFG = 3.
- The untagged set consists of all VLANs:
REW:PORT:TAG_CFG.TAG_CFG = 0.
- The untagged set consists of one VLAN <VLAN>:
REW:PORT:TAG_CFG.TAG_CFG = 1.
REW:PORT:PORT_VLAN_CFG.PORT_VLAN = <VLAN>.

Optionally, frames received as priority-tagged frames (VLAN = 0) can also be transmitted as untagged (REW:PORT:TAG_CFG.TAG_CFG=2).

Port-Based VLAN Example

Situation:

Ports 0 and 1 are isolated from ports 2 and 3 using port-based VLANs. Ports 0 and 1 are assigned port VLAN 1 and ports 2 and 3 port VLAN 2. All frames in the network are untagged.

Resolution:

```
# Port module configuration of ports 0 - 1.
# Configure the ports to always use the port VLAN.
VLAN_CFG.VLAN_AWARE_ENA = 0
# Allow only untagged frames.
DROP_CFG.DROP_UNTAGGED_ENA = 0
DROP_CFG.DROP_PRIO_C_TAGGED = 1
DROP_CFG.DROP_C_TAGGED = 1
# Configure the port VLAN to 1.
VLAN_CFG.VLAN_VLAN = 1

# Port module configuration of ports 2 - 3.
# Same as for ports 0-1, except that the port VLAN is set to 2.
VLAN_CFG.VLAN_VLAN = 2
# Analyzer configuration.
# Configure VLAN 1 to contain ports 0-1.
VLANTIDX.INDEX = 1
VLANTIDX.VLAN_PRIV_VLAN = 0
VLANTIDX.VLAN_MIRROR = 0 (don't care, for this example)
VLANTIDX.VLAN_LEARN_DISABLE = 0
VLANTIDX.VLAN_SRC_CHK = 1
VLANACCESS.VLAN_PORT_MASK = 0x03
VLANACCESS.VLAN_TBL_CMD = 2
```

```
# Configure VLAN 2 to contain ports 2-3.
VLANTIDX.INDEX = 2
VLANTIDX.VLAN_PRIV_VLAN = 0
VLANTIDX.VLAN_MIRROR = 0 (don't care, for this example)
VLANTIDX.VLAN_LEARN_DISABLE = 0
VLANTIDX.VLAN_SRC_CHK = 1
VLANACCESS.VLAN_PORT_MASK = 0x0C
VLANACCESS.VLAN_TBL_CMD = 2
```

4.3.3 PROVIDER BRIDGES AND Q-IN-Q OPERATION

The following table lists the port module configurations for provider bridge VLAN operation.

TABLE 4-15: PORT MODULE CONFIGURATIONS FOR PROVIDER BRIDGE VLAN OPERATION

Register/Register Field	Description	Replication
MAC_TAGS_CFG	Allow single tagged frames to be 4 bytes longer and double-tagged frames to be 8 bytes longer than the length configured in MAC_MAXLEN_CFG.	Per port

The following table lists the port module configurations for provider bridge VLAN operation.

TABLE 4-16: SYSTEM CONFIGURATIONS FOR PROVIDER BRIDGE VLAN OPERATION

Register/Register Field	Description	Replication
VLAN_E-TYPE_CFG.VLAN_S_TAG_E-TYPE_VAL	TPID for S-tagged frames. EtherType 0x88A8 and the configurable value VLAN_E-TYPE_CFG.VLAN_S_TAG_ETYPE_VAL are identified as the S-tag identifier.	Per port

The following table lists the analyzer configurations for provider bridge VLAN operation.

TABLE 4-17: ANALYZER CONFIGURATIONS FOR PROVIDER BRIDGE VLAN OPERATION

Register/Register Field	Description	Replication
DROP_CFG.DROP_UN-TAGGED_ENA	Discard untagged frames.	Per port
DROP_CFG.DROP_S_TAGGED_ENA	Discard VLAN S-tagged frames.	Per port
DROP_CFG.DROP_PRIORITY_S_TAGGED_ENA	Discard priority S-tagged frames.	Per port
VLAN_CFG.VLAN_AWARE_ENA	Use incoming VLAN tags in VLAN classification.	Per port
VLAN_CFG.VLAN_POP_CNT	Remove VLAN tags from frames in the rewriter.	Per port
VLAN_CFG.VLAN_TAG_TYPE	Tag type for untagged frames (Customer tag or service tag).	Per port
VLAN_CFG.VLAN_INNER_TAG_ENA	Use inner tag for VLAN classification instead of outer tag.	Per port
VLAN_CFG.VLAN_DEI VLAN_CFG.VLAN_PCP VLAN_CFG.VLAN_VID	Ingress port VLAN configuration.	Per port
VLANACCESS	VLAN table command. For indirect access to configuration of the 4096 VLANs.	None
VLANTIDX	VLAN table index. For indirect access to configuration of the 4096 VLANs.	None

The VSC7416-01 device supports the standard provider bridge features in IEEE 802.1ad (Provider Bridges). Features related to provider bridges are:

- Support for multiple tag headers (EtherTypes 0x8100, 0x88A8, and a programmable value are recognized as tag header EtherTypes)
- Pushing and popping of up to two VLAN tags
- Selective VLAN classification using either inner or outer VLAN tag
- Translating VLAN tag headers at ingress and/or at egress (using the IS1 and ES0 TCAMs)
- Enabling or disabling learning per VLAN

The following section discusses briefly how to configure these different features in the switch.

The VSC7416-01 device supports multiple VLAN tags. They can be used in MAN applications as a provider bridge, aggregating traffic from numerous independent customer LANs into the MAN space. One of the purposes of the provider bridge is to recognize and use VLAN tags so that the VLANs in the MAN space can be used independent of the customers' VLANs. This is accomplished by adding a VLAN tag with a MAN-related VID for frames entering the MAN. When leaving the MAN, the tag is stripped, and the original VLAN tag with the customer-related VID is again available. This provides a tunneling mechanism to connect remote customer VLANs through a common MAN space without interfering with the VLAN tags. All tags use EtherType 0x8100 for customer tags and EtherType 0x88A8, or a programmable value, for service provider tags.

In cases where a given service VLAN only has two member ports on the switch, the learning can be disabled for the particular VLAN (VLANTIDX.VLAN_LEARN_DISABLE) and can rely on flooding as the forwarding mechanism between the two ports. This way, the MAC table requirements are reduced.

MAN Access Switch Example

Situation:


The following is an example of setting up the device as a MAN access switch with the following requirements:

- Customer ports are aggregated into a network port for tunneling through the MAN to access remote VLANs.
- Local switching between ports of the different customers must be eliminated.
- Frames must be label-switched from network port to correct customer port without need for MAC address learning.

FIGURE 4-1: MAN ACCESS SWITCH SETUP

Frames in This Segment

Service Provider Tag (Outer Tag)		Customer Tag (Inner Tag)		Description
EtherType	VID	EtherType	VID	
0x88A8	1	0x8100	1	Frames to/from customer 1's VLAN 1
0x88A8	1	0x8100	118	Frames to/from customer 1's VLAN 118
0x88A8	1	0x8100	0	Priority-tagged frames to/from customer 1
0x88A8	2	0x8100	1	Frames to/from customer 2's VLAN 1
0x88A8	2	0x8100	4	Frames to/from customer 2's VLAN 4
0x88A8	2	N/A	N/A	Untagged frames to/from customer 2



Frames in This Segment

Customer Tag		Description
EtherType	VID	
0x8100	1	Frames in customer 1's VLAN 1
0x8100	118	Frames in customer 1's VLAN 118
0x8100	0	Customer 1's priority-tagged frames

Frames in This Segment

Customer Tag		Description
EtherType	VID	
0x8100	1	Frames in customer 2's VLAN 1
0x8100	4	Frames in customer 2's VLAN 4
N/A	N/A	Customer 2's untagged frames

This example is typically accomplished by letting each customer port have a unique port VID (PVID), which is used in the outer VLAN tag (the service provider tag). In the MAN, the VID directly indicates the customer port from which the frame is received or the customer port to which the frame is going.

A customer port is VLAN-unaware and classifies to a port-based VLAN. In the egress direction of the customer port, frames are transmitted untagged, which facilitates the stripping of the outer tag. That is, the provider tag is stripped, but the customer tag is kept. The port must allow frames with a maximum size of 1522 bytes.

Resolution:

```
# Configuration of customer 1's port (port 8).
# Allow for a single VLAN tag in the length check and set the maximum length
without VLAN
# tag to 1518 bytes.
MAC_TAGS_CFG.VLAN_LEN_AWR_ENA = 1
MAC_TAGS_CFG.VLAN_AWAR_ENA = 1
MAC_MAXLEN_CFG.MAX_LEN = 1518
# Configure the port to leave any incoming tags in the frame and to ignore any
# incoming VLAN tags in the VLAN classification. The port VID is always used
in the
# VLAN classification.
```

```
VLAN_CFG.VLAN_POP_CNT = 0
VLAN_CFG.VLAN_AWARE_ENA = 0
# Allow both C-tagged and untagged frames coming in to the device to also
support customer traffic not using VLANs to be carried across the MAN.
DROP_CFG.DROP_UNTAGGED_ENA = 0
DROP_CFG.DROP_C_TAGGED = 0
DROP_CFG.DROP_PRIO_C_TAGGED = 0
DROP_CFG.DROP_S_TAGGED = 1
DROP_CFG.DROP_PRIO_S_TAGGED = 1
# Use service provider tagging when frames from this port exit the switch.
# (EtherType 0x88A8).
VLAN_CFG.VLANTAG_TYPE = 1
# Configure the port VID to 1.
VLAN_CFG.VLAN_VID = 1
# Configure the egress side of the port to not insert tags.
# (The service provider tags are stripped in the ingress side of the MAN port).
TAG_CFG.TAG_CFG = 0
# Configuration of customer 2's port (port 9).
# Same as for customer 1's port (port 8), except that the port VID is set to 2.
VLAN_CFG.VLAN_VID = 2
# Configuration of the network port (port 16).
# MAN traffic in transit between network ports is supported by configuring all
network
# ports as follows:
# Allow for two VLAN tags in the length check and set the max length without
# VLAN tags to 1518 bytes.
MAC_TAGS_CFG.VLAN_LEN_AWR_ENA = 1
MAC_TAGS_CFG.VLAN_AWAR_ENA = 1
MAC_TAGS_CFG.PB_ENA = 1
MAC_MAXLEN_CFG.MAX_LEN = 1518
# Configure the port to use incoming VLAN tags in the VLAN classification,
# and to remove the first (outer) VLAN tag (the service tag) from incoming
frames.
VLAN_CFG.VLAN_POP_CNT = 1
VLAN_CFG.VLAN_AWARE_ENA = 1
# Allow only S-tagged frames.
DROP_CFG.DROP_UNTAGGED_ENA = 1
DROP_CFG.DROP_C_TAGGED = 1
DROP_CFG.DROP_PRIO_C_TAGGED = 1
DROP_CFG.DROP_S_TAGGED = 0
DROP_CFG.DROP_PRIO_S_TAGGED = 0
# The tag type is unused on the network port
VLAN_CFG.VLANTAG_TYPE = 0
# Configure the egress side of the port to insert tags.
TAG_CFG.TAG_CFG = 1
# Common configuration in the analyzer.
# Configure VLAN 1 to contain customer 1's port (port 8) and the network port
# (port 16). Disable learning in VLAN 1. Ingress filtering is don't care for
port
# based VLANs.
VLANTIDX.INDEX = 1
VLANTIDX.VLAN_PRIV_VLAN = 0 (don't care, for this example)
VLANTIDX.VLAN_MIRROR = 0 (don't care, for this example)
VLANTIDX.VLAN_LEARN_DISABLE = 1
VLANTIDX.VLAN_SRC_CHK = 0 (don't care, for this example)
VLANACCESS.VLAN_PORT_MASK = 0x00010100
VLANACCESS.VLAN_TBL_CMD = 2
# Configure VLAN 2 to contain customer 2's port (port 9) and the network port
# (port 16). Disable learning in VLAN 2. Ingress filtering is don't-care for
```



```

port
# based VLANs.
VLANTIDX.INDEX = 2
VLANTIDX.VLAN_PRIV_VLAN = 0 (don't care, for this example)
VLANTIDX.VLAN_MIRROR = 0 (don't care, for this example)
VLANTIDX.VLAN_LEARN_DISABLE = 1
VLANTIDX.VLAN_SRC_CHK = 0 (don't care, for this example)
VLANACCESS.VLAN_PORT_MASK = 0x00010200
VLANACCESS.VLAN_TBL_CMD = 2

```

4.3.4 PRIVATE VLANS

The following table lists the analyzer configuration registers for private VLAN support.

TABLE 4-18: PRIVATE VLAN CONFIGURATION REGISTERS

Register	Description	Replication
VLANACCESS	VLAN table command. For indirect access to configuration of the 4096 VLANs.	None
VLANTIDX	VLAN table index. For indirect access to configuration of the 4096 VLANs.	None
ISOLATED_PORTS	VLAN port mask indicating isolated ports in private VLANs.	None
COMMUNITY_PORTS	VLAN port mask indicating community ports in private VLANs.	None

When a VLAN is configured to be a private VLAN, communication between ports within that VLAN can be prevented. Two application examples are:

- Customers connected to an ISP can be members of the same VLAN, but they are not allowed to communicate with each other within that VLAN.
- Servers in a farm of web servers in a Demilitarized Zone (DMZ) are allowed to communicate with the outside world and with database servers on the inside segment, but are not allowed to communicate with each other

For private VLANs to be applied, the switch must first be configured for standard VLAN operation. For more information, see [Section 4.3.2, Standard VLAN Operation](#). When this is in place, one or more of the configured VLANs can be configured as private VLANs. Ports in a private VLAN fall into one of three groups:

- Promiscuous ports
Ports from which traffic can be forwarded to all ports in the private VLAN

Ports that can receive traffic from all ports in the private VLAN
- Community Ports
Ports from which traffic can only be forwarded to community and promiscuous ports in the private VLAN

Ports that can receive traffic from only community and promiscuous ports in the private VLAN
- Isolated ports
Ports from which traffic can only be forwarded to promiscuous ports in the private VLAN

Ports that can receive traffic from only promiscuous ports in the private VLAN

The configuration of promiscuous, community, and isolated ports applies to all private VLANs.

The forwarding of frames classified to a private VLAN happens:

- When traffic comes in on a promiscuous port in a private VLAN, the VLAN mask from the VLAN table is applied.
- When traffic comes in on a community port, the ISOLATED_PORT mask is applied in addition to the VLAN mask from the VLAN table.
- When traffic comes in on an isolated port, the ISOLATED_PORT mask and the COMMUNITY_PORT mask are applied in addition to the VLAN mask from the VLAN table.


ISP Example

Situation:

Customers A, B, and C are connected to the same switch at the ISP. Customers A and B are allowed to communicate with each other, as well as the ISP. Customer C can only communicate with the ISP. VLAN 1 is the private VLAN that isolates Customers A, B from C. Traffic on VLAN 1 coming in from the ISP (port 8) uses the VLAN mask in the VLAN table. Traffic on VLAN 1 from customer A or B has the ISOLATED_PORTS mask applied in addition to the mask from the VLAN table, with the result that traffic from customer A and B is not forwarded to customers C. Traffic on VLAN 1 from customer C has the ISOLATED_PORTS mask and the COMMUNITY_PORTS mask applied in addition to the mask from the VLAN table, with the result that traffic from customer C is not forwarded to customers A and B.

The following illustration shows the desired setup.

FIGURE 4-2: ISP EXAMPLE FOR PRIVATE VLAN



Resolution:


```
# It is assumed that Port VID and tag handling for VLAN 1 is already
# configured according to the description in Standard VLAN Operation.
# Configure VLAN 1 as a private VLAN in the VLAN table by performing these
steps:
# - Point to VLAN 1.
# - Set it as private.
# - Disable mirroring of the VLAN (not important for the example).
# - Enable learning within the VLAN (not important for the example).
# - Disable source check within the VLAN (not important for the example).
# - Include ports 0, 1, 2, and 8 in the VLAN mask.
# Insert the entry into the VLAN table.
VLANTIDX.INDEX = 1
VLANTIDX.VLAN_PRIV_VLAN = 1
VLANTIDX.VLAN_MIRROR = 0 (don't care, for this example)
VLANTIDX.VLAN_LEARN_DISABLE = 0 (don't care, for this example)
VLANTIDX.VLAN_SRC_CHK = 0 (don't care, for this example)
VLANACCESS.VLAN_PORT_MASK = 0x00000107
VLANACCESS.VLAN_TBL_CMD = 2
# Configure the private VLAN mask so that port 8 is a promiscuous
# port, ports 0 and 1 are community ports, and port 2 is an isolated port.
ISOLATED_PORTS.ISOL_PORTS = 0x00000103
COMMUNITY_PORTS.COMM_PORTS = 0x00000104
```

DMZ Example**Situation:**

VLAN 17 is a private VLAN that isolates Server1 and Server2. Traffic on VLAN 17 coming from the internal or the external router (ports 8 and 9) uses the VLAN mask in the VLAN table. Traffic on VLAN 17 from Server1 and Server2 (ports 14 and 15) has the ISOLATED_PORTS applied in addition to the mask from the VLAN table, with the result that traffic from Server1 is not forwarded to Server2 and visa versa.

The following illustration shows the desired setup.

FIGURE 4-3: DMZ EXAMPLE FOR PRIVATE VLAN



Resolution:

```
# It is assumed that Port VID and tag handling for VLAN 17 is already
# configured according to the description in Standard VLAN Operation.
# Configure VLAN 17 as a private VLAN in the VLAN table by performing these
# steps:
# - Point to VLAN 17.
# - Set it as private.
# - Disable mirroring of the VLAN (not important for the example).
# - Enable learning within the VLAN (not important for the example).
# - Disable source check within the VLAN (not important for the example).
# - Include ports 8, 9, 14, and 15 in the VLAN mask.
# - Insert the entry into the VLAN table.
VLANTIDX.INDEX = 17
VLANTIDX.VLAN_PRIV_VLAN = 1
VLANTIDX.VLAN_MIRROR = 0 (don't care, for this example)
VLANTIDX.VLAN_LEARN_DISABLE = 0 (don't care, for this example)
VLANTIDX.VLAN_SRC_CHK = 0 (don't care, for this example)
VLANACCESS.VLAN_PORT_MASK = 0x0000C300
VLANACCESS.VLAN_TBL_CMD = 2
# Configure the private VLAN mask so that ports 8 and 9 are promiscuous
# ports.
ISOLATED_PORTS.ISOL_PORTS = 0x00000300
```

4.3.5 ASYMMETRIC VLANS

Asymmetric VLANs use the same configuration registers as for standard VLAN operation. For more information about standard VLAN operation, see [Section 4.3.2, Standard VLAN Operation](#).


Asymmetric VLANs can be used to prevent communication between hosts in a network. This behavior is similar to what can be obtained by using private VLANs. For more information, see [Section 4.3.4, Private VLANs](#).

Situation:

A server and two hosts are connected to a switch. Communication between the hosts and the server should be allowed, but the hosts are not allowed to communicate directly. All traffic between the server and the hosts is untagged. Host 1 is connected to port 9, host 2 to port 10, and the server to port 8.

The host-1 port gets port VID 1 and the host-2 port gets port VID 2. The server port is a member of both VLANs 1 and 2. The server port gets port VID 3, and the two host ports are members of VLAN 3, as shown in the following illustration.

FIGURE 4-4: ASYMMETRIC VLANS



Resolution:

```
# Analyzer configurations common for ports 8, 9, and 10.
# Allow only untagged frames.
DROP_CFG.DROP_UNTAGGED_ENA = 0
DROP_CFG.DROP_C_TAGGED_ENA = 1
DROP_CFG.DROP_PRIO_C_TAGGED_ENA = 1
# As tagged frames are dropped all frames are classified to the port VID.
VLAN_CFG.VLAN_AWARE_ENA = 0 (don't care, for this example)
# Configure the egress side of the port to not insert tags.
TAG_CFG.TAG_CFG = 0
# Analyzer configuration specific for port 8. Set the port VID to 3.
VLAN_CFG.VLAN_VID = 3
VLAN_CFG.VLAN_DEI = 0 (don't care, for this example)
# Analyzer configuration specific for port 9. Set the port VID to 1.
VLAN_CFG.VLAN_VID = 1
VLAN_CFG.VLAN_DEI = 0 (don't care, for this example)

# Analyzer configuration specific for port 10. Set the port VID to 2.
VLAN_CFG.VLAN_VID = 2
VLAN_CFG.VLAN_DEI = 0 (don't care, for this example)
# Analyzer configuration common to all ports.
# Configure VLAN 1 to contain port 8.
VLANTIDX.INDEX = 1
VLANTIDX.VLAN_PRIV_VLAN = 0
VLANTIDX.VLAN_MIRROR = 0 (don't care, for this example)
VLANTIDX.VLAN_LEARN_DISABLE = 0
```

```
VLANTIDX.VLAN_SRC_CHK = 0
VLANACCESS.VLAN_PORT_MASK = 0x00000100
VLANACCESS.VLAN_TBL_CMD = 2
# Configure VLAN 2 to contain port 8.
VLANTIDX.INDEX = 2
VLANTIDX.VLAN_PRIV_VLAN = 0
VLANTIDX.VLAN_MIRROR = 0 (don't care, for this example)
VLANTIDX.VLAN_LEARN_DISABLE = 0
VLANTIDX.VLAN_SRC_CHK = 0
VLANACCESS.VLAN_PORT_MASK = 0x00000100
VLANACCESS.VLAN_TBL_CMD = 2
# Configure VLAN 3 to contain ports 9 and 10.
VLANTIDX.INDEX = 3
VLANTIDX.VLAN_PRIV_VLAN = 0
VLANTIDX.VLAN_MIRROR = 0 (don't care, for this example)
VLANTIDX.VLAN_LEARN_DISABLE = 0
VLANTIDX.VLAN_SRC_CHK = 0
VLANACCESS.VLAN_PORT_MASK = 0x00000600
VLANACCESS.VLAN_TBL_CMD = 2
```

4.3.6 SPANNING TREE PROTOCOLS

This section provides information about Rapid Spanning Tree Protocol (RSTP) support and Multiple Spanning Tree Protocol (MSTP) support. The device also supports legacy Spanning Tree Protocol (STP). STP was obsoleted by RSTP in IEEE 802.1D and is not described in this document.

It is assumed that only LAN ports connected to the switch core participate in the spanning tree protocol. This implies that BPDUs are terminated by the switch core.

4.3.6.1 Rapid Spanning Tree Protocol

The following table lists the analyzer configuration registers for Rapid Spanning Tree Protocol (RSTP) operation.

TABLE 4-19: ANALYZER CONFIGURATIONS FOR RSTP SUPPORT


Register/Register Field	Description	Replication
PGID[80-91]	Source masks used for ingress filtering.	Per port
PGID[64-79]	Aggregation masks that can be used for egress filtering for RSTP.	16
PORT_CFG.LEARN_ENA	Enable learning per port.	Per port
CPU_FWD_BPDU_CFG	Enable redirection of frames with reserved BPDU DMAC addresses.	Per port per address
CPUQ_8021_CFG.CPUQ_BPDU_VAL	CPU extraction queue for redirected BPDU frames.	Per address

To eliminate potential loops in a network, the Rapid Spanning Tree Protocol in IEEE 802.1D creates a single path between any two bridges in a network, adding stability and predictability to the network. The protocol is implemented by assigning states to all ports. Each state controls a port's functionality, limiting its ability to receive and transmit frames and learn addresses.

Establishing a spanning tree is done through the exchange of BPDUs between bridge entities. BPDUs are frequently exchanged between neighboring bridges. These frames are identified by the Bridge protocol address range (DMAC = 01-80-C2-00-00-0x).

When there is a change in the network topology, the protocol reconfigures the port states.

FIGURE 4-5: SPANNING TREE EXAMPLE



The following table lists the Rapid Spanning Tree port state properties.

TABLE 4-20: RSTP PORT STATE PROPERTIES

State	BPDU Reception	BPDU Generation	Frame Forwarding	SMAC Learning
Discarding	Yes	Yes	No	No
Learning	Yes	Yes	No	Yes
Forwarding	Yes	Yes	Yes	Yes

The legacy STP states disabled, blocking, and listening correspond to the discarding state of RSTP.

All frames with a Bridge protocol address must be redirected to the CPU. This is configured in CPU_FWD_BPDU_CFG. BPDUs are forwarded to the CPU irrespective of the port's RSTP state. CPUQ_8021_CFG.CPUQ_BPDU_VAL can be used to configure in which CPU extraction queue the BPDUs are placed. BPDU generation is done through frame injection from the CPU.

Frame forwarding is controlled through ingress filtering and egress filtering. Ingress filtering can be done by using the source masks (PGID[80-91]), and egress filtering can be done by using the aggregation masks (PGID[64-79]). Forwarding can be disabled for ports not in the Forwarding state by clearing their source masks and excluding them from all aggregation masks. The use of the aggregation masks for egress filtering does not preclude the combination of link aggregation and RSTP support. All ports in a link aggregation group that are not in the Forwarding state must be disabled in all aggregation masks. For link aggregated ports in the Forwarding state, the aggregation masks must be configured for link aggregation (such as when RSTP is not supported.)

Learning can be enabled per port with the PORT_CFG.LEARN_ENA.

The following table provides an overview of the port state configurations for port p.

TABLE 4-21: RSTP PORT STATE CONFIGURATION FOR PORT P

State	CPU_FWD_BPDU_CFG[p].BPDU_REDIR_ENA[0]	PGID[80+p]	PGID[64-79], All 16 Masks, Bit p	PORT_CFG[p].LEARN_ENA
Discarding	1	0	0	0
Learning	1	0	0	1
Forwarding	1	1 except for bit p	1	1

RSTP Example

Situation:

Port 0 is in the RSTP Discarding state. Port 2 is in the RSTP Learning state. Port 3 is in the RSTP Forwarding state. All other ports on the switch are unused.

Resolution:

```
# Get Spanning Tree Protocol BDPUs to CPU extraction queue 0 for port 0, 2,
and 3.
CPU_FWD_BPDU_CFG[0].BPDU_REDIR_ENA[0] = 1
CPU_FWD_BPDU_CFG[2].BPDU_REDIR_ENA[0] = 1
CPU_FWD_BPDU_CFG[3].BPDU_REDIR_ENA[0] = 1
CPUQ_8021_CFG.CPUQ_BPDU_VAL[0] = 0
# Configure the source mask for port 0 (Discarding state).
PGID[80] = 0x00
# Configure the source mask for port 2 (Learning state).
PGID[82] = 0x00
# Configure the source mask for port 3 (Forwarding state).
PGID[83] = 0x77
# Configure the aggregation masks to only allow forwarding to port 3
# (Forwarding state).
PGID[64-79] = 0x08
# Configure the learn mask to only allow learning on ports
# 2 (Learning state) and 3 (Forwarding state).
PORT_CFG[0].LEARN_ENA = 0
PORT_CFG[2].LEARN_ENA = 1
PORT_CFG[3].LEARN_ENA = 1
```

Multiple Spanning Tree Protocol


The following table lists the analyzer configuration registers for Multiple Spanning Tree Protocol (MSTP) operation.

TABLE 4-22: ANALYZER CONFIGURATIONS FOR MSTP SUPPORT

Register/Register Field	Description	Replication
VLANACCESS.VLAN_SRC_CHK	Per-VLAN ingress filtering enable. Part of VLAN table command for indirect access to configuration of the 4095 VLANs.	None
VLANMASK	Per-port VLAN ingress filtering enable.	None
ADVLEARN.VLAN_CHK	Disable learning for frames discarded due to VLAN membership source port filtering.	None
PORT_CFG.LEARN_ENA	Enable learning per port.	Per port
CPU_FWD_BPDU_CFG	Enable redirection of frames with reserved BPDU DMAC addresses.	Per port per address
CPUQ_8021_CFG.CPUQ_BP-DU_VAL	CPU extraction queue for redirected BPDU frames.	Per address

The Multiple Spanning Tree Protocol (MSTP) in IEEE 802.1Q increases network use, relative to RSTP, by creating multiple spanning trees that VLANs can map to independently, rather than having only one path between bridges common for all VLANs. The multiple spanning trees are created by assigning different bridge identifiers for each spanning tree. Mapping the VLANs to spanning trees is done arbitrarily.

FIGURE 4-6: MULTIPLE SPANNING TREE EXAMPLE



The Learning state is not supported for MSTP. However, this has limited impact, because when the port is taken to the Forwarding state, learning is done at wire-speed, and, as a result, the SMAC learn delay is less important. MSTP is supported for all VLANs.

The following table lists the multiple spanning tree port state properties.

TABLE 4-23: MSTP PORT STATE PROPERTIES

State per VLAN	BPDU Reception	BPDU Generation	Frame Forwarding	SMAC Learning
Discarding	Yes	Yes	No	No
Learning (not supported)	Yes	Yes	No	Yes
Forwarding	Yes	Yes	Yes	Yes

To enable the MSTP port states:

- Ensure that the switch is VLAN-aware. For more information, see [Section 4.3.2, Standard VLAN Operation](#).
- Set the ADVLEARN.VLAN_CHK bit to prevent learning of frames discarded due to VLAN ingress filtering.
- Configure all ports as defined for the forwarding state of the RSTP port. For more information, see [Table 4-21](#).

Port states per VLAN are hereafter solely configured through the VLAN masks as listed in the following table for port *p* and VLAN *v*.

TABLE 4-24: MSTP PORT STATE CONFIGURATION FOR PORT *P* AND VLAN *V*

State	VLAN_ACCESS.VLAN_SRC_CHKVLAN <i>v</i>	VLAN_ACCESS.VLAN_PORT_MASK Bit <i>p</i> , VLAN <i>v</i>
Discarding	1	0
Learning	Not supported	Not supported
Forwarding	1	1

As an alternative to setting the VLANACCESS.VLAN_SRC_CHK bit in all VLAN entries in the VLAN table, VLAN ingress filtering can be enabled globally for all VLANs on a per port basis through VLANMASK.

For all multiple spanning tree instances, BPDUs are forwarded to the CPU irrespective of the port states.

MSTP Example

Situation:

Ports 10 and 11 are both members of VLANs 20 and 21. Two spanning trees are used:

- Spanning tree for VLAN 20, where both ports 10 and 11 are in the Forwarding state
- Spanning tree for VLAN 21, where port 10 is in the Discarding state and port 11 is in the Forwarding state

All other ports on the switch are unused.

Resolution:

```
# Get all BDPUs to CPU queue 0.
CPU_FWD_BPDU_CFG[*].BPDU_REDIR_ENA[0] = 1
CPUQ_8021_CFG.CPUQ_BPDU_VAL[0] = 0
# Enable learning on all ports. The VLAN table controls forwarding and
learning.
DEV::PORT_CFG.LEARN_ENA = 1
# Disable learning of VLAN membership source port filtered frames.
ADVLEARN.VLAN_CHK = 1
# Configure VLAN 20 for ports 10 and 11 in Forwarding state.
VLANTIDX.INDEX = 20
VLANTIDX.VLAN_PRIV_VLAN = 0
VLANTIDX.VLAN_MIRROR = 0 (don't care, for this example)
VLANTIDX.VLAN_LEARN_DISABLE = 0
VLANTIDX.VLAN_SRC_CHK = 1
VLANACCESS.VLAN_PORT_MASK = 0x00000C00
VLANACCESS.VLAN_TBL_CMD = 2
```

```
# Configure VLAN 21 for port 10 in Discarding state and port 11 in Forwarding
state.
VLANTIDX.INDEX = 21
VLANTIDX.VLAN_PRIV_VLAN = 0
VLANTIDX.VLAN_MIRROR = 0 (don't care, for this example)
VLANTIDX.VLAN_LEARN_DISABLE = 0
VLANTIDX.VLAN_SRC_CHK = 1
VLANACCESS.VLAN_PORT_MASK = 0x00000800
VLANACCESS.VLAN_TBL_CMD = 2
```

4.3.7 IEEE 802.1X: NETWORK ACCESS CONTROL

IEEE 802.1X, Port-Based Network Access Control, provides a standard for authenticating and authorizing devices attached to a LAN port.

Generally, IEEE 802.1X is port-based; however, the device also supports MAC-based network access control.

This section provides information about the configuration settings for port-based and MAC-based network access control.

4.3.7.1 Port-Based Network Access Control

The following table lists the configuration settings required for port-based network access control.

TABLE 4-25: CONFIGURATIONS FOR PORT-BASED NETWORK ACCESS CONTROL

Register/Register Field	Description / Value	Replication
ANA::CPU_FWD_BPDU_CFG.BPDU_REDIR_ENA[3]	Must be set to 1 to redirect frames with destination MAC addresses 01-80-C2-00-00-03 to the CPU Port Module. IEEE 802.1X uses MAC address 01-80-C2-00-00-03.	Per port
ANA::CPUQ_8021_CFG.CPUQ_BPDU_VAL[3]	Queue to which authentication BPDUs are redirected.	None
ANA::PGID[64-79]	When a port is not yet authenticated, any forwarding of frames to the port can be disabled by clearing the port's bit in all 16 aggregation masks. After authenticated, these bits must be set.	16
ANA::PGID[80-91]	Source masks. When a port is not yet authenticated, any forwarding of frames received on the port must be disabled. This can be done by setting the ANA::PGID[80+port] to all-zeros. After authentication, the port's source mask must be set back to its normal value.	Per port

The configuration settings required for port-based network access control enable the following functionality:

- Redirects frames with DMAC 01-80-C2-00-00-03 to CPU, even if the port is not yet authenticated.
- Stops forwarding of frames to ports that are not yet authenticated. This is configured in ANA::PGID[64-79].
- Stops forwarding of frames received on ports that are not yet authenticated. This is configured in ANA::PGID[80-91].

4.3.7.2 MAC-Based Authentication with Secure CPU-Based Learning

The following table lists the configuration settings required for MAC-based network access control with secure CPU-based learning.

TABLE 4-26: CONFIGURATIONS FOR MAC-BASED NETWORK ACCESS CONTROL WITH SECURE CPU-BASED LEARNING CONFIGURATIONS

Register/Register Field	Description / Value	Replication
ANA:PORT:CPU_FWD_BP-DU_CFG.BPDU_REDIR_ENA[3]	Must be set to 1 to redirect frames with destination MAC addresses 01-80-C2-00-00-03 to the CPU Port Module. IEEE 802.1X uses MAC address 01-80-C2-00-00-03.	Per port
ANA::CPUQ_8021_CFG.CPUQ_BPDU_VAL[3]	Queue to which authentication BPDUs are redirected.	None
ANA:PORT:PORT_CFG.LEARN_ENA ANA:PORT:PORT_CFG.LEARNCPU ANA:PORT:PORT_CFG.LEARN-DROP ANA:PORT:PORT_CFG.LEARN-AUTO	Must be set to support secure CPU-based learning. See Section 4.3.1.2, Address Learning . PORT_CFG.LEARN_ENA = 1 PORT_CFG.LEARNCPU = 1 PORT_CFG.LEARNDROP = 1 PORT_CFG.LEARNAUTO = 0	Per port

The MAC-based network access control with secure CPU-based learning enables the following functionality:

- Redirects frames with DMAC 01-80-C2-00-00-03 to CPU.
- Only frames from known, authenticated MAC addresses are forwarded to other ports.
- Frames from unknown MAC addresses are redirected to CPU for authentication. After the address is authenticated, the CPU must insert an entry in the MAC table. The authentication process may be initiated from the CPU when receiving learn frames.

4.3.7.3 MAC-Based Authentication with No Learning

The following table lists the configuration settings required for MAC-based network access control with no learning.

TABLE 4-27: CONFIGURATIONS FOR MAC-BASED NETWORK ACCESS CONTROL WITH NO LEARNING

Register/Register Field	Description / Value	Replication
ANA:PORT:CPU_FWD_BP-DU_CFG.BPDU_REDIR_ENA[3]	Must be set to 1 to redirect frames with destination MAC addresses 01-80-C2-00-00-03 to the CPU Port Module. IEEE 802.1X uses MAC address 01-80-C2-00-00-03.	Per port
ANA::CPUQ_8021_CFG.CPUQ_BPDU_VAL[3]	Queue to which authentication BPDUs are redirected.	None
ANA:PORT:PORT_CFG.LEARN_ENA ANA:PORT:PORT_CFG.LEARNCPU ANA:PORT:PORT_CFG.LEARN-DROP ANA:PORT:PORT_CFG.LEARN-AUTO	Must be set to support no learning. See Section 4.3.1.2, Address Learning . PORT_CFG.LEARN_ENA = 1 PORT_CFG.LEARNCPU = 1 PORT_CFG.LEARNDROP = 1 PORT_CFG.LEARNAUTO = 0	None

The MAC-based network access control with no learning enables the following functionality:

- Frames with DMAC 01-80-C2-00-00-03 are redirected to CPU. Unauthenticated and unauthorized devices must

initiate an 802.1X session by sending 802.1X BPDUs (MAC address: 01-80-C2-00-00-03). After the address is authenticated, the CPU must insert an entry in the MAC table.

- Only frames from known, authenticated MAC addresses are forwarded to other ports.
- Frames from unknown MAC addresses are discarded and the CPU can therefore not initiate the authentication process.

4.3.8 LINK AGGREGATION

Link aggregation bundles multiple ports (member ports) together into a single logical link. It is primarily used to increase available bandwidth without introducing loops in the network and to improve resilience against faults. A link aggregation group (LAG) can be established with individual links being dynamically added or removed. This enables bandwidth to be incrementally scaled based on changing requirements. A link aggregation group can be quickly reconfigured if faults are identified.

Frames destined for a LAG are sent on only one of the LAG's member ports. The member port on which a frame is forwarded is determined by a 4-bit aggregation code (AC) that is calculated for the frame.

The aggregation code ensures that frames belonging to the same frame flow (for example, a TCP connection) are always forwarded on the same LAG member port. For that reason, reordering of frames within a flow is not possible. The aggregation code is based on the following information:

- SMAC
- DMAC
- Source and destination IPv4 address.
- Source and destination TCP/UDP ports for IPv4 packets
- Source and destination TCP/UDP ports for IPv6 packets
- IPv6 Flow Label

For best traffic distribution among the LAG member ports, enable all contributions to the aggregation code.

Each LAG can consist of up to 16 member ports. Any quantity of LAGs may be configured for the device (only limited by the quantity of ports on the device.) To configure a proper traffic distribution, the ports within a LAG must use the same link speed.

A port cannot be a member of multiple LAGs.

4.3.8.1 Link Aggregation Configuration

The following table lists the registers associated with link aggregation groups.

TABLE 4-28: LINK AGGREGATION GROUP CONFIGURATION REGISTERS

Register/Register Field	Description / Value	Replication
ANA::PGID[0 – 63]	Destination mask	64
ANA::PGID[80 – 91]	Source mask.	Per port
ANA::PGID[64 – 79]	Aggregation mask.	16
ANA::PORT_CFG.PORTID_VAL	Logical port number. Must be set to the same value for all ports that are part of a given LAG; for example, the lowest port number that is a member of the LAG.	Per port
ANA::AGGR_CFG.AC_IP6_FLOW_LBL_ENA	Use IPv6 flow label when calculating AC. Configure identically for all ports. Recommended value is 1.	None
ANA::AGGR_CFG.AC_SIPDIP_ENA	Use IPv4 source and destination IP address when calculating aggregation code. Configure identically for all ports. Recommended value is 1.	None
ANA::AGGR_CFG.AC_TCPUDP_PORT_ENA	Use IPv4 TCP/UDP port when calculating aggregation code. Configure identically for all ports. Recommended value is 1.	None

TABLE 4-28: LINK AGGREGATION GROUP CONFIGURATION REGISTERS (CONTINUED)

Register/Register Field	Description / Value	Replication
ANA:: AGGR_CFG. AC_DMACE_ENA	Use destination MAC address when calculating aggregation code. Configure identically for all ports. Recommended value is 1.	None
ANA:: AGGR_CFG. AC_SMACE_ENA	Use source MAC address when calculating aggregation code. Configure identically for all ports. Recommended value is 1.	None
ANA:: AGGR_CFG. AC_RND_ENA	Use random aggregation code. Recommended value is 0.	None

To set up a link aggregation group, the following destination masks, source masks, and aggregation masks must be configured:

- **Destination Masks: ANA::PGID[0-63].** For each of the member ports, the corresponding destination mask must be configured to include all member ports of the LAG.
- **Source Masks: ANA::PGID[80-91].** The source masks must be configured to avoid flooding frames that are received at one member port back to another member port of the LAG. As a result, the source masks for each of the member ports must be configured to exclude all of the LAG's member ports.
- **Aggregation Masks: ANA::PGID[64-79].** The aggregation masks must be configured to ensure that when a frame is destined for the LAG, it gets forwarded to exactly one of the LAG's member ports. Also, the distribution of traffic between member ports is determined by this configuration.

The following illustration shows an example of a LAG configuration.

FIGURE 4-7: LINK AGGREGATION EXAMPLE

Configuration, Switch A


Aggregation Masks		Destination Masks	
AC	Port mask	DEST	Port mask
0x0:	0b11111001101	0:	0b00000110011
0x1:	0b11111001110	1:	0b00000110011
0x2:	0b11111011100	4:	0b00000110011
0x3:	0b11111011100	5:	0b00000110011
0x4:	0b11111001101		
0x5:	0b11111001110		
0x6:	0b11111011100		
0x7:	0b11111011100		
0x8:	0b11111001101		
0x9:	0b11111001110		
0xA:	0b11111011100		
0xB:	0b11111011100		
0xC:	0b11111001101		
0xD:	0b11111001110		
0xE:	0b11111011100		
0xF:	0b11111011100		

Source Masks	
Port	Port mask
0:	0b11111001100
1:	0b11111001100
4:	0b11111001100
5:	0b11111001100


```

PORT_CFG[0].PORTID_VAL = 0
PORT_CFG[1].PORTID_VAL = 0
PORT_CFG[4].PORTID_VAL = 0
PORT_CFG[5].PORTID_VAL = 0

```



Configuration, Switch B

Aggregation Masks		Destination Masks	
AC	Port mask	DEST	Port mask
0x0:	0b11100011111	4:	0b00011110000
0x1:	0b11100101111	5:	0b00011110000
0x2:	0b11101001111	6:	0b00011110000
0x3:	0b11110001111	7:	0b00011110000
0x4:	0b11100011111		
0x5:	0b11100101111		
0x6:	0b11101001111		
0x7:	0b11110001111		
0x8:	0b11100011111		
0x9:	0b11100101111		
0xA:	0b11101001111		
0xB:	0b11110001111		
0xC:	0b11100011111		
0xD:	0b11100101111		
0xE:	0b11101001111		
0xF:	0b11110001111		

Source Masks	
Port	Port mask
4:	0b11100001111
5:	0b11100001111
6:	0b11100001111
7:	0b11100001111


```

PORT_CFG[4].PORTID_VAL = 4
PORT_CFG[5].PORTID_VAL = 4
PORT_CFG[6].PORTID_VAL = 4
PORT_CFG[7].PORTID_VAL = 4

```

In this example, ports 0, 1, 4, and 5 of switch A are configured as a LAG. These ports are connected to 4 ports (4, 5, 6, 7) of switch B, providing an aggregated bandwidth of 4 Gbps between the two switches.

The aggregation masks for switch A are configured such that frames (destined for the LAG) are distributed on the member ports as follows:

- Port 0 if frame's aggregation code (AC) is 0x0, 0x4, 0x8, 0xC
- Port 1 if frame's aggregation code (AC) is 0x1, 0x5, 0x9, 0xD
- Port 4 if frame's aggregation code (AC) is 0x2, 0x6, 0xA, 0xE
- Port 5 if frame's aggregation code (AC) is 0x3, 0x7, 0xB, 0xF

4.3.8.2 Link Aggregation Control Protocol (LACP)

LACP allows switches connected to each other to automatically discover if any ports are member of the same LAG.

To implement LACP, any LACP frames must be redirected to the CPU. Such frames are identified by the DMAC being equal to 01-80-C2-00-00-02 (Slow Protocols Multicast address).

The following table lists the registers associated with configuring the redirection of LACP frames to the CPU.

TABLE 4-29: CONFIGURATION REGISTERS FOR LACP FRAME REDIRECTION TO THE CPU

Register/Register Field	Description / Value	Replication
ANA::CPU_FWD_BP-DU_CFG.BPDU_REDIR_ENA[2]	Must be set to 1.	Per port

4.3.9 SIMPLE NETWORK MANAGEMENT PROTOCOL (SNMP)

This section provides information about the port module registers and the analyzer registers for SNMP operation.

The following table lists the system registers for SNMP operation.

TABLE 4-30: SYSTEM REGISTERS FOR SNMP SUPPORT

Register	Description	Replication
CNT	The value of the counter. For more information about how to read counters, see Section 2.4, Statistics .	None

The following table lists the analyzer registers for SNMP support.

TABLE 4-31: ANALYZER REGISTERS FOR SNMP SUPPORT

Register	Description	Replication
MACACCESS	Command register for indirect MAC table access. Supports GET_NEXT command.	None
MACHDATA	High part of data word when accessing MAC table.	None
MACLDATA	Low part of data word when accessing MAC table.	None
MACTINDX	Index for direct-mode access to MAC table.	None

For SNMP support according to IETF RFC 1157, use the following features:

- RMON counters
- MAC table GET_NEXT function

For more information about the supported RMON counters, see [Section 4.2.4, Port Counters](#).

For more information about the MAC table GET_NEXT function, see [Table 94](#).

4.3.10 MIRRORING

To debug network problems, selected traffic can be copied, or mirrored, to a mirror port where a frame analyzer can be attached to analyze the frame flow.

The traffic to be copied to the mirror port can be selected as follows:

- All frames received on a given port (also known as ingress mirroring)
- All frames transmitted on a given port (also known as egress mirroring)
- Frames selected through configured VCAP entries
- All frames classified to specific VIDs
- All frames sent to the CPU (may be useful for software debugging)
- Frames where the source MAC address is to be learned (also known as learn frame), which may be useful for software debugging

The mirror port may be any port on the device, including the CPU.

4.3.10.1 Mirroring Configuration


The following table lists configuration registers associated with mirroring.

TABLE 4-32: CONFIGURATION REGISTERS FOR MIRRORING

Register/Register Field	Description / Value	Replication
ANA::PORT_CFG.SRC_MIRROR_ENA	If set, all frames received on this port are mirrored to the port set configured in MIRRORPORTS, that is, ingress mirroring.	Per port
ANA::EMIRRORMASK	Frames forwarded to ports in this mask are mirrored to the port set configured in MIRRORPORTS, that is, egress mirroring.	Per port
ANA::VLANTIDX.VLAN_MIRROR	If set, all frames classified to this VLAN are mirrored to the port set configured in MIRRORPORTS.	One per VID
ANA::AGENCTRL.MIRROR_CPU	Frames destined for the CPU extraction queues are also forwarded to the port set configured in MIRRORPORTS.	None
ANA::MIRRORPORTS	The mirror ports. Usually only one mirror port is configured, that is, only one bit is set in this mask.	None
ANA::CPUQ_CFG.CPUQ_MIRROR	CPU extraction queue used, if CPU is included in MIRRORPORTS.	None
ANA::ADVLEARN.LEARN_MIRROR	Learn frames are also forwarded to ports marked in MIRRORPORTS.	None
VCAP Registers	Configuration of VCAP entries, for example, to trigger copy to mirror port. For more information, see Section 2.7.4, VCAP IS2 .	Per VCAP entry

The following illustration shows a port mirroring example.

FIGURE 4-8: PORT MIRRORING EXAMPLE



All traffic to and from the server on port 3 (the probe port) is mirrored to port 2 (the mirror port). Note that the mirror port may become congested, because both the Rx frames and Tx frames on the probe port become Tx frames on the mirror port. The following mirror configuration is required:

```
ANA::PORT_CFG[3].SRC_MIRROR_ENA = 1
ANA::EMIRRORMASK[3] = 1
ANA::MIRRORPORTS = 0x0000004
```

In addition to the mirror configuration settings, the egress configuration of the mirror port (port 2) must be configured identically to the egress configuration of the probe port (port 3). This is to ensure that VLAN tagging and DSCP remarking at the mirror port is performed consistently with that of the probe port, such that the frame copies at the mirror port are identical to the original frames on the probe port.

Multiple mirror conditions, such as mirror multiple probe ports, VLANs, and so on, can be enabled concurrently to the same mirror port. However, in such configurations, it may not be possible to configure the egress part of the mirror port to perform tagging and DSCP remarking consistent with that of the original frame.

4.4 IGMP and MLD Snooping

This section provides information about the features and configurations related to Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) snooping.

By default, Layer-3 multicast data traffic is flooded in a Layer-2 network in the broadcast domain spanned by the VLAN. This causes unnecessary traffic in the network and extra processing of unsolicited frames in hosts not listening to the multicast traffic. IGMP and MLD snooping enables a Layer-2 switch to listen to IGMP and MLD conversations between host and routers. The switch can then prune multicast traffic from ports that do not have a multicast listener, and as a result, do not need a copy of the multicast frame. This is done by managing the multicast group addresses and the associated port masks.

IGMP is used to manage IPv4 multicast memberships, and MLD is used to manage IPv6 multicast memberships.

The device supports IGMPv2/v3 and MLDv1/v2. IGMPv2 and MLDv1 use any-source multicasting (ASM), where the multicast listener joins a group and can receive the multicast traffic from any source. IGMPv3 and MLDv2 introduce source-specific multicasting (SSM), where both source and group are specified by the multicast listener when joining a group.

The support in the device is two-fold:

- Control plane: IGMP and MLD frames are redirected to the CPU. This enables the CPU to listen to the queries and reports.
- Data plane: By monitoring the multicast group registrations and de-registrations signaled through the IGMP and MLD frames, the CPU can setup multicast group addresses and associated ports.

4.4.1 IGMP AND MLD SNOOPING CONFIGURATION

To implement IGMP and MLD snooping, any IGMP or MLD frames must be redirected to the CPU. For information about by the conditions by which such frames are identified, see [Section 2.6.7, CPU Forwarding Determination](#). IGMP and MLD frames can be independently snooped and assigned individual CPU extraction queues.

The following table lists the registers associated with configuring the redirection of IGMP and MLD frames to the CPU.

TABLE 4-33: CONFIGURATION REGISTERS FOR IGMP AND MLD FRAME REDIRECTION TO CPU

Register/Register Field	Description / Value	Replication
ANA::CPU_F-WD_CFG.IGMP_REDIR_ENA	Must be set to 1 to redirect IGMP frames to the CPU.	Per port
ANA::CPU_F-WD_CFG.MLD_REDIR_ENA	Must be set to 1 to redirect MLD frames to the CPU.	Per port
ANA::CPUQ_CFG.CPUQ_IGMP	CPU extraction queue for IGMP frames.	None
ANA::CPUQ_CFG.CPUQ_MLD	CPU extraction queue for MLD frames.	None

4.4.2 IP MULTICAST FORWARDING CONFIGURATION

The following table lists the registers associated with configuring the multicast group addresses and the associated ports.

TABLE 4-34: IP MULTICAST CONFIGURATION REGISTERS

Register/Register Field	Description / Value	Replication
MACHDATA	MAC address and VID when accessing the MAC table.	None
MACLDATA	MAC address when accessing the MAC table.	None
MACTINDX	Direct address into the MAC table for direct read and write.	None

TABLE 4-34: IP MULTICAST CONFIGURATION REGISTERS (CONTINUED)

Register/Register Field	Description / Value	Replication
MACACCESS	Flags and command when accessing the MAC table.	None
MACTOPTIONS	Flags when accessing the MAC table	None
FLOODING_IPMC	Index into the PGID table used for flooding of IPv4/6 multicast control and data frames.	None
PGID[63:0]	Destination and flooding masks table.	64
IS1_ACTION.FID_SEL	Specifies the use of IS1_ACTION.FID_VAL for the DMAC lookup, the SMAC lookup, or for both lookups.	Per IS1 entry
IS1_ACTION.FID_VAL	FID value.	Per IS1 entry

IPv4 and IPv6 multicast group addresses are programmed in the MAC table as IPv4 and IPv6 multicast entries. For more information, see [Section 2.8.1, MAC Table](#). The entry in the MAC table also holds the set of egress ports associated with the group address.

By default, programming an IPv4 or IPv6 multicast entry in the MAC table makes it an any-source multicast, because the actual source IP address is insignificant with respect to forwarding.

To create source-specific IPv4 or IPv6 multicast entries, the Filter Identifier (FID) action in VCAP IS1 can be used, which enables creation of specific FIDs per source IP address. Entries in IS1 can contain the full IPv4 or IPv6 source address. Multiple MAC table entries holding the same IPv4 or IPv6 multicast group address but different FIDs can then be created. This effectively enables source-specific multicasting.

The switch provides full control of flooding of unknown IP multicast frames. For more information, see [Table 108](#). Generally, an IGMP and MLD snooping switch disables flooding of unknown multicast frames, except to ports connecting to multicast routers. Note that unknown IPv4 multicast control frames should be flooded to all ports, because IPv4 is not as strict as IPv6 in terms of registration for IP multicast groups.

4.5 Quality of Service (QoS)

This section discusses features and configurations related to QoS.

The device includes a number of features related to providing low-latency guaranteed services to critical network traffic such as voice and video in contrast to best-effort traffic such as web traffic and file transfers.

All incoming frames are classified to a QoS class, which is used in the queue system when assigning resources, in the arbitration from ingress to egress queues and in the egress scheduler when selecting the next frame for transmission. The device provides two methods for classifying to a QoS class and for remarking priority information in the frame: Basic and Advanced classification.

Basic QoS classification enables predefined schemes for handling Priority Code Points (PCP), Drop Eligible Indicator (DEI), and Differentiated Service Code Points (DSCP):

- QoS classification based on PCP and DEI for tagged frames. The mapping table from PCP and DEI to QoS class is programmable per port.
- QoS classification based on DSCP values. Can optionally use only trusted DSCP values. The mapping table from DSCP value to QoS class is common between all ports.
- The device has the option to work as a DS boundary node connecting two DS domains together by translating incoming/outgoing DSCP values for selected ports.
- The DSCP values can optionally be remarked based on the frame's classified QoS class.
- For untagged or non-IP frames, a default per-port QoS class is programmable.

Advanced QoS classification uses the VCAP IS1, which provides a flexible classification:

- A large range of higher layer protocol fields (Layer 2 through Layer 4) are available for rule matching.
- The IS1 action vector returns a QoS class, and translations of PCP, DEI, and DSCP values are also possible.
- Through programming of entries in IS1, QoS rules can be made as specific as needed. For example; per source MAC address, per TCP/UDP destination port number, or combination of both.

For more information about advanced QoS classification using the VCAP IS1, see [Section 4.6.2, Ingress Control Lists](#).

4.5.1 BASIC QOS CONFIGURATION

The following table lists the registers associated with configuring basic QoS.

TABLE 4-35: BASIC QOS CONFIGURATION REGISTERS

Register	Description	Replication
ANA:PORT:QOS_CFG	QoS and DSCP configuration	Per port
ANA:PORT:QOS_P-CP_DEI_MAP_CFG:	Mapping of DEI and PCP to QoS class	Per port
ANA::DSCP_CFG	DSCP configuration	Per DSCP

Situation:

Assume a configuration with the following requirements:

- All frames with DSCP=7 must get QoS class 7.
- All frames with DSCP=8 must get QoS class 5.
- DSCP = 9 is untrusted and all frames with DSCP=9 should be treated as a non-IP frame.
- VLAN-tagged frames with PCP=7 must get QoS class 7
- All other IP frames must get QoS class 1.
- All other non-IP frames must get QoS class 0.

Solution:

```
# Program overall QoS configuration
QOS_CFG.QOS_DSCP_ENA = 1
QOS_CFG.QOS_PCP_ENA = 1

# Program DSCP trust configuration ("*" = 0 through 63)
DSCP_CFG[*].DSCP_TRUST_ENA = 1
DSCP_CFG[9].DSCP_TRUST_ENA = 0

# Program DSCP QoS configuration ("*" = 0 through 63)
DSCP_CFG[*].QOS_DSCP_VAL = 1
DSCP_CFG[7].QOS_DSCP_VAL = 7
DSCP_CFG[8].QOS_DSCP_VAL = 5

# Program PCP QoS configuration ("*" = 0 through 15)
# Note: both 7 and 15 are programmed in order to don't care DEI
QOS_PCP_DEI_MAP_CFG[*] = 0
QOS_PCP_DEI_MAP_CFG[7] = 7
QOS_PCP_DEI_MAP_CFG[15] = 7

# Program default QoS class for non-IP, non-tagged frames.
QOS_CFG.QOS_DEFAULT_VAL = 0
```

4.5.2 IPV4 AND IPV6 DSCP REMARKING

IPv4 and IPv6 packets include a 6-bit Differentiated Services Code Point (DSCP), which switches and routers can use to determine the QoS class of a frame. With a proper value in the DSCP field, packets can be prioritized consistently throughout the network. Compared to QoS classification based on user priority, classification based on DSCP provides two main advantages

- DSCP field is already present in all packets (assuming all traffic is IPv4/IPv6).
- DSCP value is preserved during routing and is therefore better suited for end-to-end QoS signaling.

Some hosts may be able to send packets with an appropriate value in the DSCP field, whereas other hosts may not provide an appropriate value in the DSCP field.

For packets without an appropriate value in the DSCP field, the device can be configured to write a new DSCP value into the frame, based on the QoS class of the frame. For example, the device may have determined the QoS class based on the VLAN tag priority information (PCP and DEI). After the packet is transmitted by the egress port, the DSCP field can be rewritten with a value based on the QoS class of the frame. Any subsequent routers or switches can then be easily prioritize the frame, based on the rewritten DSCP value.

The DSCP rewriting functionality available in the device provides flexible, per-ingress port and per-DSCP-value configuration of whether frames should be subject to DSCP rewrite. If it is determined at the ingress port that the DSCP value should be rewritten and to which value, this is then signaled to the egress ports, where the actual change of the DSCP field is done.

In addition, the IS1 can be programmed to return a DSCP value as part of the action vector. This value overrules the potential DSCP value coming out of the DSCP rewrite functionality described previously. A DSCP value from either the basic classification or the advanced IS1 classification obey the same egress rules for the actual DSCP remarking.

4.5.2.1 DSCP Remarking Configuration

The following table lists the configuration registers associated with DSCP remarking.

TABLE 4-36: CONFIGURATION REGISTERS FOR DSCP REMARKING

Register/Register Field	Description / Value	Replication
ANA:PORT:DSCP_REWR_CFG	Two-bit DSCP rewrite mode per ingress port. 0x0: No DSCP rewrite. 0x1: Rewrite only if the frame's current DSCP value is zero. 0x2: Rewrite only if the frame's current DSCP value is enabled for remarking in ANA::DSCP_CFG.DSCP_REWR_ENA. 0x3: Rewrite DSCP of all frames, regardless of current DSCP value.	Per ingress port
ANA::DSCP_CFG.DSCP_REWR_ENA	Enables specific DSCP values for rewrite for ports with DSCP rewrite mode set to 0x2.	Per DSCP
ANA::DSCP_REWR_CFG.DSCP_QOS_REWR_VAL	Maps the frame's DP level and QoS class to a DSCP value.	Per DP level and per QoS class
REW::DSCP_CFG.DSCP_REWR_CFG	Enables DSCP rewrite for egress port.	Per egress port
REW::DSCP_REMAP_CFG	Remap table of DSCP values.	None

The configuration related to the ingress port controls whether a frame is to be remarked. For each ingress port, a DSCP rewrite mode is configured in ANA:PORT:DSCP_REWR_CFG. This register defines the four different modes as follows:

- 0x0: No DSCP rewrite, that is, never change the received DSCP value.
- 0x1: Rewrite if DSCP is zero. This may be useful if a DSCP value of zero indicates that the host has not written any value to the DSCP field.
- 0x2: Rewrite selected DSCP values. In ANA::DSCP_CFG.DSCP_REWR_ENA specific DSCP values can be selected for rewrite, for example, if only certain DSCP values are allowed in the network.
- 0x3: Rewrite all DSCP values.

After a frame is selected for DSCP rewrite, based on the configuration for the ingress port, the new DSCP value is determined by mapping the QoS class and DP level to a new DSCP value (ANA::DSCP_REWR_CFG.DSCP_QOS_REWR_VAL).

This DSCP value is overruled by IS1 if a hit in IS1 returns an action vector with DSCP_ENA set.

The resulting DSCP value is forwarded to the Rewriter at the egress port, which determines whether to actually write the new DSCP value into the frame (REW::DSCP_CFG.DSCP_REWR_CFG). Optionally, the DSCP value may be translated before written into the frame (REW::DSCP_REMAP_CFG) for applications where the switch acts as an DS boundary node.

When an IPv4 DSCP is rewritten, the IP header checksum is updated accordingly.

4.5.3 VOICE OVER IP (VOIP)

This section provides information about QoS in applications with Voice over IP (VoIP).

In a typical workgroup switch application with VoIP phones, both workstations and VoIP phones are connected to the switch. A workstation can be connected through a VoIP phone. Traffic from the workstation is usually untagged, whereas traffic from the VoIP phone may or may not be tagged. The QoS classification mechanism applied on the access port depends on the capabilities of the VoIP phone; these capabilities vary from phone to phone. With different VoIP phone models in the network, different access ports require different QoS classification mechanisms. The access switch can perform QoS classification, depending on the VoIP phone model, to achieve consistent VoIP QoS across the network.

Voice traffic can be identified in the following ways.

- **Source MAC address (OUI):** Most vendors use a dedicated OUI for VoIP phones.
- **EtherType:** Legacy phones may use a special EtherType for VoIP.
- **VID:** A special VID used for voice traffic.
- **UDP Port Range:** Voice traffic often uses a well-known port range for the Real-time Transport Protocol (RTP).
- **DSCP or ToS Precedence:** Many phones can set the DSCP value or the ToS precedence bits.
- **Priority Code Point:** Many phones send VLAN tagged frames and can set the priority code point.

All of these identification methods are supported by QoS classification through IS1. They can be used to determine the VoIP traffic's QoS class when entering the switch. For more information about the IS1, see [Section 2.7.3, VCAP IS1](#).

To ensure consistent QoS across the network, frames can be remarked on the uplink port. Priority Code Points and DSCP values can be remarked based on the QoS class determined by the QCLs. For more information about Priority Code Point and DSCP remarking, see [Section 2.12.1, VLAN Editing](#), and [Section 4.5.2, IPv4 and IPv6 DSCP Remark-ing](#).

Traffic received on the uplink port can usually rely on simple DSCP or PCP QoS classification.

4.6 VCAP Applications

This section provides information about Versatile Content Aware Processor (VCAP) applications for QoS classification, source IP guarding, and access control.

The following table shows the different control lists that the VCAP can be used to build.

TABLE 4-37: CONTROL LISTS AND APPLICATION

Control List	Description
Ingress control lists (ICLs)	QoS classification VLAN classification and translation policy association group classification
IPv4 source guarding control lists (S4CLs)	IPv4 source guarding
IPv6 source guarding control lists (S6CLs)	IPv6 source guarding
Access control lists (ACLs)	Access control
Egress control lists (ECLs)	Tagging and egress translations

4.6.1 NOTATION FOR CONTROL LISTS ENTRIES

Setting up a control list typically requires a large amount of register configurations. To maintain the overview of the VCAP functionality, the following control list notations are used. The register configurations are not listed. For more information about the VCAP configurations, see [Section 3.7, VCAP](#).

The notation used is:

```
entry_number vcap entry_type {entry_field=value}  
→ {action_field=value}
```

Each control entry in the notation consists of:

- The entry number specifying the TCAM address for the specific TCAM
- The VCAP used (IS1, IS2, ES0)
- The entry type (for instance NORMAL or MAC_ETYPE).

- Zero, one, or more entry fields with specified values. If no value is supplied, it is assumed that the value is 1.
- The action (indicated with →)
- Zero, one, or more action fields with specified values. If no value is supplied, it is assumed that the value is 1.

All entry fields not listed in the entry part of the control entry are set to don't care.

All action fields not listed in the action part of the control entry are set to zero.

Default actions are special, because they do not have an entry type and a pattern to match:

```
default vcap (first|second) port=value
→ {action_field=value}
```

The notation is illustrated by the following examples.

Example 1:

An example of an ACL entry:

```
255 is2 ipv4_other first igr_port_mask=(1<<7) sip=10.10.12.134
→
```

This ACL entry is located in entry number 255. It is matched for the first lookup, and it is part of the port ACL for port 7. The type is `ipv4_other`, and the action is not to change the normal flow for frames with SIP = 10.10.12.134.

Example 2:

Policy ACL A can include a monitoring rule that disables forwarding and learning of all incoming IPv4 traffic, but redirects a copy to CPU extraction queue number 3 using the hit-me-once filter. The hit-me-once filter enables the CPU to control when it ready to accept a new frame. The rule would look like this:

```
254 is2 ipv4_other first pag=A
→ hit_me_once cpu_qu_num = 3
```

Example 3:

This example shows an ACE that allows forwarding and learning of ARP requests from port 7, if the source IP address is 10.10.12.134. The ACL entry also performs ARP sanity checks that frames must pass to match. The checks include checking that it is a Layer-2 broadcast, that the hardware address space is Ethernet, that the protocol address space is IP, that the MAC address and IP address lengths are correct, and that the sender hardware address (SMAC) matches the SMAC of the frame.

```
253 is2 arp first igr_port_mask=(1<<7) l2_bc opcode=arp_request
sip=10.10.12.134
arp_addr_space_ok arp_proto_space_ok arp_len_ok
arp_sender_match
→
```

Example 4:

If the default action from first lookup for port 7 is to discard all traffic, the following notation is used:

```
default is2 first port=7
→mask_mode=1 port_mask=0x0
```

4.6.2 INGRESS CONTROL LISTS

The following table lists the registers associated with advanced QoS configuration through Ingress Control Lists.

TABLE 4-38: ADVANCED QOS CONFIGURATION REGISTER OVERVIEW

Register	Description	Replication
ANA:PORT:QOS_CFG	QoS configuration	Per port

Situation:

Assume a configuration with the following requirements:

- All frames with DSCP = 7 must get QoS class 2.
- All frames with TCP/UDP port numbers in the range 0 – 1023 must get QoS class 3, except frames with TCP/UDP port 25, which must get QoS class 1.
- All other frames must get QoS class 0.

Solution:

The resulting QoS Control List looks like this:

```
255 isl normal first etype_len ip_snap dscp = 7
→ qos_ena=1, qos_val = 2
254 isl normal first etype_len ip_snap l4_sport = 25
→ qos_ena=1, qos_val = 1
253 isl normal first etype_len ip_snap etype = 25
→ qos_ena=1, qos_val = 1
252 isl normal first etype_len ip_snap l4_sport=(key:0, mask: 0x3FF)
→ qos_ena=1, qos_val = 3
251 isl normal first etype_len ip_snap etype = (key: 0, mask: 0x3FF)
→ qos_ena=1, qos_val = 3
```

ANA:PORT:QOS_CFG.QOS_DEFAULT_VAL = 0.

4.6.3 ACCESS CONTROL LISTS

The examples operate with three levels of ACLs:

- Port ACLs
- Policy ACLs
- Switch ACLs

The port ACLs are specific to a single port or a group of ports that form a link aggregation group. For example, a port ACL can be used for source IP filtering, locking a specific source IP address to a port. For more information about this example, see [Section 4.6.4.1, Restrictive SIP Filter Using IS2](#).

The policy ACLs are shared for a group of ports that must have the same policy applied. For example, there could be one policy for ports through which workstations access the network and another policy for ports to which servers are connected.

The switch ACLs apply to all ports of the switch. They specify some general rules that apply to all traffic passing through the switch. The rules can still be rather specific, for example, covering a specific VLAN or a specific IP address.

In the examples, the resulting ACL can include one port ACL, one policy ACL, and the switch ACL. This is determined by the way the ingress port mask (IGR_PORT_MASK) and the policy association group (PAG) are used. For information about IGR_PORT_MASK and PAG, see [Section 2.7.4, VCAP IS2](#). There are several ways to use the 8-bit PAG, but in this section, all eight bits are used to point out a policy ACL. The IGR_PORT_MASK points out the port ACL. This permits one port ACL per port and a total of 256 policy ACLs. Note that ports may share the same port ACL and a port by don't caring bits in the port ACL's IGR_PORT_MASK.

Each port has a default PAG assigned to it. The IS1 VCAP can be used to change the value of the PAG based on specific protocol fields matched in the IS1 lookup. The resulting PAG is used in the IS2 VCAP lookup and is matched against the PAG field of the ACL entries.

For an ACL entry in the IS2 VCAP, the PAG and IGR_PORT_MASK use this notation:

PAG = PolicyACL_ID

IGR_PORT_MASK = 1<<PortACL_ID

Notes The "<<" operator is the bitwise left shift operator. It shifts the left operand bit-wise to the left the number of positions specified by the right operand.

The IGR_PORT_MASK is a mask so the port number is left-shifted to create the mask.

For an ACL entry that is part of a port ACL for port 8, the PAG would be (*) and IGR_PORT_MASK would be (1<<8) = 0x100. The asterisk is a wildcard, which means that the PolicyACL_ID is a don't-care. For an ACL entry that is part of policy ACL A, the PAG would be (A) and the IGR_PORT_MASK would be (*). In this case, the PortACL_ID is a don't-care.


If, for example, port 8 must have policy A applied, the PAG assigned to port 8 is (A). Using this PAG value, the following ACLs match the lookup:

- Port ACL for port 8 with PAG = (*) and IGR_PORT_MASK = (1<<8)
- Policy ACL A with PAG = (A) and IGR_PORT_MASK = (*)
- Switch ACL with PAG = (*) and IGR_PORT_MASK = (*)

The ordering of the port ACL, the policy ACL, and switch ACL in the resulting ACL follows the ordering in the TCAM. In the following illustration, the switch ACL has the highest priority, followed by the policy ACL A, and finally, the port ACL for port 8.

The resulting ingress ACL in the example is made up of the ingress ACL entries in the switch ACL, the policy ACL A, the port ACL for port 8, and the default action for port 8. The VCAP also does a second lookup, for which the resulting ACL has a common default action as the last rule.

FIGURE 4-9: RESULTING ACL FOR LOOKUP WITH PAG = (A) AND IGR_PORT_MASK = (1<<8)



4.6.4 SOURCE IP FILTER (SIP FILTER)

The VCAP enables filtering of source IP (SIP) addresses on a port also known as source IP guarding. This can be used to only allow IP traffic from a specific SIP to enter the switch on a given port. Doing this can prevent the following denial of service (DoS) attacks: LAND attack, SMURF attack, SYN flood attack, Martian attack, and Ping attack.

4.6.4.1 Restrictive SIP Filter Using IS2

A restrictive SIP filter can be applied per port in networks where only IP traffic is allowed. The filter locks a specific SIP to the port and only permits ARP frames and IPv4 frames with the specified SIP to enter the switch on the given port.

For monitoring purposes, it is possible to permit IPv4 frames with other SIPs than the SIP locked to the port. The action is to redirect to the CPU, and the amount of traffic can be reduced by using the hit-me-once feature. The ACL entry for this can be part of a policy ACL for all ports on which the SIP filter is applied.

The port ACL has the following options:

- Permit IPv4 with trusted SIP
- Permit ARP with trusted SIP passing ARP sanity checks
- Permit all IPv4 — CPU redirect with hit-me-once filter (for monitoring)
- Default port action — discard all traffic

Situation:

Apply the restrictive SIP filter on port 7 with SIP 10.10.12.134.

Resolution:

The resulting ACL for port 7 looks like this:

```
255 is2 ipv4_tcp_udp first igr_port_mask=(1<<7) sip=10.10.12.134
→
254 is2 ipv4_other first igr_port_mask=(1<<7) sip=10.10.12.134
→
253 is2 arp first igr_port_mask=(1<<7) l2_bc opcode=arp_request
sip=10.10.12.134
arp_addr_space_ok arp_proto_space_ok arp_len_ok
arp_sender_match
→
252 is2 ipv4_tcp_udp first pag=A
→ hit_me_once cpu_queue=3
251 is2 ipv4_other first pag=A
→ hit_me_once cpu_queue=3
default is2 first port=11
→mask_mode=1 port_mask=0x0
```

Applying this SIP filter requires to two entries per port plus three common entries.

4.6.4.2 Restrictive SIP Filter Using IS1 and IS2

The same filter as listed above can be achieved using the host_match actions from IS1.

Situation:

Apply the restrictive SIP filter on port 7 with SIP 10.10.12.134.

Resolution:

The resulting ACL for port 7 looks like this:

IS1:

```
255 is1 smac_sip4 igr_port=7 sip=10.10.12.134
→ host_match
```

IS2:

```
255 is2 ip4_tcp_udp first host_match=1
→
254 is2 ip4_other first host_match=1
→
253 is2 arp first igr_port_mask=(1<<7) l2_bc opcode=arp_request
sip=10.10.12.134
arp_addr_space_ok arp_proto_space_ok arp_len_ok
arp_sender_match
→
252 is2 ipv4_tcp_udp first pag=A
→ hit_me_once cpu_queue=3
251 is2 ipv4_other first pag=A
→ hit_me_once cpu_queue=3
default is2 first port=7
→mask_mode=1 port_mask=0x0
```

Applying this SIP filter requires to one entry in IS1 per port and five common entries in IS2.

4.6.4.3 Less Restrictive SIP Filter Using IS2

For networks in which non-IP protocols are allowed, for example IPX and ARP, a less restrictive SIP filter can be applied with the following port ACL:

- Permit IPv4 with trusted SIP
- Discard all IPv4
- Default port action; Permit all traffic (non-IPv4, because all IPv4 traffic is covered by the ACL entries from other two items)

For monitoring purposes, the “Discard all IPv4” ACL can be changed to perform CPU redirect. This allows the CPU to monitor all incoming IPv4 frames with source IP addresses different from the trusted SIP, but without allowing these frames to be forwarded to other ports.

Situation:

Apply the less restrictive SIP filter on port 8 with source IP address 10.10.12.134, and monitor any IPv4 traffic with unauthorized source IP addresses with hit-me-once filtering to CPU extraction queue number 2. The monitoring rule is part of policy ACL A that is applied to all user ports.

Resolution:

The resulting ingress ACL for port 8 looks like this:

```
255 is2 ipv4_tcp_udp first igr_port_mask=(1<<8) sip=10.10.12.134
→
254 is2 ipv4_other first igr_port_mask=(1<<8) sip=10.10.12.134
→
63 is2 ipv4_tcp_udp first pag=A
→ hit_me_once cpu_queue=2
62 is2 ipv4_other first pag=A
→ hit_me_once cpu_queue=2
default is2 first port=10
→
```

Applying this SIP filter requires two entries per port plus two common entries.

4.6.5 DHCP APPLICATION

A DHCP application can be supported using one policy ACL for the user ports and another policy ACL for the DHCP server ports.

On the user ports, the DHCP requests must be snooped to be able to automatically reset the SIP filters that are applied per port. DHCP replies should be prevented from being forwarded from user ports. For monitoring purposes, such illegal replies are redirected to the CPU.

On the DHCP server ports, DHCP replies are snooped to be able to automatically update the SIP filter for the user port where the reply goes.

In addition, an egress rule is needed to prevent forwarding of all DHCP requests to user ports.

Situation:

Policy ACL A is used for the user port DHCP policy, and policy ACL B is used for the DHCP server policy. The server ports are ports 8 and 9.

Snoop DHCP requests from user ports in CPU extraction queue 1, using policer 0 to protect the CPU. DHCP replies from the servers are snooped in queue 2, and are also subject to policing with policer 0. The illegal DHCP replies from user ports are redirected to queue 3 using the hit-me-once filter.

Resolution:

The PAG assigned to the user ports is (A). The PAG assigned to the DHCP server ports (8 and 9) is (B).

The following shows the ACL entries for the DHCP application:

```
255 is2 ipv4_tcp_udp protocol=udp
sport=bootp_client dport=bootp_server
→ mask_mode=1 port_mask=0x0000300
63 is2 ipv4_tcp_udp first pag=A protocol=udp
sport=bootp_client dport=bootp_server
→ cpu_copy_ena cpu_queue=1 police_ena police_idx=0
```

```
62 is2 ipv4_tcp_udp first pag=A protocol=udp
sport=bootp_server dport=bootp_client
→ hit_me_once cpu_queue=3
31 is2 ipv4_tcp_udp first pag=B protocol=udp
sport=bootp_server dport=bootp_client
→ cpu_copy_ena cpu_queue=2 police_ena police_idx=0
default is2 first
→ mask_mode=1 port_mask=0x0
default is2 second
→
```

Regardless of the number of ports covered, four ACL entries are used: one in the switch ACL, two in policy ACL A, and one in policy ACL B.

4.6.6 ARP FILTERING

The VCAP support two useful ARP filters:

- Policing ARP requests to the switch's IP address to mitigate DoS attacks by ARP flooding
- Performing general ARP sanity checks

Because these are general rules, it is sensible to make them part of the switch ACL.

Situation:

Discard all ARP frames that do not pass the ARP sanity checks. Police ARP requests to the switch's IP address 10.10.12.1 using ACL policer 2. ACL policer 2 is configured to allow 16 frames per second, and the frames are copied to CPU extraction queue 0.

RARP is not allowed in the network.

Resolution:

To do ARP filtering in the switch ACL, perform the filtering for the switch's IP address first, then allow all ARP frames passing the sanity checks, and finally, discard all remaining ARP frames. This is illustrated by the following:

```
255 is2 arp first l2_bc opcode=arp_request
dip=10.10.12.1
arp_addr_space_ok arp_proto_space_ok arp_len_ok
arp_sender_match
→ cpu_copy_ena cpu_queue=0 police_ena police_idx=255
254 is2 arp first l2_bc opcode=(arp_request or arp_reply)
arp_addr_space_ok arp_proto_space_ok arp_len_ok
arp_sender_match
→
253 is2 arp
→mask_mode=1 port_mask=0x0
```

The ACL policer configuration for policer 255 is done as follows:

```
# Set the base unit to 1 frame per second, enable the policer, and set the rate
to 16 frames per second and a burst of 1 frame:
ANA:POL[255]:POL_MODE_CFG.FRM_MODE = 1
ANA:POL[255]:POL_PIR_CFG.PIR_RATE = 16
ANA:POL[255]:POL_PIR_CFG.PIR_BURST = 3
```

Three ACL entries are used, irrespective of the number of ports covered.

4.6.7 PING POLICING

The network can easily be protected against ping attacks using a switch ACL rule that applies an ACL policer to all ping packets.

Situation:

Allow no more than 128 ping packets per second to be forwarded through the switch by means of ACL policer 15. Ping packets in excess of 128 frames per second are discarded.

Resolution:

Ping packets are ICMP frames with ICMP Type = Echo Request. Echo Request is specified by the first byte of the ICMP frame being 0x08. The rest of the ICMP frame is don't-care. ICMP frames are carried in IPv4 frames with the protocol value 0x01.

The resulting switch ACL entry is as follows:

```
127 is2 ipv4_other first protocol=icmp ip4_payload_high=0x8*  
→ police_ena police_idx=15
```

ACL policer 15 in the policer pool is configured to 128 frames per second like this:

```
ANA:POL[15]:POL_MODE_CFG.FRM_MODE = 1  
ANA:POL[15]:POL_PIR_CFG.PIR_RATE = 128  
ANA:POL[15]:POL_PIR_CFG.PIR_BURST = 1
```

One ACE is used, regardless of the number of ports covered.

4.6.8 TCP SYN POLICING

A server in the network can be protected against TCP SYN DoS attacks by policing TCP connection requests to the server's IP address.

Situation:

Allow no more than 128 new TCP connections per second to the server with IP address 10.10.12.99. Use ACL policer 5.

Resolution:

TCP connection requests are TCP frames with the SYN flag set. The resulting switch ACL entry is as follows:

```
127 is2 ipv4_tcp_udp first protocol=tcp  
dip=10.10.12.99  
syn  
→ police_ena police_idx=5
```

ACL policer 5 in the policer pool is configured to 128 frames per second by the following:

```
ANA:POL[5]:POL_MODE_CFG.FRM_MODE = 1  
ANA:POL[5]:POL_PIR_CFG.PIR_RATE = 128  
ANA:POL[5]:POL_PIR_CFG.PIR_BURST = 1
```


One ACE is used, regardless of the number of ports covered.

4.7 CPU Extraction and Injection

This section provides information about how the CPU extracts and injects frames to and from the switch core.

The following illustration shows the CPU port module used for injection and extraction.

FIGURE 4-10: CPU EXTRACTION AND INJECTION



The switch core forwards CPU extracted frames to eight CPU extraction queues. Each of these queue is then mapped to one of two CPU Extraction Groups. For each extraction group there is a scheduler (strict or round robin) which selects between the CPU extraction queues mapped to the same group.

When injecting frames, there are two CPU Injection Groups available where for instance one can be used for the Frame DMA and one can be used for manually injected frames. Both CPU injection groups have access to the switch core. However, within the switch-core, CPU injected frames are all seen as coming from port 11 in terms of queue system and analyzer configuration.

4.7.1 FORWARDING TO CPU

Several mechanisms can be used to trigger redirection or copying of frames to the CPU. They are listed in the following table.

TABLE 4-39: CONFIGURATIONS FOR REDIRECTING OR COPYING FRAMES TO THE CPU

Frame Type	Configuration (Including Selection of Extraction Queue)	Copy or Redirect
IEEE 802.1D Reserved Range DMAC = 01-80-C2-00-00-0x	ANA:PORT:CPU_FWD_BPDU_CFG ANA::CPUQ_8021_CFG.CPUQ_BPDU_VAL	Copy or redirect
IEEE 802.1D Allbridge DMAC = 01-80-C2-00-00-10	ANA:PORT: CPU_FWD_CFG.CPU_ALL- BRIDGE_REDIRENA ANA::CPUQ_CFG.CPUQ_ALLBRIDGE	Copy or redirect
IEEE 802.1D GARP Range DMAC = 01-80-C2-00-00-2x	ANA:PORT:CPU_FWD_GARP_CFG ANA::CPUQ_8021_CFG.CPUQ_GARP_VAL	Copy or redirect
IEEE 802.1D CCM/Link Trace Range DMAC = 01-80-C2-00-00-3x	ANA:PORT:CPU_FWD_CCM_CFG ANA::CPUQ_8021_CFG.CPUQ_CCM_VAL	Copy or redirect
IGMP (IPv4)	ANA:PORT: CPU_FWD_CFG.CPU_IGMP_REDIRENA ANA::CPUQ_CFG.CPUQ_IGMP	Redirect
IP Multicast Control (IPv4)	ANA:PORT: CPU_FWD_CFG.CPU_IPMC_CTRL_- COPY_ENA ANA::CPUQ_CFG.CPUQ_IPMC_CTRL	Copy

TABLE 4-39: CONFIGURATIONS FOR REDIRECTING OR COPYING FRAMES TO THE CPU

Frame Type	Configuration (Including Selection of Extraction Queue)	Copy or Redirect
MLD (IPv6)	ANA:PORT: CPU_FWD_CFG.CPU_MLD_REDIR_ENA ANA::CPUQ_CFG.CPUQ_MLD	Redirect
Versatile Register Access Protocol (VRAP)	ANA:PORT: CPU_FWD_CFG.CPU_VRAP_REDIR_ENA ANA::CPUQ_CFG2.CPUQ_VRAP	Redirect
CPU-based learning	ANA:PORT:PORT_CFG.LEARNCPU ANA::CPUQ_CFG.CPUQ_LRN	Copy
CPU-based learning of locked MAC table entries seen on a new port	ANA:PORT: PORT_CFG.LOCKED_PORTMOVE_CPU ANA::CPUQ_CFG.CPUQ_LOCKED_PORTMOVE	
CPU-based learning of frames exceeding learn limit in MAC table	ANA:PORT:PORT_CFG.LIMIT_CPU ANA::CPUQ_CFG.CPUQ_LRN	
MAC table match using MAC table	ANA::MACACCESS.MAC_CPU_COPY ANA::CPUQ_CFG.CPUQ_MAC_COPY	Copy
MAC table match using PGID table	ANA::MACACCESS.DEST_IDX ANA::PGID.PGID (bit 11) ANA::PGID.CPUQ_DST_PGID	Redirect or copy
Flooded frames	ANA::MACACCESS.DEST_IDX ANA::PGID.PGID (bit 11) ANA::PGID.CPUQ_DST_PGID	Redirect or copy
Any frame received on selected ports	ANA:PORT:CPU_SRC_COPY_ENA ANA::CPUQ_CFG.CPUQ_SRC_COPY	Copy
Mirroring	ANA::MIRRORPORTS (bit 11) ANA::CPUQ_CFG.CPUQ_MIRROR For more information about mirroring, see Section 4.3.10, Mirroring .	Copy
VCAP IS2 rules	For more information about IS2, see Section 2.7.4, VCAP IS2 .	Redirect or copy
SFlow	ANA::CPUQ_CFG.CPUQ_SFLOW For more information about SFlow, see Section 3.8.4.9, sFlow Sampling .	Copy

4.7.2 FRAME EXTRACTION

The CPU receives frames through the eight CPU extraction queues in the CPU port module. The eight queues are using resources (memory and frame descriptor pointers) from the shared queue system and are subject to the thresholds and congestion rules programmed for the CPU port (port 11) and the shared queue system in general.

The CPU can read frames from the CPU extraction queues in two ways:

- Reading registers in the CPU port module. For more information, see [Section 2.14.1, Frame Extraction](#).
- FDMA from CPU port module to RAM. For more information, see [Section 4.7, Frame DMA](#).

In addition, the VRAP engine may attach to one of the CPU extraction groups and use this to receive and process all extracted VRAP requests.

The switch core may place the 20-byte long CPU extraction header before the DMAC or after the SMAC (SYS::PORT_MODE.INCL_XTR_HDR). The CPU extraction header contains relevant side band information about the frame such as the frame's classification result (VLAN tag information, DSCP, QoS class, Rx timestamp) and the reason for sending the frame to the CPU. For more information about the contents of the CPU extraction header, see [Section TABLE 2-140:, CPU Extraction Header](#).

4.7.3 FRAME INJECTION

The CPU can inject frames through the two CPU injection groups. The injection queues use resources (memory and frame descriptor pointers) from the shared queue system and are subject to the thresholds and congestion rules programmed for the CPU port (port 11) and the shared queue system in general.

The CPU can write frames to the CPU injection groups in two ways:

- Registers access to the CPU port module. For more information, see [Section 4.7, CPU Extraction and Injection](#).
- FDMA to CPU port module. For more information, see [Section 4.7, Frame DMA](#).

In addition, the VRAP engine may attach to one of the CPU injection groups and transmit and use this for transmitting VRAP replies.

The first 20 bytes of a frame written into a CPU group is an injection header containing relevant side band information about how the frame must be processed by the switch core (SYS::PORT_MODE.INCL_INJ_HDR). For more information, see [Table 144](#).

4.7.4 FRAME EXTRACTION AND INJECTION USING AN EXTERNAL CPU

The following table lists the configuration registers associated with using an external CPU.

TABLE 4-40: CONFIGURATION REGISTERS WHEN USING AN EXTERNAL CPU

Register/Register Field	Description / Value	Replication
QSYS::EXT_CPU_CFG.EXT_CPU_PORT	Port number where external CPU is connected.	None
QSYS::EXT_CPU_CFG.EXT_CPUQ_MSK	Configures which CPU extraction queues are sent to the external CPU.	None
SYS::PORT_MODE.INCL_XTR_HDR	Enables the insertion of the CPU extraction header in egress frames.	Per port
SYS::PORT_MODE.INCL_INJ_HDR	Enables ingress port to look for CPU injection header in incoming frames.	Per port

An external CPU can connect up to any front port module and use the Ethernet interface for extracting and injecting frames into the switch core.

Note If an external CPU is connected by means of the serial interface or PCIe interface, the frame extraction and injection is performed as described in [Section 4.7.2, Frame Extraction](#) and [Section 4.7.3, Frame Injection](#).

When injecting or extracting frames, the CPU injection or extraction header is placed before the DMAC with an optional prefix. For more information about the prefixing, see [Section 2.14.3, Network Processor Interface \(NPI\)](#). When injecting frames, the CPU injection header controls whether a frame is processed by the analyzer or forwarded directly to the destination set specified in the injection header.

An internal and external CPU may coexist in a dual CPU system where the two CPUs handles different run-time protocols. When extracting CPU frames, it is selectable which CPU extraction queues are connected to the external CPU and which remain connected to the internal CPU (SYS::EXT_CPU_CFG.EXT_CPUQ_MSK). If a frame is forwarded to the CPU for more than one reason (for example, a BPDU which is also a learn frame), the frame can be forwarded to both the internal CPU extraction queues and to the external CPU.

4.8 MPLS Operation

This section provides information about MPLS operation and how multiple block-level functions work together to implement MPLS operations.

It is assumed the reader is familiar with MPLS, MPLS-TP, and MEF technologies.

4.9 Ethernet over MPLS Service

VSC7416-01 supports connection-oriented delivery of Ethernet over MPLS as a point-point or point-multipoint (P2MP) Pseudowire. VSC7416-01 also supports bridging of Ethernet over MPLS as a Virtual Private LAN Service (VPLS), which can be E-LAN or E-TREE.

Provider Edge (PE) equipment must implement both Ethernet-to-MPLS operations and MPLS-to-Ethernet operations. These are also known as Pseudowire Label Edge Router (LER) operations.


Provider equipment that forwards MPLS packets using MPLS label information must implement MPLS to MPLS operations. These are known as Label Switch Router (LSR) operations. VSC7416-01 supports point-point and P2MP MPLS LSR operations.

Each of these operations (Ethernet to MPLS, MPLS to Ethernet, and MPLS to MPLS) requires an Arrival Service Point (ASP) and a Departure Service Point (DSP) for each port that is part of the service. Some examples are shown below.

Ethernet Pseudowire Example

In this example, a bi-directional MEF E-LINE EVC connects customer site 1 to customer site 2 over a service provider's MPLS network. VSC7416-01 can implement these LER and LSR functions.

FIGURE 4-11: ETHERNET PSEUDOWIRE EXAMPLE




Each LER is an end point for an Ethernet over MPLS pseudowire, which is delivered to the MPLS network over an MPLS LSP. So each LER is also an end point for an MPLS LSP. As part of this bi-directional EVC, each LER implements one Ethernet to MPLS operation, and one MPLS to Ethernet operation.

The LSRs are only aware of the MPLS LSP; they have no knowledge of the LSP contents, which in this example carries one pseudowire, but could carry any type and number of MPLS services including pseudowires and IP/MPLS. As part of this bi-directional LSP, each LSR implements two MPLS to MPLS operations.

H-VPLS Example

In this example, a bi-directional MEF E-TREE EVC aggregates mobile data from five base stations to the radio network controller/gateway. Hierarchical VPLS (H-VPLS) is used over a provider's MPLS network. VSC7416-01 can implement these H-VPLS Edge functions (called MTU-s in VPLS standards), where multiple Ethernet UNI interfaces are associated with one Attachment Circuit (AC).

FIGURE 4-12: H-VPLS EXAMPLE



This type of EVC allows a single multipoint service to connect many base stations to one controller/gateway. In this example, the base stations must not communicate to each other, so an E-TREE (also called Rooted Multipoint or RMP) EVC is used.

If the base stations were allowed to communicate to each other without restriction, an E-LAN EVC would have been used instead of the E-TREE EVC.

Each Virtual Switch Instance (VSI) in this example implements one Ethernet to MPLS operation per UNI port and one MPLS to Ethernet operation for the network port.

4.9.1 LER OPERATIONS: ETHERNET TO MPLS

The following blocks are essential to LER operation in the Ethernet to MPLS direction. For more information, see [Table 1-6](#).

- Frame classification, using the IS1 TCAM for service classification
- Path protection
- L2 forwarding
- Rewriter

4.9.1.1 Frame Classification

Ethernet to MPLS LER operation must use IS1 service classification, and is the same as Ethernet service classification. The result is that the frame is mapped to an Arrival Service Point indicated by ISDX, a classified C-domain VID (C-VID), a COS, and a DP value. The term C-domain means the bridging domain used by customer Ethernet frames.

To support Ethernet to MPLS processing, the ISDX provides a portmask used for E-LINE and E-TREE services and for mapping to a protection group.

All service frames belonging to the same EVC are classified to a C-VID dedicated to that EVC. Regardless if it is an E-LINE EVC or a VPLS service EVC (E-LAN or E-TREE), a dedicated C-VID is assigned. The C-VID identifies the internal VLAN resource used by the service.

COS and DP determine QoS handling and egress remarking.

4.9.1.2 Path Protection Switching

Each frame can map to one Path Protection group. In the Ethernet to MPLS direction, the arriving Ethernet frame may map to:

- A 1:1 protected pseudowire, using service-level protection. No Path Protection Group resources are consumed, as service-level protection is implemented by moving one service at a time.
- A 1:1 protected path LSP using a Path Protect group. All services or LSPs that map to a 1:1 protected Path LSP can be moved between the active and standby Path LSP in one operation, with unique encapsulations used for active and standby.
- Fast Reroute Facility Bypass Protect group. All services or LSPs that are facility-protected can be moved to/from a Bypass LSP in one operation, with unique encapsulations used for Bypass and non-Bypass. In this case, generally the pseudowire is mapped to an LSP and the LSP is facility-protected using a Path Protect group.

Ethernet to MPLS Path Protection functions are implemented in the following blocks:

- IS1 provides a per-ISDX Path Protect Group mapping
- The analyzer contains a per-Path state (either active or standby), as well as a per-ISDX alternate ISDX to use as ISDX_Protect. Ethernet frames are forwarded to the active MPLS Path using one of these two ISDX values. Failover control is provided in the Analyzer.
- The rewriter provides a per-ISDX discard control which is used to selectively filter hardware-generated OAM from the standby connection. Service OAM on a path-protected pseudowire must only travel down the active LSP, but Path OAM may need to travel down both the active and standby LSPs. Service OAM failover control is provided in the Rewriter.
- These are examples of upstream path protection.

4.9.1.3 L2 Forwarding

Arriving Ethernet frames were mapped to an ISDX and C-VID by Frame Classification.

For bridged services (VPLS), the L2 Forwarding block implements standard C-VID bridging functions, including learning based on MAC Source Addresses and forwarding based on MAC Destination Addresses. The result of C-VID bridging is a forwarding portmask.

L2 Forwarding is the combination of the ISDX-specific portmask and the C-VID bridging portmask.

Services with no bridging (such as E-LINE or configured P2MP) should be mapped to C-VID. These services will be forwarded purely based on the portmask configured with each ISDX, and will not use the MAC table for learning or forwarding.

Bridged E-LAN and E-TREE services use one C-VID per service, and will forward based on MAC address and a portset restricted by that C-VID. Forwarding can be further restricted based on arrival frame ISDX (for example, E-TREE).

4.9.1.4 Rewriter

In the Ethernet to MPLS direction, the Rewriter maps Arrival Service Point to Departure Service Point (one DSP per egress frame copy) using a lookup in the ES0 TCAM, and performs the following operations.

- Push/pop/translate up to two customer frame VLAN tags as described for Ethernet processing. For more information, see [Section 4.3.3, Provider Bridges and Q-in-Q Operation](#).
- Push a pseudowire Control Word (optional), an MPLS pseudowire Label, and up to 2 more MPLS LSP Labels. The TTL value in each label is a static value configured by software in the Rewriter ENCAP_ID table.
- Push an MPLS Link Layer encapsulation (Ethernet MAC header), including up to one VLAN tag. This is also part of the Rewriter ENCAP_ID table.
- Select the TC bit settings (MPLS labels) based on classified COS and DP. These values are pre-configured in the Rewriter by software for each supported COS and DP value.

For an E-LSP, the TC bit settings will vary based on both COS and DP.

For an L-LSP, the TC bit settings will vary based only on DP.

Each pushed label may independently be an E-LSP or an L-LSP.

- PCP/DEI bit settings (MPLS Link Layer) based on classified COS and DP.

- Modify inner DSCP/PCP/DEI as described for Ethernet processing. For more information, see [Section 4.5, Quality of Service \(QoS\)](#).

4.9.2 LER OPERATIONS: MPLS TO ETHERNET

The following blocks are essential to LER operation in the MPLS to Ethernet direction.

- MPLS and frame classification, using the IS0 and IS1 TCAMs
- Path protection switching
- L2 forwarding
- Rewriter

4.9.2.1 MPLS and Frame Classification

MPLS to Ethernet LER operation uses IS0 for MPLS Link Layer (MLL) and MPLS Label classification and IS1 for classifying client Ethernet frames that were delivered over MPLS.

If an arriving frame contains MPLS and the port has MPLS processing enabled, MPLS classification is performed. The first lookup is in IS0 and uses MPLS Link Layer fields. The second lookup is also in IS0 and uses the MPLS Labels. The next lookup(s) are in IS1 and use client Ethernet fields. MPLS lookups can fail due to TCAM miss, and the frame is then handled as an Ethernet frame using IS1.

The result of this classification is that the frame is mapped to an Arrival Service Point indicated by ISDX, a classified B-VID based on the MLL lookup, a classified C-VID based on the customer Ethernet lookup, a COS, and a DP value.

The MPLS Label classification indicates which MPLS Label contains the TC bits to use for this service. For an E-LSP, the indicated TC bits provide both COS and DP values. For an L-LSP, the COS value comes directly from the ISDX and only the DP value comes from the TC bits. Each label may independently be an E-LSP or an L-LSP.

The MPLS Label classification separately indicates which MPLS Label contains the TTL bits to use for this service. For an LER service end point, the TTL bits are not checked for expiration.

To support MPLS to Ethernet processing, the ISDX provides a portmask used for E-LINE and E-TREE services and mapping to a Path Protect group.

The classified B-VID is only used for Ethernet-based forwarding using MLL in the B-domain, and not for MPLS-based forwarding. For more information, see [Section 4.9.6, G.8032 Ethernet Ring Carrying MPLS](#).

The classified C-VID is primarily used by VPLS services (E-LAN or E-TREE), and indicates the internal VLAN resources to be used by this service. C-VID can also be used by E-LINE services if multiple E-LINES are mapped to one pseudowire.

COS and DP determine QoS handling and egress remarking.

4.9.2.2 Path Protection Switching

Each frame maps to at most one Path Protect group. In the MPLS to Ethernet direction, the arriving MPLS frame may belong to either the active or standby copy of a 1:1-protected pseudowire or LSP. A frame arriving on an active pseudowire or LSP is processed, whereas a frame arriving on a standby pseudowire or LSP is discarded.

The arriving MPLS frame may also belong to a facility-protected service, either in Bypass or non-Bypass state.

MPLS to Ethernet Path Protection functions are implemented in the following blocks:

- IS0 provides a per-ISDX Path Protect Group mapping
- The analyzer contains a per-Path state (either active or standby). MPLS frames arriving on the active path are forwarded, while MPLS frames arriving on the inactive path are discarded by the Analyzer. Failover control is provided in the analyzer.

These are examples of downstream path protection.

4.9.2.3 L2 Forwarding

Arriving Ethernet frames were mapped to an ISDX and a C-VID by MPLS and Frame Classification. The L2 Forwarding block implements standard bridging functions on the internal/customer Ethernet header for bridged services, including learning based on MAC Source Addresses and forwarding based on MAC Destination Addresses.

L2 Forwarding is the combination of the ISDX-specific portmask and the C-VID bridging portmask. L2 Forwarding may also be influenced by the B-VID bridging portmask, see [Section 4.9.6, G.8032 Ethernet Ring Carrying MPLS](#).

A service (EVC) with no MAC address learning (such as a non-learning E-LINE or configured P2MP) is also assigned a C-VID resource. Services assigned this C-VID are forwarded using constrained flooding and the portmask configured with each ISDX. Thereby asymmetric forwarding is possible within the EVC also without learning on the EVC.

Bridged E-LAN and E-TREE services use one C-VID per service and will forward based on MAC address and a portset restricted by that C-VID. Forwarding can be further restricted based on arrival frame ISDX (for example, E-TREE).

4.9.2.4 Rewriter

In the MPLS to Ethernet direction, the rewriter maps Arrival Service Point to Departure Service Point (one DSP per egress frame copy) using a lookup in the ES0 TCAM, and performs these operations:

- Pop the MPLS Link Layer encapsulation (Ethernet MAC header) including up to one VLAN tag
- Pop up to three MPLS LSP and pseudowire Labels and a pseudowire Control Word (if present).
- Push/pop/translate up to two customer frame VLAN tags as described for Ethernet processing.
- Modify inner DSCP/PCP/DEI as described for Ethernet processing.

For more information about Ethernet processing, see [Section 4.5, Quality of Service \(QoS\)](#).

4.9.3 LSR OPERATIONS: MPLS TO MPLS

The following blocks are essential to LSR operations:

- MPLS and Frame Classification, using the IS0 and IS1 TCAMs
- Path Protection
- L2 Forwarding
- Rewriter

4.9.3.1 MPLS Classification

MPLS to MPLS LSR operation uses IS0 for MPLS Link Layer (MLL) and MPLS Label classification. Standard LSR operation performs no classification on the client payload being carried by MPLS.

If an arriving frame contains MPLS and the port has MPLS processing enabled, MPLS classification is performed. The first lookup is in IS0 uses MPLS Link Layer fields. The second lookup is also in IS0 and uses the MPLS Labels. MPLS lookups can fail due to TCAM miss, and the frame is then handled as an Ethernet frame using IS1.

The result of this classification is that the frame is mapped to an Arrival Service Point indicated by ISDX, a classified B-VID based on the MLL lookup, a COS, and a DP value.

The MPLS Label classification indicates which MPLS Label contains the TC bits to use for this service (normally the top label on an LSR, unless the LSR packet arrives over another LSP). For an E-LSP, the indicated TC bits provide both COS and DP values. For an L-LSP, the COS value comes directly from the ISDX and only the DP value comes from the TC bits.

The MPLS Label classification separately indicates which MPLS Label contains the TTL bits to use for this service (normally the top label on an LSR, unless the LSR packet arrives over another LSP). For an LSR service, the TTL bits are first checked to be sure the TTL has not expired. The TTL value is then decremented and passed with the frame to be used by the Rewriter.

In support of MPLS to MPLS processing, the ISDX provides a portmask used for point-point and P2MP LSR services and mapping to a Path Protect group.

The classified B-VID is only used for Ethernet-based forwarding using MLL in the B-domain, and not for MPLS-based forwarding. For more information, see [Section 4.9.6, G.8032 Ethernet Ring Carrying MPLS](#).

COS and DP determine QoS handling and egress remarking.

4.9.3.2 Path Protection Switching

Each frame maps to one protection group. The arriving MPLS frame may belong to either the active or standby copy of a 1:1-protected LSP. A frame arriving on an active LSP is processed, whereas a frame arriving on a standby LSP is discarded.

The arriving MPLS frame may also belong to a facility-protected service, either in Bypass or non-Bypass.

MPLS to MPLS Path Protection functions are implemented in the following blocks:

- IS0 provides a per-ISDX Path Protect Group mapping
- The Analyzer contains a per-Path state (either active or standby). MPLS frames arriving on the active path are for-

warded, while MPLS frames arriving on the inactive path are discarded by the Analyzer. Failover control is provided in the Analyzer.

- Path Protection switching may be configured in either the upstream or downstream direction on each LSP, but not both on the same LSP.

4.9.3.3 L2 Forwarding

Arriving MPLS frames were mapped to an ISDX and a B-VID by MPLS Classification. The L2 Forwarding block does not implement any bridging functions with standard LSR services.

L2 Forwarding is therefore based on the ISDX-specific portmask.

4.9.3.4 Rewriter

The Rewriter maps Arrival Service Point to Departure Service Point (one DSP per egress frame copy) using a lookup in the ES0 TCAM, and performs these operations:

- Pop the MPLS Link Layer encapsulation (Ethernet MAC header) including up to one VLAN tag
- Pop up to two MPLS LSP Labels if the LSR service arrived over another LSP terminated at this LSR.
- Push up to two MPLS LSP Labels if the LSR service departs over another LSP initiated at this LSR. The TTL value in each pushed label is a static value configured by software.
- Swap one MPLS LSP or pseudowire Label as part of the LSR operation. Normally this is the top label of the arriving packet, unless the LSR packet arrives over another LSP that is terminated in this LSR. It is important to note that a swap operation is NOT the same as a pop followed by a push:
- With an MPLS Label push operation TTL, TC, and S bits are all configured by software in the Rewriter
- With an MPLS Label swap operation, TTL, TC and S bits come from MPLS Classifier processing
- Push an MPLS Link Layer encapsulation (Ethernet MAC header), including up to one VLAN tag
- Select the TC bit settings (pushed MPLS labels) and PCP/DEI bit settings (MPLS Link Layer) based on classified COS and DP.
- For an E-LSP, the TC bit settings of pushed MPLS Labels will vary based on both COS and DP
- For an L-LSP, the TC bit settings of pushed MPLS Labels will vary based only on DP
- Each pushed label may independently be an E-LSP or an L-LSP

4.9.4 MPLS AND PSEUDOWIRE OAM

VSC7416-01 supports detection of MPLS and Pseudowire OAM. Detected OAM is extracted to the CPU for processing or redirected to an external MPLS OAM engine. The following types of MPLS and Pseudowire OAM are supported.

- Pseudowire VCCV type 1: OAM is indicated by presence of an ACH following the lowest label. VCCV type 1 can only be used on pseudowires with the control word option.
- Pseudowire VCCV type 2: OAM is indicated by presence of a Router Alert Label above the pseudowire label. As a more general OAM mechanism, VSC7416-01 provides a configurable forwarding behavior for when the MPLS Classification detects a specified Reserved Label within scope. For each of the 16 Reserved Labels, VSC7416-01 provides a configurable CPU extract queue or the option to discard the frame.
- Pseudowire VCCV type 3: OAM is indicated by TTL=1 on the pseudowire or lowest label.
- MPLS-TP OAM: MPLS-TP OAM is supported on Segments (links), Path LSPs, and Pseudowires (VCCV type 4). OAM is indicated by reaching the GAL (Reserved Label value 13) as the lowest label.

VSC7416-01 supports detection and hardware processing of Ethernet OAM received over an MPLS pseudowire.

4.9.4.1 Using an External MPLS OAM Engine

VSC7416-01 supports the following features for use with an external MPLS OAM engine if attached to the NPI port:

- Detection of the above OAM types, intelligent Down-MEP and Up-MEP extraction
- Extracted OAM includes arrival port and classified service in extraction header
- Extracted OAM can be mapped to its own ASP for statistics and policing
- Intelligent injection of OAM from an external engine into Down-MEP and Up-MEP locations
- Injected OAM may use injection header to bypass classification and directly control ISDX, COS, tag/remark options, and destination port set
- Injected OAM can be mapped to a DSP per destination port for statistics and per-port encapsulation of multicast OAM

4.9.5 IEEE 1588 OVER MPLS

MPLS supplies important traffic engineering, provisioning control, and protection capabilities enabling PTP to meet the tight LTE base station requirements of $< 1.5 \mu\text{s}$ phase error across the network. Some applications require even much tighter phase alignments.

When using IEEE 1588 to deliver phase synchronization, the 1588 slave (LTE base station) must determine its one-way latency from the 1588 master or BC with very high precision ($<< 1.5 \mu\text{s}$). It does so by sending delay request from slave back to master or BC, and the master or BC responds with delay response. The slave then has an accurate measure of the round-trip delay and assumes the one-way delay is exactly half of the round-trip delay.

In most networks, the delay in each direction would not be identical due to queuing delay. Further, without MPLS there may be no way to guarantee that upstream and downstream measurements follow the same path through the network (asymmetry).

Transparent clock is used to negate the effect of asymmetric queuing delays inside each box.

MPLS is used to establish a bidirectional connection between slave and master/BC, where it is guaranteed that upstream and downstream directions of this connection follow the same path through the network.

It is therefore critical to meeting LTE phase requirements that VSC7416-01 supports timestamping of PTP frames within MPLS, especially as part of a transparent clock function. VSC7416-01 supports a transparent clock all in hardware (no software interaction), for the following PTP over MPLS encapsulations:

- PTP/UDP/IPV4/MPLS (IP/MPLS)
- PTP/UDP/IPV6/MPLS (IP/MPLS)
- PTP/UDP/IPV4/Ethernet/PW/MPLS (PW)
- PTP/UDP/IPV6/Ethernet/PW/MPLS (PW)
- PTP/Ethernet/PW/MPLS (PW)

4.9.6 G.8032 ETHERNET RING CARRYING MPLS

VSC7416-01 supports the use of Ethernet-layer ring per G.8032, with MPLS functionality supported for frames added to the ring and dropped from the ring.

This configuration supports the following functions:

- Pure Ethernet customer service mapped to G.8032 Ethernet ring.
- Ethernet Pseudowire Provider Edge (LER) service, transported over G.8032 Ethernet ring.
- MPLS LSR service, transported over G.8032 Ethernet ring.
- G.8032 Ethernet Ring port functions

Frames arriving from a customer port and forwarded to a ring port use standard processing as described elsewhere.

Non-MPLS frames arriving from the ring port are processed using standard Ethernet processing.

MPLS frames arriving from a ring port use specialized processing, as described here.

If MPLS is enabled on the ring port, MPLS frames are classified using IS0. In the first IS0 lookup (MLL), it is determined whether to forward this frame to another ring port or to take this frame off the ring for local processing/forwarding. It is also possible to forward a copy to a ring port and also take one or more local copies (Ethernet Drop and Continue). In the event of an IS0 TCAM miss, the frame is forwarded as Ethernet.

The IS0 MLL controls the Ethernet forwarding around the Ethernet ring. MAC table forwarding using MLL fields is not used in this configuration.

The IS0 MLBS entries fully control the MPLS forwarding, which can be to any non-ring or ring port.

IS1 entries can be used to control pseudowire client frame forwarding based on client Ethernet fields. MAC table forwarding using client Ethernet fields is supported in this configuration.

5.0 REGISTERS

This section provides information about the programming interface, register maps, register descriptions, and register tables of the device.

In writing to registers with reserved bits, use a read-modify-write technique, where the entire register is read, but with only the user bits to be changed are modified. Do not change the values of the register and bits marked as reserved. Their read state should not be considered static or unchanging. Unspecified registers and bits must be written to 0 and can be ignored when read.

The first-level table lists all register targets and associated base addresses. The second-level table lists registers groups and offsets within targets, and the third-level tables list registers within the register groups.

Both register groups and registers may be replicated (repeated) a number of times. The repeat-count and the distance between two repetitions are listed in the Instances and Address Spacing column. Spacing is omitted if there is only one instance. The Offset within Target and the Offset within Register Group columns hold the offset of the first instance of the register group or register.

To calculate the absolute address of a given register, multiply the register group's replication number by the register group's address spacing, and add the result to the register group's offset within the target. Then multiply the register's replication number with the register's address spacing, and add the result to the register's offset within the register group. Add these two numbers to the absolute address of the target in question.

TABLE 5-1: LIST OF TARGETS AND BASE ADDRESSES

Target Name	Base Address	Description	Details
ANA	0x71900000	Analyzer Configuration	Page 321
DEV[0]	0x711E0000	Port Configuration	Page 367
DEV[1]	0x711F0000	Port Configuration	Page 367
DEV[2]	0x71200000	Port Configuration	Page 367
DEV[3]	0x71210000	Port Configuration	Page 367
DEV[8]	0x71260000	Port Configuration	Page 367
DEV[9]	0x71270000	Port Configuration	Page 367
DEV[10]	0x71280000	Port Configuration	Page 367
DEVCPU_GCB	0x71070000	CPU device configuration	Page 385
DEVCPU_ORG	0x71000000	CPU device origin	Page 414
DEVCPU_QS	0x71080000	CPU device queue system	Page 420
ES0	0x71040000	VCAP ES0 Configuration	Page 490
HSIO	0x710A0000	High Speed I/O SerDes Configuration	Page 426
ICPU_CFG	0x70000000	VCore configuration	Page 445
IS0	0x710B0000	VCAP IS0 Configuration	Page 490
IS1	0x71050000	VCAP IS1 Configuration	Page 490
IS2	0x71060000	VCAP IS2 Configuration	Page 490
OAM_MEP	0x710C0000	Versatile OAM MEP Processor (VOP) block	Page 499
PCIE	0x70101000	PCIe endpoint configuration space	Page 545
QSYS	0x71800000	Queue System Configuration	Page 551
REW	0x71030000	Rewriter Configuration	Page 569
SBA	0x70110000	Shared bus arbiter	Page 574
SYS	0x71010000	Switching Engine Configuration	Page 576
TWI	0x70100400	Two-Wire Interface controller	Page 590
UART	0x70100000	UART controller	Page 608
UART2	0x70100800	UART controller	Page 608

5.1 ANA

TABLE 5-2: REGISTER GROUPS IN ANA

Register Group Name	Offset within Target	Instances and Address Spacing	Description	Details
ANA	0x00014000	1	Common analyzer registers	Page 321
PGID	0x00009C00	1	Port Group Identifier table	Page 330
ANA_TABLES	0x00009B00	1	Analyzer tables	Page 332
MSTI_STATE	0x00009E00	1	Multiple Spanning Tree states	Page 338
OAM_UPM_LM_CNT	0x00009000	1	OAM Up-MEP loss measurement counters	Page 339
MPLS_CL	0x000140B4	1	Configuration of the MPLS-TP classifier	Page 340
MPLS_TC_MAP	0x00014180	8 0x00000020	TC to (DP, QoS) mapping tables	Page 343
PORT	0x0000C000	64 0x00000100	Classifier configuration per port	Page 344
PFC	0x00009800	11 0x00000040	Priority-based flow control configuration	Page 355
IPT	0x0000A000	1024 0x00000008	Protection table per ISDX	Page 356
PPT	0x00009AC0	1	Path group protection table per path	Page 357
FID_MAP	0x00010000	1	FID mapping table per VID	Page 358
COMMON	0x00014280	1	Common configurations for classifier	Page 359
POL	0x00000000	448 0x00000020	Policer configuration	Page 364
POL_MISC	0x00009B80	1	Policer flow control configuration	Page 366

5.1.1 ANA:ANA

Parent: [ANA](#)

Instances: 1

TABLE 5-3: REGISTERS IN ANA

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
ADVLEARN	0x00000000	1	Advanced learning setup	Page 322
VLANMASK	0x00000004	1	VLAN source port mask	Page 322
PORT_B_DOMAIN	0x00000008	1	B-domain setup	Page 323
ANAGEFIL	0x0000000C	1	Aging filter	Page 323
ANEVENTS	0x00000010	1	Event sticky bits	Page 324
STORMLIMIT_BURST	0x00000014	1	Storm policer burst	Page 325
STORMLIMIT_CFG	0x00000018	4 0x00000004	Storm policer configuration per storm policer	Page 325

TABLE 5-3: REGISTERS IN ANA (CONTINUED)

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
ISOLATED_PORTS	0x00000028	1	Private VLAN mask for isolated ports	Page 326
COMMUNITY_PORTS	0x0000002C	1	Private VLAN mask for community ports	Page 326
AUTOAGE	0x00000030	1	Auto-age timer	Page 327
MACTOPTIONS	0x00000034	1	MAC table options	Page 327
LEARNDISC	0x00000038	1	Learn discard counter	Page 328
AGENCTRL	0x0000003C	1	Analyzer configuration	Page 328
MIRRORPORTS	0x00000040	1	Mirror target ports	Page 329
EMIRRORPORTS	0x00000044	1	Egress mirror mask	Page 329
FLOODING	0x00000048	1	Standard flooding configuration	Page 329
FLOODING_IPMC	0x0000004C	1	Flooding configuration for IP multicasts	Page 329
SFLOW_CFG	0x00000050	12 0x00000004	SFlow sampling configuration per switch port	Page 330
PORT_MODE	0x00000080	13 0x00000004	Port configuration per device port	Page 330

5.1.1.1 ANA:ANA:ADVLEARN

Parent: [ANA:ANA](#)

Instances: 1

TABLE 5-4: FIELDS IN ADVLEARN

Field Name	Bit	Access	Description	Default
VLAN_CHK	11	R/W	If this bit is set, a frame discarded because of VLAN ingress filtering is not subject to learning. VLAN ingress filtering is controlled by the VLAN_SRC_CHK flag in the VLAN table (see VLANACCESS register) or the VLANMASK register.	0x0
LEARN_MIRROR	10:0	R/W	Learn frames are also forwarded to ports marked in this mask.	0x000

5.1.1.2 ANA:ANA:VLANMASK

Parent: [ANA:ANA](#)

Instances: 1

TABLE 5-5: FIELDS IN VLANMASK

Field Name	Bit	Access	Description	Default
VLANMASK	11:0	R/W	Mask for requiring VLAN ingress filtering. If the bit for the frame's physical ingress port is set in this mask, then the port must be member of ingress frame's VLAN (VLANACCESS.VLAN_PORT_MASK), otherwise the frame is discarded.	0x000

5.1.1.3 ANA:ANA:PORT_B_DOMAIN

Parent: [ANA:ANA](#)

Instances: 1

TABLE 5-6: FIELDS IN PORT_B_DOMAIN

Field Name	Bit	Access	Description	Default
PORT_B_DOMAIN	11:0	R/W	Ports marked in this mask are B-ports. Cleared ports are C-ports. The mask influences the VLAN mask. If ANA::PORT_CFG.USE_B_DOM_TBL is set for the ingress port profile, then the VLAN mask from the VLAN table is AND'ed with the B-port mask. If ANA::PORT_CFG.USE_B_DOM_TBL is cleared for the ingress port profile, then the VLAN mask for service frames (ISDX>0) is AND'ed with the C-port mask and OR'ed with the B-port mask, and the VLAN mask for non-service frames (ISDX=0) is the VLAN mask AND'ed with the C-port mask.	0x000

5.1.1.4 ANA:ANA:ANAGEFIL

Parent: [ANA:ANA](#)

Instances: 1

This register sets up which entries are touched by an aging operation (manual as well as automatic aging).

In this way, it is possible to have different aging periods in each VLAN and to have quick removal of entries on specific ports.

The register also affects the GET_NEXT MAC table command. When using the register to control the behavior of GET_NEXT, it is recommended to disable automatic aging while executing the GET_NEXT command.

TABLE 5-7: FIELDS IN ANAGEFIL

Field Name	Bit	Access	Description	Default
B_DOM_EN	22	R/W	If set, only MAC table entries with B-domain flag B_DOM matching B_DOM_VAL are aged. This is not used with GET_NEXT MAC table searching.	0x0
B_DOM_VAL	21	R/W	B-domain flag used in selective aging.	0x0

TABLE 5-7: FIELDS IN ANAGEFIL (CONTINUED)

Field Name	Bit	Access	Description	Default
AGE_LOCKED	20	R/W	Select entries to age. If cleared, unlocked entries are aged and potentially removed. If set, locked entries are aged but not removed.	0x0
PID_EN	19	R/W	If set, only MAC table entries with a destination index matching PID_VAL are aged or searched with GET_NEXT.	0x0
PID_VAL	18:14	R/W	Destination index used in selective aging or MAC table searching.	0x00
VID_EN	13	R/W	If set, only MAC table entries with a VID matching VID_VAL are aged or searched with GET_NEXT.	0x0
VID_VAL	12:0	R/W	VID used in selective aging or MAC table searching.	0x0000

5.1.1.5 ANA:ANA:ANEVENTS

Parent: [ANA:ANA](#)

Instances: 1

TABLE 5-8: FIELDS IN ANEVENTS

Field Name	Bit	Access	Description	Default
MSTI_DROP	27	Sticky	A frame was discarded due to blocking MSTI.	0x0
AUTOAGE	24	Sticky	An AUTOAGE run was performed.	0x0
STORM_DROP	22	Sticky	A frame was discarded, because it exceeded the flooding storm limitations configured in STORMLIMIT.	0x0
LEARN_DROP	21	Sticky	A frame was discarded, because it was subject to learning, and the DropMode flag was set in ADVLEARN.	0x0
AGED_ENTRY	20	Sticky	An entry was removed at CPU Learn, or CPU requested an aging process.	0x0
CPU_LEARN_FAILED	19	Sticky	A learn operation failed due to hash table depletion. CPU-based learning only.	0x0
AUTO_LEARN_FAILED	18	Sticky	A learn operation of incoming source MAC address failed due to hash table depletion. Hardware-based learning only.	0x0
LEARN_REMOVE	17	Sticky	An entry was removed when learning a new source MAC address.	0x0
AUTO_LEARNED	16	Sticky	An entry was learned from an incoming frame. Hardware-based learning only.	0x0
AUTO_MOVED	15	Sticky	A station was moved to another port.	0x0
CLASSIFIED_DROP	13	Sticky	A frame was not forwarded due to classification (such as BPDUs).	0x0
CLASSIFIED_COPY	12	Sticky	A frame was copied to the CPU due to classification.	0x0
VLAN_DISCARD	11	Sticky	A frame was discarded due to lack of VLAN membership on source port.	0x0

TABLE 5-8: FIELDS IN ANEVENTS (CONTINUED)

Field Name	Bit	Access	Description	Default
FWD_DISCARD	10	Sticky	A frame was discarded due to missing forwarding state on source port.	0x0
MULTICAST_FLOOD	9	Sticky	A frame was flooded with multicast flooding mask.	0x0
UNICAST_FLOOD	8	Sticky	A frame was flooded with unicast flooding mask.	0x0
DEST_KNOWN	7	Sticky	A frame was forwarded with known destination MAC address.	0x0
BUCKET3_MATCH	6	Sticky	A destination was found in hash table bucket 3.	0x0
BUCKET2_MATCH	5	Sticky	A destination was found in hash table bucket 2.	0x0
BUCKET1_MATCH	4	Sticky	A destination was found in hash table bucket 1.	0x0
BUCKET0_MATCH	3	Sticky	A destination was found in hash table bucket 0.	0x0
CPU_OPERATION	2	Sticky	A CPU-initiated operation on the MAC or VLAN table was processed. Default is 1 due to auto-initialization of the MAC and VLAN table.	0x1
DMAC_LOOKUP	1	Sticky	A destination address was looked up in the MAC table.	0x0
SMAC_LOOKUP	0	Sticky	A source address was looked up in the MAC table.	0x0

5.1.1.6 ANA:ANA:STORMLIMIT_BURST

Parent: [ANA:ANA](#)

Instances: 1

TABLE 5-9: FIELDS IN STORMLIMIT_BURST

Field Name	Bit	Access	Description	Default
STORM_BURST	3:0	R/W	Allowed number of frames in a burst is $2^{**}STORM\_BURST$. The maximum allowed burst is 4096 frames, which corresponds to $STORM\_BURST = 12$. The $STORM\_BURST$ is common for all storm policers.	0x0

5.1.1.7 ANA:ANA:STORMLIMIT_CFG

Parent: [ANA:ANA](#)

Instances: 4

0: UC storm policer

1: BC storm policer

2: MC policer

3: Learn policer

TABLE 5-10: FIELDS IN STORMLIMIT_CFG

Field Name	Bit	Access	Description	Default
STORM_RATE	6:3	R/W	Allowed rate of storm policer is 2**STORM_UNIT frames/sec or kiloframes/sec (see STORM_UNIT). The maximum allowed rate is 1024 kiloframes/sec, which corresponds to STORM_RATE = 10 with STORM_UNIT set to 0.	0x0
STORM_UNIT	2	R/W	If set, the base unit for the storm policer is 1 frame per second. If cleared, the base unit is 1 kiloframes per second.	0x0
STORM_MODE	1:0	R/W	Mode of operation for storm policer. 0: Disabled 1: Police CPU destination only. 2: Police front port destinations only. 3: Police both CPU and front port destinations.	0x0

5.1.1.8 ANA:ANA:ISOLATED_PORTS

Parent: [ANA:ANA](#)

Instances: 1

TABLE 5-11: FIELDS IN ISOLATED_PORTS

Field Name	Bit	Access	Description	Default
ISOL_PORTS	11:0	R/W	This mask is used in private VLANs applications. Promiscuous and community ports must be set and isolated ports must be cleared. For frames classified to a private VLAN (see the VLAN_PRIV_VLAN field in VLAN table), the resulting VLAN mask is calculated as follows: - Frames received on a promiscuous port use the VLAN mask directly. - Frames received on a community port use the VLAN mask AND'ed with the ISOL_PORTS. - Frames received on a isolated port use the VLAN mask AND'ed with the COMM_PORTS AND'ed with the ISOL_PORTS. For frames classified to a non-private VLAN, this mask is not used.	0xFFFF

5.1.1.9 ANA:ANA:COMMUNITY_PORTS

Parent: [ANA:ANA](#)

Instances: 1

TABLE 5-12: FIELDS IN COMMUNITY_PORTS

Field Name	Bit	Access	Description	Default
COMM_PORTS	11:0	R/W	This mask is used in private VLANs applications. Promiscuous and isolated ports must be set and community ports must be cleared. See ISOLATED_PORTS.ISOL_PORTS for details.	0xFFFF

5.1.1.10 ANA:ANA:AUTOAGE

Parent: [ANA:ANA](#)

Instances: 1

TABLE 5-13: FIELDS IN AUTOAGE

Field Name	Bit	Access	Description	Default
AGE_FAST	21	R/W	Sets the unit of AGE_PERIOD to 13.1 us. AGE_PERIOD must be a minimum of 3 when using the FAST option.	0x0
AGE_PERIOD	20:1	R/W	Time in seconds between automatic aging of a MAC table entry. Setting AGE_PERIOD to zero effectively disables automatic aging. An inactive unlocked MAC table entry is aged after 2*AGE_PERIOD.	0x00000
AUTOAGE_LOCKED	0	R/W	Enable setting of the AGED_FLAG for locked entries in the MAC table when doing an age scan. Locked entries are never removed but the AGED_FLAG can be used to indicate activity for the MAC address.	0x0

5.1.1.11 ANA:ANA:MACTOPTIONS

Parent: [ANA:ANA](#)

Instances: 1

TABLE 5-14: FIELDS IN MACTOPTIONS

Field Name	Bit	Access	Description	Default
REDUCED_TABLE	1	R/W	When set, the MAC table is reduced to 256 entries (64 rows of 4 entries).	0x0
SHADOW	0	R/W	Enable MAC table shadow registers. The SHADOW bit affects the behavior of the READ command in MACACCESS.MAC_TABLE_CMD: With the shadow bit set, reading bucket 0 causes the remaining 3 buckets in the row to be stored in "shadow registers". Following read accesses to bucket 1-3 return the content of the shadow registers. This is useful when reading a MAC table, which can change while being read.	0x0

VSC7416-01

5.1.1.12 ANA:ANA:LEARNDISC

Parent: [ANA:ANA](#)

Instances: 1

The total number of MAC table entries that have been or would have been learned, but have been discarded due to a lack of storage space.

TABLE 5-15: FIELDS IN LEARNDISC

Field Name	Bit	Access	Description	Default
LEARNDISC	31:0	R/W	Number of discarded learn requests due to MAC table overflow (collisions or MAC table entry limits).	0x00000000

5.1.1.13 ANA:ANA:AGENCTRL

Parent: [ANA:ANA](#)

Instances: 1

TABLE 5-16: FIELDS IN AGENCTRL

Field Name	Bit	Access	Description	Default
IGNORE_DMACE_FLAGS	11	R/W	Do not react to flags found in the DMACE entry or the corresponding flags for flooded frames (FLOOD_IGNORE_VLAN).	0x0
IGNORE_SMACE_FLAGS	10	R/W	Do not react to flags found in the SMACE entry. Note, the IGNORE_VLAN flag is not checked for SMACE entries.	0x0
FLOOD_SPECIAL	9	R/W	Flood frames using the lowest 12 bits of DMACE as destination port mask. This is only added for testing purposes.	0x0
FLOOD_IGNORE_VLAN	8	R/W	VLAN mask is not applied to flooded frames.	0x0
MIRROR_CPU	7	R/W	Frames destined for the CPU extraction queues are also forwarded to the port set configured in MIRRORPORTS.	0x0
LEARN_CPU_COPY	6	R/W	If set, auto-learned stations get the CPU_COPY flag set in the MAC table entry.	0x0
LEARN_SRC_KILL	5	R/W	If set, auto-learned stations get the SRC_KILL flag set in the MAC table entry.	0x0
LEARN_IGNORE_VLAN	4	R/W	If set, auto-learned stations get the IGNORE_VLAN flag set in the MAC table entry.	0x0
CPU_CPU_KILL_ENA	3	R/W	If set, CPU injected frames are never sent back to the CPU.	0x1
GREEN_COUNT_MODE	2	R/W	Counter mode for the Rx priority counters for green frames (C_RX_GREEN_PRIO_x) 0: Count octets 1: Count frames	0x1
YELLOW_COUNT_MODE	1	R/W	Counter mode for the Rx priority counters for yellow frames (C_RX_YELLOW_PRIO_x) 0: Count octets 1: Count frames	0x1

TABLE 5-16: FIELDS IN AGENCTRL (CONTINUED)

Field Name	Bit	Access	Description	Default
RED_COUNT_MODE	0	R/W	Counter mode for the Rx priority counters for red frames (C_RX_RED_PRIO_x) 0: Count octets 1: Count frames	0x1

5.1.1.14 ANA:ANA:MIRRORPORTS

Parent: [ANA:ANA](#)

Instances: 1

TABLE 5-17: FIELDS IN MIRRORPORTS

Field Name	Bit	Access	Description	Default
MIRRORPORTS	11:0	R/W	Ports set in this mask receive a mirror copy. If CPU is included in mask (most significant bit set), then the frame is copied to CPU extraction queue CPUQ_CFG.CPUQ_MIRROR.	0x000

5.1.1.15 ANA:ANA:EMIRRORPORTS

Parent: [ANA:ANA](#)

Instances: 1

TABLE 5-18: FIELDS IN EMIRRORPORTS

Field Name	Bit	Access	Description	Default
EMIRRORPORTS	11:0	R/W	Frames forwarded to ports in this mask are mirrored to the port set configured in MIRRORPORTS (i.e. egress port mirroring).	0x000

5.1.1.16 ANA:ANA:FLOODING

Parent: [ANA:ANA](#)

Instances: 1

TABLE 5-19: FIELDS IN FLOODING

Field Name	Bit	Access	Description	Default
FLD_UNICAST	17:12	R/W	Set the PGID mask to use when flooding unknown unicast frames.	0x3F
FLD_BROADCAST	11:6	R/W	Set the PGID mask to use when flooding unknown broadcast frames.	0x3F
FLD_MULTICAST	5:0	R/W	Set the PGID mask to use when flooding unknown multicast frames (except IP multicasts).	0x3F

5.1.1.17 ANA:ANA:FLOODING_IPMC

Parent: [ANA:ANA](#)

Instances: 1

TABLE 5-20: FIELDS IN FLOODING_IPMC

Field Name	Bit	Access	Description	Default
FLD_MC4_CTRL	23:18	R/W	Set the PGID mask to use when flooding unknown IPv4 Multicast Control frames.	0x3F
FLD_MC4_DATA	17:12	R/W	Set the PGID mask to use when flooding unknown IPv4 Multicast Data frames.	0x3F
FLD_MC6_CTRL	11:6	R/W	Set the PGID mask to use when flooding unknown IPv6 Multicast Control frames.	0x3F
FLD_MC6_DATA	5:0	R/W	Set the PGID mask to use when flooding unknown IPv6 Multicast Data frames.	0x3F

5.1.1.18 ANA:ANA:SFLOW_CFG

Parent: [ANA:ANA](#)

Instances: 12

TABLE 5-21: FIELDS IN SFLOW_CFG

Field Name	Bit	Access	Description	Default
SF_RATE	13:2	R/W	Probability of a frame being SFLOW sampled. Unit is 1/4096. A value of 0 makes 1/4096 of the candidates being forwarded to the SFLOW CPU extraction queue. A values of 4095 makes all candidates being forwarded.	0x000
SF_SAMPLE_RX	1	R/W	Enable SFLOW sampling of frames received on this port.	0x0
SF_SAMPLE_TX	0	R/W	Enable SFLOW sampling of frames transmitted on this port.	0x0

5.1.1.19 ANA:ANA:PORT_MODE

Parent: [ANA:ANA](#)

Instances: 13

These configurations exists per front port and for two CPU ports.

TABLE 5-22: FIELDS IN PORT_MODE

Field Name	Bit	Access	Description	Default
RESERVED	2:1	R/W	Must be set to its default.	0x2
L3_PARSE_CFG	0	R/W	Enable frame analysis on Layer-3 and Layer-4 protocol information. If cleared, all frames are seen as non-IP and are handled accordingly. This affects all blocks using IP information such as classification, TCAM lookups, IP flooding and forwarding, and DSCP rewriting.	0x1

5.1.2 ANA:PGID

Parent: [ANA](#)

Instances: 1

TABLE 5-23: REGISTERS IN PGID

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
PGID	0x00000000	92 0x00000004	Port Group Identifiers	Page 331

5.1.2.1 ANA:PGID:PGID

Parent: [ANA:PGID](#)

Instances: 92

Three port masks are applied to all frames, allowing transmission to a port if the corresponding bit is set in all masks.

0-63: A mask is applied based on destination analysis

64-79: A mask is applied based on aggregation analysis

80-91: A mask is applied based on source port analysis

Destination analysis:

There are 64 destination masks in total. By default, the first 11 port masks only have the bit corresponding to their port number set. These masks should not be changed, except for aggregation.

The remaining destination masks are set to 0 by default and are available for use for Layer-2 multicasts and flooding (See FLOODING and FLOODING_IPMC).

Aggregation analysis:

The aggregation port masks are used to select only one port within each aggregation group. These 16 masks must be setup to select only one port in each aggregated port group.

For ports, which are not part of any aggregation group, the corresponding bits in all 16 masks must be set.

I.e. if no aggregation is configured, all masks must be set to all-ones.

The aggregation mask used for the forwarding of a given frame is selected by the frame's aggregation code (see AGGRCTRL).

Source port analysis:

The source port masks are used to prevent frames from being looped back to the ports on which they were received, and must be updated according to the

aggregation configuration. A frame that is received on port n, uses mask 80+n as a mask to filter out destination ports to avoid loopback, or to facilitate port grouping (port-based VLANs). The default values are that all bits are set except for the index number.

TABLE 5-24: FIELDS IN PGID

Field Name	Bit	Access	Description	Default
PGID	11:0	R/W	When a mask is chosen, bit N must be set for the frame to be transmitted on port N.	0xFFFF
CPUQ_DST_PGID	29:27	R/W	CPU extraction queue used when CPU port is enabled in PGID. Only applicable for the destination analysis.	0x0

5.1.3 ANA:ANA_TABLES

Parent: [ANA](#)

Instances: 1

TABLE 5-25: REGISTERS IN ANA_TABLES

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
ANMOVED	0x00000030	1	Station move logger	Page 332
MACHDATA	0x00000034	1	MAC address high	Page 332
MACLDATA	0x00000038	1	MAC address low	Page 333
MACACCESS	0x0000003C	1	MAC table command	Page 333
MACTINDX	0x00000040	1	MAC table index	Page 335
VLANACCESS	0x00000044	1	VLAN table command	Page 335
VLANTIDX	0x00000048	1	VLAN table index	Page 336
ISDXACCESS	0x0000004C	1	ISDX table command	Page 336
ISDXTIDX	0x00000050	1	ISDX table index	Page 337
ENTRYLIM	0x00000000	12 0x00000004	MAC table entry limits per switch port	Page 337
PTP_ID_HIGH	0x00000054	1	PTP identifiers 63-32	Page 338
PTP_ID_LOW	0x00000058	1	PTP identifiers 31-0	Page 338

5.1.3.1 ANA:ANA_TABLES:ANMOVED

Parent: [ANA:ANA_TABLES](#)

Instances: 1

TABLE 5-26: FIELDS IN ANMOVED

Field Name	Bit	Access	Description	Default
ANMOVED	11:0	R/W	This mask detects port moves in the MAC table. When a station is learned on a new port while already learned on another port, the bit corresponding to the new port is set in this mask. This mask can be used to detect topology problems in the network, where stations are learned on multiple ports repeatedly. If some bits in this register get asserted repeatedly, the ports can be shut down, or management warnings can be issued.	0x000

5.1.3.2 ANA:ANA_TABLES:MACHDATA

Parent: [ANA:ANA_TABLES](#)

Instances: 1

TABLE 5-27: FIELDS IN MACHDATA

Field Name	Bit	Access	Description	Default
VID	28:16	R/W	VID used in MAC table operations through MACACCESS. For read operations, the VID value is returned in this field.	0x0000
MACHDATA	15:0	R/W	Most significant 16 MAC address bits used in MAC table operations through MACACCESS.	0x0000

5.1.3.3 ANA:ANA_TABLES:MACLDATA

Parent: [ANA:ANA_TABLES](#)

Instances: 1

TABLE 5-28: FIELDS IN MACLDATA

Field Name	Bit	Access	Description	Default
MACLDATA	31:0	R/W	Lower 32 MAC address bits used in MAC table operations through MACACCESS.	0x00000000

5.1.3.4 ANA:ANA_TABLES:MACACCESS

Parent: [ANA:ANA_TABLES](#)

Instances: 1

This register is used for updating or reading the MAC table from the CPU.

The command (MAC_TABLE_CMD) selects between different operations and uses the following encoding:

000 - IDLE:

The previous operation has completed.

001 - LEARN:

Insert/learn new entry in MAC table. Position given by (MAC, VID) in MACHDATA and MACLDATA.

010 - FORGET:

Delete/unlearn entry given by (MAC, VID) in MACHDATA and MACLDATA.

Both locked and unlocked entries are deleted.

011 - AGE:

Start an age scan on the MAC table.

100 - GET_NEXT:

Get the smallest entry in the MAC table numerically larger than the (MAC, VID) specified in MACHDATA and MACLDATA. The VID and MAC are evaluated as a 60-bit number with the VID being most significant.

101 - INIT:

Table is initialized (completely cleared).

110 - READ:

VSC7416-01

The READ command is divided into two modes: Direct mode and indirect mode.

Direct mode (read):

With MACACCESS.VALID cleared, the entry pointed to by MACTINDX.INDEX (row) and MACTINDX.BUCKET (column) is read.

Indirect mode (lookup):

With MACACCESS.VALID set, the entry pointed to by (MAC, VID) in the MACHDATA and MACLDATA is read.

111 - WRITE

Write entry. Address of the entry is specified in MACTINDX.INDEX (row) and MACTINDX.BUCKET (column).

An existing entry (locked or unlocked) is overwritten.

The MAC_TABLE_CMD must be IDLE before a new command can be issued.

The AGE and CLEAR commands run for approximately 50 us. The other commands execute immediately.

The flags IGNORE_VLAN and MAC_CPU_COPY are ignored for DMAC lookup if AGENCTRL.IGNORE_DMAC_FLAGS is set.

The flags SRC_KILL and MAC_CPU_COPY are ignored for SMAC lookup if AGENCTRL.IGNORE_SMAC_FLAGS is set.

TABLE 5-29: FIELDS IN MACACCESS

Field Name	Bit	Access	Description	Default
B_DOM	18	R/W	If set, MAC table entry is only applicable in the B-domain. Ports with ANA::PORT_CFG:USE_B_DOM_TBL set are B-ports.	0x0
MAC_CPU_COPY	15	R/W	Frames matching this entry are copied to the CPU extraction queue CPUQ_CFG.CPUQ_MAC_COPY. Applies to both SMAC and DMAC lookup.	0x0
SRC_KILL	14	R/W	Frames matching this entry are discarded. Applies only to the SMAC lookup. For discarding frames based on the DMAC lookup a NULL PGID mask can be used.	0x0
IGNORE_VLAN	13	R/W	The VLAN mask is ignored for this destination. Applies only to DMAC lookup.	0x0
AGED_FLAG	12	R/W	This flag is set on every aging run. Entry is removed if flag is already set. The flag is cleared when the entry is target for a SMAC lookup. Locked entries will not be removed. Bit is for IPv6 Multicast used for port 25.	0x0
VALID	11	R/W	Entry is valid.	0x0
ENTRY_TYPE	10:9	R/W	Type of entry: 0: Normal entry eligible for aging. 1: Locked entry. Entry is not removed by aging. 2: IPv4 Multicast entry. Entry is not removed by aging. 3: IPv6 Multicast entry. Entry is not removed by aging.	0x0

TABLE 5-29: FIELDS IN MACACCESS (CONTINUED)

Field Name	Bit	Access	Description	Default
DEST_IDX	8:3	R/W	Index for the destination masks table (PGID). For unicasts, this is a number from 0-EXB_PORT_CNT_MINUS_ONE.	0x00
MAC_TABLE_CMD	2:0	R/W	MAC Table Command. See below.	0x0

5.1.3.5 ANA:ANA_TABLES:MACTINDX

Parent: [ANA:ANA_TABLES](#)

Instances: 1

TABLE 5-30: FIELDS IN MACTINDX

Field Name	Bit	Access	Description	Default
BUCKET	12:11	R/W	Selects one of the four MAC table entries in a row. The row is addressed with the INDEX field.	0x0
M_INDEX	10:0	R/W	The index selects one of the 2048 MAC table rows. Within a row the entry is addressed by the BUCKET field	0x000

5.1.3.6 ANA:ANA_TABLES:VLANACCESS

Parent: [ANA:ANA_TABLES](#)

Instances: 1

The VLAN_TBL_CMD field of this register is used for updating and reading the VLAN table. The command (VLAN_TBL_CMD) selects between different operations and uses the following encoding:

00 - IDLE:

The previous operation has completed.

01 - READ:

The VLAN table entry set in VLANTIDX.INDEX is returned in VLANACCESS.VLAN_PORT_MASK and the VLAN flags in VLANTIDX.

10 - WRITE:

The VLAN table entry pointed to by VLANTIDX.INDEX is updated with VLANACCESS.VLAN_PORT_MASK and the VLAN flags in VLANTIDX.

11 - INIT:

The VLAN table is initialized to default values (all ports are members of all VLANs).

The VLAN_TBL_CMD must be IDLE before a new command can be issued. The INIT command run for approximately 50 us whereas the other commands execute immediately. When an operation has completed, VLAN_TBL_CMD changes to IDLE.

TABLE 5-31: FIELDS IN VLANACCESS

Field Name	Bit	Access	Description	Default
VLAN_PORT_MASK	13:2	R/W	Frames classified to this VLAN can only be sent to ports in this mask. Note that the CPU port module is always member of all VLANs and its VLAN membership can therefore not be configured through this mask.	0x7FF
VLAN_TBL_CMD	1:0	R/W	VLAN Table Command.	0x0

5.1.3.7 ANA:ANA_TABLES:VLANTIDX

Parent: [ANA:ANA_TABLES](#)

Instances: 1

TABLE 5-32: FIELDS IN VLANTIDX

Field Name	Bit	Access	Description	Default
VLAN_PRIV_VLAN	15	R/W	If set, a VLAN is a private VLAN. See PRIV_VLAN_MASK for details.	0x0
VLAN_LEARN_DISABLED	14	R/W	Disable learning for this VLAN.	0x0
VLAN_MIRROR	13	R/W	If set, all frames classified to this VLAN are mirrored to the port set configured in MIRRORPORTS.	0x0
VLAN_SRC_CHK	12	R/W	If set, VLAN ingress filtering is enabled for this VLAN. If set, a frame's ingress port must be member of the frame's VLAN, otherwise the frame is discarded.	0x0
V_INDEX	11:0	R/W	Index used to select VLAN table entry for read/write operations (see VLANACCESS). This value equals the VID.	0x000

5.1.3.8 ANA:ANA_TABLES:ISDXACCESS

Parent: [ANA:ANA_TABLES](#)

Instances: 1

The ISDX_TBL_CMD field of this register is used for updating and reading the ISDX table. The command (ISDX_TBL_CMD) selects between different operations and uses the following encoding:

00 - IDLE:

The previous operation has completed.

01 - READ:

The ISDX table entry set in ISDXTIDX.INDEX is returned in ISDXACCESS.ISDX_PORT_MASK and ISDXTIDX.

10 - WRITE:

The ISDX table entry pointed to by ISDXTIDX.INDEX is updated with ISDXACCESS.ISDX_PORT_MASK and ISDX-TIDX

11 - INIT:

The ISDX table is initialized to default values (all ports are members of all ISDXs).

The ISDX_TBL_CMD must be IDLE before a new command can be issued. The INIT command run for approximately 50 us whereas the other commands execute immediately. When an operation has completed, ISDX_TBL_CMD changes to IDLE.

TABLE 5-33: FIELDS IN ISDXACCESS

Field Name	Bit	Access	Description	Default
ISDX_PORT_MASK	13:2	R/W	Frames classified to this ISDX can only be sent to ports in this mask. Note that the CPU port module is always member of all ISDXs and its ISDX membership can therefore not be configured through this mask.	0xFFFF
ISDX_TBL_CMD	1:0	R/W	ISDX Table Command.	0x0

5.1.3.9 ANA:ANA_TABLES:ISDXTIDX

Parent: [ANA:ANA_TABLES](#)

Instances: 1

TABLE 5-34: FIELDS IN ISDXTIDX

Field Name	Bit	Access	Description	Default
ISDX_MSTI	20:15	R/W	Index of associated MSTP instance, see ANA::MSTI_STATE.	0x00
ISDX_ES0_KEY_ENA	14	R/W	Enable use of ISDX value in ES0 lookup in the rewriter. Otherwise a VID lookup is made.	0x1
ISDX_SDLBI	30:21	R/W	Index of service DLB policer. If ISDX_SDLBI is zero, service policing is disabled for this service and frames classified to the service can be policed by a QoS policer instead, see ANA:PORT:POL_CFG.QUEUE_POL_ENA. 0: Service policing is disabled. 1-1023: The service policer pointed to by ISDX_SDLBI is used for frames classified to the ISDX.	0x000
ISDX_INDEX	9:0	R/W	Index used to select ISDX table entry for read/write operations (see ISDXACCESS).	0x000
ISDX_FORCE_ENA	10	R/W	If set, the VLAN mask and the destination mask are not used in the forwarding decision.	0x0

5.1.3.10 ANA:ANA_TABLES:ENTRYLIM

Parent: [ANA:ANA_TABLES](#)

Instances: 12

TABLE 5-35: FIELDS IN ENTRYLIM

Field Name	Bit	Access	Description	Default
ENTRYLIM	17:14	R/W	Maximum number of unlocked entries in the MAC table learned on this port. Locked entries and IPMC entries do not obey this limit. Both auto-learned and unlocked CPU-learned entries obey this limit. 0: 1 entry 1: 2 entries n: 2**n entries >12: 8192 entries	0xD
ENTRYSTAT	13:0	R/W	Current number of unlocked MAC table entries learned on this port.	0x0000

5.1.3.11 ANA:ANA_TABLES:PTP_ID_HIGH

Parent: [ANA:ANA_TABLES](#)

Instances: 1

TABLE 5-36: FIELDS IN PTP_ID_HIGH

Field Name	Bit	Access	Description	Default
PTP_ID_HIGH	31:0	R/W	Bit vector with current use of PTP timestamp identifiers 32 through 63. Timestamp identifier is 63 is reserved for signaling that no identifiers are available. A timestamp identifier is released by setting the corresponding bit. Bit 0: Timestamp identifier 32 ... Bit 31: Timestamp identifier 63.	0x00000000

5.1.3.12 ANA:ANA_TABLES:PTP_ID_LOW

Parent: [ANA:ANA_TABLES](#)

Instances: 1

TABLE 5-37: FIELDS IN PTP_ID_LOW

Field Name	Bit	Access	Description	Default
PTP_ID_LOW	31:0	R/W	Bit vector with current use of PTP timestamp identifiers 0 through 31. A timestamp identifier is released by setting the corresponding bit. Bit 0: Timestamp identifier 0 ... Bit 31: Timestamp identifier 31.	0x0000000F

5.1.4 ANA:MSTI_STATE

Parent: [ANA](#)

Instances: 1

TABLE 5-38: REGISTERS IN MSTI_STATE

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
MSTI_STATE	0x00000000	128 0x00000004	Multiple spanning tree states	Page 339

5.1.4.1 ANA:MSTI_STATE:MSTI_STATE

Parent: [ANA:MSTI_STATE](#)

Instances: 128

Configuration of forwarding states per MSTi.

The table is looked up twice per frame:

- 1) Entries 0 through 63 are looked up using the frame's FID as key. If the FID is larger than 63, then entry 0 is looked up.
- 2) Entries 64 through 127 are looked up using the ISDX_MSTI from the ISDX table. The key is ISDX_MSTI+64.

TABLE 5-39: FIELDS IN MSTI_STATE

Field Name	Bit	Access	Description	Default
MSTI_FWD_ENA	10:0	R/W	Enable/disable MSTP forwarding per port. Ports in MSTP forwarding state must be enabled in this port mask. If a port is disabled in this mask, frames received on the port are not forwarded, and frames are not forwarded to the port.	0x7FF

5.1.5 ANA:OAM_UPM_LM_CNT

Parent: [ANA](#)

Instances: 1

TABLE 5-40: REGISTERS IN OAM_UPM_LM_CNT

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
OAM_UPM_LM_CNT	0x00000000	512 0x00000004	OAM Up-MEP loss measurement counters	Page 339

5.1.5.1 ANA:OAM_UPM_LM_CNT:OAM_UPM_LM_CNT

Parent: [ANA:OAM_UPM_LM_CNT](#)

Instances: 512

TABLE 5-41: FIELDS IN OAM_UPM_LM_CNT

Field Name	Bit	Access	Description	Default
OAM_UPM_LM_CNT_VAL	31:0	R/W	Up-MEP OAM service Tx loss measurement counters.	0x00000000

5.1.6 ANA:MPLS_CL

Parent: [ANA](#)

Instances: 1

TABLE 5-42: REGISTERS IN MPLS_CL

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
OAM_CONF	0x00000000	1	OAM configuration	Page 340
FRAME_EXTRACT	0x00000004	7 0x00000004	Frame extraction of MPLS-TP error frames	Page 341
RSVD_LABELS_ENA	0x00000020	16 0x00000004	Configuration of reserved labels	Page 341
SEGMENT_OAM	0x00000060	13 0x00000004	Configure support for segment OAM on physical ports.	Page 342
EXCP_VPROFILE_CONF	0x00000094	1	Virtual port profile configuration for MPLS-TP exception frames	Page 342

5.1.6.1 ANA:MPLS_CL:OAM_CONF

Parent: [ANA:MPLS_CL](#)

Instances: 1

TABLE 5-43: FIELDS IN OAM_CONF

Field Name	Bit	Access	Description	Default
CPU_OAM_QUEUE	9:7	R/W	CPU extraction queue for the following MPLS OAM frames: - VCCV1 OAM - VCCV2 RAL OAM - VCCV3 OAM - GAL MEP LSP OAM - GAL MIP LSP OAM	0x0
VCCV2_LABEL	6:3	R/W	Reserved label used for PW VCCV2 OAM channel. Default is to use reserved label = 1 (Router Alert).	0x1
VCCV2_ENA	2	R/W	Enable VCCV2 OAM signalling. If enabled, frames with reserved label VCCV2_LABEL as outer label are allowed an MLBS lookup using the label following the VCCV2 label. If disabled, such frames are treated as segment reserved label frames. '0': Disable VCCV2 OAM signalling '1': Enable VCCV2 OAM signalling	0x0

TABLE 5-43: FIELDS IN OAM_CONF (CONTINUED)

Field Name	Bit	Access	Description	Default
OAM_TTL_CHECK	1	R/W	Enable TTL expiry check for OAM frames. If a frame with TTL < 2 is classified as OAM, it is extracted as an TTL expiry error instead of OAM, see FRAME_EXTRACT. 0: Do not perform TTL expiry check for OAM frames. 1: Perform TTL expiry check for OAM frames.	0x1
RESERVED	0	R/W	Must be set to its default.	0x1

5.1.6.2 ANA:MPLS_CL:FRAME_EXTRACT**Parent:** [ANA:MPLS_CL](#)**Instances:** 7

For each of the MLBS processing errors, it is possible to redirect the frame to a configurable CPU extraction queue:

- 0: TTL expiry
- 1: Expected CW
- 2: Expected ACh
- 3: Illegal segment OAM
- 4: Unexpected S-bit clear
- 5: Unknown reserved label
- 6: MLBS miss

TABLE 5-44: FIELDS IN FRAME_EXTRACT

Field Name	Bit	Access	Description	Default
EXTRACT_ENA	15:3	R/W	Decides if MPLS-TP error frames are redirected to the CPU or discarded. For bit x: 0: Discard MPLS-TP error frames received on port x 1: Redirect MPLS-TP error frames to CPU received on port x	0x0000
EXTRACT_QUEUE	2:0	R/W	Selects the queue to which MPLS-TP error frames are redirected, if enabled by EXTRACT_ENA.	0x0

5.1.6.3 ANA:MPLS_CL:RSVD_LABELS_ENA**Parent:** [ANA:MPLS_CL](#)**Instances:** 16

This register allows enabling each of the 16 MPLS reserved labels (label value < 16) for exception handling.

If a label is enabled, it is copied to the CPU extraction queue configured in RSVD_CPU_QUEUE.

The reserved labels are assigned the following meaning by RFC 3032:

Label 0: IPv4 Explicit NULL Label

Label 1: Router Alert Label (RAL)

Label 2: IPv6 Explicit NULL Label

Label 3: Reserved (Implicit NULL Label - Never in stack)

Label 4: (Reserved - Not yet assigned)
Label 5: (Reserved - Not yet assigned)
Label 6: (Reserved - Not yet assigned)
Label 7: (Reserved - Not yet assigned)
Label 8: (Reserved - Not yet assigned)
Label 9: (Reserved - Not yet assigned)
Label 10: (Reserved - Not yet assigned)
Label 11: (Reserved - Not yet assigned)
Label 12: (Reserved - Not yet assigned)
Label 13: Generic Alert Label (GAL)
Label 14: OAM Alert Label (OAL)
Label 15: (Reserved - Not yet assigned)

TABLE 5-45: FIELDS IN RSVD_LABELS_ENA

Field Name	Bit	Access	Description	Default
RSVD_LBL_ENA	3	R/W	Enable this reserved label for OAM. If enabled for OAM, the MPLS frame is forwarded to the CPU if the reserved label is found within the scope in the label stack.	0x0
RSVD_CPU_QUEUE	2:0	R/W	CPU extraction queue for frames with a reserved label.	0x0

5.1.6.4 ANA:MPLS_CL:SEGMENT_OAM

Parent: [ANA:MPLS_CL](#)

Instances: 13

TABLE 5-46: FIELDS IN SEGMENT_OAM

Field Name	Bit	Access	Description	Default
SEGMENT_OAM_ENA	8	R/W	Enable segment OAM on this physical port.	0x0
SEGMENT_OAM_LBL	7:4	R/W	Configures the reserved label used for segment OAM on this segment. Only the lower 4 bits must be configured because the top 16 bits are zero per definition.	0x0
SEGMENT_OAM_BOS	3	R/W	Enable check that the reserved label used for segment OAM on this port must be bottom of stack.	0x0
SEGMENT_OAM_QUEUE	2:0	R/W	CPU extraction queue for segment OAM frames.	0x0

5.1.6.5 ANA:MPLS_CL:EXCP_VPROFILE_CONF

Parent: [ANA:MPLS_CL](#)

Instances: 1

Configures the virtual port profile to be used for MPLS-TP exception frames.

TABLE 5-47: FIELDS IN EXCP_VPROFILE_CONF

Field Name	Bit	Access	Description	Default
MLBS_MISS_SRC	27	R/W	Selects the source of the virtual port profile for frames that do not have an MLBS hit in the MPLS-TP classifier. 0: Use MLBS_MISS_VPROFILE 1: Use MLL_action.VPROFILE	0x0
MLBS_MISS_VPROFILE	26:21	R/W	Virtual port profile for MLBS miss frames. The use of this field is enabled by MLBS_MISS_SRC.	0x00
OAM_SRC	20	R/W	Selects the source of the virtual port profile for MPLS-TP OAM frames terminated in this device. Segment OAM frames: 0: Use OAM_VPROFILE 1: Physical ingress port number. Other OAM frames: 0: Use OAM_VPROFILE 1: Use the MLBS_action.VPROFILE	0x0
OAM_VPROFILE	19:14	R/W	Virtual port profile for MPLS-TP OAM frames terminated in this device. The use of this field is enabled by OAM_SRC.	0x00
RSVD_LBL_SRC	13	R/W	Selects the source of the virtual port profile for MPLS-TP frames for which the classification exposes a reserved label with handling enabled in RSVD_LABELS_ENA. Segment reserved label frames (label is outer-most label): 0: Use RSVD_LBL_VPROFILE 1: Physical ingress port number. Other reserved label frames: 0: Use RSVD_LBL_VPROFILE 1: Use the MLBS_action.VPROFILE	0x0
RSVD_LBL_VPROFILE	12:7	R/W	Virtual port profile for MPLS-TP frames for which the classification exposes a reserved label with handling enabled in RSVD_LABELS_ENA. The use of this field is enabled by RSVD_LBL_SRC.	0x00
MLBS_ERR_SRC	6	R/W	Selects the source of the virtual port profile for MLBS error frames. 0: Use MLBS_ERR_VPROFILE 1: Use the MLL_action.VPROFILE	0x0
MLBS_ERR_VPROFILE	5:0	R/W	Virtual port profile for MLBS error frames. The use of this field is enabled by MLBS_ERR_SRC.	0x00

5.1.7 ANA:MPLS_TC_MAP

Parent: [ANA](#)

Instances: 8

TABLE 5-48: REGISTERS IN MPLS_TC_MAP

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
TC_MAP_TBL	0x00000000	8 0x00000004	TC mapping tables	Page 344

5.1.7.1 ANA:MPLS_TC_MAP:TC_MAP_TBL

Parent: [ANA:MPLS_TC_MAP](#)

Instances: 8

Incoming TC bits can be mapped to internal drop precedence level and QoS class for E-LSP.

The TC_MAP action out of the IS0 MLBS lookup selects which table to use.

TABLE 5-49: FIELDS IN TC_MAP_TBL

Field Name	Bit	Access	Description	Default
TC_MAP_QOS	3:1	R/W	For MPLS data frames, map the MPLS data frame's TC to a QoS class: QoS class = ANA:MPLS_TC_MAP[MLBS_action.TC_MAP]:TC_MAP_TBL[TC].TC_MAP_QOS. This value is used as the virtual port's default QoS class and overrules the setting in ANA::QOS_CFG.QOS_DEFAULT_VAL.	0x0
TC_MAP_DP	0	R/W	For MPLS data frames, map the MPLS data frame's TC to a drop precedence level: DP level = ANA:MPLS_TC_MAP[MLBS_action.TC_MAP]:TC_MAP_TBL[TC].TC_MAP_DP. This value is used as the virtual port's default DP level and overrules the setting in ANA::QOS_CFG.DP_DEFAULT_VAL.	0x0

5.1.8 ANA:PORT

Parent: [ANA](#)

Instances: 64

TABLE 5-50: REGISTERS IN PORT

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
VLAN_CFG	0x00000000	1	Port VLAN configuration	Page 345
DROP_CFG	0x00000004	1	VLAN acceptance filtering	Page 346
QOS_CFG	0x00000008	1	QoS and DSCP configuration	Page 346
VCAP_CFG	0x0000000C	1	VCAP configuration	Page 347
VCAP_S1_KEY_CFG	0x00000010	3 0x00000004	VCAP S1 key configuration per S1 lookup	Page 348
VCAP_S2_CFG	0x0000001C	1	VCAP S2 configuration	Page 348

TABLE 5-50: REGISTERS IN PORT (CONTINUED)

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
QOS_PCP_DEI_MAP_CFG	0x00000020	16 0x00000004	Mapping of DEI and PCP to QoS class and drop precedence level	Page 349
CPU_FWD_CFG	0x00000060	1	CPU forwarding of special protocols	Page 349
CPU_FWD_BPDU_CFG	0x00000064	1	CPU forwarding of BPDU frames	Page 350
CPU_FWD_GARP_CFG	0x00000068	1	CPU forwarding of GARP frames	Page 351
CPU_FWD_CCM_CFG	0x0000006C	1	CPU forwarding of CCM/Link trace frames	Page 351
PORT_CFG	0x00000070	1	Special port configuration	Page 352
POL_CFG	0x00000074	1	Policer selection	Page 353
PTP_CFG	0x00000078	1	Timing protocol configuration	Page 354
PTP_DLY1_CFG	0x0000007C	1	PTP ingress delay 1 configuration	Page 355
PTP_DLY2_CFG	0x00000080	1	PTP ingress delay 2 configuration	Page 355

5.1.8.1 ANA:PORT:VLAN_CFG

Parent: [ANA:PORT](#)

Instances: 1

TABLE 5-51: FIELDS IN VLAN_CFG

Field Name	Bit	Access	Description	Default
VLAN_AWARE_ENA	20	R/W	Enable VLAN awareness. If set, Q-tag headers are processed during the basic VLAN classification. If cleared, Q-tag headers are ignored during the basic VLAN classification.	0x0
VLAN_POP_CNT	19:18	R/W	Number of tag headers to remove from ingress frame. 0: Keep all tags. 1: Pop up to 1 tag (outer tag if available). 2: Pop up to 2 tags (outer and inner tag if available). 3: Disable rewriting of VLAN tags and DSCP value. The rewriter can still update the FCS.	0x0
VLAN_INNER_TAG_ENA	17	R/W	Set if the inner Q-tag must be used instead of the outer Q-tag. If the received frame is single tagged, the outer tag is used. This bit influences the VLAN acceptance filter (DROP_CFG), the basic VLAN classification (VLAN_CFG), and the basic QoS classification (QOS_CFG).	0x0

TABLE 5-51: FIELDS IN VLAN_CFG (CONTINUED)

Field Name	Bit	Access	Description	Default
VLAN_TAG_TYPE	16	R/W	Tag Protocol Identifier type for port-based VLAN. 0: C-tag (EtherType = 0x8100) 1: S-tag (EtherType = 0x88A8 or configurable value (VLAN_ETYPE_CFG))	0x0
VLAN_DEI	15	R/W	DEI value for port-based VLAN.	0x0
VLAN_PCP	14:12	R/W	PCP value for port-based VLAN.	0x0
VLAN_VID	11:0	R/W	VID value for port-based VLAN.	0x000

5.1.8.2 ANA:PORT:DROP_CFG

Parent: [ANA:PORT](#)

Instances: 1

TABLE 5-52: FIELDS IN DROP_CFG

Field Name	Bit	Access	Description	Default
DROP_UNTAGGED_ENA	6	R/W	Drop untagged frames.	0x0
DROP_S_TAGGED_ENA	5	R/W	Drop S-tagged frames (VID different from 0 and EtherType = 0x88A8 or configurable value (VLAN_ETYPE_CFG)).	0x0
DROP_C_TAGGED_ENA	4	R/W	Drop C-tagged frames (VID different from 0 and EtherType = 0x8100).	0x0
DROP_PRIO_S_TAGGED_ENA	3	R/W	Drop S-tagged frames (VID=0 and EtherType = 0x88A8 or configurable value (VLAN_ETYPE_CFG)).	0x0
DROP_PRIO_C_TAGGED_ENA	2	R/W	Drop priority C-tagged frames (VID=0 and EtherType = 0x8100).	0x0
DROP_NULL_MAC_ENA	1	R/W	Drop frames with source or destination MAC address equal to 0x000000000000.	0x0
DROP_MC_SMAC_ENA	0	R/W	Drop frames with multicast source MAC address.	0x0

5.1.8.3 ANA:PORT:QOS_CFG

Parent: [ANA:PORT](#)

Instances: 1

TABLE 5-53: FIELDS IN QOS_CFG

Field Name	Bit	Access	Description	Default
DP_DEFAULT_VAL	8	R/W	Default drop precedence level.	0x0
QOS_DEFAULT_VAL	7:5	R/W	Default QoS class.	0x0
QOS_DSCP_ENA	4	R/W	If set, the DP level and QoS class can be based on DSCP values.	0x0
QOS_PCP_ENA	3	R/W	If set, DP level and QoS class can be based on the PCP and DEI bits for tagged frames.	0x0

TABLE 5-53: FIELDS IN QOS_CFG (CONTINUED)

Field Name	Bit	Access	Description	Default
DSCP_TRANSLATE_ENA	2	R/W	Set if the DSCP value must be translated before using the DSCP value. If set, the translated DSCP value is given from DSCP_CFG[DSCP].DSCP_TRANSLATE_VAL.	0x0
DSCP_REWR_CFG	1:0	R/W	Configure which DSCP values to rewrite based on DP level and QoS class. If the DSCP value is to be rewritten, then the new DSCP = DSCP_REWR_CFG[8*DP level + QoS class].DSCP_QOS_REWR_VAL. 0: Rewrite none. 1: Rewrite if DSCP=0 2: Rewrite for selected values configured in DSCP_CFG[DSCP].DSCP_REWR_ENA. 3: Rewrite all.	0x0

5.1.8.4 ANA:PORT:VCAP_CFG

Parent: [ANA:PORT](#)

Instances: 1

TABLE 5-54: FIELDS IN VCAP_CFG

Field Name	Bit	Access	Description	Default
S1_ENA	14	R/W	If S1 is enabled, each frame received on this port is processed and matched against the entries in the S1 TCAM. Each frame results in three lookups.	0x0
S1_DMAC_DIP_ENA	13:11	R/W	Set if the destination MAC address and the destination IP address must be used in S1 key instead of the source MAC address and the source IP address. One bit per lookup. Only applicable to S1 keys S1_NORMAL and S1_NORMAL_IP6.	0x0
S1_VLAN_INNER_TAG_ENA	10:8	R/W	Set if the inner Q-tag must be passed on to the S1 TCAM instead of the outer Q-tag. For single tagged frames, the outer tag is used. For untagged frames, the port VLAN is used. This bit influences the TPID, VID, PCP, and DEI input to the S1 key generation. One bit per lookup. Only applicable to S1_NORMAL.	0x0
OAM_DETECT_CFG	7:6	R/W	Configures the detection of OAM frames. 0: Do not detect OAM; frame is treated as data even if EtherType is 0x8902. 1: Detect as OAM if frame is untagged and EtherType is 0x8902. 2: Detect as OAM if frame is single-tagged and EtherType is 0x8902. 3: Detect as OAM if frame is double-tagged and EtherType is 0x8902. The configuration of OAM detection can be changed by IS1 action OAM_DETECT_SEL.	0x0
PAG_VAL	5:0	R/W	Default PAG value used as input to S2. The PAG value can be changed by S1 actions.	0x00

VSC7416-01

5.1.8.5 ANA:PORT:VCAP_S1_KEY_CFG

Parent: [ANA:PORT](#)

Instances: 3

TABLE 5-55: FIELDS IN VCAP_S1_KEY_CFG

Field Name	Bit	Access	Description	Default
S1_KEY_IP6_CFG	6:4	R/W	Selects key per lookup in S1 for IPv6 frames. 0: Use key S1_NORMAL 1: Use key S1_7TUPLE 2: Use key S1_5TUPLE_IP4 3: Use key S1_NORMAL_IP6 4: Use key S1_5TUPLE_IP6 5: Use key S1_DBL_VID	0x0
S1_KEY_IP4_CFG	3:2	R/W	Selects key per lookup in S1 for IPv4 frames. 0: Use key S1_NORMAL 1: Use key S1_7TUPLE 2: Use key S1_5TUPLE_IP4 3: Use key S1_DBL_VID	0x0
S1_KEY_OTHER_CFG	1:0	R/W	Selects key per lookup in S1 for non-IP frames. 0: Use key S1_NORMAL 1: Use key S1_7TUPLE 2: Use key S1_DBL_VID	0x0

5.1.8.6 ANA:PORT:VCAP_S2_CFG

Parent: [ANA:PORT](#)

Instances: 1

TABLE 5-56: FIELDS IN VCAP_S2_CFG

Field Name	Bit	Access	Description	Default
S2_ENA	14	R/W	If S2 is enabled, each frame received on this port is processed and matched against the entries in the S2 TCAM. Each frame results in two lookups to determine both an ingress and an egress action.	0x0
S2_SNAP_DIS	13:12	R/W	If set, MAC_SNAP frames received on this port are treated as MAC_LLIC frames when matching in S2. Bit 0 controls the first lookup and bit 1 controls the second lookup.	0x0
S2_ARP_DIS	11:10	R/W	If set, MAC_ARP frames received on this port are treated as MAC_ETYPE frames when matching in S2. Bit 0 controls the first lookup and bit 1 controls the second lookup.	0x0
S2_IP_TCPUDP_DIS	9:8	R/W	If set, IP_TCPUDP frames received on this port are treated as MAC_ETYPE frames when matching in S2. Bit 0 controls the first lookup and bit 1 controls the second lookup.	0x0

TABLE 5-56: FIELDS IN VCAP_S2_CFG (CONTINUED)

Field Name	Bit	Access	Description	Default
S2_IP_OTHER_DIS	7:6	R/W	If set, IP_OTHER frames received on this port are treated as MAC_ETYPE frames when matching in S2. Bit 0 controls the first lookup and bit 1 controls the second lookup.	0x0
S2_IP6_CFG	5:2	R/W	S2_IP6_CFG controls the key generation for IPv6 frames. Bits 1:0 control the first lookup and bits 3:2 control the second lookup. 0: IPv6 frames are matched against IP6_TCP_UDP or IP6_OTHER entries 1: IPv6 frames are matched against IP6_STD entries 2: IPv6 frames are matched against IP4_TCP_UDP or IP4_OTHER entries 3: IPv6 frames are matched against MAC_ETYPE entries	0x1
S2_OAM_DIS	1:0	R/W	If set, OAM frames received on this port are treated as MAC_ETYPE frames when matching in S2. Bit 0 controls the first lookup and bit 1 controls the second lookup.	0x0

5.1.8.7 ANA:PORT:QOS_PCP_DEI_MAP_CFG

Parent: [ANA:PORT](#)

Instances: 16

TABLE 5-57: FIELDS IN QOS_PCP_DEI_MAP_CFG

Field Name	Bit	Access	Description	Default
DP_PCP_DEI_VAL	3	R/W	Map the frame's PCP and DEI values to a drop precedence level. DP level = QOS_PCP_DEI_MAP_CFG[index].DP_PCP_DEI_VAL, where index = 8*DEI + PCP. Only applicable to tagged frames. The use of Inner or outer tag can be selected using VLAN_CFG.VLAN_INNER_TAG_ENA.	0x0
QOS_PCP_DEI_VAL	2:0	R/W	Map the frame's PCP and DEI values to a QoS class. QoS class = QOS_PCP_DEI_MAP_CFG[index].QOS_PCP_DEI_VAL, where index = 8*DEI + PCP. Only applicable to tagged frames. The use of inner or outer tag can be selected using VLAN_CFG.VLAN_INNER_TAG_ENA.	0x0

5.1.8.8 ANA:PORT:CPU_FWD_CFG

Parent: [ANA:PORT](#)

Instances: 1

TABLE 5-58: FIELDS IN CPU_FWD_CFG

Field Name	Bit	Access	Description	Default
CPU_VRAP_REDIR_ENA	7	R/W	If set, VRAP frames are redirected to the CPU extraction queue given by CPUQ_CFG2.CPUQ_VRAP.	0x0
CPU_MLD_REDIR_ENA	6	R/W	If set, MLD frames are redirected to the CPU.	0x0
CPU_IGMP_REDIR_ENA	5	R/W	If set, IGMP frames are redirected to the CPU.	0x0
CPU_IPMC_CTRL_COPY_ENA	4	R/W	If set, IPv4 multicast control frames (destination IP address in the range 224.0.0.x) are copied to the CPU.	0x0
CPU_SRC_COPY_ENA	3	R/W	If set, all frames received on this port are copied to the CPU extraction queue given by CPUQ_CFG.CPUQ_SRC_COPY.	0x0
CPU_ALLBRIDGE_DROP_ENA	2	R/W	If set, All LANs Bridge Management Group Address frames (DMAC = 01-80-C2-00-00-10) are not forwarded to any front ports. Together with CPU_ALLBRIDGE_REDIR_ENA, CPU_ALLBRIDGE_DROP_ENA controls the forwarding: CPU_ALLBRIDGE_DROP_ENA=0, CPU_ALLBRIDGE_REDIR_ENA=0: No action CPU_ALLBRIDGE_DROP_ENA=0, CPU_ALLBRIDGE_REDIR_ENA=1: Redirect to CPU CPU_ALLBRIDGE_DROP_ENA=1, CPU_ALLBRIDGE_REDIR_ENA=0: Discard CPU_ALLBRIDGE_DROP_ENA=1, CPU_ALLBRIDGE_REDIR_ENA=1: Copy to CPU	0x0
CPU_ALLBRIDGE_REDIR_ENA	1	R/W	If set, All LANs Bridge Management Group Address frames (DMAC = 01-80-C2-00-00-10) are redirected to the CPU. See also CPU_ALLBRIDGE_DROP_ENA.	0x0
CPU_OAM_ENA	0	R/W	If set, frames redirected to the cpu will still be added to the OAM statistics.	0x0

5.1.8.9 ANA:PORT:CPU_FWD_BPDU_CFG

Parent: [ANA:PORT](#)

Instances: 1

TABLE 5-59: FIELDS IN CPU_FWD_BPDU_CFG

Field Name	Bit	Access	Description	Default
BPDU_DROP_ENA	31:16	R/W	If bit x is set, BPDU frame (DMAC = 01-80-C2-00-00-0x) is not forwarded to any front ports. Together with BPDU_REDIR_ENA, BPDU_DROP_ENA controls the forwarding of BPDU frames: BPDU_DROP_ENA=0, BPDU_REDIR_ENA=0: No action BPDU_DROP_ENA=0, BPDU_REDIR_ENA=1: Redirect to CPU BPDU_DROP_ENA=1, BPDU_REDIR_ENA=0: Discard BPDU_DROP_ENA=1, BPDU_REDIR_ENA=1: Copy to CPU	0x0000
BPDU_REDIR_ENA	15:0	R/W	If bit x is set, BPDU frame (DMAC = 01-80-C2-00-00-0x) is redirected to the CPU. See also BPDU_DROP_ENA.	0x0000

5.1.8.10 ANA:PORT:CPU_FWD_GARP_CFG

Parent: [ANA:PORT](#)

Instances: 1

TABLE 5-60: FIELDS IN CPU_FWD_GARP_CFG

Field Name	Bit	Access	Description	Default
GARP_DROP_ENA	31:16	R/W	If bit x is set, GARP frame (DMAC = 01-80-C2-00-00-2x) is not forwarded to any front port. Together with GARP_REDIR_ENA, GARP_DROP_ENA controls the forwarding of GARP frames: GARP_DROP_ENA=0, GARP_REDIR_ENA=0: No action GARP_DROP_ENA=0, GARP_REDIR_ENA=1: Redirect to CPU GARP_DROP_ENA=1, GARP_REDIR_ENA=0: Discard GARP_DROP_ENA=1, GARP_REDIR_ENA=1: Copy to CPU	0x0000
GARP_REDIR_ENA	15:0	R/W	If bit x is set, GARP frame (DMAC = 01-80-C2-00-00-2x) is redirected to the CPU. See also GARP_DROP_ENA.	0x0000

5.1.8.11 ANA:PORT:CPU_FWD_CCM_CFG

Parent: [ANA:PORT](#)

Instances: 1

TABLE 5-61: FIELDS IN CPU_FWD_CCM_CFG

Field Name	Bit	Access	Description	Default
CCM_DROP_ENA	31:16	R/W	If bit x is set, CCM/Link trace frame (DMAC = 01-80-C2-00-00-3x) is not forwarded to any front port. Together with CCM_REDIR_ENA, CCM_DROP_ENA controls the forwarding of CCM/Link trace frames: CCM_DROP_ENA=0, CCM_REDIR_ENA=0: No action CCM_DROP_ENA=0, CCM_REDIR_ENA=1: Redirect to CPU CCM_DROP_ENA=1, CCM_REDIR_ENA=0: Discard CCM_DROP_ENA=1, CCM_REDIR_ENA=1: Copy to CPU	0x0000
CCM_REDIR_ENA	15:0	R/W	If bit x is set, CCM/Link trace frame (DMAC = 01-80-C2-00-00-3x) is redirected to the CPU. See also CCM_DROP_ENA.	0x0000

5.1.8.12 ANA:PORT:PORT_CFG

Parent: [ANA:PORT](#)

Instances: 1

TABLE 5-62: FIELDS IN PORT_CFG

Field Name	Bit	Access	Description	Default
SRC_MIRROR_ENA	15	R/W	If set, all frames received on this port are mirrored to the port set configured in MIRRORPORTS (ie. ingress mirroring). For egress mirroring, see EMIRRORMASK.	0x0
LIMIT_DROP	14	R/W	If set, learn frames on an ingress port, which has exceeded the maximum number of MAC table entries are discarded. Forwarding to CPU is still allowed. Note that if LEARN_ENA is cleared, then the LIMIT_DROP is ignored.	0x0
LIMIT_CPU	13	R/W	If set, learn frames on an ingress port, which has exceeded the maximum number of MAC table entries are copied to the CPU extraction queue specified in CPUQ_CFG.CPUQ_LRN. Note that if LEARN_ENA is cleared, then the LIMIT_CPU is ignored.	0x0
LOCKED_PORTMOVE_DROP	12	R/W	If set, incoming frames triggering a port move for a locked entry in the MAC table received on this port are discarded. Forwarding to CPU is still allowed. Note that if LEARN_ENA is cleared, then the LOCKED_PORTMOVE_DROP is ignored.	0x0

TABLE 5-62: FIELDS IN PORT_CFG (CONTINUED)

Field Name	Bit	Access	Description	Default
LOCKED_PORTMOVE_CPU	11	R/W	If set, incoming frames triggering a port move for a locked MAC table entry received on this port are copied to the CPU extraction queue specified in CPUQ_CFG.CPUQ_LOCKED_PORTMOVE. Note that if LEARN_ENA is cleared, then the LOCKED_PORTMOVE_CPU is ignored.	0x0
LEARNDROP	10	R/W	If set, incoming learn frames received on this port are discarded. Forwarding to CPU is still allowed. Note that if LEARN_ENA is cleared, then the LEARNDROP is ignored.	0x0
LEARNCPU	9	R/W	If set, incoming learn frames received on this port are copied to the CPU extraction queue specified in AGENTRL.CPUQ_LRN. Note that if LEARN_ENA is cleared, then the LEARNCPU is ignored.	0x0
LEARNAUTO	8	R/W	If set, incoming learn frames received on this port are auto learned. Note that if LEARN_ENA is cleared, then the LEARNAUTO is ignored.	0x1
LEARN_ENA	7	R/W	Enable learning for frames received on this port. If cleared, learning is skipped and any configuration settings in LEARNAUTO, LEARNCPU, LEARNDROP is ignored.	0x1
RECV_ENA	6	R/W	Enable forwarding from this port based on VLAN and destination MAC address. If cleared, the VLAN and MAC tables are not looked up and result is set to 0. Learning and CPU forwarding are still possible.	0x1
PORTID_VAL	5:2	R/W	Logical port number for front port. If port is not a member of a LLAG, then PORTID must be set to the physical port number. If port is a member of a LLAG, then PORTID must be set to the common PORTID_VAL used for all member ports of the LLAG. The value must not exceed the number of physical ports on the device.	0x0
USE_B_DOM_TBL	1	R/W	Select B-domain table for frames received on this ingress port. MAC table is looked up and learning updated in the B-domain, not affecting or hitting stations learned on C-ports.	0x0
LSR_MODE	0	R/W	Assume LSR operation on this port. Rewriting will be enabled after popped encapsulation.	0x0

5.1.8.13 ANA:PORT:POL_CFG

Parent: [ANA:PORT](#)

Instances: 1

TABLE 5-63: FIELDS IN POL_CFG

Field Name	Bit	Access	Description	Default
PORT_POL_ENA	17	R/W	Enable port policing for frames using this virtual port profile. Port policing on physical port P uses policer P.	0x0
QUEUE_POL_ENA	16:9	R/W	Bit mask, where bit<n> enables policing of frames classified to QoS class n on this port. Policing of QoS class Q on physical port P uses policer 32+P*8+Q. Only applicable to non-service-classified frames (ISDX=0) or service-classified frames with disabled service policer (ISDX-TIDX.ISDX_SDLBI=0).	0x00
POL_ORDER	8:0	R/W	Each frame is checked against three policers: 1) port, 2) QoS or service, and 3) VCAP. In this register, a bit set makes updating of a policer dependant on the result from another. Bit<n+3*m> set means: Policer state <n> is checked before policer <m> is updated. Bit0: Port policer must be open in order to update port policer with frame Bit1: QoS or service policer must be open in order to update port policer with frame Bit2: VCAP policer must be open in order to update port policer with frame Bit3: Port policer must be open in order to update QoS or service policer with frame Bit4: QoS or service policer must be open in order to update QoS or service policer with frame Bit5: VCAP policer must be open in order to update QoS or service policer with frame Bit6: Port policer must be open in order to update VCAP policer with frame Bit7: QoS or service policer must be open in order to update VCAP policer with frame Bit8: VCAP policer must be open in order to update VCAP policer with frame	0x1FF

5.1.8.14 ANA:PORT:PTP_CFG

Parent: [ANA:PORT](#)

Instances: 1

TABLE 5-64: FIELDS IN PTP_CFG

Field Name	Bit	Access	Description	Default
PTP_BACKPLANE_MODE	0	R/W	Backplane mode for PTP. If set, the receive timestamp is read from the PTP header (4 reserved bytes at offset 16) instead of from the ingress timer. Ingress delays, PTP_DLY1_CFG and PTP_DLY2_CFG, are not applied in backplane mode.	0x0

5.1.8.15 ANA:PORT:PTP_DLY1_CFG

Parent: [ANA:PORT](#)

Instances: 1

TABLE 5-65: FIELDS IN PTP_DLY1_CFG

Field Name	Bit	Access	Description	Default
USE_QOS_AS_PCP_ENA	31	R/W	If set, the frame's classified PCP out of the basic classifier is set to the frame's classified QoS class. This overrides normal PCP classification in the basic classifier. The PCP classification can still be changed by IS1 actions PCP_ENA and PCP_VAL.	0x0
PTP_DLY1_VAL	30:0	R/W	Delay optionally subtracted from the frame's receive timestamp. This can for instance be a peer-to-peer link delay or a delay due to asymmetry. The subtraction is enabled through IS2 action REW_OP for one-step PTP. Unit is ns. Value is signed.	0x00000000

5.1.8.16 ANA:PORT:PTP_DLY2_CFG

Parent: [ANA:PORT](#)

Instances: 1

TABLE 5-66: FIELDS IN PTP_DLY2_CFG

Field Name	Bit	Access	Description	Default
PTP_DLY2_VAL	31:0	R/W	Delay optionally subtracted from the frame's receive timestamp. This can for instance be a peer-to-peer link delay or a delay due to asymmetry. The subtraction is enabled through IS2 action REW_OP for one-step PTP. This is not applicable if the ingress port is in backplane mode. Unit is ns. Value is signed.	0x00000000

5.1.9 ANA:PFC

Parent: [ANA](#)

Instances: 11

TABLE 5-67: REGISTERS IN PFC

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
PFC_CFG	0x00000000	1	Priority-based flow control configuration	Page 356

5.1.9.1 ANA:PFC:PFC_CFG

Parent: [ANA:PFC](#)

Instances: 1

TABLE 5-68: FIELDS IN PFC_CFG

Field Name	Bit	Access	Description	Default
RX_PFC_ENA	9:2	R/W	Enable PFC per priority. Bit n enables PFC on priority n.	0x00
FC_LINK_SPEED	1:0	R/W	Configures the link speed. This is used to evaluate the time specifications in incoming pause frames. 0: 2500 Mbps 1: 1000 Mbps 2: 100 Mbps 3: 10 Mbps	0x1

5.1.10 ANA:IPT

Parent: [ANA](#)

Instances: 1024

TABLE 5-69: REGISTERS IN IPT

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
OAM_MEP_CFG	0x00000000	1	Configuration of OAM MEP	Page 356
IPT	0x00000004	1	ISDX protection table	Page 357

5.1.10.1 ANA:IPT:OAM_MEP_CFG

Parent: [ANA:IPT](#)

Instances: 1

TABLE 5-70: FIELDS IN OAM_MEP_CFG

Field Name	Bit	Access	Description	Default
MEP_IDX_P	12:7	R/W	Configures the associated protected OAM MEP for service protection.	0x00
MEP_IDX	6:1	R/W	Configures the associated OAM MEP.	0x00
MEP_IDX_ENA	0	R/W	Use the associated OAM MEP.	0x0

5.1.10.2 ANA:IPT:IPT

Parent: [ANA:IPT](#)

Instances: 1

TABLE 5-71: FIELDS IN IPT

Field Name	Bit	Access	Description	Default
IPT_CFG	18:17	R/W	Protection configuration for ISDX. 0: No protection 1: Replace classified ISDX with ISDX_P if the ISDX's path group uses the protected path. 2: Working path: Discard the frame if the ISDX's path group uses the protected path. 3: Protected path: Discard the frame if the ISDX's path group uses the working path.	0x0
ISDX_P	16:7	R/W	Replace the frame's classified ISDX with ISDX_P if the ISDX's path group is upstream-protected and is using the protect path. The new ISDX_P is used in the ISDX table, in VCAP IS2, VCAP ES0, and when updating ingress service counters.	0x000
PPT_IDX	6:0	R/W	Path protection group. This index points to the path protection entry in the path protection table used for this ISDX. Only indexes 0 through 75 are valid.	0x00

5.1.11 ANA:PPT

Parent: [ANA](#)

Instances: 1

TABLE 5-72: REGISTERS IN PPT

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
PPT	0x00000000	3 0x00000004	Path protection table	Page 357

5.1.11.1 ANA:PPT:PPT

Parent: [ANA:PPT](#)

Instances: 3

First register in this group maps to path protection group indexes 31:0.

Second register in this group maps to path protection group indexes 63:32.

Third register in this group maps to path protection group indexes 75:64.

TABLE 5-73: FIELDS IN PPT

Field Name	Bit	Access	Description	Default
PATH_STAT	31:0	R/W	Path protection status with one entry per path protection group. The path protection group is selected by IPT.PPT_IDX from the ISDX protection table. Bit n in PPT[0].PATH_CFG returns the status for path protection group n, n <= 31. Bit n in PPT[1].PATH_CFG returns the status for path protection group n+32, 32 <= n <= 63. Bit n in PPT[2].PATH_CFG returns the status for path protection group n+64, 64 <= n <= 75. Each bit is encoded the following way: 0: Path protection group uses working path. 1: Path protection group uses protected path.	0x00000000

5.1.12 ANA:FID_MAP

Parent: [ANA](#)

Instances: 1

TABLE 5-74: REGISTERS IN FID_MAP

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
FID_MAP	0x00000000	4096 0x00000004	Map VIDs to FIDs	Page 358

5.1.12.1 ANA:FID_MAP:FID_MAP

Parent: [ANA:FID_MAP](#)

Instances: 4096

FID mapping table per VID.

The table is looked up with the frame's VID as key. The returned FID is used in MAC table lookups, MAC table learning, and in MSTP forwarding (ANA::MSTI_STATE).

TABLE 5-75: FIELDS IN FID_MAP

Field Name	Bit	Access	Description	Default
FID_C_VAL	11:6	R/W	FID used for frames received on C-ports (ANA:PORT:PORT_CFG.USE_B_DOM_TA BLE=0). If FID_C_VAL=0 then FID_C_VAL=VID is used instead.	0x00

TABLE 5-75: FIELDS IN FID_MAP (CONTINUED)

Field Name	Bit	Access	Description	Default
FID_B_VAL	5:0	R/W	FID used for frames received on B-ports (ANA:PORT_CFG.USE_B_DOM_TA_BLE=1). If FID_B_VAL=0 then FID_B_VAL=VID is used instead.	0x00

5.1.13 ANA:COMMON

Parent: [ANA](#)

Instances: 1

TABLE 5-76: REGISTERS IN COMMON

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
AGGR_CFG	0x00000000	1	Aggregation code generation	Page 359
CPUQ_CFG	0x00000004	1	CPU extraction queue configuration	Page 360
CPUQ_CFG2	0x00000008	1	CPU extraction queue configuration 2	Page 360
CPUQ_8021_CFG	0x0000000C	16 0x00000004	CPU extraction queue per address of BPDU, GARP, and CCM frames.	Page 361
DSCP_CFG	0x0000004C	64 0x00000004	DSCP configuration per DSCP value.	Page 361
DSCP_REWR_CFG	0x0000014C	16 0x00000004	DSCP rewrite values per DP level and QoS class	Page 361
VCAP_RNG_TYPE_CFG	0x0000018C	8 0x00000004	VCAP range checkers	Page 362
VCAP_RNG_VAL_CFG	0x000001AC	8 0x00000004	Range configuration per range checker	Page 362
VRAP_CFG	0x000001CC	1	VRAP classifier configuration	Page 362
VRAP_HDR_DATA	0x000001D0	1	VRAP data	Page 363
VRAP_HDR_MASK	0x000001D4	1	VRAP mask	Page 363
DISCARD_CFG	0x000001D8	1	Various options for discard filters	Page 363
FID_CFG	0x000001DC	1	FID selector configuration	Page 364

5.1.13.1 ANA:COMMON:AGGR_CFG

Parent: [ANA:COMMON](#)

Instances: 1

TABLE 5-77: FIELDS IN AGGR_CFG

Field Name	Bit	Access	Description	Default
AC_RND_ENA	7	R/W	Use pseudo random number for aggregation code. Override other contributions.	0x0
AC_DMACH_ENA	6	R/W	Use the lower 12 bits of the destination MAC address for aggregation code.	0x0

TABLE 5-77: FIELDS IN AGGR_CFG (CONTINUED)

Field Name	Bit	Access	Description	Default
AC_SMAC_ENA	5	R/W	Use the lower 12 bits of the source MAC address for aggregation code.	0x0
AC_IP6_FLOW_LBL_ENA	4	R/W	Use the 20-bit IPv6 flow label for aggregation code.	0x0
AC_IP6_TCPUDP_ENA	3	R/W	Use least significant 8 bits of both source port and destination port of IPv6 frames for aggregation code.	0x0
AC_IP4_SIPDIP_ENA	2	R/W	Use least significant 8 bits of both source IP address and destination IP address of IPv4 frames for aggregation code.	0x0
AC_IP4_TCPUDP_ENA	1	R/W	Use least significant 8 bits of both source port and destination port of IPv4 frames for aggregation code.	0x0
AC_ISDX_ENA	0	R/W	Use the ISDX for aggregation code.	0x0

5.1.13.2 ANA:COMMON:CPUQ_CFG

Parent: [ANA:COMMON](#)

Instances: 1

TABLE 5-78: FIELDS IN CPUQ_CFG

Field Name	Bit	Access	Description	Default
CPUQ_MLD	29:27	R/W	CPU extraction queue used for MLD frames.	0x0
CPUQ_IGMP	26:24	R/W	CPU extraction queue used for IGMP frames.	0x0
CPUQ_IPMC_CTRL	23:21	R/W	CPU extraction queue used for IPv4 multi-cast control frames.	0x0
CPUQ_ALLBRIDGE	20:18	R/W	CPU extraction queue used for All LANs Bridge Management Group Address frames (DMAC = 01-80-C2-00-00-10).	0x0
CPUQ_LOCKED_PORTMOVE	17:15	R/W	CPU extraction queue for frames triggering a port move for a locked MAC table entry.	0x0
CPUQ_SRC_COPY	14:12	R/W	CPU extraction queue for frames copied due to CPU_SRC_COPY_ENA	0x0
CPUQ_MAC_COPY	11:9	R/W	CPU extraction queue for frames copied due to CPU_COPY return by MAC table lookup	0x0
CPUQ_LRN	8:6	R/W	CPU extraction queue for frames copied due to learned or moved stations.	0x0
CPUQ_MIRROR	5:3	R/W	CPU extraction queue for frames copied due to mirroring to the CPU.	0x0
CPUQ_SFLOW	2:0	R/W	CPU extraction queue for frames copied due to SFLOW sampling.	0x0

5.1.13.3 ANA:COMMON:CPUQ_CFG2

Parent: [ANA:COMMON](#)

Instances: 1

TABLE 5-79: FIELDS IN CPUQ_CFG2

Field Name	Bit	Access	Description	Default
CPUQ_VRAP	2:0	R/W	CPU extraction queue used for VRAP frames.	0x0

5.1.13.4 ANA:COMMON:CPUQ_8021_CFG

Parent: [ANA:COMMON](#)

Instances: 16

The register instance number corresponds to the address of the extracted frame. For instance: CPUQ_8021_CFG[4].CPUQ_BPDU_VAL is the CPU extraction queue used for BPDU frames with address 01-80-C2-00-00-04.

TABLE 5-80: FIELDS IN CPUQ_8021_CFG

Field Name	Bit	Access	Description	Default
CPUQ_BPDU_VAL	8:6	R/W	CPU extraction queue used for BPDU frames.	0x0
CPUQ_GARP_VAL	5:3	R/W	CPU extraction queue used for GARP frames.	0x0
CPUQ_CCM_VAL	2:0	R/W	CPU extraction queue used for CCM/Link trace frames.	0x0

5.1.13.5 ANA:COMMON:DSCP_CFG

Parent: [ANA:COMMON](#)

Instances: 64

TABLE 5-81: FIELDS IN DSCP_CFG

Field Name	Bit	Access	Description	Default
DP_DSCP_VAL	11	R/W	Maps the frame's DSCP value to a drop precedence level. This is enabled in QOS_CFG.QOS_DSCP_ENA.	0x0
QOS_DSCP_VAL	10:8	R/W	Maps the frame's DSCP value to a QoS class. This is enabled in QOS_CFG.QOS_DSCP_ENA.	0x0
DSCP_TRANSLATE_VAL	7:2	R/W	Translated DSCP value triggered if DSCP translation is set for port (QOS_CFG[port].DSCP_TRANSLATE_ENA)	0x00
DSCP_TRUST_ENA	1	R/W	Must be set for a DSCP value if the DSCP value is to be used for QoS classification.	0x0
DSCP_REWR_ENA	0	R/W	Set if the DSCP value is selected to be rewritten. This is controlled in QOS_CFG.DSCP_REWR_CFG.	0x0

5.1.13.6 ANA:COMMON:DSCP_REWR_CFG

Parent: [ANA:COMMON](#)

Instances: 16

TABLE 5-82: FIELDS IN DSCP_REWR_CFG

Field Name	Bit	Access	Description	Default
DSCP_QOS_REWR_VAL	5:0	R/W	Map the frame's DP level and QoS class to a DSCP value. $DSCP = DSCP\_REWR\_CFG[8*DP\ level + QoS\ class].DSCP\_QOS\_REWR\_VAL$. This is controlled in <code>QOS_CFG.DSCP_REWR_CFG</code> and <code>DSCP_CFG.DSCP_REWR_ENA</code> .	0x00

5.1.13.7 ANA:COMMON:VCAP_RNG_TYPE_CFG

Parent: [ANA:COMMON](#)

Instances: 8

TABLE 5-83: FIELDS IN VCAP_RNG_TYPE_CFG

Field Name	Bit	Access	Description	Default
VCAP_RNG_CFG	2:0	R/W	0: Idle 1: TCP/UDP destination port is matched against range 2: TCP/UDP source port is matched against range 3: TCP/UDP source and destination ports are matched against range. Match if either source or destination port is within range. 4: VID is matched against range (S1: VID in frame, S2: classified VID) 5: DSCP value is matched against range 6: Reserved 7: Reserved	0x0

5.1.13.8 ANA:COMMON:VCAP_RNG_VAL_CFG

Parent: [ANA:COMMON](#)

Instances: 8

TABLE 5-84: FIELDS IN VCAP_RNG_VAL_CFG

Field Name	Bit	Access	Description	Default
VCAP_RNG_MIN_VAL	31:16	R/W	Lower value. Value is included in range.	0x0000
VCAP_RNG_MAX_VAL	15:0	R/W	Upper value. Value is included in range.	0x0000

5.1.13.9 ANA:COMMON:VRAP_CFG

Parent: [ANA:COMMON](#)

Instances: 1

TABLE 5-85: FIELDS IN VRAP_CFG

Field Name	Bit	Access	Description	Default
VRAP_VLAN_AWARE_ENA	12	R/W	If set, VRAP frames must be single VLAN tagged and the frame's VID must match ANA::VRAP_CFG.VRAP_VID. If cleared, VRAP frames must be untagged.	0x0
VRAP_VID	11:0	R/W	VID value for VRAP frames.	0x000

5.1.13.10 ANA:COMMON:VRAP_HDR_DATA

Parent: [ANA:COMMON](#)

Instances: 1

TABLE 5-86: FIELDS IN VRAP_HDR_DATA

Field Name	Bit	Access	Description	Default
VRAP_HDR_DATA	31:0	R/W	The frame's VRAP header (4 bytes after EPID) is matched against VRAP_HDR_DATA, except for bits don't cared by VRAP_HDR_MASK.	0x10000000

5.1.13.11 ANA:COMMON:VRAP_HDR_MASK

Parent: [ANA:COMMON](#)

Instances: 1

TABLE 5-87: FIELDS IN VRAP_HDR_MASK

Field Name	Bit	Access	Description	Default
VRAP_HDR_MASK	31:0	R/W	Bits set in VRAP_HDR_MASK don't care the equivalent bits in VRAP_HDR_DATA.	0x07FFFFFF

5.1.13.12 ANA:COMMON:DISCARD_CFG

Parent: [ANA:COMMON](#)

Instances: 1

TABLE 5-88: FIELDS IN DISCARD_CFG

Field Name	Bit	Access	Description	Default
DROP_TAGGING_ISDX0	3	R/W	Frames discarded due to wrong tagging will be classified to ISDX 0.	0x1
DROP_CTRLPROT_ISDX0	2	R/W	Frames discarded with L2/L3 protocol filter will be classified to ISDX 0.	0x1
DROP_TAGGING_S2_ENA	1	R/W	Frames discarded due to wrong tagging will still allow S2 lookup.	0x1
DROP_CTRLPROT_S2_ENA	0	R/W	Frames discarded with L2/L3 protocol filter will still allow S2 lookup.	0x1

5.1.13.13 ANA:COMMON:FID_CFG

Parent: [ANA:COMMON](#)

Instances: 1

TABLE 5-89: FIELDS IN FID_CFG

Field Name	Bit	Access	Description	Default
VID_MC_ENA	0	R/W	If set, the frame's classified VID is used instead of the FID for multicast frames when hashing into the MAC table for the DMAC lookup. The IS1 action FID_SEL can still overrule this setting. The SMAC lookup is unchanged by this configuration.	0x0

5.1.14 ANA:POL

Parent: [ANA](#)

Instances: 448

General purpose policers selected by port configuration and ACL actions

TABLE 5-90: REGISTERS IN POL

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
POL_PIR_CFG	0x00000000	1	Excess configuration	Page 364
POL_CIR_CFG	0x00000004	1	Committed configuration	Page 364
POL_MODE_CFG	0x00000008	1	Common configuration for this policer	Page 365
POL_PIR_STATE	0x0000000C	1	Excess state of the policer	Page 365
POL_CIR_STATE	0x00000010	1	Committed state of the policer	Page 366

5.1.14.1 ANA:POL:POL_PIR_CFG

Parent: [ANA:POL](#)

Instances: 1

TABLE 5-91: FIELDS IN POL_PIR_CFG

Field Name	Bit	Access	Description	Default
PIR_RATE	20:6	R/W	Excess information rate (EIR). Unit is 33.3 kbps.	0x0000
PIR_BURST	5:0	R/W	Excess burst size (EBS). Maximum allowed value is 60. Unit is 4 kilobytes.	0x00

5.1.14.2 ANA:POL:POL_CIR_CFG

Parent: [ANA:POL](#)

Instances: 1

TABLE 5-92: FIELDS IN POL_CIR_CFG

Field Name	Bit	Access	Description	Default
CIR_RATE	20:6	R/W	Committed information rate (CIR). Unit is 33.3 kbps.	0x0000
CIR_BURST	5:0	R/W	Committed burst size (CBS). Maximum allowed value is 60. Unit is 4 kilobytes.	0x00

5.1.14.3 ANA:POL:POL_MODE_CFG

Parent: [ANA:POL](#)

Instances: 1

TABLE 5-93: FIELDS IN POL_MODE_CFG

Field Name	Bit	Access	Description	Default
IPG_SIZE	9:5	R/W	Size of interframe gap to add to each frame if line rate policing is chosen in FRM_MODE.	0x14
FRM_MODE	4:3	R/W	Accounting mode of the policer. 0: Line rate. Measure bytes in frames including IPG_SIZE. 1: Data rate. Measure bytes in frames excluding IPG. 2: Frame rate. Measure frames with rate unit = 33.3 frames per second and burst unit = 32.8 frames. 3: Frame rate. Measure frames with rate unit = 0.33 frame per second burst unit = 0.3 frames.	0x0
DLB_COUPLED	2	R/W	Coupling flag. If enabled, yellow frames are allowed to use CIR as well as EIR. If disabled, yellow frames are only allowed to use CIR. 0: Coupling is disabled. 1: Coupling is enabled.	0x0
CIR_ENA	1	R/W	Enable dual leaky bucket mode. If disabled, the policer operates as a single leaky bucket policer governed by POL_PIR_CFG. If enabled, the policer operates as a dual leaky bucket policer governed by POL_PIR_CFG and POL_CIR_CFG.	0x0
RESERVED	0	R/W	Must be set to its default.	0x1

5.1.14.4 ANA:POL:POL_PIR_STATE

Parent: [ANA:POL](#)

Instances: 1

TABLE 5-94: FIELDS IN POL_PIR_STATE

Field Name	Bit	Access	Description	Default
PIR_LVL	21:0	R/W	Current fill level of the excess bucket. Unit is 0.5 bits.	0x000000

VSC7416-01

5.1.14.5 ANA:POL:POL_CIR_STATE

Parent: [ANA:POL](#)

Instances: 1

TABLE 5-95: FIELDS IN POL_CIR_STATE

Field Name	Bit	Access	Description	Default
CIR_LVL	21:0	R/W	Current fill level of the committed bucket. Unit is 0.5 bits.	0x000000

5.1.15 ANA:POL_MISC

Parent: [ANA](#)

Instances: 1

TABLE 5-96: REGISTERS IN POL_MISC

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
POL_FLOWC	0x00000000	27 0x00000004	Flow control configuration per port	Page 366
POL_HYST	0x0000006C	1	Set flow control hysteresis	Page 366

5.1.15.1 ANA:POL_MISC:POL_FLOWC

Parent: [ANA:POL_MISC](#)

Instances: 27

TABLE 5-97: FIELDS IN POL_FLOWC

Field Name	Bit	Access	Description	Default
POL_FLOWC	0	R/W	Use MAC flow control for lowering ingress rate 0: Standard policing. Frames are discarded when the rate is exceeded. 1: Flow control policing. Policer instructs the MAC to issue pause frames when the rate is exceeded.	0x0

5.1.15.2 ANA:POL_MISC:POL_HYST

Parent: [ANA:POL_MISC](#)

Instances: 1

TABLE 5-98: FIELDS IN POL_HYST

Field Name	Bit	Access	Description	Default
POL_FC_HYST	9:4	R/W	Set hysteresis for when to re-open a bucket after the burst capacity has been used. Unit is 1 kilobytes. This applies to policer in flow control mode (POL_FLOWC=1).	0x02

TABLE 5-98: FIELDS IN POL_HYST (CONTINUED)

Field Name	Bit	Access	Description	Default
POL_DROP_HYST	3:0	R/W	Set hysteresis for when to re-open a bucket after the burst capacity has been used. Unit is 2 kilobytes. This applies to policer in drop mode (POL_FLOWC=0).	0x0

5.2 DEV

TABLE 5-99: REGISTER GROUPS IN DEV

Register Group Name	Offset within Target	Instances and Address Spacing	Description	Details
PORT_MODE	0x00000000	1	Port configuration	Page 367
MAC_CFG_STATUS	0x00000010	1	MAC configuration and status	Page 369
PCS1G_CFG_STATUS	0x0000003C	1	PCS 1G Configuration Status Registers	Page 374
PCS1G_TSTPAT_CFG_STATUS	0x00000080	1	PCS1G Testpattern Configuration and Status Registers	Page 381
PCS_FX100_CONFIGURATION	0x00000088	1	PCS FX100 Configuration Registers	Page 382
PCS_FX100_STATUS	0x0000008C	1	PCS FX100 Status Registers	Page 383

5.2.1 DEV:PORT_MODE

Parent: [DEV](#)

Instances: 1

TABLE 5-100: REGISTERS IN PORT_MODE

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
CLOCK_CFG	0x00000000	1	Clock and reset configuration	Page 367
PORT_MISC	0x00000004	1	Forwarding configuration of special frame types	Page 368
EEE_CFG	0x0000000C	1	Control Energy Efficient Ethernet operation.	Page 368

5.2.1.1 DEV:PORT_MODE:CLOCK_CFG

Parent: [DEV:PORT_MODE](#)

Instances: 1

TABLE 5-101: FIELDS IN CLOCK_CFG

Field Name	Bit	Access	Description	Default
MAC_TX_RST	7	R/W	Reset the MAC Tx clock domain.	0x1
MAC_RX_RST	6	R/W	Reset the MAC Rx clock domain.	0x1
PCS_TX_RST	5	R/W	Reset the PCS Tx clock domain.	0x1
PCS_RX_RST	4	R/W	Reset the PCS Rx clock domain.	0x1

TABLE 5-101: FIELDS IN CLOCK_CFG (CONTINUED)

Field Name	Bit	Access	Description	Default
PORT_RST	3	R/W		0x1
RESERVED	2	R/W	Must be set to its default.	0x1
LINK_SPEED	1:0	R/W	Selects the link speed. 0: No link 1: 1000/2500 Mbps 2: 100 Mbps 3: 10 Mbps	0x0

5.2.1.2 DEV:PORT_MODE:PORT_MISC

Parent: [DEV:PORT_MODE](#)

Instances: 1

TABLE 5-102: FIELDS IN PORT_MISC

Field Name	Bit	Access	Description	Default
FWD_ERROR_ENA	4	R/W	Forward frames with errors signaled by the MAC.	0x0
FWD_PAUSE_ENA	3	R/W	Forward pause frames (EtherType = 0x8808, opcode = 0x0001). The reaction to incoming pause frames is controlled independently of FWD_PAUSE_ENA.	0x0
FWD_CTRL_ENA	2	R/W	Forward MAC control frames excluding pause frames (EtherType = 0x8808, opcode different from 0x0001).	0x0

5.2.1.3 DEV:PORT_MODE:EEE_CFG

Parent: [DEV:PORT_MODE](#)

Instances: 1

TABLE 5-103: FIELDS IN EEE_CFG

Field Name	Bit	Access	Description	Default
EEE_ENA	22	R/W	Enable EEE operation on the port. A port enters the low power mode when no egress queues have data ready. The port is activated when one of the following conditions is true: - A queue has been non-empty for EEE_TIMER_AGE. - A queue has more than EEE_HIGH_FRAMES frames pending. - A queue has more than EEE_HIGH_BYTES bytes pending. - A queue is marked as a fast queue, and has data pending.	0x0

TABLE 5-103: FIELDS IN EEE_CFG (CONTINUED)

Field Name	Bit	Access	Description	Default
EEE_TIMER_AGE	21:15	R/W	Maximum time frames in any queue must wait before the port is activated. The default value corresponds to 48 us. Time = $4^{**}(\text{EEE\_TIMER\_AGE}/16) * (\text{EEE\_TIMER\_AGE} \bmod 16)$ microseconds	0x23
EEE_TIMER_WAKEUP	14:8	R/W	Time from the egress port is activated until frame transmission is restarted. Default value corresponds to 16 us. Time = $4^{**}(\text{EEE\_TIMER\_WAKEUP}/16) * (\text{EEE\_TIMER\_WAKEUP} \bmod 16)$ microseconds	0x14
EEE_TIMER_HOLDOFF	7:1	R/W	When all queues are empty, the port is kept active until this time has passed. Default value corresponds to 5 us. Time = $4^{**}(\text{EEE\_TIMER\_HOLDOFF}/16) * (\text{EEE\_TIMER\_HOLDOFF} \bmod 16)$ microseconds	0x05
PORT_LPI	0	R/O	Status bit indicating whether port is in low-power-idle due to the LPI algorithm (EEE_CFG). If set, transmissions are held back.	0x0

5.2.2 DEV:MAC_CFG_STATUS

Parent: [DEV](#)

Instances: 1

The 1G/2.5G MAC module contains configuration and status registers related to the MAC module of the 1G and 2.5G ports.

TABLE 5-104: REGISTERS IN MAC_CFG_STATUS

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
MAC_ENA_CFG	0x00000000	1	Mode Configuration Register	Page 370
MAC_MODE_CFG	0x00000004	1	Mode Configuration Register	Page 370
MAC_MAXLEN_CFG	0x00000008	1	Max Length Configuration Register	Page 370
MAC_TAGS_CFG	0x0000000C	1	VLAN / Service tag configuration register	Page 370
MAC_ADV_CHK_CFG	0x00000010	1	Advanced Check Feature Configuration Register	Page 371
MAC_IFG_CFG	0x00000014	1	Inter Frame Gap Configuration Register	Page 371
MAC_HDX_CFG	0x00000018	1	Half Duplex Configuration Register	Page 372
MAC_FC_MAC_LOW_CFG	0x00000020	1	MAC Flow Control Configuration Register	Page 373
MAC_FC_MAC_HIGH_CFG	0x00000024	1	MAC Flow Control Configuration Register	Page 373
MAC_STICKY	0x00000028	1	Sticky Bit Register	Page 373

5.2.2.1 DEV:MAC_CFG_STATUS:MAC_ENA_CFG

Parent: [DEV:MAC_CFG_STATUS](#)

Instances: 1

TABLE 5-105: FIELDS IN MAC_ENA_CFG

Field Name	Bit	Access	Description	Default
RX_ENA	4	R/W	Receiver Module Enable. '0': Receiver Module Disabled '1': Receiver Module Enabled	0x0
TX_ENA	0	R/W	Transmitter Module Enable. '0': Transmitter Module Disabled '1': Transmitter Module Enabled	0x0

5.2.2.2 DEV:MAC_CFG_STATUS:MAC_MODE_CFG

Parent: [DEV:MAC_CFG_STATUS](#)

Instances: 1

TABLE 5-106: FIELDS IN MAC_MODE_CFG

Field Name	Bit	Access	Description	Default
GIGA_MODE_ENA	4	R/W	Enables 1 Gbps mode. '0': 10/100 Mbps mode '1': 1 Gbps mode. Note: FDX_ENA must also be set.	0x1
FDX_ENA	0	R/W	Enables Full Duplex: '0': Half Duplex '1': Full duplex. Note: Full duplex MUST be selected if GIGA_MODE is enabled.	0x1

5.2.2.3 DEV:MAC_CFG_STATUS:MAC_MAXLEN_CFG

Parent: [DEV:MAC_CFG_STATUS](#)

Instances: 1

TABLE 5-107: FIELDS IN MAC_MAXLEN_CFG

Field Name	Bit	Access	Description	Default
MAX_LEN	15:0	R/W	When set, single tagged frames are allowed to be 4 bytes longer than the MAC_MAXLEN_CFG configuration and double tagged frames are allowed to be 8 bytes longer. Single tagged frames are adjusted if VLAN_AWR_ENA is also set. Double tagged frames are adjusted if both VLAN_AWR_ENA and VLAN_DBL_AWR_ENA are set.	0x05EE

5.2.2.4 DEV:MAC_CFG_STATUS:MAC_TAGS_CFG

Parent: [DEV:MAC_CFG_STATUS](#)

Instances: 1

The MAC can be configured to accept 0, 1 and 2 tags and the TAG value can be user-defined.

TABLE 5-108: FIELDS IN MAC_TAGS_CFG

Field Name	Bit	Access	Description	Default
TAG_ID	31:16	R/W	This field defines which EtherTypes are recognized as a VLAN TPID - besides 0x8100. The value is used for all tag positions. I.e. a double tagged frame can have the following tag values: (TAG1,TAG2): (0x8100, 0x8100) (0x8100, TAG_ID) (TAG_ID, 0x8100) or (TAG_ID, TAG_ID) Single tagged frame can have the following TPID values: 0x8100 or TAG_ID.	0x8100
VLAN_DBL_AWR_ENA	1	R/W	If set, double tagged frames are subject to length adjustments (VLAN_LEN_AWR_ENA). VLAN_AWR_ENA must be set when VLAN_DBL_AWR_ENA is set. '0': The MAC does not look for inner tags. '1': The MAC accepts inner tags with TPID=0x8100 or TAG_ID.	0x0
VLAN_AWR_ENA	0	R/W	If set, single tagged frames are subject to length adjustments (VLAN_LEN_AWR_ENA). '0': The MAC does not look for any tags. '1': The MAC accepts outer tags with TPID=0x8100 or TAG_ID.	0x0
VLAN_LEN_AWR_ENA	2	R/W	When set, single tagged frames are allowed to be 4 bytes longer than the MAC_MAX_LEN_CFG configuration and double tagged frames are allowed to be 8 bytes longer. Single tagged frames are adjusted if VLAN_AWR_ENA is also set. Double tagged frames are adjusted if both VLAN_AWR_ENA and VLAN_DBL_AWR_ENA are set.	0x1

5.2.2.5 DEV:MAC_CFG_STATUS:MAC_ADV_CHK_CFG

Parent: [DEV:MAC_CFG_STATUS](#)

Instances: 1

TABLE 5-109: FIELDS IN MAC_ADV_CHK_CFG

Field Name	Bit	Access	Description	Default
LEN_DROP_ENA	0	R/W	Length Drop Enable: Configures the Receive Module to drop frames in reference to in-range and out-of-range errors: '0': Length Drop Disabled '1': Length Drop Enabled.	0x0

5.2.2.6 DEV:MAC_CFG_STATUS:MAC_IFG_CFG

Parent: [DEV:MAC_CFG_STATUS](#)

Instances: 1

TABLE 5-110: FIELDS IN MAC_IFG_CFG

Field Name	Bit	Access	Description	Default
TX_IFG	12:8	R/W	Used to adjust the duration of the inter-frame gap in the Tx direction and must be set according to the speed and duplex settings. 10/100 Mbps, HDX, FDX: 0x19, 0x13 1000 Mbps: 0x07.	0x07
RX_IFG2	7:4	R/W	Used to adjust the duration of the second part of the inter-frame gap in the Rx direction and must be set according to the speed and duplex settings. 10/100 Mbps, HDX, FDX: 0x8, 0xB 1000 Mbps: 0x1.	0x1
RX_IFG1	3:0	R/W	Used to adjust the duration of the first part of the inter-frame gap in the Rx direction and must be set according to the speed settings. 10/100 Mbps: 0x7 1000 Mbps: 0x5.	0x5

5.2.2.7 DEV:MAC_CFG_STATUS:MAC_HDX_CFG

Parent: [DEV:MAC_CFG_STATUS](#)

Instances: 1

TABLE 5-111: FIELDS IN MAC_HDX_CFG

Field Name	Bit	Access	Description	Default
WEXC_DIS	24	R/W	Determines whether the MAC backs off after an excessive collision has occurred. If set, back off is disabled after excessive collisions. '0': Back off after excessive collisions '1': Don't back off after excessive collisions	0x0
SEED	23:16	R/W	Seed value loaded into the PRBS of the MAC. Used to prevent excessive collision events.	0x00
SEED_LOAD	12	R/W	Load SEED value into PRNG register. A SEED value is loaded into the PRNG register of the MAC, when SEED_LOAD is asserted. After a load, the SEED_LOAD must be deasserted. '0': Do not load SEED value '1': Load SEED value.	0x0

TABLE 5-111: FIELDS IN MAC_HDX_CFG (CONTINUED)

Field Name	Bit	Access	Description	Default
RETRY_AFTER_EXC_COL_ENA	8	R/W	This bit is used to setup the MAC to retransmit a frame after an early collision even though 16 (or more) early collisions have occurred. '0': A frame is discarded and counted as an excessive collision if 16 collisions occur for this frame. '1': The MAC retransmits a frame after an early collision, regardless of the number of previous early collisions. The backoff sequence is reset after every 16 collisions.	0x0
LATE_COL_POS	6:0	R/W	Adjustment of early/late collision boundary: This bitgroup is used to adjust the MAC so that a collision on a shared transmission medium before bit 512 is handled as an early collision, whereas a collision after bit 512 is handled as a late collision, i.e. no retransmission is performed.	0x43

5.2.2.8 DEV:MAC_CFG_STATUS:MAC_FC_MAC_LOW_CFG

Parent: [DEV:MAC_CFG_STATUS](#)

Instances: 1

TABLE 5-112: FIELDS IN MAC_FC_MAC_LOW_CFG

Field Name	Bit	Access	Description	Default
MAC_LOW	23:0	R/W	Lower three bytes in the SMAC in generated flow control frames. 0xNNN: Lower three DMAC bytes	0x000000

5.2.2.9 DEV:MAC_CFG_STATUS:MAC_FC_MAC_HIGH_CFG

Parent: [DEV:MAC_CFG_STATUS](#)

Instances: 1

TABLE 5-113: FIELDS IN MAC_FC_MAC_HIGH_CFG

Field Name	Bit	Access	Description	Default
MAC_HIGH	23:0	R/W	Higher three bytes in the SMAC in generated flow control frames. 0xNNN: Higher three DMAC bytes	0x000000

5.2.2.10 DEV:MAC_CFG_STATUS:MAC_STICKY

Parent: [DEV:MAC_CFG_STATUS](#)

Instances: 1

Clear the sticky bits by writing a '0' in the relevant bitgroups (writing a '1' sets the bit)!.

TABLE 5-114: FIELDS IN MAC_STICKY

Field Name	Bit	Access	Description	Default
RX_IPG_SHRINK_STICKY	9	Sticky	Sticky bit indicating that an inter packet gap shrink was detected (IPG < 12 bytes).	0x0
RX_PREAM_SHRINK_STICKY	8	Sticky	Sticky bit indicating that a preamble shrink was detected (preamble < 8 bytes). '0': no preamble shrink was detected '1': a preamble shrink was detected one or more times Bit is cleared by writing a '1' to this position.	0x0
RX_CARRIER_EXT_STICKY	7	Sticky	Sticky bit indicating that a carrier extend was detected. '0': no carrier extend was detected '1': one or more carrier extends were detected Bit is cleared by writing a '1' to this position.	0x0
RX_CARRIER_EXT_ERR_STICKY	6	Sticky	Sticky bit indicating that a carrier extend error was detected. '0': no carrier extend error was detected '1': one or more carrier extend errors were detected Bit is cleared by writing a '1' to this position.	0x0
RX_JUNK_STICKY	5	Sticky	Sticky bit indicating that junk was received (bytes not recognized as a frame). '0': no junk was received '1': junk was received one or more times Bit is cleared by writing a '1' to this position.	0x0
TX_RETRANSMIT_STICKY	4	Sticky	Sticky bit indicating that the transmit MAC asked the host for a frame retransmission. '0': no tx retransmission was initiated '1': one or more tx retransmissions were initiated Bit is cleared by writing a '1' to this position.	0x0
TX_JAM_STICKY	3	Sticky	Sticky bit indicating that the transmit host issued a jamming signal. '0': the transmit host issued no jamming signal '1': the transmit host issued one or more jamming signals Bit is cleared by writing a '1' to this position.	0x0
TX_FIFO_OFLW_STICKY	2	Sticky	Sticky bit indicating that the MAC transmit FIFO has overrun.	0x0
TX_FRM_LEN_OVR_STICKY	1	Sticky	Sticky bit indicating that the transmit frame length has overrun. I.e. a frame longer than 64K occurred. '0': no tx frame length error occurred '1': one or more tx frames length errors occurred Bit is cleared by writing a '1' to this position.	0x0
TX_ABORT_STICKY	0	Sticky	Sticky bit indicating that the transmit host initiated abort was executed.	0x0

5.2.3 DEV:PCS1G_CFG_STATUS

Parent: [DEV](#)

Instances: 1

Configuration and status register set for PCS1G

TABLE 5-115: REGISTERS IN PCS1G_CFG_STATUS

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
PCS1G_CFG	0x00000000	1	PCS1G Configuration	Page 375
PCS1G_MODE_CFG	0x00000004	1	PCS1G Mode Configuration	Page 376
PCS1G_SD_CFG	0x00000008	1	PCS1G Signal Detect Configuration	Page 376
PCS1G_ANEG_CFG	0x0000000C	1	PCS1G Aneg Configuration	Page 376
PCS1G_ANEG_NP_CFG	0x00000010	1	PCS1G Aneg Next Page Configuration	Page 377
PCS1G_LB_CFG	0x00000014	1	PCS1G Loopback Configuration	Page 377
PCS1G_ANEG_STATUS	0x00000020	1	PCS1G ANEG Status Register	Page 378
PCS1G_ANEG_NP_STATUS	0x00000024	1	PCS1G Aneg Next Page Status Register	Page 378
PCS1G_LINK_STATUS	0x00000028	1	PCS1G link status	Page 378
PCS1G_LINK_DOWN_CNT	0x0000002C	1	PCS1G link down counter	Page 379
PCS1G_STICKY	0x00000030	1	PCS1G sticky register	Page 379
PCS1G_DEBUG_STATUS	0x00000034	1	PCS1G debug status	Page 380
PCS1G_LPI_CFG	0x00000038	1	PCS1G Low Power Idle Configuration	Page 380
PCS1G_LPI_WAKE_ERROR_CNT	0x0000003C	1	PCS1G wake error counter	Page 380
PCS1G_LPI_STATUS	0x00000040	1	PCS1G Low Power Idle Status	Page 380

5.2.3.1 DEV:PCS1G_CFG_STATUS:PCS1G_CFG

Parent: [DEV:PCS1G_CFG_STATUS](#)

Instances: 1

PCS1G main configuration register

TABLE 5-116: FIELDS IN PCS1G_CFG

Field Name	Bit	Access	Description	Default
LINK_STATUS_TYPE	4	R/W	Set type of link_status indication at CPU-System 0: Sync_status (from PCS synchronization state machine) 1: Bit 15 of PCS1G_ANEG_STATUS.lp_adv_ability (Link up/down)	0x0
PCS_ENA	0	R/W	PCS enable 0: Disable PCS 1: Enable PCS	0x0

5.2.3.2 DEV:PCS1G_CFG_STATUS:PCS1G_MODE_CFG

Parent: [DEV:PCS1G_CFG_STATUS](#)

Instances: 1

PCS1G mode configuration

TABLE 5-117: FIELDS IN PCS1G_MODE_CFG

Field Name	Bit	Access	Description	Default
UNIDIR_MODE_ENA	4	R/W	Unidirectional mode enable. Implementation of 802.3, Clause 66. When asserted, this enables MAC to transmit data independent of the state of the receive link. 0: Unidirectional mode disabled 1: Unidirectional mode enabled	0x0
SGMII_MODE_ENA	0	R/W	Selection of PCS operation 0: PCS is used in SERDES mode 1: PCS is used in SGMII mode. Configuration bit PCS1G_ANEG_CFG.SW_RE-SOLVE_ENA must be set additionally	0x1

5.2.3.3 DEV:PCS1G_CFG_STATUS:PCS1G_SD_CFG

Parent: [DEV:PCS1G_CFG_STATUS](#)

Instances: 1

PCS1G signal_detect configuration

TABLE 5-118: FIELDS IN PCS1G_SD_CFG

Field Name	Bit	Access	Description	Default
SD_SEL	8	R/W	Signal detect selection (select input for internal signal_detect line) 0: Select signal_detect line from hardmacro 1: Select external signal_detect line	0x0
SD_POL	4	R/W	Signal detect polarity: The signal level on signal_detect input pin must be equal to SD_POL to indicate signal detection (SD_ENA must be set) 0: Signal Detect input pin must be '0' to indicate a signal detection 1: Signal Detect input pin must be '1' to indicate a signal detection	0x1
SD_ENA	0	R/W	Signal Detect Enable 0: The Signal Detect input pin is ignored. The PCS assumes an active Signal Detect at all times 1: The Signal Detect input pin is used to determine if a signal is detected	0x1

5.2.3.4 DEV:PCS1G_CFG_STATUS:PCS1G_ANEG_CFG

Parent: [DEV:PCS1G_CFG_STATUS](#)

Instances: 1

PCS1G Auto-negotiation configuration register

TABLE 5-119: FIELDS IN PCS1G_ANEG_CFG

Field Name	Bit	Access	Description	Default
ADV_ABILITY	31:16	R/W	Advertised Ability Register: Holds the capabilities of the device as described IEEE 802.3, Clause 37. If SGMII mode is selected (PCS1G_MODE_CFG.SGMII_MODE_ENA = 1), SW_RESOLVE_ENA must be set.	0x0000
SW_RESOLVE_ENA	8	R/W	Software Resolve Abilities 0: If Auto Negotiation fails (no matching HD or FD capabilities) the link is disabled 1: The result of an Auto Negotiation is ignored - the link can be setup via software. This bit must be set in SGMII mode.	0x0
ANEG_RESTART_ONE_SHOT	1	One-shot	Auto Negotiation Restart 0: No action 1: Restart Auto Negotiation	0x0
ANEG_ENA	0	R/W	Auto Negotiation Enable 0: Auto Negotiation Disabled 1: Auto Negotiation Enabled	0x0

5.2.3.5 DEV:PCS1G_CFG_STATUS:PCS1G_ANEG_NP_CFG

Parent: [DEV:PCS1G_CFG_STATUS](#)

Instances: 1

PCS1G Auto-negotiation configuration register for next-page function

TABLE 5-120: FIELDS IN PCS1G_ANEG_NP_CFG

Field Name	Bit	Access	Description	Default
NP_TX	31:16	R/W	Next page register: Holds the next-page information as described in IEEE 802.3, Clause 37	0x0000
NP_LOADED_ONE_SHOT	0	One-shot	Next page loaded 0: next page is free and can be loaded 1: next page register has been filled (to be set after np_tx has been filled)	0x0

5.2.3.6 DEV:PCS1G_CFG_STATUS:PCS1G_LB_CFG

Parent: [DEV:PCS1G_CFG_STATUS](#)

Instances: 1

PCS1G Loop-Back configuration register

TABLE 5-121: FIELDS IN PCS1G_LB_CFG

Field Name	Bit	Access	Description	Default
TBI_HOST_LB_ENA	0	R/W	Loops data in PCS (TBI side) from egress direction to ingress direction. The Rx clock is automatically set equal to the Tx clock 0: TBI Loopback Disabled 1: TBI Loopback Enabled	0x0

5.2.3.7 DEV:PCS1G_CFG_STATUS:PCS1G_ANEG_STATUS

Parent: [DEV:PCS1G_CFG_STATUS](#)

Instances: 1

PCS1G Auto-negotiation status register

TABLE 5-122: FIELDS IN PCS1G_ANEG_STATUS

Field Name	Bit	Access	Description	Default
LP_ADV_ABILITY	31:16	R/O	Advertised abilities from link partner as described in IEEE 802.3, Clause 37	0x0000
PR	4	R/O	Resolve priority 0: ANEG is in progress 1: ANEG nearly finished - priority can be resolved (via software)	0x0
PAGE_RX_STICKY	3	Sticky	Status indicating whether a new page has been received. 0: No new page received 1: New page received Bit is cleared by writing a 1 to this position.	0x0
ANEG_COMPLETE	0	R/O	Auto Negotiation Complete 0: No Auto Negotiation has been completed 1: Indicates that an Auto Negotiation has completed successfully	0x0

5.2.3.8 DEV:PCS1G_CFG_STATUS:PCS1G_ANEG_NP_STATUS

Parent: [DEV:PCS1G_CFG_STATUS](#)

Instances: 1

PCS1G Auto-negotiation next page status register

TABLE 5-123: FIELDS IN PCS1G_ANEG_NP_STATUS

Field Name	Bit	Access	Description	Default
LP_NP_RX	31:16	R/O	Next page ability register from link partner as described in IEEE 802.3, Clause 37	0x0000

5.2.3.9 DEV:PCS1G_CFG_STATUS:PCS1G_LINK_STATUS

Parent: [DEV:PCS1G_CFG_STATUS](#)

Instances: 1

PCS1G link status register

TABLE 5-124: FIELDS IN PCS1G_LINK_STATUS

Field Name	Bit	Access	Description	Default
SIGNAL_DETECT	8	R/O	Indicates whether or not the selected Signal Detect input line is asserted 0: No signal detected 1: Signal detected	0x0

TABLE 5-124: FIELDS IN PCS1G_LINK_STATUS (CONTINUED)

Field Name	Bit	Access	Description	Default
LINK_STATUS	4	R/O	Indicates whether the link is up or down. A link is up when ANEG state machine is in state LINK_OK or AN_DISABLE_LINK_OK 0: Link down 1: Link up	0x0
SYNC_STATUS	0	R/O	Indicates if PCS has successfully synchronized 0: PCS is out of sync 1: PCS has synchronized	0x0

5.2.3.10 DEV:PCS1G_CFG_STATUS:PCS1G_LINK_DOWN_CNT

Parent: [DEV:PCS1G_CFG_STATUS](#)

Instances: 1

PCS1G link down counter register

TABLE 5-125: FIELDS IN PCS1G_LINK_DOWN_CNT

Field Name	Bit	Access	Description	Default
LINK_DOWN_CNT	7:0	R/W	Link Down Counter. A counter that counts the number of times a link has been down. The counter does not saturate at 255 and is only cleared when writing 0 to the register	0x00

5.2.3.11 DEV:PCS1G_CFG_STATUS:PCS1G_STICKY

Parent: [DEV:PCS1G_CFG_STATUS](#)

Instances: 1

PCS1G status register for sticky bits

TABLE 5-126: FIELDS IN PCS1G_STICKY

Field Name	Bit	Access	Description	Default
LINK_DOWN_STICKY	4	Sticky	The sticky bit is set when the link has been down - i.e. if the ANEG state machine has not been in the AN_DISABLE_LINK_OK or LINK_OK state for one or more clock cycles. This occurs if e.g. ANEG is restarted or for example if signal-detect or synchronization has been lost for more than 10 ms (1.6 ms in SGMII mode). By setting the UDLT bit, the required down time can be reduced to 9,77 us (1.56 us) 0: Link is up 1: Link has been down Bit is cleared by writing a 1 to this position.	0x0
OUT_OF_SYNC_STICKY	0	Sticky	Sticky bit indicating if PCS synchronization has been lost 0: Synchronization has not been lost at any time 1: Synchronization has been lost for one or more clock cycles Bit is cleared by writing a 1 to this position.	0x0

5.2.3.12 DEV:PCS1G_CFG_STATUS:PCS1G_DEBUG_STATUS

Parent: [DEV:PCS1G_CFG_STATUS](#)

Instances: 1

PCS1G debug status register

TABLE 5-127: FIELDS IN PCS1G_DEBUG_STATUS

Field Name	Bit	Access	Description	Default
XMIT_MODE	13:12	R/O	Indicates the mode of the TBI 00: Idle mode 01: Configuration mode 10: Reserved 11: Data mode	0x0

5.2.3.13 DEV:PCS1G_CFG_STATUS:PCS1G_LPI_CFG

Parent: [DEV:PCS1G_CFG_STATUS](#)

Instances: 1

Configuration register for Low Power Idle (Energy Efficient Ethernet)

TABLE 5-128: FIELDS IN PCS1G_LPI_CFG

Field Name	Bit	Access	Description	Default
RESERVED	20	R/W	Must be set to its default.	0x1
LPI_RX_WTIM	5:4	R/W	Max wake-up time before link_fail 00: 10 us 01: 13 us 10: 17 us 11: 20 us	0x3
TX_ASSERT_LPIDLE	0	R/W	Assert Low-Power Idle (LPI) in transmit mode 0: Disable LPI transmission 1: Enable LPI transmission	0x0

5.2.3.14 DEV:PCS1G_CFG_STATUS:PCS1G_LPI_WAKE_ERROR_CNT

Parent: [DEV:PCS1G_CFG_STATUS](#)

Instances: 1

PCS1G Low Power Idle wake error counter (Energy Efficient Ethernet)

TABLE 5-129: FIELDS IN PCS1G_LPI_WAKE_ERROR_CNT

Field Name	Bit	Access	Description	Default
WAKE_ERROR_CNT	15:0	R/W	Wake Error Counter. A counter that is incremented when the link partner does not send wake-up burst in due time. The counter saturates at 65535 and is cleared when writing 0 to the register	0x0000

5.2.3.15 DEV:PCS1G_CFG_STATUS:PCS1G_LPI_STATUS

Parent: [DEV:PCS1G_CFG_STATUS](#)

Instances: 1

Status register for Low Power Idle (Energy Efficient Ethernet)

TABLE 5-130: FIELDS IN PCS1G_LPI_STATUS

Field Name	Bit	Access	Description	Default
RX_LPI_FAIL	16	R/O	Receiver has failed to recover from Low-Power Idle mode 0: No failure 1: Failed to recover from LPI mode	0x0
RX_LPI_EVENT_STICKY	12	Sticky	Receiver Low-Power idle occurrence 0: No LPI symbols received 1: Receiver has received LPI symbols Bit is cleared by writing a 1 to this position.	0x0
RX_QUIET	9	R/O	Receiver Low-Power Quiet mode 0: Receiver not in quiet mode 1: Receiver is in quiet mode	0x0
RX_LPI_MODE	8	R/O	Receiver Low-Power Idle mode 0: Receiver not in low power idle mode 1: Receiver is in low power idle mode	0x0
TX_LPI_EVENT_STICKY	4	Sticky	Transmitter Low-Power idle occurrence 0: No LPI symbols transmitted 1: Transmitter has transmitted LPI symbols Bit is cleared by writing a 1 to this position.	0x0
TX_QUIET	1	R/O	Transmitter Low-Power Quiet mode 0: Transmitter not in quiet mode 1: Transmitter is in quiet mode	0x0
TX_LPI_MODE	0	R/O	Transmitter Low-Power Idle mode 0: Transmitter not in low power idle mode 1: Transmitter is in low power idle mode	0x0

5.2.4 DEV:PCS1G_TSTPAT_CFG_STATUS

Parent: [DEV](#)

Instances: 1

PCS1G testpattern configuration and status register set

TABLE 5-131: REGISTERS IN PCS1G_TSTPAT_CFG_STATUS

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
PCS1G_TSTPAT_MODE_CFG	0x00000000	1	PCS1G TSTPAT MODE CFG	Page 381
PCS1G_TSTPAT_STATUS	0x00000004	1	PCS1G TSTPAT STATUS	Page 382

5.2.4.1 DEV:PCS1G_TSTPAT_CFG_STATUS:PCS1G_TSTPAT_MODE_CFG

Parent: [DEV:PCS1G_TSTPAT_CFG_STATUS](#)

Instances: 1

PCS1G testpattern mode configuration register (Frame based pattern 4 and 5 might be not available depending on chip type)

TABLE 5-132: FIELDS IN PCS1G_TSTPAT_MODE_CFG

Field Name	Bit	Access	Description	Default
JTP_SEL	2:0	R/W	Jitter Test Pattern Select: Enables and selects the jitter test pattern to be transmitted. The jitter test patterns are according to the IEEE 802.3, Annex 36A 0: Disable transmission of test patterns 1: High frequency test pattern - repeated transmission of D21.5 code group 2: Low frequency test pattern - repeated transmission of K28.7 code group 3: Mixed frequency test pattern - repeated transmission of K28.5 code group 4: Long continuous random test pattern (packet length is 1524 bytes) 5: Short continuous random test pattern (packet length is 360 bytes)	0x0

5.2.4.2 DEV:PCS1G_TSTPAT_CFG_STATUS:PCS1G_TSTPAT_STATUS

Parent: [DEV:PCS1G_TSTPAT_CFG_STATUS](#)

Instances: 1

PCS1G testpattern status register

TABLE 5-133: FIELDS IN PCS1G_TSTPAT_STATUS

Field Name	Bit	Access	Description	Default
JTP_ERR_CNT	15:8	R/W	Jitter Test Pattern Error Counter. Due to re-sync measures it might happen that single errors are not counted (applies for 2.5gpbs mode). The counter saturates at 255 and is only cleared when writing 0 to the register	0x00
JTP_ERR	4	R/O	Jitter Test Pattern Error 0: Jitter pattern checker has found no error 1: Jitter pattern checker has found an error	0x0
JTP_LOCK	0	R/O	Jitter Test Pattern Lock 0: Jitter pattern checker has not locked 1: Jitter pattern checker has locked	0x0

5.2.5 DEV:PCS_FX100_CONFIGURATION

Parent: [DEV](#)

Instances: 1

Configuration register set for PCS 100Base-FX logic

TABLE 5-134: REGISTERS IN PCS_FX100_CONFIGURATION

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
PCS_FX100_CFG	0x00000000	1	PCS 100Base FX Configuration	Page 383

5.2.5.1 DEV:PCS_FX100_CONFIGURATION:PCS_FX100_CFG

Parent: DEV:PCS_FX100_CONFIGURATION

Instances: 1

Configuration bit groups for 100Base-FX PCS

TABLE 5-135: FIELDS IN PCS_FX100_CFG

Field Name	Bit	Access	Description	Default
SD_SEL	26	R/W	Signal detect selection (select input for internal signal_detect line) 0: Select signal_detect line from hardmacro 1: Select external signal_detect line	0x0
SD_POL	25	R/W	Signal detect polarity: The signal level on signal_detect input pin must be equal to SD_POL to indicate signal detection (SD_ENA must be set). Use '1' when SD_SEL is set to hardmacro. 0: Signal Detect input pin must be '0' to indicate a signal detection 1: Signal Detect input pin must be '1' to indicate a signal detection	0x1
SD_ENA	24	R/W	Signal Detect Enable 0: The Signal Detect input pin is ignored. The PCS assumes an active Signal Detect at all times 1: The Signal Detect input pin is used to determine if a signal is detected	0x1
RESERVED	15:12	R/W	Must be set to its default.	0x4
LINKHYSTTIMER	7:4	R/W	Link hysteresis timer configuration. The hysteresis time lasts [linkhysttimer] * 65536 ns + 2320 ns. If linkhysttime is set to 5, the hysteresis lasts the minimum time of 330 us as specified in IEEE802.3 - 24.3.3.4.	0x5
UNIDIR_MODE_ENA	3	R/W	Unidirectional mode enable. Implementation of 802.3 clause 66. When asserted, this enables MAC to transmit data independent of the state of the receive link. 0: Unidirectional mode disabled 1: Unidirectional mode enabled	0x0
FEFCHK_ENA	2	R/W	Far-End Fault (FEF) detection enable 0: Disable FEF detection 1 Enable FEF detection	0x1
FEFGEN_ENA	1	R/W	Far-End Fault (FEF) generation enable 0: Disable FEF generation 1 Enable FEF generation	0x1
PCS_ENA	0	R/W	PCS enable 0: Disable PCS 1: Enable PCS	0x0

5.2.6 DEV:PCS_FX100_STATUS

Parent: DEV

Instances: 1

Status register set for PCS 100Base-FX logic

TABLE 5-136: REGISTERS IN PCS_FX100_STATUS

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
PCS_FX100_STATUS	0x00000000	1	PCS 100Base FX Status	Page 384

5.2.6.1 DEV:PCS_FX100_STATUS:PCS_FX100_STATUS

Parent: [DEV:PCS_FX100_STATUS](#)

Instances: 1

Status bit groups for 100Base-FX PCS. Note: If sigdet_cfg != "00" is selected status signal "signal_detect" shows the internal signal_detect value is gated with the status of rx toggle-rate control circuitry.

TABLE 5-137: FIELDS IN PCS_FX100_STATUS

Field Name	Bit	Access	Description	Default
PCS_ERROR_STICKY	7	Sticky	PCS error has occurred 1: RX_ER was high while RX_DV active 0: No RX_ER indication found while RX_DV active Bit is cleared by writing a 1 to this position.	0x0
FEF_FOUND_STICKY	6	Sticky	Far-end Fault state has occurred 1: A Far-End Fault has been detected 0: No Far-End Fault occurred Bit is cleared by writing a 1 to this position.	0x0
SSD_ERROR_STICKY	5	Sticky	Stream Start Delimiter error occurred 1: A Start-of-Stream Delimiter error has been detected 0: No SSD error occurred Bit is cleared by writing a 1 to this position.	0x0
SYNC_LOST_STICKY	4	Sticky	Synchronization lost 1: Synchronization lost 0: No sync lost occurred Bit is cleared by writing a 1 to this position.	0x0
FEF_STATUS	2	R/O	Current status of Far-end Fault detection state 1: Link currently in fault state 0: Link is in normal state	0x0
SIGNAL_DETECT	1	R/O	Current status of selected signal_detect input line 1: Proper signal detected 0: No proper signal found	0x0
SYNC_STATUS	0	R/O	Status of synchronization 1: Link established 0: No link found	0x0

5.3 DEVCPU_GCB

TABLE 5-138: REGISTER GROUPS IN DEVCPU_GCB

Register Group Name	Offset within Target	Instances and Address Spacing	Description	Details
CHIP_REGS	0x00000000	1	Software/hardware interaction	Page 385
VCORE_ACCESS	0x00000020	1	VCore SBA access	Page 386
GPIO	0x00000034	1	GPIO configuration/status	Page 390
MIIM	0x0000005C	2 0x00000024	MIIM master controller	Page 393
MIIM_READ_SCAN	0x000000A4	1	MIIM master controller scan results	Page 396
TEMP_SENSOR	0x000000AC	1	Temperature sensor control	Page 397
SIO_CTRL	0x000000B4	1	Serial IO control configuration	Page 398
FAN_CFG	0x00000174	1	Fan controller configuration	Page 402
FAN_STAT	0x00000178	1	Fan controller status	Page 403
PTP_CFG	0x0000017C	1	PTP controller configuration	Page 404
PTP_STAT	0x000001B4	1	PTP controller status	Page 408
MEMITGR	0x000001C4	1	Memory integrity monitor	Page 409
VRAP	0x000001DC	1	VRAP controller	Page 413

5.3.1 DEVCPU_GCB:CHIP_REGS

Parent: [DEVCPU_GCB](#)

Instances: 1

TABLE 5-139: REGISTERS IN CHIP_REGS

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
CHIP_ID	0x00000000	1	Chip ID register	Page 385
GPR	0x00000004	1	General purpose register	Page 386
SOFT_RST	0x00000008	1	Reset control register	Page 386
HW_STAT	0x00000010	1	Various status indications	Page 386

5.3.1.1 DEVCPU_GCB:CHIP_REGS:CHIP_ID

Parent: [DEVCPU_GCB:CHIP_REGS](#)

Instances: 1

JTAG encoded Chip ID register.

TABLE 5-140: FIELDS IN CHIP_ID

Field Name	Bit	Access	Description	Default
REV_ID	31:28	R/O	Revision ID.	0x1
PART_ID	27:12	R/O	Part ID.	0x7416
MFG_ID	11:1	R/O	Manufacturer ID.	0x074

TABLE 5-140: FIELDS IN CHIP_ID (CONTINUED)

Field Name	Bit	Access	Description	Default
ONE	0	R/O	Always 1.	0x1

5.3.1.2 DEVCPU_GCB:CHIP_REGS:GPR

Parent: [DEVCPU_GCB:CHIP_REGS](#)

Instances: 1

TABLE 5-141: FIELDS IN GPR

Field Name	Bit	Access	Description	Default
GPR	31:0	R/W	General purpose register for software development.	0x00000000

5.3.1.3 DEVCPU_GCB:CHIP_REGS:SOFT_RST

Parent: [DEVCPU_GCB:CHIP_REGS](#)

Instances: 1

TABLE 5-142: FIELDS IN SOFT_RST

Field Name	Bit	Access	Description	Default
SOFT_CHIP_RST	0	R/W	Set this field to reset the whole chip. This field is automatically cleared by the reset. Note: It is possible for the VCore System to protect itself from this soft-reset, for more info see ICPU_CFG::RESET.CORE_RST_PROTECT inside the VCore register space.	0x0

5.3.1.4 DEVCPU_GCB:CHIP_REGS:HW_STAT

Parent: [DEVCPU_GCB:CHIP_REGS](#)

Instances: 1

TABLE 5-143: FIELDS IN HW_STAT

Field Name	Bit	Access	Description	Default
MEM_FAIL	0	R/O	This field is set if a hardware fail has been detected in any of the memories during startup-initialization of the chip. This field is valid after release of reset.	0x0

5.3.2 DEVCPU_GCB:VCORE_ACCESS

Parent: [DEVCPU_GCB](#)

Instances: 1

TABLE 5-144: REGISTERS IN VCore_ACCESS

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
VA_CTRL	0x00000000	1	Control register for VCore accesses	Page 387
VA_ADDR	0x00000004	1	Address register for VCore accesses	Page 387
VA_DATA	0x00000008	1	Data register for VCore accesses	Page 388
VA_DATA_INCR	0x0000000C	1	Data register for VCore accesses (w. auto increment of address)	Page 389
VA_DATA_INERT	0x00000010	1	Data register for VCore accesses (will not initiate access)	Page 390

5.3.2.1 DEVCPU_GCB:VCore_ACCESS:VA_CTRL

Parent: [DEVCPU_GCB:VCore_ACCESS](#)

Instances: 1

TABLE 5-145: FIELDS IN VA_CTRL

Field Name	Bit	Access	Description	Default
VA_ERR	3:2	R/O	If the VCore access logic detects an error this field is set based on the nature of the error. This is a read-only field which is cleared by the VCore access logic when a new access is (successfully) accepted. 0: No errors detected. 1: SBA not ready when accessed. 2: SBA reported error. 3: New data or address written during active access.	0x0
VA_BUSY_RD	1	R/O	This field is set to the value of DEVCPU_GCB::VA_CTRL.VA_BUSY whenever one of the data registers DEVCPU_GCB::VA_DATA, DEVCPU_GCB::VA_DATA_INCR, or DEVCPU_GCB::VA_DATA_RO is read. By examining this field it is possible to determine if DEVCPU_GCB::VA_CTRL.VA_BUSY was set at the time a read from one of these registers was performed.	0x0
VA_BUSY	0	R/O	This field is set by hardware when an access into VCore domain is started, and cleared when the access is done.	0x0

5.3.2.2 DEVCPU_GCB:VCore_ACCESS:VA_ADDR

Parent: [DEVCPU_GCB:VCore_ACCESS](#)

Instances: 1

TABLE 5-146: FIELDS IN VA_ADDR

Field Name	Bit	Access	Description	Default
VA_ADDR	31:0	R/W	The address to access in the VCore domain, all addresses must be 32-bit aligned (i.e. the two least significant bit must always be 0). When accesses are initiated using the DEVCPU_GCB::VA_DATA_INCR register, then this field is auto-incremented by 4 at the end of the transfer. The memory region of the VCore that maps to switch-core registers cannot be accessed by using these registers.	0x00000000

5.3.2.3 DEVCPU_GCB:VCORE_ACCESS:VA_DATA

Parent: [DEVCPU_GCB:VCORE_ACCESS](#)

Instances: 1

The VA_DATA, VA_DATA_INCR, and VA_DATA_INERT registers are used for indirect access into the VCore domain. The functionality of the VA_DATA_INCR and VA_DATA_INERT registers are similar to this register - but with minor exceptions. These exceptions are fleshed out in the description of the respective registers.

TABLE 5-147: FIELDS IN VA_DATA

Field Name	Bit	Access	Description	Default
VA_DATA	31:0	R/W	Reading or writing from/to this field initiates accesses into the VCore domain. While an access is ongoing (DEVCPU_GCB::VA_CTRL.VA_BUSY is set) this field may not be written. It is possible to read this field while an access is ongoing, but the data returned will be 0x88888888. When writing to this field; a write into the VCore domain is initiated to the address specified in the DEVCPU_GCB::VA_ADDR register, with the data that was written to this field. Only 32-bit writes are supported. This field may not be written to until the DEVCPU_GCB::VA_CTRL.VA_BUSY indicates that no accesses is ongoing. When reading from this field; a read from the VCore domain is initiated from the address specified in the DEVCPU_GCB::VA_ADDR register. Important: The data that is returned from reading this field (and stating an access) is not the result of the newly initiated read, instead the data from the last access is returned. The result of the newly initiated read access will be ready once the DEVCPU_GCB::VA_CTRL.VA_BUSY field shows that the access is done. Note: When the result of a read-access is read from this field (the second read), a new access will automatically be initiated. This is desirable when reading a series of addresses from VCore domain. If a new access is not desirable, then the result should be read from the DEVCPU_GCB::VA_DATA_INERT register instead of this field.	0x00000000

5.3.2.4 DEVCPU_GCB:VCORE_ACCESS:VA_DATA_INCR

Parent: [DEVCPU_GCB:VCORE_ACCESS](#)

Instances: 1

TABLE 5-148: FIELDS IN VA_DATA_INCR

Field Name	Bit	Access	Description	Default
VA_DATA_INCR	31:0	R/W	This field behaves in the same way as DEVCPU_GCB::VA_DATA.VA_DATA. Except when an access is initiated by using this field (either read or write); the address register (DEVCPU_GCB::VA_ADDR) is automatically incremented by 4 at the end of the access, i.e. when DEVCPU_GCB::VA_CTRL.VA_BUSY is deasserted.	0x00000000

5.3.2.5 DEVCPU_GCB:VCORE_ACCESS:VA_DATA_INERT

Parent: [DEVCPU_GCB:VCORE_ACCESS](#)

Instances: 1

TABLE 5-149: FIELDS IN VA_DATA_INERT

Field Name	Bit	Access	Description	Default
VA_DATA_INERT	31:0	R/W	This field behaves in the same way as DEVCPU_GCB::VA_DATA.VA_DATA. Except accesses (read or write) does not initiate VCore accesses. Writing to this register just overwrites the value currently held by all of the data registers (DEVCPU_GCB::VA_DATA, DEVCPU_GCB::VA_DATA_INCR, and DEVCPU_GCB::VA_DATA_INERT).	0x00000000

5.3.3 DEVCPU_GCB:GPIO

Parent: [DEVCPU_GCB](#)

Instances: 1

Each register in this group contains one field with one bit per GPIO pin. Bit 0 in each field corresponds to GPIO0, bit 1 to GPIO1, and so on.

TABLE 5-150: REGISTERS IN GPIO

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
GPIO_OUT_SET	0x00000000	1	Atomic set of GPIO output value	Page 390
GPIO_OUT_CLR	0x00000004	1	Atomic clear of GPIO output value	Page 391
GPIO_OUT	0x00000008	1	GPIO output value	Page 391
GPIO_IN	0x0000000C	1	Current value at GPIO pins	Page 391
GPIO_OE	0x00000010	1	Direction of GPIOs (output enable)	Page 391
GPIO_INTR	0x00000014	1	Interrupt (GPIO input value change event)	Page 392
GPIO_INTR_ENA	0x00000018	1	GPIO interrupt enable	Page 392
GPIO_INTR_IDENT	0x0000001C	1	Currently interrupting GPIOs	Page 392
GPIO_ALT	0x00000020	2 0x00000004	GPIO alternate functions	Page 392

5.3.3.1 DEVCPU_GCB:GPIO:GPIO_OUT_SET

Parent: [DEVCPU_GCB:GPIO](#)

Instances: 1

TABLE 5-151: FIELDS IN GPIO_OUT_SET

Field Name	Bit	Access	Description	Default
G_OUT_SET	31:0	One-shot	Set bits in this register to atomically set corresponding bits in DEVCPU_GCB::GPIO_OUT. This register return 0 on read.	0x00000000

5.3.3.2 DEVCPU_GCB:GPIO:GPIO_OUT_CLR

Parent: [DEVCPU_GCB:GPIO](#)

Instances: 1

TABLE 5-152: FIELDS IN GPIO_OUT_CLR

Field Name	Bit	Access	Description	Default
G_OUT_CLR	31:0	One-shot	Set bits in this register to atomically clear corresponding bits in DEVCPU_GCB::GPIO_OUT. This register return 0 on read.	0x00000000

5.3.3.3 DEVCPU_GCB:GPIO:GPIO_OUT

Parent: [DEVCPU_GCB:GPIO](#)

Instances: 1

TABLE 5-153: FIELDS IN GPIO_OUT

Field Name	Bit	Access	Description	Default
G_OUT	31:0	R/W	Controls the value on the GPIO pins enabled for output (via the GPIO_OE register). Atomic (safe) modifications to the contents of this register can be performed by using the DEVCPU_GCB::GPIO_OUT_SET and DEVCPU_GCB::GPIO_OUT_CLR registers.	0x00000000

5.3.3.4 DEVCPU_GCB:GPIO:GPIO_IN

Parent: [DEVCPU_GCB:GPIO](#)

Instances: 1

TABLE 5-154: FIELDS IN GPIO_IN

Field Name	Bit	Access	Description	Default
G_IN	31:0	R/O	GPIO input register. Reflects the current state of the corresponding GPIO pins.	0x00000000

5.3.3.5 DEVCPU_GCB:GPIO:GPIO_OE

Parent: [DEVCPU_GCB:GPIO](#)

Instances: 1

TABLE 5-155: FIELDS IN GPIO_OE

Field Name	Bit	Access	Description	Default
G_OE	31:0	R/W	Configures the direction of the GPIO pins. 0: Input 1: Output	0x00000000

5.3.3.6 DEVCPU_GCB:GPIO:GPIO_INTR

Parent: [DEVCPU_GCB:GPIO](#)

Instances: 1

TABLE 5-156: FIELDS IN GPIO_INTR

Field Name	Bit	Access	Description	Default
G_INTR	31:0	Sticky	Indicates whether a GPIO input has changed since last clear. 0: No change 1: GPIO has changed	0x00000000

5.3.3.7 DEVCPU_GCB:GPIO:GPIO_INTR_ENA

Parent: [DEVCPU_GCB:GPIO](#)

Instances: 1

TABLE 5-157: FIELDS IN GPIO_INTR_ENA

Field Name	Bit	Access	Description	Default
G_INTR_ENA	31:0	R/W	Enables individual GPIO pins for interrupt.	0x00000000

5.3.3.8 DEVCPU_GCB:GPIO:GPIO_INTR_IDENT

Parent: [DEVCPU_GCB:GPIO](#)

Instances: 1

TABLE 5-158: FIELDS IN GPIO_INTR_IDENT

Field Name	Bit	Access	Description	Default
G_INTR_IDENT	31:0	R/O	Shows which GPIO sources that are currently interrupting. This field is the result of an AND-operation between the DEVCPU_GCB::GPIO_INTR and the DEVCPU_GCB::GPIO_INTR_ENA registers.	0x00000000

5.3.3.9 DEVCPU_GCB:GPIO:GPIO_ALT

Parent: [DEVCPU_GCB:GPIO](#)

Instances: 2

TABLE 5-159: FIELDS IN GPIO_ALT

Field Name	Bit	Access	Description	Default
G_ALT	31:0	R/W	Configures alternate functions for individual GPIOs. See datasheet for information on possible alternate functions. The LSB of the GPIO encoding is placed in replication 0, MSB is placed in replication 1. For example; to encode Alternate mode 2 for GPIO[n] write DEVCPU_GCB::GPIO_ALT[0][n] = 0 and DEVCPU_GCB::GPIO_ALT[1][n] = 1. Note: This register is only reset by the device's reset input, i.e. it is not affected by soft reset! 0: GPIO mode 1: Alternate mode 1 2: Alternate mode 2 3: Alternate mode 3	0x00000000

5.3.4 DEVCPU_GCB:MIIM

Parent: [DEVCPU_GCB](#)

Instances: 2

TABLE 5-160: REGISTERS IN MIIM

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
MII_STATUS	0x00000000	1	MIIM status	Page 393
MII_CMD	0x00000008	1	MIIM command	Page 394
MII_DATA	0x0000000C	1	MIIM reply data	Page 395
MII_CFG	0x00000010	1	MIIM configuration	Page 395
MII_SCAN_0	0x00000014	1	MIIM scan 0	Page 395
MII_SCAN_1	0x00000018	1	MIIM scan 1	Page 396
MII_SCAN_LAST_RSLTS	0x0000001C	1	MIIM results	Page 396
MII_SCAN_LAST_RSLTS_VLD	0x00000020	1	MIIM results valid	Page 396

5.3.4.1 DEVCPU_GCB:MIIM:MII_STATUS

Parent: [DEVCPU_GCB:MIIM](#)

Instances: 1

TABLE 5-161: FIELDS IN MII_STATUS

Field Name	Bit	Access	Description	Default
MIIM_STAT_BUSY	3	R/O	Indicates the current state of the MIIM controller. When read operations are done (no longer busy), then read data is available via the DEVCPU_GCB::MII_DATA register. 0: MIIM controller is in idle state 1: MIIM controller is busy performing MIIM cmd (either read or write cmd)	0x0

TABLE 5-161: FIELDS IN MII_STATUS (CONTINUED)

Field Name	Bit	Access	Description	Default
MIIM_STAT_OPR_PEND	2	R/O	The MIIM controller has a CMD fifo of depth one. When this field is 0, then it is safe to write another MIIM command to the MIIM controller. 0: Read or write not pending 1: Read or write pending	0x0
MIIM_STAT_PENDING_RD	1	R/O	Indicates whether a read operation via the MIIM interface is in progress or not. 0: Read not in progress 1: Read in progress	0x0
MIIM_STAT_PENDING_WR	0	R/O	Indicates whether a write operation via the MIIM interface is in progress or not. 0: Write not in progress 1: Write in progress	0x0
MIIM_SCAN_COMPLETE	4	R/O	Signals if all PHYs have been scanned (with auto scan) at least once. 0: Auto scan has not scanned all PHYs 1: PHYs has been scanned at least once	0x0

5.3.4.2 DEVCPU_GCB:MIIM:MII_CMD

Parent: [DEVCPU_GCB:MIIM](#)

Instances: 1

TABLE 5-162: FIELDS IN MII_CMD

Field Name	Bit	Access	Description	Default
MIIM_CMD_VLD	31	One-shot	Must be set for starting a new PHY access. This bit is automatically cleared. 0: Write to this register is ignored 1: Write to this register is processed	0x0
MIIM_CMD_PHYAD	29:25	R/W	Indicates the addressed PHY number.	0x00
MIIM_CMD_REGAD	24:20	R/W	Indicates the addressed of the register within the PHY that shall be accessed.	0x00
MIIM_CMD_WRDATA	19:4	R/W	Data to be written in the PHY register.	0x0000
MIIM_CMD_SINGLE_SCAN	3	R/W	Select if scanning of the PHY shall be done once, or scanning should be done continuously. 0: Do continuously PHY scanning 1: Stop once all PHY have been scanned	0x0
MIIM_CMD_OPR_FIELD	2:1	R/W	Indicates type of operation. Clause 22: 1: Write 2: Read Clause 45: 0: Address 1: Write 2: Read inc 3: Read	0x0

TABLE 5-162: FIELDS IN MII_CMD (CONTINUED)

Field Name	Bit	Access	Description	Default
MIIM_CMD_SCAN	0	R/W	Indicates whether automatic scanning of PHY registers is enabled. When enabled, the PHY-number for each automatic read is continuously round-robined from DEVCPU_GCB::PHY_ADDR_LOW through DEVCPU_GCB::PHY_ADDR_HIGH. This function is started upon a read operation. Scan shall be disabled before reconfiguring the MIIM controller. 0: Disabled 1: Enabled	0x0

5.3.4.3 DEVCPU_GCB:MIIM:MII_DATA

Parent: [DEVCPU_GCB:MIIM](#)

Instances: 1

TABLE 5-163: FIELDS IN MII_DATA

Field Name	Bit	Access	Description	Default
MIIM_DATA_SUCCESS	17:16	R/O	Indicates whether a read operation failed or succeeded. 0: OK 3: Error	0x0
MIIM_DATA_RDDATA	15:0	R/O	Data read from PHY register.	0x0000

5.3.4.4 DEVCPU_GCB:MIIM:MII_CFG

Parent: [DEVCPU_GCB:MIIM](#)

Instances: 1

TABLE 5-164: FIELDS IN MII_CFG

Field Name	Bit	Access	Description	Default
MIIM_CFG_PRESCALE	7:0	R/W	Configures the MIIM clock frequency. This is computed as $\text{system\_clk}/(2^{*(1+X)})$, where X is the value written to this register. Note : Setting X to 0 is invalid and will result in the same frequency as setting X to 1.	0x32
MIIM_ST_CFG_FIELD	10:9	R/W	The ST (start-of-frame) field of the MIIM frame format adopts the value of this field. This must be configured for either clause 22 or 45 MIIM operation. 0: Clause 45 1: Clause 22 Other values are reserved.	0x1

5.3.4.5 DEVCPU_GCB:MIIM:MII_SCAN_0

Parent: [DEVCPU_GCB:MIIM](#)

Instances: 1

TABLE 5-165: FIELDS IN MII_SCAN_0

Field Name	Bit	Access	Description	Default
MIIM_SCAN_PHYADHI	9:5	R/W	Indicates the high PHY number to scan during automatic scanning.	0x00
MIIM_SCAN_PHYADLO	4:0	R/W	Indicates the low PHY number to scan during automatic scanning.	0x00

5.3.4.6 DEVCPU_GCB:MIIM:MII_SCAN_1

Parent: [DEVCPU_GCB:MIIM](#)

Instances: 1

TABLE 5-166: FIELDS IN MII_SCAN_1

Field Name	Bit	Access	Description	Default
MIIM_SCAN_MASK	31:16	R/W	Indicates the mask for comparing the PHY registers during automatic scan.	0x0000
MIIM_SCAN_EXPECT	15:0	R/W	Indicates the expected value for comparing the PHY registers during automatic scan.	0x0000

5.3.4.7 DEVCPU_GCB:MIIM:MII_SCAN_LAST_RSLTS

Parent: [DEVCPU_GCB:MIIM](#)

Instances: 1

TABLE 5-167: FIELDS IN MII_SCAN_LAST_RSLTS

Field Name	Bit	Access	Description	Default
MIIM_LAST_RSLT	31:0	R/O	Indicates for each PHY if a PHY register has matched the expected value (with mask). This register reflects the value of the last reading of the phy register. 0: Mismatch 1: Match	0x00000000

5.3.4.8 DEVCPU_GCB:MIIM:MII_SCAN_LAST_RSLTS_VLD

Parent: [DEVCPU_GCB:MIIM](#)

Instances: 1

TABLE 5-168: FIELDS IN MII_SCAN_LAST_RSLTS_VLD

Field Name	Bit	Access	Description	Default
MIIM_LAST_RSLT_VLD	31:0	R/O	Indicates for each PHY if a PHY register matched are valid or not. 0: Scan result not valid 1: Scan result valid	0x00000000

5.3.5 DEVCPU_GCB:MIIM_READ_SCAN

Parent: [DEVCPU_GCB](#)

Instances: 1

TABLE 5-169: REGISTERS IN MIIM_READ_SCAN

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
MII_SCAN_RSLTS_STICKY	0x00000000	2 0x00000004	MIIM results sticky	Page 397

5.3.5.1 DEVCPU_GCB:MIIM_READ_SCAN:MII_SCAN_RSLTS_STICKY

Parent: [DEVCPU_GCB:MIIM_READ_SCAN](#)

Instances: 2

TABLE 5-170: FIELDS IN MII_SCAN_RSLTS_STICKY

Field Name	Bit	Access	Description	Default
MIIM_SCAN_RSLTS_STICKY	31:0	R/O	Indicates for each PHY if a PHY register has had a mismatch of the expected value (with mask) since last reading of MIIM_SCAN_RSLTS_STICKY. Result is sticky, and result will indicate if there has been a mismatch since the last reading of this register. Upon reading this register it is reset to 0xFFFFFFFF. 0 : Mismatch 1 : Match.	0x00000000

5.3.6 DEVCPU_GCB:TEMP_SENSOR

Parent: [DEVCPU_GCB](#)

Instances: 1

TABLE 5-171: REGISTERS IN TEMP_SENSOR

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
TEMP_SENSOR_CTRL	0x00000000	1	Temperature sensor control	Page 397
TEMP_SENSOR_DATA	0x00000004	1	Temperature sensor data	Page 398

5.3.6.1 DEVCPU_GCB:TEMP_SENSOR:TEMP_SENSOR_CTRL

Parent: [DEVCPU_GCB:TEMP_SENSOR](#)

Instances: 1

TABLE 5-172: FIELDS IN TEMP_SENSOR_CTRL

Field Name	Bit	Access	Description	Default
TEMP_SENSOR_POWER_DOWN	1	R/W	Set to power down the temperature sensor. 0: Normal operation 1: Power down mode	0x1

TABLE 5-172: FIELDS IN TEMP_SENSOR_CTRL (CONTINUED)

Field Name	Bit	Access	Description	Default
TEMP_SENSOR_RESET_N	0	R/W	Set to release temperature sensor from reset. 0: Reset state 1: Normal operation	0x0

5.3.6.2 DEVCPU_GCB:TEMP_SENSOR:TEMP_SENSOR_DATA

Parent: [DEVCPU_GCB:TEMP_SENSOR](#)

Instances: 1

TABLE 5-173: FIELDS IN TEMP_SENSOR_DATA

Field Name	Bit	Access	Description	Default
TEMP_SENSOR_DATA_VALID	8	R/O	Set when data from the temperature sensor is valid. It will take some time after enabling the temperature sensor until it returns valid data.	0x0
TEMP_SENSOR_DATA	7:0	R/O	Temperature data readout, this field is valid when DEVCPU_GCB::TEMP_SENSOR_DATA.TEMP_SENSOR_DATA_VALID is set. Once valid data is returned then it will continually be updated until the temperature sensor is put into power down or reset (see DEVCPU_GCB::TEMP_SENSOR_CTRL for more information). This value is read directly from another clock domain, keep reading field value until same value is returned on two successive reads. $TEMP(degC) = 135.3degC - 0.65degC * TEMP\_SENSOR\_DATA$	0x00

5.3.7 DEVCPU_GCB:SIO_CTRL

Parent: [DEVCPU_GCB](#)

Instances: 1

TABLE 5-174: REGISTERS IN SIO_CTRL

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
SIO_INPUT_DATA	0x00000000	4 0x00000004	Last value at SGPIO pins	Page 399
SIO_INT_POL	0x00000010	4 0x00000004	SGPIO interrupt polarity	Page 399
SIO_PORT_INT_ENA	0x00000020	1	SGPIO interrupt enable per port	Page 399
SIO_PORT_CONFIG	0x00000024	32 0x00000004	Output configuration	Page 399
SIO_PORT_ENABLE	0x000000A4	1	Port enable	Page 400
SIO_CONFIG	0x000000A8	1	General configurations	Page 400
SIO_CLOCK	0x000000AC	1	SGPIO shift clock frequency	Page 402

TABLE 5-174: REGISTERS IN SIO_CTRL (CONTINUED)

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
SIO_INT_REG	0x000000B0	4 0x00000004	Currently interrupting SGPIOs	Page 402

5.3.7.1 DEVCPU_GCB:SIO_CTRL:SIO_INPUT_DATA

Parent: [DEVCPU_GCB:SIO_CTRL](#)

Instances: 4

TABLE 5-175: FIELDS IN SIO_INPUT_DATA

Field Name	Bit	Access	Description	Default
S_IN	31:0	R/O	Serial input data. Each replication, N, holds bit N from all ports - bit N from port M is mapped to replication N bit M. Values of disabled gpios are undefined.	0x00000000

5.3.7.2 DEVCPU_GCB:SIO_CTRL:SIO_INT_POL

Parent: [DEVCPU_GCB:SIO_CTRL](#)

Instances: 4

TABLE 5-176: FIELDS IN SIO_INT_POL

Field Name	Bit	Access	Description	Default
INT_POL	31:0	R/W	Interrupt polarity. Bit n from all ports. This register defines at which logic value an interrupt is generated. For bit0 this register is also used to define the polarity of the "loss of signal" output. 0: interrupt and "loss of signal" is active high 1: interrupt and "loss of signal" is active low	0x00000000

5.3.7.3 DEVCPU_GCB:SIO_CTRL:SIO_PORT_INT_ENA

Parent: [DEVCPU_GCB:SIO_CTRL](#)

Instances: 1

TABLE 5-177: FIELDS IN SIO_PORT_INT_ENA

Field Name	Bit	Access	Description	Default
INT_ENA	31:0	R/W	Interrupt enable vector with one enable bit for each port. Port order: (portN down to port0) 0: Interrupt is disabled for the port 1: Interrupt is enabled for the port	0x00000000

5.3.7.4 DEVCPU_GCB:SIO_CTRL:SIO_PORT_CONFIG

Parent: [DEVCPU_GCB:SIO_CTRL](#)

Instances: 32

TABLE 5-178: FIELDS IN SIO_PORT_CONFIG

Field Name	Bit	Access	Description	Default
BIT_SOURCE	11:0	R/W	Output source select for the four outputs from each port. The source select is encoded using three bits for each output bit: Output bit0 is defined by (2 down to 0), output bit1 is defined by (5 down to 3), output bit2 is defined by (8 down to 6), and output bit3 is defined by (11 down to 9). 0: Forced 0 1: Forced 1 2: Blink mode 0 3: Blink mode 1 4: Link activity blink mode 0 5: Link activity blink mode 1 6: Link activity blink mode 0 inversed polarity 7: Link activity blink mode 1 inversed polarity	0x000

5.3.7.5 DEVCPU_GCB:SIO_CTRL:SIO_PORT_ENABLE

Parent: [DEVCPU_GCB:SIO_CTRL](#)

Instances: 1

TABLE 5-179: FIELDS IN SIO_PORT_ENABLE

Field Name	Bit	Access	Description	Default
P_ENA	31:0	R/W	Port enable vector with one enable bit for each port. Port order: (portN down to port0) 0: Port is disabled 1: Port is enabled	0x00000000

5.3.7.6 DEVCPU_GCB:SIO_CTRL:SIO_CONFIG

Parent: [DEVCPU_GCB:SIO_CTRL](#)

Instances: 1

TABLE 5-180: FIELDS IN SIO_CONFIG

Field Name	Bit	Access	Description	Default
SIO_BMODE_1	21:20	R/W	Configuration for blink mode 1. Supports three different blink modes and a "burst toggle" mode in which blink mode 1 will alternate for each burst. 0: Blink freq appr. 20Hz 1: Blink freq appr. 10Hz 2: Blink freq appr. 5Hz 3: Burst toggle	0x0
SIO_BMODE_0	19:18	R/W	Configuration of blink mode 0. Supports four different blink modes. 0: Blink freq appr. 20Hz 1: Blink freq appr. 10Hz 2: Blink freq appr. 5Hz 3: Blink freq appr. 2.5Hz	0x0

TABLE 5-180: FIELDS IN SIO_CONFIG (CONTINUED)

Field Name	Bit	Access	Description	Default
SIO_BLINK_RESET	17	R/W	Reset the blink counters. Used to synchronize the blink modes between different chips. 0: Blink counter is running 1: Blink counter is reset until sio_blink_reset is unset again	0x0
SIO_INT_ENA	16:13	R/W	Bit interrupt enable. Enables interrupts for the four gpios in a port. Is applied to all ports. 0: Interrupt is disabled for bit n for all ports 1: Interrupt is enabled for bit n for all ports	0x0
SIO_BURST_GAP_DIS	12	R/W	Set to disable burst gap.	0x0
SIO_BURST_GAP	11:7	R/W	Configures the length of burst gap in steps of approx. 1 ms. Burst gap can be disabled by setting DEVCPU_GCB::SIO_CONFIG.SIO_BURST_GAP_DIS. 0: 1.05 ms burst gap 1: 2.10 ms burst gap ... 31: 33.55 ms burst gap	0x00
SIO_SINGLE_SHOT	6	One-shot	Use this to output a single burst. Will be cleared by hardware when the burst has finished.	0x0
SIO_AUTO_REPEAT	5	R/W	Use this to output repeated bursts interleaved with burst gaps. Must be manually reset again to stop output of bursts.	0x0
SIO_LD_POLARITY	4	R/W	Polarity of the "Ld" signal 0: load signal is active low 1: load signal is active high	0x0
SIO_PORT_WIDTH	3:2	R/W	Number of SGPIOs pr. port. 0: 1 gpio pr. port 1: 2 gpios pr. port 2: 3 gpios pr. port 3: 4 gpios pr. port	0x0
SIO_REVERSE_OUTPUT	1	R/W	Reverse the output bitstream. The default order of the output bit stream is (displayed in transmitted order): (portN bit3, portN bit2, ..., port0 bit1, port0 bit0) The reverse order of the output bit stream is (displayed in transmitted order): (port0 bit0, port0 bit1, ..., portN bit2, portN bit3) 0: Do not reverse 1: Reverse	0x0
SIO_REVERSE_INPUT	0	R/W	Reverse the input bitstream. The default order of the input bit stream is (displayed in received order): (port0 bit0, port0 bit1, ..., portN bit2, portN bit3) The reverse order of the input bit stream is (displayed in received order): (portN bit3, portN bit2, ..., port0 bit1, port0 bit0) 0: Do not reverse 1: Reverse	0x0

VSC7416-01

5.3.7.7 DEVCPU_GCB:SIO_CTRL:SIO_CLOCK

Parent: [DEVCPU_GCB:SIO_CTRL](#)

Instances: 1

TABLE 5-181: FIELDS IN SIO_CLOCK

Field Name	Bit	Access	Description	Default
SIO_CLK_FREQ	11:0	R/W	SIO controller clock frequency. Divides the 156.25 MHz system clk with value of this field. E.g. the system clk is 156.25 MHz and this field is set to 10, the output frequency will be 15.625 MHz. 0: Disable clock 1: Reserved, do not use Others : Clock divider value.	0x000

5.3.7.8 DEVCPU_GCB:SIO_CTRL:SIO_INT_REG

Parent: [DEVCPU_GCB:SIO_CTRL](#)

Instances: 4

TABLE 5-182: FIELDS IN SIO_INT_REG

Field Name	Bit	Access	Description	Default
INT_REG	31:0	Sticky	Interrupt register. Bit n from all ports. Disabled gpios are always '0'. Note: The SGPIO interrupt is level sensitive, these sticky bits are set when the corresponding SGPIO value is shifted into the device. Bit order (portM bit-n down to portM bit-0). 0: No interrupt for given gpio 1: Interrupt for given gpio	0x00000000

5.3.8 DEVCPU_GCB:FAN_CFG

Parent: [DEVCPU_GCB](#)

Instances: 1

TABLE 5-183: REGISTERS IN FAN_CFG

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
FAN_CFG	0x00000000	1	Fan controller configuration	Page 402

5.3.8.1 DEVCPU_GCB:FAN_CFG:FAN_CFG

Parent: [DEVCPU_GCB:FAN_CFG](#)

Instances: 1

TABLE 5-184: FIELDS IN FAN_CFG

Field Name	Bit	Access	Description	Default
DUTY_CYCLE	23:16	R/W	Define the duty cycle 0x00: Always "off" 0xFF: Always "on"	0x00
PWM_FREQ	6:4	R/W	Set the frequency of the PWM output 0: 25 kHz 1: 120 Hz 2: 100 Hz 3: 80 Hz 4: 60 Hz 5: 40 Hz 6: 20 Hz 7: 10 Hz	0x0
INV_POL	3	R/W	Define the polarity of the PWM output. 0: PWM is logic 1 when "on" 1: PWM is logic 0 when "on"	0x0
GATE_ENA	2	R/W	Enable gating of the TACH input by the PWM output so that only TACH pulses received when PWM is "on" are counted. 0: Disabled 1: Enabled	0x0
PWM_OPEN_COL_ENA	1	R/W	Configure the PWM output to be open collector	0x0
FAN_STAT_CFG	0	R/W	Configure behavior of TACH input tick counter, see DEVCPU_GCB::FAN_CNT for more information.	0x0

5.3.9 DEVCPU_GCB:FAN_STAT

Parent: [DEVCPU_GCB](#)

Instances: 1

TABLE 5-185: REGISTERS IN FAN_STAT

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
FAN_CNT	0x00000000	1	TACH counter	Page 403

5.3.9.1 DEVCPU_GCB:FAN_STAT:FAN_CNT

Parent: [DEVCPU_GCB:FAN_STAT](#)

Instances: 1

TABLE 5-186: FIELDS IN FAN_CNT

Field Name	Bit	Access	Description	Default
FAN_CNT	15:0	R/O	Counts the number of TACH input ticks. If DEVCPU_GCB::FAN_CFG.FAN_STAT_CFG is set then this is a wrapping counter that shows the total number of registered TACH ticks. If DEVCPU_GCB::FAN_CFG.FAN_STAT_CFG is cleared then this counter is updated once every second with the number of TACH ticks registered during the last second.	0x0000

5.3.10 DEVCPU_GCB:PTP_CFG

Parent: [DEVCPU_GCB](#)

Instances: 1

TABLE 5-187: REGISTERS IN PTP_CFG

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
PTP_MISC_CFG	0x00000000	1	Misc PTP configurations	Page 404
PTP_UPPER_LIMIT_CFG	0x00000004	1	Master counter upper limit	Page 405
PTP_UPPER_LIMIT_1_TIME_ADJ_CFG	0x00000008	1	Master counter upper limit one time adjust	Page 405
PTP_SYNC_INTR_ENA_CFG	0x0000000C	1	Sync Interrupt enable	Page 406
GEN_EXT_CLK_HIGH_PERIOD_CFG	0x00000010	2 0x00000004	Generated external clock high period	Page 406
GEN_EXT_CLK_LOW_PERIOD_CFG	0x00000018	2 0x00000004	Generated external clock low period	Page 407
GEN_EXT_CLK_CFG	0x00000020	2 0x00000004	External clock to internal master synchronization configuration	Page 407
CLK_ADJ_CFG	0x00000028	1	Generated clock adjustment configuration	Page 407
CLK_ADJ_FRQ	0x0000002C	1	Generated clock frequency adjustment	Page 408
PTP_EXT_SYNC_TIME_CFG	0x00000030	2 0x00000004	External sync time configuration	Page 408

5.3.10.1 DEVCPU_GCB:PTP_CFG:PTP_MISC_CFG

Parent: [DEVCPU_GCB:PTP_CFG](#)

Instances: 1

TABLE 5-188: FIELDS IN PTP_MISC_CFG

Field Name	Bit	Access	Description	Default
EXT_SYNC_OUTP_INV	14:13	R/W	Inversion of external sync output. 0: External sync output is not inverted 1: External sync output is inverted	0x0

TABLE 5-188: FIELDS IN PTP_MISC_CFG (CONTINUED)

Field Name	Bit	Access	Description	Default
EXT_SYNC_OUTP_SEL	12:11	R/W	GPIO selection of external sync output. 0: Use generated external clock as external sync output on GPIO 1588 PPS0/1 1: Use Master counter reset as external sync output on GPIO 1588 PPS0/1	0x0
EXT_SYNC_OUTP_ENA	10:9	R/W	External sync output enable. 0: External sync output is disabled 1: External sync output is enabled	0x0
EXT_SYNC_INP_INV	8:7	R/W	Inversion of external sync input. 0: External sync input is not inverted 1: External sync input is inverted	0x0
EXT_SYNC_INP_ENA	6:5	R/W	External sync input enable. 0: External sync input is disabled 1: External sync input is enabled	0x0
EXT_SYNC_CAP_ENA	4:3	R/W	Configure if a synchronization input pulse stores current Master counter value in EXT_SYNC_CURRENT_TIME_STAT. 0: Current Master counter value is not captured 1: Current Master counter value is captured	0x0
EXT_SYNC_ENA	2:1	R/W	Enable synchronization to external sync. If set, the Master counter is reset to the value set in PTP_EXT_SYNC_TIME_CFG. 0: Sync on external signal is disabled 1: Sync on external signal is enabled	0x0
PTP_ENA	0	R/W	Enable master counter. 0: Master counter disabled 1: Master counter enabled	0x0

5.3.10.2 DEVCPU_GCB:PTP_CFG:PTP_UPPER_LIMIT_CFG

Parent: [DEVCPU_GCB:PTP_CFG](#)

Instances: 1

TABLE 5-189: FIELDS IN PTP_UPPER_LIMIT_CFG

Field Name	Bit	Access	Description	Default
PTP_UPPER_LIMIT	29:0	R/W	Counter value where the Master counter should be reset. Unit is ns.	0x3B9ACA00

5.3.10.3 DEVCPU_GCB:PTP_CFG:PTP_UPPER_LIMIT_1_TIME_ADJ_CFG

Parent: [DEVCPU_GCB:PTP_CFG](#)

Instances: 1

TABLE 5-190: FIELDS IN PTP_UPPER_LIMIT_1_TIME_ADJ_CFG

Field Name	Bit	Access	Description	Default
PTP_UPPER_LIMIT_1_TIME_ADJ_SHOT	30	One-shot	One time enable for PTP_UPPER_LIMIT_1_TIME_ADJ 0: Normal operation 1: Timer is adjusted by usage of PTP_UPPER_LIMIT_1_TIME_ADJ Bit is cleared by HW	0x0
PTP_UPPER_LIMIT_1_TIME_ADJ	29:0	R/W	Counter value where the Master counter should be reset. Unit is ns.	0x3B9ACA00

5.3.10.4 DEVCPU_GCB:PTP_CFG:PTP_SYNC_INTR_ENA_CFG

Parent: [DEVCPU_GCB:PTP_CFG](#)

Instances: 1

TABLE 5-191: FIELDS IN PTP_SYNC_INTR_ENA_CFG

Field Name	Bit	Access	Description	Default
EXT_SYNC_CURRENT_TIME_ENA	2:1	R/W	Interrupt mask. Masks interrupt generation when a synchronization pulse is received on external sync input pin. 0: Interrupt is not generated 1: Interrupt is generated	0x0
SYNC_STAT_ENA	0	R/W	Interrupt mask. Masks interrupt generation when Master Timer generates a synchronization pulse. 0: Interrupt is not generated 1: Interrupt is generated	0x0

5.3.10.5 DEVCPU_GCB:PTP_CFG:GEN_EXT_CLK_HIGH_PERIOD_CFG

Parent: [DEVCPU_GCB:PTP_CFG](#)

Instances: 2

TABLE 5-192: FIELDS IN GEN_EXT_CLK_HIGH_PERIOD_CFG

Field Name	Bit	Access	Description	Default
MASTER_CNT_CLK_ENA	29	R/W	If set, a 31.25 MHz external clock is generated based on the 1588 master timer. Any adjustments to the 1588 master timer are also reflected in the external clock.	0x0
GEN_EXT_CLK_HIGH_PERIOD	28:0	R/W	High period for generated external clock in system clock cycles. N: External clock signal is high for (N+1)*system_clk cycles. E.g. N=999, system clock = 250 MHz which means 4 ns clk period. High Phase is 4 us.	0x000030D4

5.3.10.6 DEVCPU_GCB:PTP_CFG:GEN_EXT_CLK_LOW_PERIOD_CFG

Parent: [DEVCPU_GCB:PTP_CFG](#)

Instances: 2

TABLE 5-193: FIELDS IN GEN_EXT_CLK_LOW_PERIOD_CFG

Field Name	Bit	Access	Description	Default
GEN_EXT_CLK_LOW_PERIOD	29:0	R/W	Low period for generated external clock in system clock cycles. N: External clock signal is low for (N+1)*system_clk cycles. E.g. N=999, system clock = 250 MHz which means 4 ns clk period. Low Phase is 4 us.	0x000030D4

5.3.10.7 DEVCPU_GCB:PTP_CFG:GEN_EXT_CLK_CFG

Parent: [DEVCPU_GCB:PTP_CFG](#)

Instances: 2

TABLE 5-194: FIELDS IN GEN_EXT_CLK_CFG

Field Name	Bit	Access	Description	Default
GEN_EXT_CLK_SYNC_ENA	2	R/W	Enable sync of generated external clock to PTP sync master. 0: Synchronization is disabled 1: Synchronization is enabled	0x0
GEN_EXT_CLK_ENA	0	R/W	Enable generated external clock. 0: Generated external clock disabled. 1: Generated external clock enabled	0x0

5.3.10.8 DEVCPU_GCB:PTP_CFG:CLK_ADJ_CFG

Parent: [DEVCPU_GCB:PTP_CFG](#)

Instances: 1

TABLE 5-195: FIELDS IN CLK_ADJ_CFG

Field Name	Bit	Access	Description	Default
CLK_ADJ_DIR	2	R/W	Clock frequency adjustment direction. 0: Positive adjustment. Every N system clock cycles, a 1 is added to the counter. => clock period is decreased, clock frequency is increased 1: Negative adjustment. Every N system clock cycles, a 1 is subtracted from the counter. => clock period is increased, clock frequency is decreased	0x0
CLK_ADJ_ENA	1	R/W	Clock frequency adjust enable. 0: Adjustment Disabled 1: Adjustment Enabled	0x0

VSC7416-01

TABLE 5-195: FIELDS IN CLK_ADJ_CFG (CONTINUED)

Field Name	Bit	Access	Description	Default
CLK_ADJ_UPD	0	R/W	Defines when the updated adjustment value and direction takes effect. 0: updated values take immediate effect 1: updated values take effect after the next sync pulse	0x0

5.3.10.9 DEVCPU_GCB:PTP_CFG:CLK_ADJ_FRQ

Parent: [DEVCPU_GCB:PTP_CFG](#)

Instances: 1

TABLE 5-196: FIELDS IN CLK_ADJ_FRQ

Field Name	Bit	Access	Description	Default
CLK_ADJ	29:0	R/W	Clock frequency adjust. N: Number of system clock cycles after which the counter for the clock must be adjusted.	0x00004E1F

5.3.10.10 DEVCPU_GCB:PTP_CFG:PTP_EXT_SYNC_TIME_CFG

Parent: [DEVCPU_GCB:PTP_CFG](#)

Instances: 2

TABLE 5-197: FIELDS IN PTP_EXT_SYNC_TIME_CFG

Field Name	Bit	Access	Description	Default
EXT_SYNC_TIME	29:0	R/W	The current time is set to the value of this field when a rising edge is seen on the associated external sync input.	0x00000000

5.3.11 DEVCPU_GCB:PTP_STAT

Parent: [DEVCPU_GCB](#)

Instances: 1

TABLE 5-198: REGISTERS IN PTP_STAT

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
PTP_CURRENT_TIME_STAT	0x00000000	1	Current PTP master timer value	Page 408
EXT_SYNC_CURRENT_TIME_STAT	0x00000004	2 0x00000004	External sync time status	Page 409
PTP_EVT_STAT	0x0000000C	1	External sync time events	Page 409

5.3.11.1 DEVCPU_GCB:PTP_STAT:PTP_CURRENT_TIME_STAT

Parent: [DEVCPU_GCB:PTP_STAT](#)

Instances: 1

TABLE 5-199: FIELDS IN PTP_CURRENT_TIME_STAT

Field Name	Bit	Access	Description	Default
PTP_CURRENT_TIME	29:0	R/O	Current master counter value. Unit is 1 ns.	0x00000000

5.3.11.2 DEVCPU_GCB:PTP_STAT:EXT_SYNC_CURRENT_TIME_STAT

Parent: [DEVCPU_GCB:PTP_STAT](#)

Instances: 2

TABLE 5-200: FIELDS IN EXT_SYNC_CURRENT_TIME_STAT

Field Name	Bit	Access	Description	Default
EXT_SYNC_CURRENT_TIME	29:0	R/O	Snapshot of current time, when a rising edge was seen in on the external sync input. Note: A new value is only captured when the associated sticky bit is not set. Current time in clock_ticks when the rising edge on the external sync input was seen. Note: This has to be adjusted by 3 clock ticks for synchronizing the signal to core clock.	0x00000000

5.3.11.3 DEVCPU_GCB:PTP_STAT:PTP_EVT_STAT

Parent: [DEVCPU_GCB:PTP_STAT](#)

Instances: 1

TABLE 5-201: FIELDS IN PTP_EVT_STAT

Field Name	Bit	Access	Description	Default
CLK_ADJ_UPD_STICKY	3	Sticky	Identifies if the adjust value update has already happened in case the adjustment is only allowed to take place at sync. If update is allowed to take place immediately the sticky bit is unused. 0: updated has not yet happened 1: updated has happened	0x0
EXT_SYNC_CURRENT_TIME_STICKY	2:1	Sticky	Sticky bit that indicates a synchronization pulse has been captured on external sync input pin. 0: No Timestamp has been captured 1: New Timestamp has been captured	0x0
SYNC_STAT	0	Sticky	Master timer has generated a synchronization pulse to the Slave Timers. 0: No master timer wrap happened 1: Master timer wrap happened	0x0

5.3.12 DEVCPU_GCB:MEMITGR

Parent: [DEVCPU_GCB](#)

Instances: 1

TABLE 5-202: REGISTERS IN MEMITGR

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
MEMITGR_CTRL	0x00000000	1	Monitor control	Page 410
MEMITGR_STAT	0x00000004	1	Monitor status	Page 410
MEMITGR_INFO	0x00000008	1	Memory indication	Page 411
MEMITGR_IDX	0x0000000C	1	Memory index	Page 412
MEMITGR_DIV	0x00000010	1	Monitor speed	Page 413

5.3.12.1 DEVCPU_GCB:MEMITGR:MEMITGR_CTRL

Parent: [DEVCPU_GCB:MEMITGR](#)

Instances: 1

TABLE 5-203: FIELDS IN MEMITGR_CTRL

Field Name	Bit	Access	Description	Default
ACTIVATE	0	One-shot	Setting this field transitions the integrity monitor between operating modes. Transitioning between modes takes time, this field remains set until the new mode is reached. During this time the monitor also reports busy DEVCPU_GCB::MEMITGR_MODE.MODE_BUSY is set). From IDLE (DEVCPU_GCB::MEMITGR_MODE.MODE_IDLE is set) the monitor can transition into either DETECT or LISTEN mode, the DETECT mode is entered if a memory reports an indication - the LISTEN mode is entered if no indications are reported. The first time after reset the monitor will not detect indications, that is; it will transition directly from IDLE to LISTEN mode. From DETECT (DEVCPU_GCB::MEMITGR_MODE.MODE_DETECT is set) the monitor can transition into either DETECT or LISTEN mode, the DETECT mode is entered if more indications are reported - the LISTEN mode is entered if no more indications are reported. From LISTEN (DEVCPU_GCB::MEMITGR_MODE.MODE_LISTEN is set) the monitor can transition into IDLE mode. Software shall not set this field when the monitor is BUSY (when DEVCPU_GCB::MEMITGR_STAT.MODE_BUSY is set.)	0x0

5.3.12.2 DEVCPU_GCB:MEMITGR:MEMITGR_STAT

Parent: [DEVCPU_GCB:MEMITGR](#)

Instances: 1

TABLE 5-204: FIELDS IN MEMITGR_STAT

Field Name	Bit	Access	Description	Default
INDICATION	4	R/O	If this field is set then there is an indication from one of the memories that needs to be analyzed. An indication is either a parity detection or an error correction. This field is only set when the monitor is in LISTEN mode (DEVCPU_GCB::MEMITGR_MODE.MODE_LISTEN is set), in all other states (including BUSY) this field returns 0.	0x0
MODE_LISTEN	3	R/O	This field is set when the monitor is in LISTEN mode, during listen mode the monitor continually check for parity/correction indications from the memories.	0x0
MODE_DETECT	2	R/O	This field is set when the monitor is in DETECT mode, during detect mode the DEVCPU_GCB::MEMITGR_INFO and DEVCPU_GCB::MEMITGR_IDX registers contains valid information about one indication.	0x0
MODE_IDLE	1	R/O	This field is set when the monitor is in IDLE mode.	0x1
MODE_BUSY	0	R/O	The busy signal is a copy of the DEVCPU_GCB::MEMITGR_CTRL.ACTIVATE field, see description of that field for more information about the different states/modes of the monitor.	0x0

5.3.12.3 DEVCPU_GCB:MEMITGR:MEMITGR_INFO

Parent: [DEVCPU_GCB:MEMITGR](#)

Instances: 1

This field is only valid when the monitor is in the DETECT (DEVCPUGCB::MEMITGR_MODE.MODE_DETECT is set) mode.

TABLE 5-205: FIELDS IN MEMITGR_INFO

Field Name	Bit	Access	Description	Default
MEM_ERR	31	R/O	This field is set if the monitor has detected a parity indication (or an unrecoverable correction).	0x0
MEM_COR	30	R/O	This field is set if the monitor has detected a correction.	0x0

TABLE 5-205: FIELDS IN MEMITGR_INFO (CONTINUED)

Field Name	Bit	Access	Description	Default
MEM_ERR_OVF	29	R/O	This field is set if the monitor has detected a parity indication (or an unrecoverable correction) for which the address has not been recorded. If DEVCPU_GCB::MEMITGR_INFO.MEM_ERR is set then there has been more than one indication, then only the address of the newest indication has been kept. If DEVCPU_GCB::MEMITGR_INFO.MEM_ERR is cleared then an indication has occurred for which the address could not be stored, this is a very rare situation that can only happen if an indication is detected just as the memory is talking to the monitor.	0x0
MEM_COR_OVF	28	R/O	This field is set if the monitor has correction indication for which the address has not been recorded. If DEVCPU_GCB::MEMITGR_INFO.MEM_ERR is set then there has also been a parity indication (or an unrecoverable correction) which takes priority over correction indications. If DEVCPU_GCB::MEMITGR_INFO.MEM_ERR is cleared and DEVCPU_GCB::MEMITGR_INFO.MEM_COR is set then there has been more than one correction indication, then only the address of the newest correction indication has been kept. If DEVCPU_GCB::MEMITGR_INFO.MEM_ERR and DEVCPU_GCB::MEMITGR_INFO.MEM_COR is both cleared then a correction indication has occurred for which the address could not be stored, this is a very rare situation that can only happen if an indication is detected just as the memory is talking to the monitor.	0x0
MEM_ADDR	27:0	R/O	This field is valid only when DEVCPU_GCB::MEMITGR.MEM_ERR or DEVCPU_GCB::MEMITGR.MEM_COR is set.	0x00000000

5.3.12.4 DEVCPU_GCB::MEMITGR::MEMITGR_IDX

Parent: [DEVCPU_GCB::MEMITGR](#)

Instances: 1

This field is only valid when the monitor is in the DETECT (DEVCPU_GCB::MEMITGR_MODE.MODE_DETECT is set) mode.

TABLE 5-206: FIELDS IN MEMITGR_IDX

Field Name	Bit	Access	Description	Default
MEM_IDX	15:0	R/O	This field contains a unique index for the memory for which info is currently provided in DEVCPU_GCB::MEMITGR_MEMINFO. Indexes are counted from 1 (not 0).	0x0000

5.3.12.5 DEVCPU_GCB:MEMITGR:MEMITGR_DIV

Parent: [DEVCPU_GCB:MEMITGR](#)

Instances: 1

TABLE 5-207: FIELDS IN MEMITGR_DIV

Field Name	Bit	Access	Description	Default
MEM_DIV	15:0	R/W	Configure divider for generating the sync-pulse to memories (controls the speed at which the monitor talks to the memories). The lower this is set the faster indications can be read out of the memories. Set to 0x40	0x00FF

5.3.13 DEVCPU_GCB:VRAP

Parent: [DEVCPU_GCB](#)

Instances: 1

TABLE 5-208: REGISTERS IN VRAP

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
VRAP_ACCESS_STAT	0x00000000	1	VRAP events	Page 413

5.3.13.1 DEVCPU_GCB:VRAP:VRAP_ACCESS_STAT

Parent: [DEVCPU_GCB:VRAP](#)

Instances: 1

TABLE 5-209: FIELDS IN VRAP_ACCESS_STAT

Field Name	Bit	Access	Description	Default
FRM_RECV_STICKY	3	Sticky	This field is set if a valid VRAP (Versatile Register Access Protocol) frame has been received.	0x0
CMD_INVALID_STICKY	2	Sticky	This field is set if an invalid command inside a valid VRAP frame has been received. The VRAP engine has ignored the command.	0x0
FRM_INVALID_STICKY	1	Sticky	This field is set if an invalid VRAP frame has been received and discarded by the VRAP-engine. Frames with a VRAP header different from V1 are considered invalid.	0x0

TABLE 5-209: FIELDS IN VRAP_ACCESS_STAT (CONTINUED)

Field Name	Bit	Access	Description	Default
REPLY_ABORT_STICKY	0	Sticky	This field is set if a VRAP reply frame has been aborted. This may happen if a protocol violation is detected during VRAP request frame processing.	0x0

5.4 DEVCPU_ORG

TABLE 5-210: REGISTER GROUPS IN DEVCPU_ORG

Register Group Name	Offset within Target	Instances and Address Spacing	Description	Details
DEVCPU_ORG	0x00000000	1	Origin registers	Page 414

5.4.1 DEVCPU_ORG:DEVCPU_ORG

Parent: [DEVCPU_ORG](#)

Instances: 1

TABLE 5-211: REGISTERS IN DEVCPU_ORG

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
IF_CTRL	0x00000000	1	Physical interface control	Page 414
IF_CFGSTAT	0x00000004	1	Physical interface configuration and status	Page 415
ORG_CFG	0x00000008	1	Origin configuration	Page 415
ERR_CNTS	0x0000000C	1	Error counters	Page 416
TIMEOUT_CFG	0x00000010	1	Timeout configuration	Page 416
GPR	0x00000014	1	General purpose register	Page 417
MAILBOX_SET	0x00000018	1	Atomic set of mailbox	Page 417
MAILBOX_CLR	0x0000001C	1	Atomic clear of mailbox	Page 417
MAILBOX	0x00000020	1	Mailbox	Page 418
SEMA_CFG	0x00000024	1	Semaphore configuration	Page 418
SEMA0	0x00000028	1	Semaphore 0	Page 418
SEMA0_OWNER	0x0000002C	1	Semaphore 0 owner	Page 419
SEMA1	0x00000030	1	Semaphore 1	Page 419
SEMA1_OWNER	0x00000034	1	Semaphore 1 owner	Page 419

5.4.1.1 DEVCPU_ORG:DEVCPU_ORG:IF_CTRL

Parent: [DEVCPU_ORG:DEVCPU_ORG](#)

Instances: 1

TABLE 5-212: FIELDS IN IF_CTRL

Field Name	Bit	Access	Description	Default
IF_CTRL	3:0	R/W	This register configures critical interface parameters, it is constructed so that it can always be written correctly no matter the current state of the interface. When initializing a physical interface, then this is the first register that must be written, the state of the interface at that point may be unknown and therefore the following scheme is required to bring the interface to a known state: When writing a 4-bit value to this field construct a 32-bit data-word as follows: a) copy the 4-bit value into bits 3:0, 11:8, 19:16, and 27:24. b) reverse the 4-bit value and copy into bits 7:4, 15:12, 23:20, and 31:28. Example: To write the value 2 to this field; the 32-bit data-word to write is "0x42424242". Bit 0 configures endianness (when applicable), 0:Little-Endian, 1:Big-Endian. Bit 1 configures bit-order (when applicable), 0:MSB-first, 1:LSB-first. Bit 2,3 are reserved and should be kept 0. For the SI interface the default value of this field is 0x1. For all other interfaces the default value is 0x0.	0x0

5.4.1.2 DEVCPU_ORG:DEVCPU_ORG:IF_CFGSTAT

Parent: [DEVCPU_ORG:DEVCPU_ORG](#)

Instances: 1

TABLE 5-213: FIELDS IN IF_CFGSTAT

Field Name	Bit	Access	Description	Default
IF_STAT	16	Sticky	SI interface: This field is set if the SI interface has read data from device before it was ready (this can happen if the SI frequency is too high or when too few padding bytes has been specified).	0x0
IF_CFG	3:0	R/W	SI interface: This is the number of padding bytes to insert before read-data is shifted out of the device. This is needed when using high serial interface frequencies.	0x0

5.4.1.3 DEVCPU_ORG:DEVCPU_ORG:ORG_CFG

Parent: [DEVCPU_ORG:DEVCPU_ORG](#)

Instances: 1

TABLE 5-214: FIELDS IN ORG_CFG

Field Name	Bit	Access	Description	Default
FAST_WR	0	R/W	Clear this field to make write accesses return status OK because they are finished before status of the access is known. All non-OK responses will be logged in DEVCPU_ORG::ERR_CNTS no matter the value of this field.	0x1

5.4.1.4 DEVCPU_ORG:DEVCPU_ORG:ERR_CNTS

Parent: [DEVCPU_ORG:DEVCPU_ORG](#)

Instances: 1

TABLE 5-215: FIELDS IN ERR_CNTS

Field Name	Bit	Access	Description	Default
ERR_TGT_BUSY	19:16	R/W	The "Target Busy" indication is triggered when an interface tries to access a target which is currently reset or if another interface is using the target.	0x0
ERR_WD_DROP_ORG	15:12	R/W	The "Watch Dog Drop Origin" indication is triggered when the origin does not receive reply from the CSR ring within a given amount of time. This cannot happen during normal operation.	0x0
ERR_WD_DROP	11:8	R/W	The "Watch Dog Drop" indication is triggered when a target is too long about processing a request. Usually requests are processed immediately but some accesses requires interaction with the core-logic, when this logic is in reset or during very heavy traffic load there is a chance of timing out in the target.	0x0
ERR_NO_ACTION	7:4	R/W	The "No Action" indication is triggered when a target is accessed with a non-existing address. In other words, the target did not contain a register at the requested address.	0x0
ERR_UTM	3:0	R/W	The "Unknown Target Module" indication is triggered when a non-existing target is requested. In other words there was no target with the requested target-id.	0x0

5.4.1.5 DEVCPU_ORG:DEVCPU_ORG:TIMEOUT_CFG

Parent: [DEVCPU_ORG:DEVCPU_ORG](#)

Instances: 1

This register is shared between all interfaces on the Origin.

TABLE 5-216: FIELDS IN TIMEOUT_CFG

Field Name	Bit	Access	Description	Default
TIMEOUT_CFG	11:0	R/W	The contents of this field controls the timeout delay for the CSR system. Setting this field to 0 disables timeout. Timeout is handled as follows: A counter that decrements continually, when reaching 0 it will wrap to the value specified by this field. When a target has been processing a request for three "wraps" the target time-out and generate a WD_-DROP indication. In the origin an Interface that has been processing a request for four "wraps" will time out and generate a WD_-DROP_ORG indication.	0xFF

5.4.1.6 DEVCPU_ORG:DEVCPU_ORG:GPR

Parent: [DEVCPU_ORG:DEVCPU_ORG](#)

Instances: 1

This register is shared between all interfaces on the Origin.

TABLE 5-217: FIELDS IN GPR

Field Name	Bit	Access	Description	Default
GPR	31:0	R/W	General purpose 32-bit registers for debug and software development. The contents of this register can always (safely) be read. However write operations from different masters (to this register), which occur at (exactly) the same time, will fail.	0x00000000

5.4.1.7 DEVCPU_ORG:DEVCPU_ORG:MAILBOX_SET

Parent: [DEVCPU_ORG:DEVCPU_ORG](#)

Instances: 1

TABLE 5-218: FIELDS IN MAILBOX_SET

Field Name	Bit	Access	Description	Default
MAILBOX_SET	31:0	One-shot	Set bits in this register to atomically set corresponding bits in the DEVCPU_ORG::MAILBOX register. This register return 0 on read.	0x00000000

5.4.1.8 DEVCPU_ORG:DEVCPU_ORG:MAILBOX_CLR

Parent: [DEVCPU_ORG:DEVCPU_ORG](#)

Instances: 1

TABLE 5-219: FIELDS IN MAILBOX_CLR

Field Name	Bit	Access	Description	Default
MAILBOX_CLR	31:0	One-shot	Set bits in this register to atomically clear corresponding bits in the DEVCPU_ORG::MAILBOX register. This register return 0 on read.	0x00000000

5.4.1.9 DEVCPU_ORG:DEVCPU_ORG:MAILBOX

Parent: [DEVCPU_ORG:DEVCPU_ORG](#)

Instances: 1

This register is shared between all interfaces on the Origin.

TABLE 5-220: FIELDS IN MAILBOX

Field Name	Bit	Access	Description	Default
MAILBOX	31:0	R/W	Mailbox register which is shared between all interfaces on the Origin. Atomic (safe) modifications to the contents of this register can be performed by using the DEVCPU_ORG::MAILBOX_CLR and DEVCPU_ORG::MAILBOX_SET registers.	0x00000000

5.4.1.10 DEVCPU_ORG:DEVCPU_ORG:SEMA_CFG

Parent: [DEVCPU_ORG:DEVCPU_ORG](#)

Instances: 1

This register is shared between all interfaces on the Origin.

TABLE 5-221: FIELDS IN SEMA_CFG

Field Name	Bit	Access	Description	Default
SEMA_INTR_POL	1:0	R/W	By default semaphore-interrupt is generated when a semaphore is free. By setting this field interrupt is generated when semaphore is taken, bit 0 corresponds to semaphore 0, bit 1 to semaphore 1. 0: Interrupt on taken semaphore 1: Interrupt on free semaphore	0x0

5.4.1.11 DEVCPU_ORG:DEVCPU_ORG:SEMA0

Parent: [DEVCPU_ORG:DEVCPU_ORG](#)

Instances: 1

TABLE 5-222: FIELDS IN SEMA0

Field Name	Bit	Access	Description	Default
SEMA0	0	R/W	General semaphore. The first interface to read this field will be granted the semaphore, read DEVCPU_ORG::SEMA0_OWNER to see if ownership of semaphore was granted. Any interface can release the semaphore by writing (any value) to this field.	0x1

5.4.1.12 DEVCPU_ORG:DEVCPU_ORG:SEMA0_OWNER

Parent: [DEVCPU_ORG:DEVCPU_ORG](#)

Instances: 1

TABLE 5-223: FIELDS IN SEMA0_OWNER

Field Name	Bit	Access	Description	Default
SEMA0_OWNER	31:0	R/O	Current owner of the semaphore. This field is a one-hot encoded vector, each bit in this vector correspond to an interface on the origin. If this field return 0, then the semaphore was free at the time of reading the register. 0: Semaphore is free. 1: VCore System owns semaphore 2: VRAP owns semaphore 4: SI owns smaphore 8: MIIM owns semaphore	0x00000000

5.4.1.13 DEVCPU_ORG:DEVCPU_ORG:SEMA1

Parent: [DEVCPU_ORG:DEVCPU_ORG](#)

Instances: 1

TABLE 5-224: FIELDS IN SEMA1

Field Name	Bit	Access	Description	Default
SEMA1	0	R/W	General semaphore. The first interface to read this field will be granted the semaphore, read DEVCPU_ORG::SEMA1_OWNER to see if ownership of semaphore was granted. Any interface can release the semaphore by writing (any value) to this field.	0x1

5.4.1.14 DEVCPU_ORG:DEVCPU_ORG:SEMA1_OWNER

Parent: [DEVCPU_ORG:DEVCPU_ORG](#)

Instances: 1

TABLE 5-225: FIELDS IN SEMA1_OWNER

Field Name	Bit	Access	Description	Default
SEMA1_OWNER	31:0	R/O	Current owner of the semaphore. This field is a one-hot encoded vector, each bit in this vector correspond to an interface on the origin. If this field return 0, then the semaphore was free at the time of reading the register. 0: Semaphore is free. 1: VCore System owns semaphore 2: VRAP owns semaphore 4: SI owns smaphore 8: MIIM owns semaphore	0x00000000

5.5 DEVCPU_QS

TABLE 5-226: REGISTER GROUPS IN DEVCPU_QS

Register Group Name	Offset within Target	Instances and Address Spacing	Description	Details
XTR	0x00000000	1	Frame extraction	Page 420
INJ	0x00000024	1	Frame injection	Page 423

5.5.1 DEVCPU_QS:XTR

Parent: [DEVCPU_QS](#)

Instances: 1

The DEVCPU_QS::XTR_GRP_CFG, DEVCPU_QS::XTR_RD, and DEVCPU_QS::FRM_PRUNING registers are replicated once per extraction group. Replication index n corresponds to extraction group number n.

TABLE 5-227: REGISTERS IN XTR

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
XTR_GRP_CFG	0x00000000	2 0x00000004	Extraction group configuration	Page 420
XTR_RD	0x00000008	2 0x00000004	Manual extraction data	Page 421
XTR_FRM_PRUNING	0x00000010	2 0x00000004	Extraction frame pruning	Page 422
XTR_FLUSH	0x00000018	1	Extraction group flush	Page 422
XTR_DATA_PRESENT	0x0000001C	1	Extraction status	Page 423

5.5.1.1 DEVCPU_QS:XTR:XTR_GRP_CFG

Parent: [DEVCPU_QS:XTR](#)

Instances: 2

TABLE 5-228: FIELDS IN XTR_GRP_CFG

Field Name	Bit	Access	Description	Default
MODE	3:2	R/W	Configures mode of the extraction group. Each extraction group can be assigned to one of three owners. Note: The VRAP block support only one context, if more than one extraction group is assigned the lowest group-number will be used. 0: VRAP block 1: Manual extraction (via DEVCPU_QS registers) 2: FDMA extraction and manual extraction via SBA registers	0x0
STATUS_WORD_POS	1	R/W	Select order of last data and status words. This field only applies to manual extraction mode (see DEVCPU_QS::XTR_GRP_CFG.MODE). 0: Status just before last data 1: Status just after last data	0x1
BYTE_SWAP	0	R/W	This field allows swapping the endianness of DEVCPU_QS::XTR_RD. Most software will want to read extraction data in network order (big-endian mode), i.e. the first byte of the destination MAC address to be placed on byte-address 0 of DEVCPU_QS::XTR_RD. In order to do this a little endian CPU must set this field, a big endian CPU must clear this field. This field only applies to manual extraction mode (see DEVCPU_QS::XTR_GRP_CFG.MODE). 0: Same endianness 1: Swap endianness	0x1

5.5.1.2 DEVCPU_QS:XTR:XTR_RD

Parent: [DEVCPU_QS:XTR](#)

Instances: 2

TABLE 5-229: FIELDS IN XTR_RD

Field Name	Bit	Access	Description	Default
DATA	31:0	R/O	Frame Data. Read from this register to obtain the next 32 bits of the frame data currently stored in the CPU queue system. Each read must check for the special values "0x8000000n", 0<=n<=7, as described by the encoding. Note: The encoding is presented in little endian format, if swapping is used (see DEVCPU_QS::XTR_GRP_CFG.BYTE_SWAP), then the special values are swapped as well. I.e. a little endian CPU using BYTE_SWAP=1 has to look for "0x0n000080" instead of "0x8000000n". The position of the unused/valid bytes follows the endianness encoding and swapping. n=0-3: EOF. Unused bytes in end-of-frame word is 'n' n=4 : EOF, but truncated n=5 : EOF Aborted. Frame invalid n=6 : Escape. Next read is packet data n=7 : Data not ready for reading out	0x00000000

5.5.1.3 DEVCPU_QS:XTR:XTR_FRM_PRUNING

Parent: [DEVCPU_QS:XTR](#)

Instances: 2

TABLE 5-230: FIELDS IN XTR_FRM_PRUNING

Field Name	Bit	Access	Description	Default
PRUNE_SIZE	7:0	R/W	Extracted frames for the corresponding queue are pruned to (PRUNE_SIZE+1) 32-bit words. Note : PRUNE_SIZE is the frame data size + optional IFH (see SYS::PORT_MODE.INCL_XTR_HDR for more information on enabling IFH for extraction frames). 0 : No pruning 1: Frames extracted are pruned to 8 bytes 2: Frames extracted are pruned to 12 bytes ... 255: Frames extracted are pruned to 1024 bytes	0x00

5.5.1.4 DEVCPU_QS:XTR:XTR_FLUSH

Parent: [DEVCPU_QS:XTR](#)

Instances: 1

TABLE 5-231: FIELDS IN XTR_FLUSH

Field Name	Bit	Access	Description	Default
FLUSH	1:0	R/W	Enable software flushing of a CPU queue. Note the queue will continue to be flushed until this field is cleared by SW. The flushing will automatically stop on frame boundary so OQS is allowed to transmit to the CPU queue during flushing. 0: No action 1: Do CPU queue flushing	0x0

5.5.1.5 DEVCPU_QS:XTR:XTR_DATA_PRESENT

Parent: [DEVCPU_QS:XTR](#)

Instances: 1

TABLE 5-232: FIELDS IN XTR_DATA_PRESENT

Field Name	Bit	Access	Description	Default
DATA_PRESENT	1:0	R/O	Shows if data is available for a specific group. It remains set until all frame data have been extracted. This field is only set if the mode of the group is set to manual extraction via DEVCPU_QS registers. 0: No frames available for this CPU queue group 1: At least one frame is available for this CPU queue group	0x0

5.5.2 DEVCPU_QS:INJ

Parent: [DEVCPU_QS](#)

Instances: 1

The DEVCPU_QS::INJ_GRP_CFG, DEVCPU_QS::INJ_WR, DEVCPU_QS::INJ_CTRL, and DEVCPU_QS::INJ_ERR registers are replicated once per injection group. Replication index n corresponds to injection group number n.

TABLE 5-233: REGISTERS IN INJ

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
INJ_GRP_CFG	0x00000000	2 0x00000004	Injection group configuration	Page 423
INJ_WR	0x00000008	2 0x00000004	Manual injection data	Page 424
INJ_CTRL	0x00000010	2 0x00000004	Injection control	Page 424
INJ_STATUS	0x00000018	1	Injection status	Page 425
INJ_ERR	0x0000001C	2 0x00000004	Injection errors	Page 426

5.5.2.1 DEVCPU_QS:INJ:INJ_GRP_CFG

Parent: [DEVCPU_QS:INJ](#)

Instances: 2

TABLE 5-234: FIELDS IN INJ_GRP_CFG

Field Name	Bit	Access	Description	Default
MODE	3:2	R/W	Configures mode of the injection group. Each injection group can be assigned to one of three owners. Note: The VRAP block support only one context, if more than one injection group is assigned the lowest group-number will be used. 0: VRAP block 1: Manual injection (via DEVCPU_QS registers) 2: FDMA injection and manual injection via SBA registers	0x0
BYTE_SWAP	0	R/W	This field allows swapping the endianness of the DEVCPU_QS::INJ_WR register. Most software will want to write injection data in network order (big-endian mode), i.e. the first byte of the destination MAC address to be placed on byte-address 0 of DEVCPU_QS::INJ_WR. In order to do this a little endian CPU must set this field, a big endian CPU must clear this field. This field only applies to manual extraction mode (see DEVCPU_QS::INJ_GRP_CFG.MODE). 0: Same endianness as CPU 1: Swap endianness	0x1

5.5.2.2 DEVCPU_QS:INJ:INJ_WR

Parent: [DEVCPU_QS:INJ](#)

Instances: 2

TABLE 5-235: FIELDS IN INJ_WR

Field Name	Bit	Access	Description	Default
DATA	31:0	R/W	Frame Write. Write to this register inject the next 32 bits of the frame data currently injected into the chip. Reading from this register returns 0.	0x00000000

5.5.2.3 DEVCPU_QS:INJ:INJ_CTRL

Parent: [DEVCPU_QS:INJ](#)

Instances: 2

TABLE 5-236: FIELDS IN INJ_CTRL

Field Name	Bit	Access	Description	Default
GAP_SIZE	22:21	R/W	Set this field to configure a minimum delay after each frame EOF when injecting to the QS. 0 though 15 additional 32-bit word delays can be inserted after eof. The number of delays is configured by setting low 2 bit in replication 0 and high two 2 in replication 1. So that the number of delays to insert is {INJ_CTRL[1].GAP_SIZE, INJ_CTRL[0].GAP_SIZE}.	0x0
ABORT	20	One-shot	Set to abort the current frame.	0x0
EOF	19	One-shot	Set to indicate that the next data written to DEVCPU_QS::INJ_WR is end-of-frame. At the same time as setting this field, also set DEVCPU_QS::INJ_CTRL.VLD_BYTES to indicate the number of valid data bytes in the end-of-frame word.	0x0
SOF	18	One-shot	Set to indicate that the next data written to DEVCPU_QS::INJ_WR is start-of-frame.	0x0
VLD_BYTES	17:16	R/W	Set to indicate how many bytes of the next data written to DEVCPU_QS::INJ_WR which are valid. This field is only used during end-of-frame words (see DEVCPU_QS::INJ_CTRL.EOF for more information). The position of the valid bytes follows the endianness encoding and swapping. 0: All bytes are valid n: 'n' byte are valid	0x0

5.5.2.4 DEVCPU_QS:INJ:INJ_STATUS

Parent: [DEVCPU_QS:INJ](#)

Instances: 1

TABLE 5-237: FIELDS IN INJ_STATUS

Field Name	Bit	Access	Description	Default
WMARK_REACHED	5:4	R/O	Before the CPU injects a frame, software may check if the input queue has reached high watermark. If watermark in the IQS has been reached this bit will be set. 0: Input queue has not reached high watermark 1: Input queue has reached high watermark, and frames injected may be dropped due to buffer overflow	0x0
FIFO_RDY	3:2	R/O	When '1' the injector group's FIFO is ready for additional data written through the DEVCPU_QS::INJ_WR register. 0: The injector group cannot accept additional data 1: The injector group is able to accept additional data	0x0

TABLE 5-237: FIELDS IN INJ_STATUS (CONTINUED)

Field Name	Bit	Access	Description	Default
INJ_IN_PROGRESS	1:0	R/O	When '1' the injector group is in the process of receiving a frame, and at least one write to INJ_WR remains before the frame is forwarded to the front ports. When '0' the injector group is waiting for an initiation of a frame injection. 0: A frame injection is not in progress 1: A frame injection is in progress	0x0

5.5.2.5 DEVCPU_QS:INJ:INJ_ERR

Parent: [DEVCPU_QS:INJ](#)

Instances: 2

TABLE 5-238: FIELDS IN INJ_ERR

Field Name	Bit	Access	Description	Default
ABORT_ERR_STICKY	1	Sticky	Set if a frame has been aborted because of double-SOF injection (missing EOF).	0x0
WR_ERR_STICKY	0	Sticky	Set in case of overflow as a result of not obeying FIFO-ready	0x0

5.6 HSIO

Register Collection for Control of Macros (SERDES1G, SERDES6G, LCPLL)

TABLE 5-239: REGISTER GROUPS IN HSIO

Register Group Name	Offset within Target	Instances and Address Spacing	Description	Details
PLL5G_CFG	0x00000000	1	PLL5G Configuration Registers	Page 427
PLL5G_STATUS	0x0000001C	1	PLL5G Status Registers	Page 428
RCOMP_STATUS	0x00000028	1	RCOMP Status Registers	Page 428
SYNC_ETH_CFG	0x0000002C	1	SYNC_ETH Configuration Registers	Page 429
MCB_PLL5G_RCOMP_CFG	0x00000038	1	MCB PLL5G_RCOMP Configuration Register	Page 430
SERDES1G_ANA_CFG	0x0000003C	1	SERDES1G Analog Configuration Registers	Page 431
SERDES1G_DIG_CFG	0x00000058	1	SERDES1G Digital Configuration Register	Page 436
SERDES1G_DIG_STATUS	0x00000070	1	SERDES1G Digital Status Register	Page 437
MCB_SERDES1G_CFG	0x00000074	1	MCB SERDES1G Configuration Register	Page 437
SERDES6G_ANA_CFG	0x00000078	1	SERDES6G Analog Configuration Registers	Page 438
SERDES6G_DIG_CFG	0x0000009C	1	SERDES6G Digital Configuration Registers	Page 443

TABLE 5-239: REGISTER GROUPS IN HSIO (CONTINUED)

Register Group Name	Offset within Target	Instances and Address Spacing	Description	Details
MCB_SERDES6G_CFG	0x000000C4	1	MCB SERDES6G Configuration Register	Page 445

5.6.1 HSIO:PLL5G_CFG

Parent: [HSIO](#)

Instances: 1

Configuration register set for PLL5G.

TABLE 5-240: REGISTERS IN PLL5G_CFG

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
PLL5G_CFG0	0x00000000	1	PLL5G Configuration 0	Page 427
PLL5G_CFG3	0x0000000C	1	PLL5G Configuration 3	Page 427

5.6.1.1 HSIO:PLL5G_CFG:PLL5G_CFG0

Parent: [HSIO:PLL5G_CFG](#)

Instances: 1

Configuration register 0 for PLL5G

TABLE 5-241: FIELDS IN PLL5G_CFG0

Field Name	Bit	Access	Description	Default
RESERVED	5:0	R/W	Must be set to its default.	0x11
CLK_DIV	11:6	R/W	Applicable to LCPLL2 only. Setting for CLK-OUT2 divider: 0-4: Reserved. 5: 250 MHz. 6-7: Reserved	0x06
RESERVED	12	R/W	Must be set to its default.	0x1
RESERVED	13	R/W	Must be set to its default.	0x1
RESERVED	14	R/W	Must be set to its default.	0x1
RESERVED	15	R/W	Must be set to its default.	0x1
RESERVED	17:16	R/W	Must be set to its default.	0x2
RESERVED	22:18	R/W	Must be set to its default.	0x0D
RESERVED	26:23	R/W	Must be set to its default.	0x7
RESERVED	28	R/W	Must be set to its default.	0x1
RESERVED	29	R/W	Must be set to its default.	0x1
RESERVED	30	R/W	Must be set to its default.	0x1

5.6.1.2 HSIO:PLL5G_CFG:PLL5G_CFG3

Parent: [HSIO:PLL5G_CFG](#)

Instances: 1

Configuration register 3 for PLL5G

TABLE 5-242: FIELDS IN PLL5G_CFG3

Field Name	Bit	Access	Description	Default
RESERVED	7:0	R/W	Must be set to its default.	0x28
RESERVED	14	R/W	Must be set to its default.	0x1
RESERVED	17	R/W	Must be set to its default.	0x1
CLK_SEL	21:19	R/W	Applicable to LCPLL2 only. Select clock output: 0-2: Reserved. 3: Enable CLKOUT2 clock. 4-7: Reserved	0x4

5.6.2 HSIO:PLL5G_STATUS

Parent: [HSIO](#)

Instances: 1

Status register set for PLL5G.

TABLE 5-243: REGISTERS IN PLL5G_STATUS

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
PLL5G_STATUS0	0x00000000	1	PLL5G Status 0	Page 428

5.6.2.1 HSIO:PLL5G_STATUS:PLL5G_STATUS0

Parent: [HSIO:PLL5G_STATUS](#)

Instances: 1

Status register 0 for the PLL5G

TABLE 5-244: FIELDS IN PLL5G_STATUS0

Field Name	Bit	Access	Description	Default
LOCK_STATUS	0	R/O	PLL lock status 0: not locked, 1: locked	0x0
READBACK_DATA	8:1	R/O	RCPLL Interface to read back internal data of the FSM.	0x00
CALIBRATION_DONE	9	R/O	RCPLL Flag that indicates that the calibration procedure has finished.	0x0
CALIBRATION_ERR	10	R/O	RCPLL Flag that indicates errors that may occur during the calibration procedure.	0x0
OUT_OF_RANGE_ERR	11	R/O	RCPLL Flag that indicates a out of range condition while NOT in calibration mode.	0x0
RANGE_LIM	12	R/O	RCPLL Flag range limiter signaling	0x0

5.6.3 HSIO:RCOMP_STATUS

Parent: [HSIO](#)

Instances: 1

Status register set for RCOMP.

TABLE 5-245: REGISTERS IN RCOMP_STATUS

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
RCOMP_STATUS	0x00000000	1	RCOMP Status	Page 429

5.6.3.1 HSIO:RCOMP_STATUS:RCOMP_STATUS

Parent: [HSIO:RCOMP_STATUS](#)

Instances: 1

Status register bits for the RCOMP

TABLE 5-246: FIELDS IN RCOMP_STATUS

Field Name	Bit	Access	Description	Default
BUSY	12	R/O	Resistor comparison activity 0: resistor measurement finished or inactive 1: resistor measurement in progress	0x0
DELTA_ALERT	7	R/O	Alarm signal if rcomp isn't best choice any-more 0: inactive 1: active	0x0
RCOMP	3:0	R/O	Measured resistor value 0: maximum resistance value 15: minimum resistance value	0x0

5.6.4 HSIO:SYNC_ETH_CFG

Parent: [HSIO](#)

Instances: 1

Configuration register set for SYNC_ETH.

TABLE 5-247: REGISTERS IN SYNC_ETH_CFG

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
SYNC_ETH_CFG0	0x00000000	1	SYNC ETH Configuration 0	Page 429
SYNC_ETH_CFG1	0x00000004	1	SYNC ETH Configuration 1	Page 430
SYNC_ETH_CFG2	0x00000008	1	SYNC ETH Configuration 2	Page 430

5.6.4.1 HSIO:SYNC_ETH_CFG:SYNC_ETH_CFG0

Parent: [HSIO:SYNC_ETH_CFG](#)

Instances: 1

Configuration register for RCVRD_CLK pad enable (1 bit per recovered clock output)

TABLE 5-248: FIELDS IN SYNC_ETH_CFG0

Field Name	Bit	Access	Description	Default
RECO_CLK_ENA	9:0	R/W	Enable recovered clock pads 0: Disable (high-impedance) 1: Enable (output recovered clock)	0x000

5.6.4.2 HSIO:SYNC_ETH_CFG:SYNC_ETH_CFG1

Parent: [HSIO:SYNC_ETH_CFG](#)

Instances: 1

Configuration register for RCVRD_CLK selection (1 bit per recovered clock output)

TABLE 5-249: FIELDS IN SYNC_ETH_CFG1

Field Name	Bit	Access	Description	Default
RECO_CLK_SEL0	19:16	R/W	select recovered clock source for RCVD_ - CLK[0] 0-9: select source 10-15: reserved	0x0
RECO_CLK_SEL	9:0	R/W	select common recovered clock source 0: Use individual clock source 1: Use common clock source	0x000

5.6.4.3 HSIO:SYNC_ETH_CFG:SYNC_ETH_CFG2

Parent: [HSIO:SYNC_ETH_CFG](#)

Instances: 1

Selection register for SYNC_ETH.

TABLE 5-250: FIELDS IN SYNC_ETH_CFG2

Field Name	Bit	Access	Description	Default
SEL_RECO_CLK_B	1:0	R/W	Select recovered clock divider B 0: No clock dividing 1: Divide clock by 5 2: Divide clock by 4 3: Reserved	0x0

5.6.5 HSIO:MCB_PLL5G_RCOMP_CFG

Parent: [HSIO](#)

Instances: 1

All PLL5G_RCOMP macros are accessed via the serial Macro Configuration Bus (MCB). Each macro is configured by one MCB slave. All MCB slaves are connected in a daisy-chain loop.

TABLE 5-251: REGISTERS IN MCB_PLL5G_RCOMP_CFG

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
MCB_PLL5G_RCOMP_AD- DR_CFG	0x00000000	1	MCB PLL5G RCOMP Address Cfg	Page 431

5.6.5.1 HSIO:MCB_PLL5G_RCOMP_CFG:MCB_PLL5G_RCOMP_ADDR_CFG

Parent: [HSIO:MCB_PLL5G_RCOMP_CFG](#)

Instances: 1

Configuration of PLL5G_RCOMP MCB slaves to be accessed

TABLE 5-252: FIELDS IN MCB_PLL5G_RCOMP_ADDR_CFG

Field Name	Bit	Access	Description	Default
PLL5G_RCOMP_WR_ONE_SHOT	31	One-shot	Initiate a write access to marked PLL5G_RCOMP slaves 0: No write operation pending 1: Initiate write to slaves (kept 1 until write operation has finished)	0x0
PLL5G_RCOMP_RD_ONE_SHOT	30	One-shot	Initiate a read access to marked PLL5G_RCOMP slaves 0: No read operation pending (read op finished after bit has been set) 1: Initiate a read access (kept 1 until read operation has finished)	0x0
PLL5G_RCOMP_ADDR	1:0	R/W	Activation vector for PLL5G_RCOMP-Slaves, one-hot coded, each bit is related to one macro, e.g. bit 0 enables/disables access to macro No. 0 0: Disable macro access via MCB 1: Enable macro access via MCB	0x1

5.6.6 HSIO:SERDES1G_ANA_CFG

Parent: [HSIO](#)

Instances: 1

Configuration register set for SERDES1G (analog parts)

TABLE 5-253: REGISTERS IN SERDES1G_ANA_CFG

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
SERDES1G_DES_CFG	0x00000000	1	SERDES1G Deserializer Cfg	Page 431
SERDES1G_IB_CFG	0x00000004	1	SERDES1G Input Buffer Cfg	Page 432
SERDES1G_OB_CFG	0x00000008	1	SERDES1G Output Buffer Cfg	Page 433
SERDES1G_SER_CFG	0x0000000C	1	SERDES1G Serializer Cfg	Page 434
SERDES1G_COMMON_CFG	0x00000010	1	SERDES1G Common Cfg	Page 435
SERDES1G_PLL_CFG	0x00000014	1	SERDES1G Pll Cfg	Page 435

5.6.6.1 HSIO:SERDES1G_ANA_CFG:SERDES1G_DES_CFG

Parent: [HSIO:SERDES1G_ANA_CFG](#)

Instances: 1

Configuration register for SERDES1G deserializer

TABLE 5-254: FIELDS IN SERDES1G_DES_CFG

Field Name	Bit	Access	Description	Default
DES_PHS_CTRL	16:13	R/W	Control of phase regulator logic. Bit 3 must be set to 0. 0: Disabled 1: Enabled with 99 ppm limit 2: Enabled with 202 ppm limit 3: Enabled with 485 ppm limit 4: Enabled if corresponding PCS is in sync with 50 ppm limit 5: Enabled if corresponding PCS is in sync with 99 ppm limit 6: Enabled if corresponding PCS is in sync with 202 ppm limit 7: Enabled if corresponding PCS is in sync with 485 ppm limit	0x6
DES_CPMD_SEL	12:11	R/W	Deserializer phase control, main CP/MD select 0: Directly from DES 1: Through hysteresis stage from DES 2: From core 3: Disabled	0x0
DES_MBTR_CTRL	10:8	R/W	Des phase control for 180 degrees deadlock block mode of operation 0: Depending on density of input pattern 1: Active until PCS has synchronized 2: Depending on density of input pattern until PCS has synchronized 3: Never 4: Always 5-7: Reserved	0x2
DES_BW_ANA	7:5	R/W	Bandwidth selection for proportional path of CDR loop. 0: No division 1: Divide by 2 2: Divide by 4 3: Divide by 8 4: Divide by 16 5: Divide by 32 6: Divide by 64 7: Divide by 128	0x6
DES_BW_HYST	3:1	R/W	Selection of time constant for integrative path of CDR loop. 0: Divide by 2 1: Divide by 4 2: Divide by 8 3: Divide by 16 4: Divide by 32 5: Divide by 64 6: Divide by 128 7: Divide by 256	0x7

5.6.6.2 HSIO:SERDES1G_ANA_CFG:SERDES1G_IB_CFG

Parent: [HSIO:SERDES1G_ANA_CFG](#)

Instances: 1

Configuration register for SERDES1G input buffer

TABLE 5-255: FIELDS IN SERDES1G_IB_CFG

Field Name	Bit	Access	Description	Default
IB_FX100_ENA	27	R/W	Switches signal detect circuit into low frequency mode, must be used in FX100 mode	0x0
RESERVED	26:24	R/W	Must be set to its default.	0x1
IB_DET_LEV	21:19	R/W	Detect thresholds: 0: 159-189mVppd 1: 138-164mVppd 2: 109-124mVppd 3: 74-89mVppd	0x0
IB_HYST_LEV	14	R/W	Input buffer hysteresis levels: 0: 59-79mV 1: 81-124mV	0x0
IB_ENA_CMV_TERM	13	R/W	Enable common mode voltage termination 0: Low termination (VDD_A x 0.7) 1: High termination (VDD_A)	0x1
IB_ENA_DC_COUPLING	12	R/W	Enable dc-coupling of input signal 0: Disable 1: Enable	0x0
IB_ENA_DETLEV	11	R/W	Enable detect level circuit 0: Disable 1: Enable	0x0
IB_ENA_HYST	10	R/W	Enable hysteresis for input signal. Hysteresis can only be enabled if DC offset compensation is disabled. 0: Disable 1: Enable	0x0
IB_ENA_OFFSET_COMP	9	R/W	Enable offset compensation of input stage. This bit must be disabled to enable hysteresis (IB_ENA_HYST). 0: Disable 1: Enable	0x1
IB_EQ_GAIN	8:6	R/W	Selects weighting between AC and DC input path: 0: Reserved 1: Reserved 2: 0dB (recommended value) 3: 1.5dB 4: 3dB 5: 6dB 6: 9dB 12.5dB	0x2
IB_SEL_CORNER_FREQ	5:4	R/W	Corner frequencies of AC path: 0: 1.3GHz 1: 1.5GHz 2: 1.6GHz 3: 1.8GHz	0x0
RESERVED	3:0	R/W	Must be set to its default.	0xB

5.6.6.3 HSIO:SERDES1G_ANA_CFG:SERDES1G_OB_CFG

Parent: [HSIO:SERDES1G_ANA_CFG](#)

VSC7416-01

Instances: 1

Configuration register for SERDES1G output buffer

TABLE 5-256: FIELDS IN SERDES1G_OB_CFG

Field Name	Bit	Access	Description	Default
OB_SLP	18:17	R/W	Slope / slew rate control: 0: 45ps 1: 85ps 2: 105ps 3: 115ps	0x3
OB_AMP_CTRL	16:13	R/W	Amplitude control in steps of 50mVppd. 0: 0.4Vppd 15: 1.1Vppd	0x0
RESERVED	12:10	R/W	Must be set to its default.	0x2
OB_VCM_CTRL	7:4	R/W	Common mode voltage control: 0: Reserved 1: 440mV 2: 480mV 3: 460mV 4: 530mV 5: 500mV 6: 570mV 7: 550mV	0x4
RESERVED	3:0	R/W	Must be set to its default.	0x1

5.6.6.4 HSIO:SERDES1G_ANA_CFG:SERDES1G_SER_CFG

Parent: [HSIO:SERDES1G_ANA_CFG](#)

Instances: 1

Configuration register for SERDES1G serializer

TABLE 5-257: FIELDS IN SERDES1G_SER_CFG

Field Name	Bit	Access	Description	Default
SER_IDLE	9	R/W	Invert output D0b for idle-mode of OB 0: Non-inverting 1: Inverting	0x0
SER_DEEMPH	8	R/W	Invert and delays (one clk cycle) output D1 for de-emphasis of OB 0: Non-inverting and non-delaying 1: Inverting and delaying	0x0
SER_ENHYS	3	R/W	Enable hysteresis for phase alignment 0: Disable hysteresis 1: Enable hysteresis	0x0
SER_BIG_WIN	2	R/W	Use wider window for phase alignment 0: Use small window for low jitter (100 to 200ps) 1: Use wide window for higher jitter (150 to 300 ps)	0x0
SER_EN_WIN	1	R/W	Enable window for phase alignment 0: Disable window 1: Enable window	0x0

TABLE 5-257: FIELDS IN SERDES1G_SER_CFG (CONTINUED)

Field Name	Bit	Access	Description	Default
SER_ENALI	0	R/W	Enable phase alignment 0: Disable phase alignment 1: Enable phase alignment	0x0

5.6.6.5 HSIO:SERDES1G_ANA_CFG:SERDES1G_COMMON_CFG

Parent: [HSIO:SERDES1G_ANA_CFG](#)

Instances: 1

Configuration register for common SERDES1G functions Note: When enabling the facility loop (ena_floop) also the phase alignment in the serializer has to be enabled and configured adequate.

TABLE 5-258: FIELDS IN SERDES1G_COMMON_CFG

Field Name	Bit	Access	Description	Default
SYS_RST	31	R/W	System reset (low active) 0: Apply reset (not self-clearing) 1: Reset released	0x0
SE_AUTO_SQUELCH_B_ENA	22	R/W	Enable auto-squelching for sync. ethernet bus B 0: Disable 1: Enable	0x0
SE_AUTO_SQUELCH_A_ENA	21	R/W	Enable auto-squelching for sync. ethernet bus A 0: Disable 1: Enable	0x0
RECO_SEL_B	20	R/W	Select recovered clock of this lane on sync. ethernet bus B 0: Lane not selected 1: Lane selected	0x0
RECO_SEL_A	19	R/W	Select recovered clock of this lane on sync. ethernet bus A 0: Lane not selected 1: Lane selected	0x0
ENA_LANE	18	R/W	Enable lane 0: Disable lane 1: Enable line	0x0
ENA_ELOOP	11	R/W	Enable equipment loop 0: Disable 1: Enable	0x0
ENA_FLOOP	10	R/W	Enable facility loop 0: Disable 1: Enable	0x0
RESERVED	0	R/W	Must be set to its default.	0x1

5.6.6.6 HSIO:SERDES1G_ANA_CFG:SERDES1G_PLL_CFG

Parent: [HSIO:SERDES1G_ANA_CFG](#)

Instances: 1

Configuration register for SERDES1G RCPLL

TABLE 5-259: FIELDS IN SERDES1G_PLL_CFG

Field Name	Bit	Access	Description	Default
RESERVED	21	R/W	Must be set to its default.	0x1
PLL_FSM_CTRL_DATA	15:8	R/W	Control data for FSM	0x00
PLL_FSM_ENA	7	R/W	Enable FSM	0x0

5.6.7 HSIO:SERDES1G_DIG_CFG

Parent: [HSIO](#)

Instances: 1

Configuration register set for SERDES1G digital BIST and DFT functions.

TABLE 5-260: REGISTERS IN SERDES1G_DIG_CFG

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
SERDES1G_MISC_CFG	0x00000014	1	SERDES1G Misc Configuration	Page 436

5.6.7.1 HSIO:SERDES1G_DIG_CFG:SERDES1G_MISC_CFG

Parent: [HSIO:SERDES1G_DIG_CFG](#)

Instances: 1

Configuration register for miscellaneous functions

TABLE 5-261: FIELDS IN SERDES1G_MISC_CFG

Field Name	Bit	Access	Description	Default
SEL_RECO_CLK	17:16	R/W	Select recovered clock divider 0: No clock dividing 1: Divide clock by 5 2: Divide clock by 4 3: Reserved	0x0
DES_100FX_CPMD_ENA	8	R/W	Enable deserializer cp/md handling for 100fx mode 0: Disable 1: Enable	0x0
RX_LPI_MODE_ENA	5	R/W	Enable RX-Low-Power feature (Power control by LPI-FSM in connected PCS) 0: Disable 1: Enable	0x0
TX_LPI_MODE_ENA	4	R/W	Enable TX-Low-Power feature (Power control by LPI-FSM in connected PCS) 0: Disable 1: Enable	0x0
RX_DATA_INV_ENA	3	R/W	Enable data inversion received from Deserializer 0: Disable 1: Enable	0x0

TABLE 5-261: FIELDS IN SERDES1G_MISC_CFG (CONTINUED)

Field Name	Bit	Access	Description	Default
TX_DATA_INV_ENA	2	R/W	Enable data inversion sent to Serializer 0: Disable 1: Enable	0x0
LANE_RST	0	R/W	Lane Reset 0: No reset 1: Reset (not self-clearing)	0x0

5.6.8 HSIO:SERDES1G_DIG_STATUS

Parent: [HSIO](#)

Instances: 1

Status register set for SERDES1G digital BIST and DFT functions.

TABLE 5-262: REGISTERS IN SERDES1G_DIG_STATUS

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
SERDES1G_DFT_STATUS	0x00000000	1	SERDES1G DFT Status	Page 437

5.6.8.1 HSIO:SERDES1G_DIG_STATUS:SERDES1G_DFT_STATUS

Parent: [HSIO:SERDES1G_DIG_STATUS](#)

Instances: 1

Status register of SERDES1G DFT functions

TABLE 5-263: FIELDS IN SERDES1G_DFT_STATUS

Field Name	Bit	Access	Description	Default
BIST_NOSYNC	2	R/O	BIST sync result 0: Synchronization successful 1: Synchronization on BIST data failed	0x0

5.6.9 HSIO:MCB_SERDES1G_CFG

Parent: [HSIO](#)

Instances: 1

All SERDES1G macros are accessed via the serial Macro Configuration Bus (MCB). Each macro is configured by one MCB slave. All MCB slaves are connected in a daisy-chain loop.

TABLE 5-264: REGISTERS IN MCB_SERDES1G_CFG

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
MCB_SERDES1G_ADDR_CFG	0x00000000	1	MCB SERDES1G Address Cfg	Page 437

5.6.9.1 HSIO:MCB_SERDES1G_CFG:MCB_SERDES1G_ADDR_CFG

Parent: [HSIO:MCB_SERDES1G_CFG](#)

Instances: 1

VSC7416-01

Configuration of SERDES1G MCB slaves to be accessed

TABLE 5-265: FIELDS IN MCB_SERDES1G_ADDR_CFG

Field Name	Bit	Access	Description	Default
SERDES1G_WR_ONE_SHOT	31	One-shot	Initiate a write access to marked SER-DES1G slaves 0: No write operation pending 1: Initiate write to slaves (kept 1 until write operation has finished)	0x0
SERDES1G_RD_ONE_SHOT	30	One-shot	Initiate a read access to marked SERDES1G slaves 0: No read operation pending (read op finished after bit has been set) 1: Initiate a read access (kept 1 until read operation has finished)	0x0
SERDES1G_ADDR	24:0	R/W	Activation vector for SERDES1G-Slaves, one-hot coded, each bit is related to one macro, e.g. bit 0 enables/disables access to macro no. 0. 0: Disable macro access via MCB 1: Enable macro access via MCB	0x1FFFFFFF

5.6.10 HSIO:SERDES6G_ANA_CFG

Parent: [HSIO](#)

Instances: 1

Configuration register set for SERDES6G (analog parts)

TABLE 5-266: REGISTERS IN SERDES6G_ANA_CFG

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
SERDES6G_DES_CFG	0x00000000	1	SERDES6G Deserializer Cfg	Page 438
SERDES6G_IB_CFG	0x00000004	1	SERDES6G Input Buffer Cfg	Page 439
SERDES6G_IB_CFG1	0x00000008	1	SERDES6G Input Buffer Cfg1	Page 440
SERDES6G_OB_CFG	0x0000000C	1	SERDES6G Output Buffer Cfg	Page 440
SERDES6G_OB_CFG1	0x00000010	1	SERDES6G Output Buffer Cfg1	Page 441
SERDES6G_SER_CFG	0x00000014	1	SERDES6G Serializer Cfg	Page 441
SERDES6G_COMMON_CFG	0x00000018	1	SERDES6G Common Cfg	Page 442
SERDES6G_PLL_CFG	0x0000001C	1	SERDES6G Pll Cfg	Page 443

5.6.10.1 HSIO:SERDES6G_ANA_CFG:SERDES6G_DES_CFG

Parent: [HSIO:SERDES6G_ANA_CFG](#)

Instances: 1

Configuration register for SERDES6G deserializer

TABLE 5-267: FIELDS IN SERDES6G_DES_CFG

Field Name	Bit	Access	Description	Default
DES_PHS_CTRL	16:13	R/W	Control of phase regulator logic. Bit 3 must always be set to 0. Optimal settings for bits 2:0 are 4 through 7; recommended setting is 6. 0: Disabled 1: Enabled with 99 ppm limit 2: Enabled with 202 ppm limit 3: Enabled with 485 ppm limit 4: Enabled if corresponding PCS is in sync with 50 ppm limit 5: Enabled if corresponding PCS is in sync with 99 ppm limit 6: Enabled if corresponding PCS is in sync with 202 ppm limit 7: Enabled if corresponding PCS is in sync with 485 ppm limit	0x6
DES_MBTR_CTRL	12:10	R/W	Des phase control for 180 degrees deadlock block mode of operation 0: Depending on density of input pattern 1: Active until PCS has synchronized 2: Depending on density of input pattern until PCS has synchronized 3: Never 4: Always 5-7: Reserved	0x2
DES_CPMD_SEL	9:8	R/W	DES phase control, main CP/MD select 0: Directly from DES 1: Through hysteresis stage from DES 2: From core 3: Disabled	0x0
DES_BW_HYST	7:5	R/W	Selection of time constant for integrative path of the CDR loop. 0: Reserved 1: Reserved 2: Reserved 3: Divide by 16 4: Divide by 32 5: Divide by 64 6: Divide by 128 7: Divide by 256	0x5
DES_BW_ANA	3:1	R/W	Bandwidth selection for proportional path of the CDR loop. 0: Reserved 1: Reserved 2: Divide by 4 3: Divide by 8 4: Divide by 16 5: Divide by 32 6: Divide by 64 7: Divide by 128	0x3

5.6.10.2 HSIO:SERDES6G_ANA_CFG:SERDES6G_IB_CFG

Parent: [HSIO:SERDES6G_ANA_CFG](#)

VSC7416-01

Instances: 1

Configuration register 0 for SERDES6G input buffer

TABLE 5-268: FIELDS IN SERDES6G_IB_CFG

Field Name	Bit	Access	Description	Default
RESERVED	30:28	R/W	Must be set to its default.	0x1
RESERVED	17:14	R/W	Must be set to its default.	0x4
RESERVED	13:10	R/W	Must be set to its default.	0xF
RESERVED	9:7	R/W	Must be set to its default.	0x5
IB_VBCOM	6:4	R/W	Level detection thresholds in steps of approximately 8mV: 0: 60mV, 7: 120mV	0x4
RESERVED	3:0	R/W	Must be set to its default.	0x2

5.6.10.3 HSIO:SERDES6G_ANA_CFG:SERDES6G_IB_CFG1

Parent: [HSIO:SERDES6G_ANA_CFG](#)

Instances: 1

Configuration register 1 for SERDES6G input buffer

TABLE 5-269: FIELDS IN SERDES6G_IB_CFG1

Field Name	Bit	Access	Description	Default
RESERVED	11:8	R/W	Must be set to its default.	0xF
IB_CTERM_ENA	5	R/W	Common mode termination 0: Disable 1: Enable	0x1
RESERVED	4	R/W	Must be set to its default.	0x1
IB_ENA_OFFSAC	3	R/W	Auto offset compensation for ac path 0: Disable 1: Enable	0x1
IB_ENA_OFFSDC	2	R/W	Auto offset compensation for dc path 0: Disable 1: Enable	0x1
IB_FX100_ENA	1	R/W	Increases timing constant for level detect circuit, must be used in FX100 mode 0: Normal speed 1: Slow speed (oversampling)	0x0

5.6.10.4 HSIO:SERDES6G_ANA_CFG:SERDES6G_OB_CFG

Parent: [HSIO:SERDES6G_ANA_CFG](#)

Instances: 1

Configuration register 0 for SERDES6G output buffer

TABLE 5-270: FIELDS IN SERDES6G_OB_CFG

Field Name	Bit	Access	Description	Default
OB_IDLE	31	R/W	PCIe support 1: idle - force to 0V differential 0: Normal mode	0x0

TABLE 5-270: FIELDS IN SERDES6G_OB_CFG (CONTINUED)

Field Name	Bit	Access	Description	Default
OB_ENA1V_MODE	30	R/W	Output buffer supply voltage 1: Set to nominal 1V 0: Set to higher voltage	0x0
OB_POL	29	R/W	Polarity of output signal 0: Normal 1: Inverted	0x1
OB_POST0	28:23	R/W	Coefficients for 1st Post Cursor (MSB is sign)	0x00
OB_POST1	22:18	R/W	Coefficients for 2nd Post Cursor (MSB is sign)	0x00
OB_PREC	17:13	R/W	Coefficients for Pre Cursor (MSB is sign)	0x00
OB_SR_H	8	R/W	Half the predriver speed, use for slew rate control 0: Disable - slew rate < 60 ps 1: Enable - slew rate > 60 ps	0x1
RESERVED	7:4	R/W	Must be set to its default.	0x1
OB_SR	3:0	R/W	Driver speed, fine adjustment of slew rate 30-60ps (if OB_SR_H = 0), 60-140ps (if OB_SR_H = 1)	0x7

5.6.10.5 HSIO:SERDES6G_ANA_CFG:SERDES6G_OB_CFG1

Parent: [HSIO:SERDES6G_ANA_CFG](#)

Instances: 1

Configuration register 1 for SERDES6G output buffer

TABLE 5-271: FIELDS IN SERDES6G_OB_CFG1

Field Name	Bit	Access	Description	Default
OB_ENA_CAS	8:6	R/W	Output skew, used for skew adjustment in SGMII mode	0x1
OB_LEV	5:0	R/W	Level of output amplitude 0: lowest level 63: highest level	0x00

5.6.10.6 HSIO:SERDES6G_ANA_CFG:SERDES6G_SER_CFG

Parent: [HSIO:SERDES6G_ANA_CFG](#)

Instances: 1

Configuration register for SERDES6G serializer

TABLE 5-272: FIELDS IN SERDES6G_SER_CFG

Field Name	Bit	Access	Description	Default
SER_ENHYS	3	R/W	Enable hysteresis for phase alignment 0: Disable hysteresis 1: Enable hysteresis	0x0
SER_EN_WIN	1	R/W	Enable window for phase alignment 0: Disable window 1: Enable window	0x0

TABLE 5-272: FIELDS IN SERDES6G_SER_CFG (CONTINUED)

Field Name	Bit	Access	Description	Default
SER_ENALI	0	R/W	Enable phase alignment 0: Disable phase alignment 1: Enable phase alignment	0x0

5.6.10.7 HSIO:SERDES6G_ANA_CFG:SERDES6G_COMMON_CFG

Parent: [HSIO:SERDES6G_ANA_CFG](#)

Instances: 1

Configuration register for common SERDES6G functions Note: When enabling the facility loop (ena_floop) also the phase alignment in the serializer has to be enabled and configured adequate.

TABLE 5-273: FIELDS IN SERDES6G_COMMON_CFG

Field Name	Bit	Access	Description	Default
SYS_RST	31	R/W	System reset (low active) 0: Apply reset (not self-clearing) 1: Reset released	0x0
SE_AUTO_SQUELCH_B_ENA	22	R/W	Enable auto-squelching for sync. ethernet bus B 0: Disable 1: Enable	0x0
SE_AUTO_SQUELCH_A_ENA	21	R/W	Enable auto-squelching for sync. ethernet bus A 0: Disable 1: Enable	0x0
RECO_SEL_B	20	R/W	Select recovered clock of this lane on sync. ethernet bus B 0: Lane not selected 1: Lane selected	0x0
RECO_SEL_A	19	R/W	Select recovered clock of this lane on sync. ethernet bus A 0: Lane not selected 1: Lane selected	0x0
ENA_LANE	18	R/W	Enable lane 0: Disable lane 1: Enable line	0x0
ENA_ELOOP	11	R/W	Enable equipment loop 0: Disable 1: Enable	0x0
ENA_FLOOP	10	R/W	Enable facility loop 0: Disable 1: Enable	0x0
HRATE	7	R/W	Enable half rate 0: Disable 1: Enable	0x0
QRATE	6	R/W	Enable quarter rate 0: Disable 1: Enable	0x1

TABLE 5-273: FIELDS IN SERDES6G_COMMON_CFG (CONTINUED)

Field Name	Bit	Access	Description	Default
IF_MODE	5:4	R/W	Interface mode 0: Reserved 1: 10-bit mode 2: Reserved 3: 20-bit mode	0x1

5.6.10.8 HSIO:SERDES6G_ANA_CFG:SERDES6G_PLL_CFG

Parent: [HSIO:SERDES6G_ANA_CFG](#)

Instances: 1

Configuration register for SERDES6G RCPLL

TABLE 5-274: FIELDS IN SERDES6G_PLL_CFG

Field Name	Bit	Access	Description	Default
PLL_DIV4	20	R/W	Enable div4 mode	0x0
PLL_ENA_ROT	18	R/W	Enable rotation	0x0
PLL_FSM_CTRL_DATA	15:8	R/W	Control data for FSM	0x3C
PLL_FSM_ENA	7	R/W	Enable FSM	0x0
PLL_ROT_DIR	2	R/W	Select rotation direction	0x0
PLL_ROT_FRQ	1	R/W	Select rotation frequency	0x0

5.6.11 HSIO:SERDES6G_DIG_CFG

Parent: [HSIO](#)

Instances: 1

Configuration register set for SERDES6G digital BIST and DFT functions.

TABLE 5-275: REGISTERS IN SERDES6G_DIG_CFG

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
SERDES6G_DIG_CFG	0x00000000	1	SERDES6G Digital Configuration register	Page 443
SERDES6G_MISC_CFG	0x0000001C	1	SERDES6G Misc Configuration	Page 444

5.6.11.1 HSIO:SERDES6G_DIG_CFG:SERDES6G_DIG_CFG

Parent: [HSIO:SERDES6G_DIG_CFG](#)

Instances: 1

Configuration register for SERDES6G digital functions

TABLE 5-276: FIELDS IN SERDES6G_DIG_CFG

Field Name	Bit	Access	Description	Default
SIGDET_AST	5:3	R/W	Signal detect assertion time 0: 0 us 1: 35 us 2: 70 us 3: 105 us 4: 140 us 5..7: reserved	0x0
SIGDET_DST	2:0	R/W	Signal detect de-assertion time 0: 0 us 1: 250 us 2: 350 us 3: 450 us 4: 550 us 5..7: reserved	0x0

5.6.11.2 HSIO:SERDES6G_DIG_CFG:SERDES6G_MISC_CFG

Parent: [HSIO:SERDES6G_DIG_CFG](#)

Instances: 1

Configuration register for miscellaneous functions

TABLE 5-277: FIELDS IN SERDES6G_MISC_CFG

Field Name	Bit	Access	Description	Default
SEL_RECO_CLK	17:16	R/W	Select recovered clock divider 0: No clock dividing 1: Divide clock by 5 2: Divide clock by 4 3: Reserved	0x0
DES_100FX_CPMD_ENA	8	R/W	Enable deserializer cp/md handling for 100fx mode 0: Disable 1: Enable	0x0
RX_LPI_MODE_ENA	5	R/W	Enable RX-Low-Power feature (Power control by LPI-FSM in connected PCS) 0: Disable 1: Enable	0x0
TX_LPI_MODE_ENA	4	R/W	Enable TX-Low-Power feature (Power control by LPI-FSM in connected PCS) 0: Disable 1: Enable	0x0
RX_DATA_INV_ENA	3	R/W	Enable data inversion received from Deserializer 0: Disable 1: Enable	0x0
TX_DATA_INV_ENA	2	R/W	Enable data inversion sent to Serializer 0: Disable 1: Enable	0x0
LANE_RST	0	R/W	Lane Reset 0: No reset 1: Reset (not self-clearing)	0x0

5.6.12 HSIO:MCB_SERDES6G_CFG

Parent: [HSIO](#)

Instances: 1

All SERDES6G macros are accessed via the serial Macro Configuration Bus (MCB). Each macro is configured by one MCB Slave. All MCB Slaves are connected in a daisy-chain loop.

TABLE 5-278: REGISTERS IN MCB_SERDES6G_CFG

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
MCB_SERDES6G_ADDR_CFG	0x00000000	1	MCB SERDES6G Address Cfg	Page 445

5.6.12.1 HSIO:MCB_SERDES6G_CFG:MCB_SERDES6G_ADDR_CFG

Parent: [HSIO:MCB_SERDES6G_CFG](#)

Instances: 1

Configuration of SERDES6G MCB Slaves to be accessed

TABLE 5-279: FIELDS IN MCB_SERDES6G_ADDR_CFG

Field Name	Bit	Access	Description	Default
SERDES6G_WR_ONE_SHOT	31	One-shot	Initiate a write access to marked SERDES6G Slaves 0: No write operation pending 1: Initiate write to slaves (kept 1 until write operation has finished)	0x0
SERDES6G_RD_ONE_SHOT	30	One-shot	Initiate a read access to marked SERDES6G Slaves 0: No read operation pending (read op finished after bit has been set) 1: Initiate a read access (kept 1 until read operation has finished)	0x0
SERDES6G_ADDR	15:0	R/W	Activation vector for SERDES6G-Slaves, one-hot coded, each bit is related to one macro, e.g. bit 0 enables/disables access to macro no. 0. 0: Disable macro access via MCB 1: Enable macro access via MCB	0xFFFF

5.7 ICPU_CFG

TABLE 5-280: REGISTER GROUPS IN ICPU_CFG

Register Group Name	Offset within Target	Instances and Address Spacing	Description	Details
CPU_SYSTEM_CTRL	0x00000000	1	CPU system configurations	Page 446
SPI_MST	0x0000003C	1	SPI boot master	Page 449
INTR	0x00000070	1	Interrupt controller	Page 450
TIMERS	0x000000DC	1	Timers	Page 458

TABLE 5-280: REGISTER GROUPS IN ICPU_CFG (CONTINUED)

Register Group Name	Offset within Target	Instances and Address Spacing	Description	Details
MEMCTRL	0x00000108	1	DDR2/3 memory controller	Page 461
TWI_DELAY	0x00000190	1	TWI hold time configuration	Page 471
TWI_SPIKE_FILTER	0x00000194	1	TWI spike filter configuration	Page 472
FDMA	0x00000198	1	Frame DMA	Page 472
PCIe	0x0000026C	1	PCIe endpoint	Page 483
MANUAL_XTRINJ	0x00004000	1	Manual extraction and injection via FDMA	Page 488

5.7.1 ICPU_CFG:CPU_SYSTEM_CTRL

Parent: [ICPU_CFG](#)

Instances: 1

TABLE 5-281: REGISTERS IN CPU_SYSTEM_CTRL

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
GPR	0x00000000	8 0x00000004	General purpose registers	Page 446
RESET	0x00000020	1	Reset settings	Page 446
GENERAL_CTRL	0x00000024	1	General control	Page 447
GENERAL_STAT	0x00000028	1	General status	Page 448

5.7.1.1 ICPU_CFG:CPU_SYSTEM_CTRL:GPR

Parent: [ICPU_CFG:CPU_SYSTEM_CTRL](#)

Instances: 8

TABLE 5-282: FIELDS IN GPR

Field Name	Bit	Access	Description	Default
GPR	31:0	R/W	General purpose 8 times 32-bit registers for software development and debug. Note: These registers are only reset by the device's reset input, i.e. they are not affected by soft reset!	0x00000000

5.7.1.2 ICPU_CFG:CPU_SYSTEM_CTRL:RESET

Parent: [ICPU_CFG:CPU_SYSTEM_CTRL](#)

Instances: 1

TABLE 5-283: FIELDS IN RESET

Field Name	Bit	Access	Description	Default
CORE_RST_CPU_ONLY	3	R/W	Set this field to enable VCore System reset protection. It is possible to protect the VCore System from soft-reset (issued via RESET:CORE_RST_FORCE) and watch-dog-timeout. When this field is set the aforementioned resets only reset the VCore CPU, not the VCore System. 0: Soft-reset and WDT-event and reset entire VCore 1: Soft-reset and WDT-event only reset the VCore CPU	0x0
CORE_RST_PROTECT	2	R/W	Set this field to enable VCore reset protection. It is possible to protect the entire VCore from chip-level soft-reset (issued via DEVCPU_GCB::SOFT_CHIP_RST.SOFT_CHIP_RST). Setting this field does not protect against hard-reset of the chip (by asserting the reset pin). 0: No reset protection 1: VCore is protected from chip-level-soft-reset	0x0
CORE_RST_FORCE	1	One-shot	Set this field to generate a soft reset for the VCore. This field will be cleared when the reset has taken effect. It is possible to protect the VCore system (everything else than the VCore CPU) from reset via RESET.CORE_RST_CPU_ONLY. 0: VCore is not reset 1: Initiate soft reset of the VCore	0x0
MEM_RST_FORCE	0	R/W	Clear this field to release the DDR2/3 memory controller from reset. 0: Memory controller is not reset 1: Memory controller is forced in reset	0x1

5.7.1.3 ICPU_CFG:CPU_SYSTEM_CTRL:GENERAL_CTRL

Parent: [ICPU_CFG:CPU_SYSTEM_CTRL](#)

Instances: 1

TABLE 5-284: FIELDS IN GENERAL_CTRL

Field Name	Bit	Access	Description	Default
CPU_MIPS_DIS	8	R/W	The default value if this field depends on strapping of the VCore. Set this field to disable the MIPS core (and instead enable 8051).	0x0
IF_MIIM_SLV_ENA	7	R/W	The default value if this field depends on strapping of the VCore. Set this field to enable the MIIM slave-interface on the GPIOs.	0x0

TABLE 5-284: FIELDS IN GENERAL_CTRL (CONTINUED)

Field Name	Bit	Access	Description	Default
IF_SI_MST_ENA	3	R/W	The default value if this field depends on strapping of the VCore. Set this field to force SI interfaces into master mode. This field must be set if the VCore is started manually. Note, while this field is set, it is not possible for an external CPU to access registers in the chip via SI.	0x0
CPU_BE_ENA	2	R/W	The default value if this field depends on strapping of the VCore. Set this field to force the MIPS VCore CPU into Big-Endian mode.	0x0
CPU_DIS	1	R/W	The default value if this field depends on strapping of the VCore. Clear this field to allow booting of the VCore CPU, while this field is set the VCore CPU is held in reset. 0: VCore CPU is allowed to boot 1: VCore CPU is forced in reset	0x1
BOOT_MODE_ENA	0	R/W	Use this field to change from Boot mode to Normal mode. In Boot mode, the reset vector of the VCore CPU maps to CS0 on the FLASH interface. When in Normal mode, this address maps instead to the DRAM Controller. The DRAM Controller must be operational before disabling Boot mode. After setting Boot mode, this register must be read back. The change in Boot mode takes effect during read. 0: The VCore memory map is in Normal mode. 1: The VCore memory map is in Boot mode.	0x1

5.7.1.4 ICPU_CFG:CPU_SYSTEM_CTRL:GENERAL_STAT

Parent: [ICPU_CFG:CPU_SYSTEM_CTRL](#)

Instances: 1

TABLE 5-285: FIELDS IN GENERAL_STAT

Field Name	Bit	Access	Description	Default
VCORE_CFG	7:4	R/O	Shows the value of the VCore_CFG strapping inputs.	0xF

TABLE 5-285: FIELDS IN GENERAL_STAT (CONTINUED)

Field Name	Bit	Access	Description	Default
REG_IF_ERR	3:1	R/O	Debug information for checking register read/write problems. This is a read-only field which can only be cleared by reset of the VCore System. 0: No errors detected. 1: Non-32bit access to VCore or SwC registers, access has been discarded, read returns 0x8888 or 0x88. 2: SwC registers not ready when accessed, access has been discarded, read returns 0x88888888. 3: SwC registers reported error, check DEVCPU_ORG::ERR_CNTS for error, read returns 0x88888888. 4: Unimplemented VCore register, access has been discarded, read returns 0x88888888.	0x0
CPU_SLEEP	0	R/O	This field is set if the VCore CPU has entered sleep mode.	0x0

5.7.2 ICPU_CFG:SPI_MST

Parent: [ICPU_CFG](#)

Instances: 1

TABLE 5-286: REGISTERS IN SPI_MST

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
SPI_MST_CFG	0x00000000	1	SPI boot master configuration	Page 449
SW_MODE	0x00000014	1	Manual control of the SPI interface	Page 450

5.7.2.1 ICPU_CFG:SPI_MST:SPI_MST_CFG

Parent: [ICPU_CFG:SPI_MST](#)

Instances: 1

TABLE 5-287: FIELDS IN SPI_MST_CFG

Field Name	Bit	Access	Description	Default
FAST_READ_ENA	10	R/W	The type of read-instruction that the SPI Controller generates for reads. 0: READ (slow read - Instruction code - 0x03) 1: FAST READ (fast read - Instruction code - 0x0B)	0x0
CS_DESELECT_TIME	9:5	R/W	The minimum number of SPI clock cycles for which the SPI chip select (SI_nEn) must be deasserted in between transfers. Typical value of this is 100 ns. Setting this field to 0 is illegal.	0x1F

TABLE 5-287: FIELDS IN SPI_MST_CFG (CONTINUED)

Field Name	Bit	Access	Description	Default
CLK_DIV	4:0	R/W	Controls the clock frequency for the SPI interface (SI_Clk). The clock frequency is VCore system clock divided by the value of this field. Setting this field to 0 or 1 value is illegal.	0x1F

5.7.2.2 ICPU_CFG:SPI_MST:SW_MODE

Parent: [ICPU_CFG:SPI_MST](#)

Instances: 1

Note: There are 4 chip selects in total, but only chip select 0 is mapped to IO-pin (SI_nEn). The rest of the SPI chip selects are available as alternate functions on GPIOs, these must be enabled in the GPIO controller before they can be controlled via this register.

TABLE 5-288: FIELDS IN SW_MODE

Field Name	Bit	Access	Description	Default
SW_PIN_CTRL_MODE	13	R/W	Set to enable software pin control mode (Bit banging), when set software has direct control of the SPI interface. This mode is used for writing into flash.	0x0
SW_SPI_SCK	12	R/W	Value to drive on SI_Clk output. This field is only used if SW_MODE.SW_PIN_CTRL_MODE is set.	0x0
SW_SPI_SCK_OE	11	R/W	Set to enable drive of SI_Clk output. This field is only used if SW_MODE.SW_PIN_CTRL_MODE is set.	0x0
SW_SPI_SDO	10	R/W	Value to drive on SI_DO output. This field is only used if SW_MODE.SW_PIN_CTRL_MODE is set.	0x0
SW_SPI_SDO_OE	9	R/W	Set to enable drive of SI_DO output. This field is only used if SW_MODE.SW_PIN_CTRL_MODE is set.	0x0
SW_SPI_CS	8:5	R/W	Value to drive on SI_nEn outputs, each bit in this field maps to a corresponding chip-select (0 though 3). This field is only used if SW_MODE.SW_PIN_CTRL_MODE is set. Note: Chip selects 1 though 3 are available as alternate GPIO functions.	0x0
SW_SPI_CS_OE	4:1	R/W	Set to enable drive of SI_nEn outputs, each bit in this field maps to a corresponding chip-select (0 though 3). This field is only used if SW_MODE.SW_PIN_CTRL_MODE is set. Note: Chip selects 1 though 3 are available as alternate GPIO functions.	0x0
SW_SPI_SDI	0	R/O	Current value of the SI_DI input.	0x0

5.7.3 ICPU_CFG:INTR

Parent: [ICPU_CFG](#)

Instances: 1

These are the VCore interrupt controller registers. The controller has 25 interrupt sources (inputs) which can each be mapped to one of 4 interrupt destinations (outputs). The controller has a dedicated block for handling 11 device interrupts; the handling of device interrupts can be considered a small interrupt controller located before the interrupt controller.

TABLE 5-289: REGISTERS IN INTR

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
INTR_RAW	0x00000000	1	Interrupt raw status	Page 451
INTR_TRIGGER	0x00000004	1	Interrupt trigger mode	Page 452
INTR_FORCE	0x00000008	1	Interrupt force sticky event	Page 452
INTR_STICKY	0x0000000C	1	Interrupt sticky status	Page 452
INTR_BYPASS	0x00000010	1	Interrupt bypass enable	Page 453
INTR_ENA	0x00000014	1	Interrupt enable	Page 453
INTR_ENA_CLR	0x00000018	1	Atomic clear of interrupt enable	Page 453
INTR_ENA_SET	0x0000001C	1	Atomic set of interrupt	Page 454
INTR_IDENT	0x00000020	1	Currently active interrupt sources	Page 454
DST_INTR_MAP	0x00000024	4 0x00000004	Mapping of source to destination interrupts	Page 454
DST_INTR_IDENT	0x00000034	4 0x00000004	Currently active interrupt sources per destination	Page 454
EXT_INTR_POL	0x00000044	1	External interrupt polarity	Page 455
EXT_INTR_DRV	0x00000048	1	External interrupt output drive mode	Page 455
EXT_INTR_DIR	0x0000004C	1	External interrupt output enable	Page 455
DEV_INTR_POL	0x00000050	1	Device interrupt polarity	Page 456
DEV_INTR_RAW	0x00000054	1	Device interrupt raw status	Page 456
DEV_INTR_TRIGGER	0x00000058	1	Device interrupt trigger mode	Page 456
DEV_INTR_STICKY	0x0000005C	1	Device interrupt sticky status	Page 457
DEV_INTR_BYPASS	0x00000060	1	Device interrupt bypass enable	Page 457
DEV_INTR_ENA	0x00000064	1	Device interrupt enable	Page 458
DEV_INTR_IDENT	0x00000068	1	Currently active device interrupts	Page 458

5.7.3.1 ICPU_CFG:INTR:INTR_RAW

Parent: [ICPU_CFG:INTR](#)

Instances: 1

TABLE 5-290: FIELDS IN INTR_RAW

Field Name	Bit	Access	Description	Default
INTR_RAW	24:0	R/O	Shows the current value of individual interrupt sources. All sources are active high (the external interrupts has been corrected for polarity as configured in ICPU_CFG::EXT_INTR_POL).	0x0000000

5.7.3.2 ICPU_CFG:INTR:INTR_TRIGGER

Parent: [ICPU_CFG:INTR](#)

Instances: 1

TABLE 5-291: FIELDS IN INTR_TRIGGER

Field Name	Bit	Access	Description	Default
INTR_TRIGGER	24:0	R/W	Configure trigger mode of individual interrupt sources. The trigger mode determines how the value of the ICPU_CFG::INTR_RAW register is transferred to the ICPU_CFG::INTR_STICKY register. For a level-triggered interrupts ICPU_CFG::INTR_STICKY is set for as long as the corresponding bit in ICPU_CFG::INTR_RAW is high - i.e. is not possible to clear a bit in ICPU_CFG::INTR_STICKY until the corresponding ICPU_CFG::INTR_RAW is zero. For a level triggered interrupt bits in ICPU_CFG::INTR_STICKY is set when the corresponding bit in ICPU_CFG::INTR_RAW changes value. 0: Interrupt is level-activated 1: Interrupt is edge-triggered	0x0000000

5.7.3.3 ICPU_CFG:INTR:INTR_FORCE

Parent: [ICPU_CFG:INTR](#)

Instances: 1

TABLE 5-292: FIELDS IN INTR_FORCE

Field Name	Bit	Access	Description	Default
INTR_FORCE	24:0	One-shot	Set to force corresponding ICPU_CFG::INTR_STICKY bits. This field may be useful during development of software interrupt handler functions.	0x0000000

5.7.3.4 ICPU_CFG:INTR:INTR_STICKY

Parent: [ICPU_CFG:INTR](#)

Instances: 1

TABLE 5-293: FIELDS IN INTR_STICKY

Field Name	Bit	Access	Description	Default
INTR_STICKY	24:0	Sticky	This register is set based on source interrupt events or by debug-force. See ICPU_CFG::INTR_TRIGGER and ICPU_CFG::INTR_FORCE for more information. Bits in this register remains set until cleared by software.	0x0000000

5.7.3.5 ICPU_CFG:INTR:INTR_BYPASS

Parent: [ICPU_CFG:INTR](#)

Instances: 1

TABLE 5-294: FIELDS IN INTR_BYPASS

Field Name	Bit	Access	Description	Default
INTR_BYPASS	24:0	R/W	This register allows bypass of ICPU_CFG::INTR_STICKY for individual interrupt sources. When an interrupt source is in bypass mode then ICPU_CFG::INTR_RAW is used instead of ICPU_CFG::INTR_STICKY. Note: Enabling bypass does not make sense for all interrupt sources. It should only be used when the corresponding interrupt is sticky at the source. For example manual extraction data available interrupts can be configured to bypass, because the interrupt will remain asserted until the available data has been extracted. Note: The device interrupt is bypassed per default, "stickyness" is already implemented by ICPU_CFG::DEV_INTR_STICKY.	0x0000001

5.7.3.6 ICPU_CFG:INTR:INTR_ENA

Parent: [ICPU_CFG:INTR](#)

Instances: 1

TABLE 5-295: FIELDS IN INTR_ENA

Field Name	Bit	Access	Description	Default
INTR_ENA	24:0	R/W	Set to enable propagation of individual interrupt sources to destinations. Atomic access to this register (needed in a multithreaded system) can be implemented by the ICPU_CFG::INTR_ENA_CLR and ICPU_CFG::INTR_ENA_SET registers.	0x0000000

5.7.3.7 ICPU_CFG:INTR:INTR_ENA_CLR

Parent: [ICPU_CFG:INTR](#)

Instances: 1

TABLE 5-296: FIELDS IN INTR_ENA_CLR

Field Name	Bit	Access	Description	Default
INTR_ENA_CLR	24:0	One-shot	Set bit(s) in this register to clear the corresponding bits in ICPU_CFG::INTR_ENA. This register can be used for atomic access to ICPU_CFG::INTR_ENA register.	0x0000000

5.7.3.8 ICPU_CFG:INTR:INTR_ENA_SET

Parent: [ICPU_CFG:INTR](#)

Instances: 1

TABLE 5-297: FIELDS IN INTR_ENA_SET

Field Name	Bit	Access	Description	Default
INTR_ENA_SET	24:0	One-shot	Set bit(s) in this register to set the corresponding bits in ICPU_CFG::INTR_ENA. This register can be used for atomic access to ICPU_CFG::INTR_ENA register.	0x0000000

5.7.3.9 ICPU_CFG:INTR:INTR_IDENT

Parent: [ICPU_CFG:INTR](#)

Instances: 1

TABLE 5-298: FIELDS IN INTR_IDENT

Field Name	Bit	Access	Description	Default
INTR_IDENT	24:0	R/O	Shows the currently active interrupt sources. For interrupt sources that are not bypassed this register is a result of AND'ing ICPU_CFG::INTR_STICKY with ICPU_CFG::INTR_ENA.	0x0000000

5.7.3.10 ICPU_CFG:INTR:DST_INTR_MAP

Parent: [ICPU_CFG:INTR](#)

Instances: 4

Replicated per destination interrupt.

TABLE 5-299: FIELDS IN DST_INTR_MAP

Field Name	Bit	Access	Description	Default
DST_INTR_MAP	24:0	R/W	Set to enable mapping of individual interrupt sources to interrupt destinations. This register is replicated once for each destination interrupt.	0x0000000

5.7.3.11 ICPU_CFG:INTR:DST_INTR_IDENT

Parent: [ICPU_CFG:INTR](#)

Instances: 4

Replicated per destination interrupt.

TABLE 5-300: FIELDS IN DST_INTR_IDENT

Field Name	Bit	Access	Description	Default
DST_INTR_IDENT	24:0	R/O	Shows the currently active interrupt sources per destination interrupt. The contents of this register is equal to ICPU_CFG::INTR_IDENT AND'ed with the corresponding ICPU_CFG::DST_INTR_MAP. If any bit is set in this register the corresponding destination interrupt is asserted.	0x0000000

5.7.3.12 ICPU_CFG:INTR:EXT_INTR_POL

Parent: [ICPU_CFG:INTR](#)

Instances: 1

TABLE 5-301: FIELDS IN EXT_INTR_POL

Field Name	Bit	Access	Description	Default
EXT_INTR_POL	1:0	R/W	Set individual bits in this register to configure polarity of the corresponding external interrupt. This setting is applies to both to input and output mode. The polarity is corrected at the edge of the chip, internally interrupts are always active-high. 0: External interrupt is active low 1: External interrupt is active high	0x0

5.7.3.13 ICPU_CFG:INTR:EXT_INTR_DRV

Parent: [ICPU_CFG:INTR](#)

Instances: 1

TABLE 5-302: FIELDS IN EXT_INTR_DRV

Field Name	Bit	Access	Description	Default
EXT_INTR_DRV	1:0	R/W	This register configures drive mode of the corresponding external interrupt. This setting only applies to external interrupts that are configured as output, see ICPU_CFG::EXT_INTR_DIR.EXT_INTR_DIR for more information. 0: Only drive external interrupt output when asserted (tristate when inactive) 1: External interrupt output is always driven	0x0

5.7.3.14 ICPU_CFG:INTR:EXT_INTR_DIR

Parent: [ICPU_CFG:INTR](#)

Instances: 1

TABLE 5-303: FIELDS IN EXT_INTR_DIR

Field Name	Bit	Access	Description	Default
EXT_INTR_DIR	1:0	R/W	This register configures the direction of external interrupts. When an external interrupt is configured for input it can be used as an interrupt source in the interrupt controller (registers ICPU_CFG:INTR_*). When configured as an output it can be used as an interrupt destination (registers ICPU_CFG::DST_*). 0: External interrupt is input 1: External interrupt is output	0x0

5.7.3.15 ICPU_CFG:INTR:DEV_INTR_POL

Parent: [ICPU_CFG:INTR](#)

Instances: 1

Each bit in this register corresponds a device-interrupt, i.e. bit 3 corresponds to device-interrupt number 3.

TABLE 5-304: FIELDS IN DEV_INTR_POL

Field Name	Bit	Access	Description	Default
DEV_INTR_POL	10:0	R/W	Set individual bits in this register to configure polarity of the corresponding device interrupt. 0: External interrupt is active low 1: External interrupt is active high	0x7FF

5.7.3.16 ICPU_CFG:INTR:DEV_INTR_RAW

Parent: [ICPU_CFG:INTR](#)

Instances: 1

TABLE 5-305: FIELDS IN DEV_INTR_RAW

Field Name	Bit	Access	Description	Default
DEV_INTR_RAW	10:0	R/O	Shows the current value of individual device interrupt sources. All sources are active high (sources have been corrected for polarity as configured in ICPU_CFG::DEV_INTR_POL).	0x000

5.7.3.17 ICPU_CFG:INTR:DEV_INTR_TRIGGER

Parent: [ICPU_CFG:INTR](#)

Instances: 1

TABLE 5-306: FIELDS IN DEV_INTR_TRIGGER

Field Name	Bit	Access	Description	Default
DEV_INTR_TRIGGER	10:0	R/W	Configure trigger mode of individual device interrupt sources. The trigger mode determines how the value of the ICPU_CFG::DEV_INTR_RAW register is transferred to the ICPU_CFG::DEV_INTR_STICKY register. For a level-triggered interrupts ICPU_CFG::DEV_INTR_STICKY is set for as long as the corresponding bit in ICPU_CFG::DEV_INTR_RAW is high - i.e. is not possible to clear a bit in ICPU_CFG::DEV_INTR_STICKY until the corresponding ICPU_CFG::DEV_INTR_RAW is zero. For a level triggered interrupt bits in ICPU_CFG::DEV_INTR_STICKY is set when the corresponding bit in ICPU_CFG::DEV_INTR_RAW changes value. 0: Interrupt is level-activated 1: Interrupt is edge-triggered	0x000

5.7.3.18 ICPU_CFG:INTR:DEV_INTR_STICKY

Parent: [ICPU_CFG:INTR](#)

Instances: 1

TABLE 5-307: FIELDS IN DEV_INTR_STICKY

Field Name	Bit	Access	Description	Default
DEV_INTR_STICKY	10:0	Sticky	This register is set based on device interrupt source events. See ICPU_CFG::DEV_INTR_TRIGGER for more information. Bits in this register remains set until cleared by software.	0x000

5.7.3.19 ICPU_CFG:INTR:DEV_INTR_BYPASS

Parent: [ICPU_CFG:INTR](#)

Instances: 1

TABLE 5-308: FIELDS IN DEV_INTR_BYPASS

Field Name	Bit	Access	Description	Default
DEV_INTR_BYPASS	10:0	R/W	This register allows bypass of ICPU_CFG::DEV_INTR_STICKY for individual device interrupt sources. When a device interrupt source is in bypass mode then ICPU_CFG::DEV_INTR_RAW is used instead of ICPU_CFG::DEV_INTR_STICKY. Please see note on bypass in ICPU_CFG::INTR_BYPASS.	0x000

5.7.3.20 ICPU_CFG:INTR:DEV_INTR_ENA

Parent: [ICPU_CFG:INTR](#)

Instances: 1

TABLE 5-309: FIELDS IN DEV_INTR_ENA

Field Name	Bit	Access	Description	Default
DEV_INTR_ENA	10:0	R/W	Set to enable propagation of individual device interrupt sources to the main interrupt controller.	0x000

5.7.3.21 ICPU_CFG:INTR:DEV_INTR_IDENT

Parent: [ICPU_CFG:INTR](#)

Instances: 1

TABLE 5-310: FIELDS IN DEV_INTR_IDENT

Field Name	Bit	Access	Description	Default
DEV_INTR_IDENT	10:0	R/O	Shows the currently active interrupt sources. For interrupt sources that are not bypassed this register is a result of AND'ing ICPU_CFG::DEV_INTR_STICKY with ICPU_CFG::DEV_INTR_ENA.	0x000

5.7.4 ICPU_CFG:TIMERS

Parent: [ICPU_CFG](#)

Instances: 1

TABLE 5-311: REGISTERS IN TIMERS

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
WDT	0x00000000	1	Watchdog timer	Page 458
TIMER_TICK_DIV	0x00000004	1	Timer tick divider	Page 459
TIMER_VALUE	0x00000008	3 0x00000004	Timer value	Page 459
TIMER_RELOAD_VALUE	0x00000014	3 0x00000004	Timer reload value	Page 460
TIMER_CTRL	0x00000020	3 0x00000004	Timer control	Page 460

5.7.4.1 ICPU_CFG:TIMERS:WDT

Parent: [ICPU_CFG:TIMERS](#)

Instances: 1

TABLE 5-312: FIELDS IN WDT

Field Name	Bit	Access	Description	Default
WDT_STATUS	9	R/O	Shows whether the last reset was caused by a watchdog timer reset. This field is updated during reset, therefore it is always valid. 0: Reset was not caused by WDT 1: Reset was caused by WDT timeout	0x0
WDT_ENABLE	8	R/W	Use this field to enable or disable the watchdog timer. When the WDT is enabled, it causes a reset after 2 seconds if it is not periodically reset. This field is only read by the WDT after a successful lock sequence (see ICPU_CFG::WDT.WDT_LOCK). 0: WDT is disabled 1: WDT is enabled	0x0
WDT_LOCK	7:0	R/W	Use this field to configure and reset the WDT. When writing 0xBE to this field immediately followed by writing 0xEF, the WDT resets and configurations are read from this register (as provided when the 0xEF is written). When the WDT is enabled, writing any value other than 0xBE or 0xEF after 0xBE is written, causes a WDT reset as if the timer had run out.	0x00

5.7.4.2 ICPU_CFG:TIMERS:TIMER_TICK_DIV

Parent: [ICPU_CFG:TIMERS](#)

Instances: 1

TABLE 5-313: FIELDS IN TIMER_TICK_DIV

Field Name	Bit	Access	Description	Default
TIMER_TICK_DIV	17:0	R/W	The timer tick generator runs from a 156.25MHz base clock. By default, the divider value generates a timer tick every 100 us (10 KHz). The timer tick is used for all of the timers (except the WDT). This field must not be set to generate a timer tick of less than 0.1 us (higher than 10 MHz). If this field is changed, it may take up to 2 ms before the timers are running stable at the new frequency. The timer tick frequency is: 156.25MHz/(TIMER_TICK_DIV+1).	0x03D08

5.7.4.3 ICPU_CFG:TIMERS:TIMER_VALUE

Parent: [ICPU_CFG:TIMERS](#)

Instances: 3

TABLE 5-314: FIELDS IN TIMER_VALUE

Field Name	Bit	Access	Description	Default
TIMER_VAL	31:0	R/W	The current value of the timer. When enabled via <code>TIMER_CTRL.TIMER_ENA</code> the timer decrements at every timer tick (see <code>TIMER_TICK_DIV</code> for more info on timer tick frequency). When the timer has reached 0, and a timer-tick is received, then an interrupt is generated. For example; If a periodic interrupt is needed every 1ms, and the timer tick is generated every 100us then the <code>TIMER_VALUE</code> (and <code>TIMER_RELOAD_VALUE</code>) must be configured to 9. By default the timer will reload from the <code>TIMER_RELOAD_VALUE</code> when interrupt is generated, and then continue decrementing from the reloaded value. It is possible to make the timer stop after generating interrupt by setting <code>TIMER_CTRL.ONE_SHOT</code> .	0x00000000

5.7.4.4 ICPU_CFG:TIMERS:TIMER_RELOAD_VALUE

Parent: [ICPU_CFG:TIMERS](#)

Instances: 3

TABLE 5-315: FIELDS IN TIMER_RELOAD_VALUE

Field Name	Bit	Access	Description	Default
RELOAD_VAL	31:0	R/W	The contents of this field are loaded into the corresponding timer (<code>TIMER_VALUE</code>) when it wraps (decrements a zero).	0x00000000

5.7.4.5 ICPU_CFG:TIMERS:TIMER_CTRL

Parent: [ICPU_CFG:TIMERS](#)

Instances: 3

TABLE 5-316: FIELDS IN TIMER_CTRL

Field Name	Bit	Access	Description	Default
ONE_SHOT_ENA	2	R/W	When set the timer will automatically disable itself after it has generated interrupt.	0x0
TIMER_ENA	1	R/W	When enabled, the corresponding timer decrements at each timer-tick. If <code>TIMER_CTRL.ONE_SHOT_ENA</code> is set this field is cleared when the timer reach 0 and interrupt is generated. 0: Timer is disabled 1: Timer is enabled	0x0

TABLE 5-316: FIELDS IN TIMER_CTRL (CONTINUED)

Field Name	Bit	Access	Description	Default
FORCE_RELOAD	0	One-shot	Set this field to force the reload of the timer, this will set the TIMER_VALUE to TIMER_RELOAD_VALUE for the corresponding timer. This field can be set at the same time as enabling the counter, in that case the counter will be reloaded and then enabled for counting.	0x0

5.7.5 ICPU_CFG:MEMCTRL

Parent: [ICPU_CFG](#)

Instances: 1

TABLE 5-317: REGISTERS IN MEMCTRL

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
MEMCTRL_CTRL	0x00000000	1	DDR2/3 control	Page 461
MEMCTRL_CFG	0x00000004	1	DDR2/3 configuration	Page 462
MEMCTRL_STAT	0x00000008	1	DDR2/3 status	Page 464
MEMCTRL_REF_PERIOD	0x0000000C	1	Refresh configuration	Page 464
MEMCTRL_ZQCAL	0x00000010	1	DDR3 ZQ-calibration	Page 465
MEMCTRL_TIMING0	0x00000014	1	Timing configuration	Page 465
MEMCTRL_TIMING1	0x00000018	1	Timing configuration	Page 466
MEMCTRL_TIMING2	0x0000001C	1	Timing configuration	Page 467
MEMCTRL_TIMING3	0x00000020	1	Timing configuration	Page 467
MEMCTRL_MR0_VAL	0x00000028	1	Mode register 0 value	Page 468
MEMCTRL_MR1_VAL	0x0000002C	1	Mode register 1 value	Page 468
MEMCTRL_MR2_VAL	0x00000030	1	Mode register 2 value	Page 468
MEMCTRL_MR3_VAL	0x00000034	1	Mode register 3 value	Page 469
MEMCTRL_TERMRES_CTRL	0x00000038	1	On-die-termination configuration	Page 469
MEMCTRL_DQS_DLY	0x00000040	2 0x00000004	DQS window configuration	Page 469
MEMPHY_CFG	0x00000050	1	SSTL configuration	Page 470
MEMPHY_ZCAL	0x00000078	1	SSTL drive-strength calibration	Page 470

5.7.5.1 ICPU_CFG:MEMCTRL:MEMCTRL_CTRL

Parent: [ICPU_CFG:MEMCTRL](#)

Instances: 1

TABLE 5-318: FIELDS IN MEMCTRL_CTRL

Field Name	Bit	Access	Description	Default
PWR_DOWN	3	R/W	Set this field to force the memory module into low power self refresh mode. The ICPU_CFG::MEMCTRL_STAT.PWR_DOWN_ACK is set when the controller has executed the command. Clear this field to bring the controller back to normal operation. Note: Before using power-down the ICPU_CFG::MEMCTRL_TIMING2.INIT_DLY must be reconfigured, see field description for more information.	0x0
MDSET	2	One-shot	Set this field to do memory register write command. The register to write is defined by ICPU_CFG::MEMCTRL_MR0_VAL[15:14], the data to write is defined by ICPU_CFG::MEMCTRL_MR0_VAL[15:0].	0x0
STALL_REF_ENA	1	R/W	Set this field to temporarily give software unhindered access to memory; two things changes in the controller: a) Refreshes are postponed until ICPU_CFG::MEMCTRL_REF_PERIOD.MAX_PEND_REF is exceeded. b) Only a single refresh will be issued when exceeding ICPU_CFG::MEMCTRL_REF_PERIOD.MAX_PEND_REF. When this field is NOT set (the normal case) the memory controller try to do refreshes during controller idle periods, also once refresh has started - all pending refreshes will be performed. Note: Interrupt routines and other high-priority tasks can use this field to ensure uninterrupted access to the memory, such routines must clear this field when memory access is no longer critical.	0x0
INITIALIZE	0	One-shot	Set this field to force the memory controller to initialize the SDRAM. This field is automatically cleared after the controller has started initialization of the SDRAM. Wait for ICPU_CFG::MEMCTRL_STAT.INIT_DONE before performing read/write operations on the controller. Note: The ICPU_CFG::MEMCTRL_CFG, ICPU_CFG::MEMCTRL_REF_PERIOD, ICPU_CFG::MEMCTRL_TIMING*, and ICPU_CFG::MEMCTRL_MR* registers must be configured appropriately before setting this field.	0x0

5.7.5.2 ICPU_CFG::MEMCTRL::MEMCTRL_CFG

Parent: [ICPU_CFG::MEMCTRL](#)

Instances: 1

TABLE 5-319: FIELDS IN MEMCTRL_CFG

Field Name	Bit	Access	Description	Default
DDR_512MBYTE_PLUS	16	R/W	Set this field to enable support for more than 512MByte SDRAM.	0x0
DDR_ECC_ERR_ENA	15	R/W	Set this field to enable propagation of ECC errors to SBA by generating bus-error event when ECC error is detected. This field may not be set unless ICPU_CFG::MEMCTRL_CFG.DDR_ECC_ENA is also set.	0x0
DDR_ECC_COR_ENA	14	R/W	Set this field to enable propagation of ECC corrections to SBA by generating bus-error event when ECC correction is detected. This field may not be set unless ICPU_CFG::MEMCTRL_CFG.DDR_ECC_ENA is also set.	0x0
DDR_ECC_ENA	13	R/W	This field enables ECC mode of the SDRAM controller. This field may only be set when ICPU_CFG::MEMCTRL_CFG.DDR_WIDTH is configured for 16bit. In ECC mode byte lane 1 is used for ECC information and byte lane 0 is used for data. When using ECC mode, then the amount of random-access memory available to an application is half of the physically attached memory.	0x0
DDR_WIDTH	12	R/W	This field configures the interface width of the SDRAM controller. If 8bit is selected then byte lane 0 must be populated by external SDRAM memory. 0: 8bit 1: 16bit	0x0
DDR_MODE	11	R/W	This field configures the operating mode of the SDRAM controller. 0: DDR2 1: DDR3	0x0
BURST_SIZE	10	R/W	The number of data-bytes that is transmitted during one burst (of the defined burst length: ICPU_CFG::MEMCTRL_CFG.BURST_LEN). 0: 8 data-bytes per burst. 1: 16 data-bytes per burst.	0x0
BURST_LEN	9	R/W	The burst size that is used by the SDRAM controller. The SDRAM must be configured with the corresponding burst size (through the ICPU_CFG::MEMCTRL_MR0_VAL register.) Note: The number of data-bytes that is transmitted during one burst must be encoded in the ICPU_CFG::MEMCTRL_CFG.BURST_SIZE field. 0 : BURST4 1 : BURST8	0x0

TABLE 5-319: FIELDS IN MEMCTRL_CFG (CONTINUED)

Field Name	Bit	Access	Description	Default
BANK_CNT	8	R/W	Number of banks in the SDRAM configuration being used. 0 : 4 banks 1 : 8 banks	0x0
MSB_ROW_ADDR	7:4	R/W	Set to 1 less than the number of row address bits for the SDRAM configuration in use.	0x0
MSB_COL_ADDR	3:0	R/W	Set to 1 less than the number of column address bits for the SDRAM configuration in use. For example; for a memory that is using column addresses {A11, A9-A0} this field must be set to 10.	0x0

5.7.5.3 ICPU_CFG:MEMCTRL:MEMCTRL_STAT

Parent: [ICPU_CFG:MEMCTRL](#)

Instances: 1

TABLE 5-320: FIELDS IN MEMCTRL_STAT

Field Name	Bit	Access	Description	Default
RDATA_MASKED	5	Sticky	Set if the controller has masked a data-set towards the SBA (too many data words received from SDRAM). This field may be set during training. If this event occurs during normal operation then the DQS window may be wrongly configured.	0x0
RDATA_DUMMY	4	Sticky	Set if the controller has inserted missing data-set towards the SBA (too few data words received from SDRAM). This field may be set during training. If this event occurs during normal operation then the DQS window may be wrongly configured.	0x0
RDATA_ECC_ERR	3	Sticky	Set if the controller is enabled for ECC and an uncorrectable error has been detected.	0x0
RDATA_ECC_COR	2	Sticky	Set if the controller is enabled for ECC and a dataset has been corrected (a correctable error has been corrected).	0x0
PWR_DOWN_ACK	1	R/O	Is set once the memory controller has put the SDRAM in low power self refresh mode (result of setting ICPU_CFG::MEMCTRL_CFG.PWD_DOWN). When this field is set SDRAM interface signal may turned off (all signals except the CKE) may be left floating.	0x0
INIT_DONE	0	R/O	This field is set after initialization of the SDRAM is done. When this field is set then read/write operations can be performed.	0x0

5.7.5.4 ICPU_CFG:MEMCTRL:MEMCTRL_REF_PERIOD

Parent: [ICPU_CFG:MEMCTRL](#)

Instances: 1

TABLE 5-321: FIELDS IN MEMCTRL_REF_PERIOD

Field Name	Bit	Access	Description	Default
MAX_PEND_REF	19:16	R/W	Maximum number of refreshes that are allowed to be outstanding at any time. If the number of outstanding refreshes exceeds this value, the memory controller will stop data accesses, and issue refreshes. If no outstanding refreshes is allowed then set this field to 0.	0x8
REF_PERIOD	15:0	R/W	Refresh interval of the SDRAM expressed in number of clock cycles. This value is calculated by dividing the average periodic refresh interval (tREFI) by the clock period.	0x0100

5.7.5.5 ICPU_CFG:MEMCTRL:MEMCTRL_ZQCAL

Parent: [ICPU_CFG:MEMCTRL](#)

Instances: 1

TABLE 5-322: FIELDS IN MEMCTRL_ZQCAL

Field Name	Bit	Access	Description	Default
ZQCAL_LONG	1	One-shot	Set this field to issue long ZQ calibration command. This field is cleared when calibration has been performed.	0x0
ZQCAL_SHORT	0	One-shot	Set this field to issue short ZQ calibration command. This field is cleared when calibration has been performed.	0x0

5.7.5.6 ICPU_CFG:MEMCTRL:MEMCTRL_TIMING0

Parent: [ICPU_CFG:MEMCTRL](#)

Instances: 1

The Following parameters are needed for configuration of the ICPU_CFG::MEMCTRL_TIMING* registers. All asynchronous delays should be converted (round up) to the corresponding number of DDR controller clock cycles. Note that DDR modules may specify a minimum number of clock cycles for some parameters:

RL=CL

BL=2 for burst4, 4 for burst8

Additional for DDR2 memories:

WL=RL-1

MD=tMRD

ID=400ns

SD=tXSRD

OW=WL-2

OR=RL-3

RP=tRP for 4-bank modules, tRPA for 8-bank modules

FAW=1 for 4-bank modules, tFAW for 8-bank modules

Additional for DDR3 memories:

VSC7416-01

WL=CWL
MD=tMOD
ID=tXPR
SD=tDLLK
OW=2
OR=2
RP=tRP
FAW=tFAW

TABLE 5-323: FIELDS IN MEMCTRL_TIMING0

Field Name	Bit	Access	Description	Default
RD_TO_WR_DLY	31:28	R/W	Delay from read to write on same chip select. RL+BL+1-WL	0x4
WR_CS_CHANGE_DLY	27:24	R/W	Delay from write to write on different chip selects. If ICPU_CFG::MEMCTRL_TERMRES_CTRL.ODT_WR_ENA==9 then increase this field by 2 (to allow ODT handover between destination devices). BL-1	0x3
RD_CS_CHANGE_DLY	23:20	R/W	Delay from read to read on different chip selects. BL	0x2
RAS_TO_PRECH_DLY	19:16	R/W	Delay from RAS to precharge. tRAS_min-1	0x0
WR_TO_PRECH_DLY	15:12	R/W	Delay from write to precharge. WL+BL+tWR-1	0x0
RD_TO_PRECH_DLY	11:8	R/W	Delay from read to precharge. BL-1	0x0
WR_DATA_XFR_DLY	7:4	R/W	Delay from write command to data. WL-1	0x0
RD_DATA_XFR_DLY	3:0	R/W	Delay from read command to data. Important; this delay is further increased by the DQS delay logic - see ICPU_CFG::MEMCTRL_DQS_DLY for more information. RL-3	0x0

5.7.5.7 ICPU_CFG:MEMCTRL:MEMCTRL_TIMING1

Parent: [ICPU_CFG:MEMCTRL](#)

Instances: 1

See ICPU_CFG::MEMCTRL_TIMING0 for description of important parameters.

TABLE 5-324: FIELDS IN MEMCTRL_TIMING1

Field Name	Bit	Access	Description	Default
RAS_TO_RAS_SAME_BANK_DLY	31:24	R/W	Delay from RAS to RAS within same bank. tRC-1	0x00
BANK8_FAW_DLY	23:16	R/W	Four bank activate period. FAW-1	0x00

TABLE 5-324: FIELDS IN MEMCTRL_TIMING1 (CONTINUED)

Field Name	Bit	Access	Description	Default
PRECH_TO_RAS_DLY	15:12	R/W	Delay from precharge to RAS. tRP-1	0x0
RAS_TO_RAS_DLY	11:8	R/W	Delay from RAS to RAS. tRRD-1	0x0
RAS_TO_CAS_DLY	7:4	R/W	Delay from RAS to CAS. tRCD-1	0x0
WR_TO_RD_DLY	3:0	R/W	Delay from write to read. WL+BL+tWTR-1	0x0

5.7.5.8 ICPU_CFG::MEMCTRL::MEMCTRL_TIMING2

Parent: [ICPU_CFG::MEMCTRL](#)

Instances: 1

See ICPU_CFG::MEMCTRL_TIMING0 for description of important parameters.

TABLE 5-325: FIELDS IN MEMCTRL_TIMING2

Field Name	Bit	Access	Description	Default
PRECH_ALL_DLY	31:28	R/W	Delay after precharge all. RP-1	0x0
MDSET_DLY	27:24	R/W	Delay after register-write. MD-1	0x0
REF_DLY	23:16	R/W	Delay after refresh. tRFC-1	0x00
INIT_DLY	15:0	R/W	Delay for initialization (see ICPU_CFG::MEMCTRL_CTRL.INITIALIZE). Before initialization: ID-1 After initialization: SD-1	0x0000

5.7.5.9 ICPU_CFG::MEMCTRL::MEMCTRL_TIMING3

Parent: [ICPU_CFG::MEMCTRL](#)

Instances: 1

See ICPU_CFG::MEMCTRL_TIMING0 for description of important parameters.

TABLE 5-326: FIELDS IN MEMCTRL_TIMING3

Field Name	Bit	Access	Description	Default
ODT_RD_DLY	15:12	R/W	Delay from read to ODT assert. External ODT assert for read commands is enabled by ICPU_CFG::MEMCTRL_TERM- RES_CTRL.ODT_RD_ENA. OR-1	0x0
ODT_WR_DLY	11:8	R/W	Delay from write to ODT assert. External ODT assert for write commands is enabled by ICPU_CFG::MEMCTRL_TERM- RES_CTRL.ODT_WR_ENA. OW-1	0x0

TABLE 5-326: FIELDS IN MEMCTRL_TIMING3 (CONTINUED)

Field Name	Bit	Access	Description	Default
LOCAL_ODT_RD_DLY	7:4	R/W	Delay from read to local read-termination activate. Important; this delay is further increased by the DQS delay logic - see ICPU_CFG::MEMCTRL_DQS_DLY for more information. Local read-termination is enabled by ICPU_CFG::MEMCTRL_TERMRES_CTRL.LOCAL_ODT_RD_ENA. RL-3	0x0
WR_TO_RD_CS_CHANGE_DLY	3:0	R/W	Delay from write to write on different chip selects. WL+tWTR-1	0x0

5.7.5.10 ICPU_CFG:MEMCTRL:MEMCTRL_MR0_VAL

Parent: [ICPU_CFG:MEMCTRL](#)

Instances: 1

TABLE 5-327: FIELDS IN MEMCTRL_MR0_VAL

Field Name	Bit	Access	Description	Default
MR0_VAL	15:0	R/W	Value to be programmed into the mode register (0) during SDRAM initialization. During initialization bit 8 (DLL Reset) of this register must be set to 0, the memory controller automatically sets this bit when required during the initialization procedure. After initialization this field is used for sending custom MDSET commands. If ICPU_CFG::MEMCTRL_CTRL.MDSET is set then the value of this field is written to the register defined by bit [15:14].	0x0000

5.7.5.11 ICPU_CFG:MEMCTRL:MEMCTRL_MR1_VAL

Parent: [ICPU_CFG:MEMCTRL](#)

Instances: 1

TABLE 5-328: FIELDS IN MEMCTRL_MR1_VAL

Field Name	Bit	Access	Description	Default
MR1_VAL	15:0	R/W	Value to be programmed into mode register 1 / extended mode register during SDRAM initialization. Bits 7 thorough 9 (OCD Calibration Program) of this register must be set to 0x7, the memory controller sets this field when required during the initialization procedure.	0x0000

5.7.5.12 ICPU_CFG:MEMCTRL:MEMCTRL_MR2_VAL

Parent: [ICPU_CFG:MEMCTRL](#)

Instances: 1

TABLE 5-329: FIELDS IN MEMCTRL_MR2_VAL

Field Name	Bit	Access	Description	Default
MR2_VAL	15:0	R/W	Value to be programmed into mode register 2 / extended mode register 2 during SDRAM initialization.	0x0000

5.7.5.13 ICPU_CFG:MEMCTRL:MEMCTRL_MR3_VAL

Parent: [ICPU_CFG:MEMCTRL](#)

Instances: 1

TABLE 5-330: FIELDS IN MEMCTRL_MR3_VAL

Field Name	Bit	Access	Description	Default
MR3_VAL	15:0	R/W	Value to be programmed into mode register 3 / extended mode register 3 during SDRAM initialization.	0x0000

5.7.5.14 ICPU_CFG:MEMCTRL:MEMCTRL_TERMRES_CTRL

Parent: [ICPU_CFG:MEMCTRL](#)

Instances: 1

TABLE 5-331: FIELDS IN MEMCTRL_TERMRES_CTRL

Field Name	Bit	Access	Description	Default
ODT_RD_EXT	11	R/W	Set this field to extend the ODT termination output by one clock during read operations.	0x0
ODT_RD_ENA	10:7	R/W	Set to enable ODT output during read operations. 0: Reading will not assert ODT 2: Assert ODT0 for CS1 read 4: Assert ODT1 for CS0 read Other values are reserved	0x0
ODT_WR_EXT	6	R/W	Set this field to extend the ODT termination output by one clock during write operations.	0x0
ODT_WR_ENA	5:2	R/W	Set to assert ODT output(s) during write operations. 0: Writing will not assert ODT 3: Assert ODT0 for any write 9: Assert ODT0 for CS0 write and ODT1 for CS1 write 12: Assert ODT1 for any write Others values are reserved	0x0
LOCAL_ODT_RD_EXT	1	R/W	Set this field to extend the local termination by one clock during read operations.	0x0
LOCAL_ODT_RD_ENA	0	R/W	Set to enable local termination during a read operation.	0x0

5.7.5.15 ICPU_CFG:MEMCTRL:MEMCTRL_DQS_DLY

Parent: [ICPU_CFG:MEMCTRL](#)

Instances: 2

VSC7416-01

This register is replicated two times, once for each Byte Lane (first replication corresponds to Byte Lane 0).

TABLE 5-332: FIELDS IN MEMCTRL_DQS_DLY

Field Name	Bit	Access	Description	Default
RESERVED	10:8	R/W	Must be set to its default.	0x3
RESERVED	7:5	R/W	Must be set to its default.	0x3
DQS_DLY	4:0	R/W	This field configures read-window delay as an offset in 1/4 clock cycles from the fixed read-delay configured in MEMCTRL_TIMING0:RD_DATA_XFR_DLY.	0x00

5.7.5.16 ICPU_CFG:MEMCTRL:MEMPHY_CFG

Parent: [ICPU_CFG:MEMCTRL](#)

Instances: 1

TABLE 5-333: FIELDS IN MEMPHY_CFG

Field Name	Bit	Access	Description	Default
PHY_FIFO_RST	7	R/W	Soft-reset to the FIFO blocks in the memory controller physical interface. Leave at default value. 0: PHY FIFOs is in working mode. 1: PHY FIFOs is forced in reset.	0x0
PHY_ODT_OE	4	R/W	Set to enable output drive of the ODT output.	0x0
PHY_CK_OE	3	R/W	Set to enable output drive of the CK/nCK and CKE outputs.	0x0
PHY_CL_OE	2	R/W	Set to enable output drive of the Command Lane outputs.	0x0
PHY_SSTL_ENA	1	R/W	Set this field to enable the SSTL mode for the memory controllers physical interfaces.	0x1
PHY_RST	0	R/W	Master reset to the memory controller physical interface. 0: PHY is in working mode. 1: PHY is forced in reset.	0x1

5.7.5.17 ICPU_CFG:MEMCTRL:MEMPHY_ZCAL

Parent: [ICPU_CFG:MEMCTRL](#)

Instances: 1

TABLE 5-334: FIELDS IN MEMPHY_ZCAL

Field Name	Bit	Access	Description	Default
ZCAL_PROG_ODT	8:5	R/W	Together with the external reference resistor this field configures the SSTL On-Die-Termination (ODT) impedance. This field must be configured prior to, or at the same time as, setting the ICPU_CFG::MEMPHY_ZCAL.ZCAL_ENA field. 2: 150ohms 3: 120ohms 5: 75ohms 7: 60ohms 8: 50ohms 11: 40ohms 13: 34ohms Other values are reserved.	0x3
ZCAL_PROG	4:1	R/W	Together with the external reference resistor this field configures the SSTL output drive-strength. This field must be configured prior to, or at the same time as, setting the ICPU_CFG::MEMPHY_ZCAL.ZCAL_ENA field. 2: 150ohms 3: 120ohms 5: 75ohms 7: 60ohms 8: 50ohms 11: 40ohms 13: 34ohms Other values are reserved.	0xB
ZCAL_ENA	0	One-shot	Set this field to start automatic SSTL output and ODT drive-strength calibration. This field is cleared when the automatic calibration has completed.	0x0

5.7.6 ICPU_CFG:TWI_DELAY

Parent: [ICPU_CFG](#)

Instances: 1

TABLE 5-335: REGISTERS IN TWI_DELAY

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
TWI_CONFIG	0x00000000	1	TWI hold time configuration	Page 471

5.7.6.1 ICPU_CFG:TWI_DELAY:TWI_CONFIG

Parent: [ICPU_CFG:TWI_DELAY](#)

Instances: 1

TABLE 5-336: FIELDS IN TWI_CONFIG

Field Name	Bit	Access	Description	Default
TWI_CNT_RELOAD	8:1	R/W	Configure the hold time delay to apply to SDA after SCK when transmitting from the device. This delay is expressed in a number of VCore System clock cycles. The delay value should be as close to 300ns as possible without going below 300ns. Set to $(300\text{ns}/4.8\text{ns} + 2) = 65$	0x00
TWI_DELAY_ENABLE	0	R/W	Set this field to enable hold time on the TWI SDA output. When enabled the TWI_CONFIG.TWI_CNT_RELOAD field determines the amount of hold time to apply to SDA.	0x0

5.7.7 ICPU_CFG:TWI_SPIKE_FILTER

Parent: [ICPU_CFG](#)

Instances: 1

TABLE 5-337: REGISTERS IN TWI_SPIKE_FILTER

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
TWI_SPIKE_FILTER_CFG	0x00000000	1	TWI spike filter configuration	Page 472

5.7.7.1 ICPU_CFG:TWI_SPIKE_FILTER:TWI_SPIKE_FILTER_CFG

Parent: [ICPU_CFG:TWI_SPIKE_FILTER](#)

Instances: 1

TABLE 5-338: FIELDS IN TWI_SPIKE_FILTER_CFG

Field Name	Bit	Access	Description	Default
SPIKE_FILTER_CFG	4:0	R/W	0: 1 period 1: 2 periods 2 : 3 periods	0x00

5.7.8 ICPU_CFG:FDMA

Parent: [ICPU_CFG](#)

Instances: 1

For registers and fields in this group, which are replicated per channel; replication-index-0 corresponds to extraction-channel 0, replication-index-1 corresponds to extraction-channel 1, and so on. The injection channels follows after the extraction channels.

TABLE 5-339: REGISTERS IN FDMA

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
FDMA_DCB_LL	0x00000000	4 0x00000004	Pointer to next DCB	Page 473
FDMA_DCB_DATAP	0x00000010	4 0x00000004	Pointer to data block	Page 474
FDMA_DCB_DATA	0x00000020	4 0x00000004	Length of data block	Page 474
FDMA_DCB_STAT	0x00000030	4 0x00000004	Status word	Page 474
FDMA_DCB_LL_PREV	0x00000040	4 0x00000004	Pointer to current DCB	Page 475
FDMA_CH_STAT	0x00000050	1	Current channel status	Page 475
FDMA_CH_SAFE	0x00000054	1	Current channel safe-status	Page 475
FDMA_CH_ACTIVATE	0x00000058	1	Activate specific channels	Page 475
FDMA_CH_DISABLE	0x0000005C	1	Disable specific channels	Page 476
FDMA_CH_FORCEDIS	0x00000060	1	Ungraceful disable of specific channels	Page 476
FDMA_CH_CNT	0x00000064	4 0x00000004	Channel counters	Page 476
FDMA_CH_INJ_TOKEN_CNT	0x00000074	2 0x00000004	Injection channel token counter	Page 477
FDMA_CH_INJ_TO- KEN_TICK_RLD	0x0000007C	2 0x00000004	Injection channel token tick counter reload value	Page 478
FDMA_CH_INJ_TO- KEN_TICK_CNT	0x00000084	2 0x00000004	Injection channel token tick counter	Page 478
FDMA_EVT_ERR	0x0000008C	1	Error event	Page 478
FDMA_EVT_ERR_CODE	0x00000090	1	Additional error information	Page 479
FDMA_INTR_LL	0x00000094	1	LLP-event	Page 479
FDMA_INTR_LL_ENA	0x00000098	1	LLP-event interrupt enable	Page 479
FDMA_INTR_FRM	0x0000009C	1	FRM-event	Page 480
FDMA_INTR_FRM_ENA	0x000000A0	1	FRM-event interrupt enable	Page 480
FDMA_INTR_SIG	0x000000A4	1	SIG-event	Page 480
FDMA_INTR_SIG_ENA	0x000000A8	1	SIG-event interrupt enable	Page 480
FDMA_INTR_ENA	0x000000AC	1	Channel interrupt enable	Page 480
FDMA_INTR_IDENT	0x000000B0	1	Currently interrupting channels	Page 481
FDMA_CH_CFG	0x000000B4	4 0x00000004	Channel specific configurations	Page 481
FDMA_GCFG	0x000000C4	1	General FDMA configurations	Page 482
FDMA_GSTAT	0x000000C8	1	General FDMA status	Page 483
FDMA_IDLECNT	0x000000CC	1	FDMA idle Counter	Page 483
FDMA_CONST	0x000000D0	1	Constants for this FDMA implementation.	Page 483

5.7.8.1 ICPU_CFG:FDMA:FDMA_DCB_LL

Parent: [ICPU_CFG:FDMA](#)

Instances: 4

TABLE 5-340: FIELDS IN FDMA_DCB_LL

Field Name	Bit	Access	Description	Default
LLP	31:0	R/W	This field is used by the FDMA for tracking lists of DCBs. This field is updated automatically when the FDMA load DCBs from memory. This field can only be modified when the channel is in safe mode, see ICPU_CFG::FDAM_CH_SAFE.CH_SAFE for more information.	0x00000000

5.7.8.2 ICPU_CFG:FDMA:FDMA_DCB_DATAP

Parent: [ICPU_CFG:FDMA](#)

Instances: 4

TABLE 5-341: FIELDS IN FDMA_DCB_DATAP

Field Name	Bit	Access	Description	Default
DATAP	31:0	R/O	For debug, current data-pointer.	0x00000000

5.7.8.3 ICPU_CFG:FDMA:FDMA_DCB_DATAL

Parent: [ICPU_CFG:FDMA](#)

Instances: 4

TABLE 5-342: FIELDS IN FDMA_DCB_DATAL

Field Name	Bit	Access	Description	Default
TOKEN	17	R/O	For debug, current token-indication.	0x0
DATAL	15:0	R/O	For debug, current data-length.	0x0000

5.7.8.4 ICPU_CFG:FDMA:FDMA_DCB_STAT

Parent: [ICPU_CFG:FDMA](#)

Instances: 4

This register is updated by the FDMA during extraction or injection. Software cannot rely on the value of this register.

TABLE 5-343: FIELDS IN FDMA_DCB_STAT

Field Name	Bit	Access	Description	Default
BLOCKO	31:20	R/O	Block offset in bytes, the value of this field is loaded from the DCB.	0x000
PD	19	R/O	Pruned/Done indication.	0x0
ABORT	18	R/O	Abort indication.	0x0
EOF	17	R/O	Set when the current DCB contains end-of-frame.	0x0
SOF	16	R/O	Set when the current DCB contains start-of-frame.	0x0

TABLE 5-343: FIELDS IN FDMA_DCB_STAT (CONTINUED)

Field Name	Bit	Access	Description	Default
BLOCKL	15:0	R/O	Block size in bytes, excluding offset (as specified in ICPU_CFG::FDMA_DCB_STAT.BLOCKO). For frames that span multiple DCBs, this field only shows the amount of data in the current DCB.	0x0000

5.7.8.5 ICPU_CFG:FDMA:FDMA_DCB_LL_PREV

Parent: [ICPU_CFG:FDMA](#)

Instances: 4

TABLE 5-344: FIELDS IN FDMA_DCB_LL_PREV

Field Name	Bit	Access	Description	Default
LL_PREV	31:2	R/O	This field holds the pointer to current DCB (the previous ICPU_CFG::FDMA_DCB_LL.DCB_LL).	0x00000000

5.7.8.6 ICPU_CFG:FDMA:FDMA_CH_STAT

Parent: [ICPU_CFG:FDMA](#)

Instances: 1

TABLE 5-345: FIELDS IN FDMA_CH_STAT

Field Name	Bit	Access	Description	Default
CH_STAT	3:0	R/O	Shows status for all FDMA channels, there is one bit per channel. 0:Disabled 1:Updating, or Active	0x0

5.7.8.7 ICPU_CFG:FDMA:FDMA_CH_SAFE

Parent: [ICPU_CFG:FDMA](#)

Instances: 1

TABLE 5-346: FIELDS IN FDMA_CH_SAFE

Field Name	Bit	Access	Description	Default
CH_SAFE	3:0	R/O	When set it is safe for software to read/modify/write ICPU_CFG::FDMA_DCB_LL.LL_P, ICPU_CFG::FDMA_CH_CNT.CH_CNT_SIG, ICPU_CFG::FDMA_CH_CNT.CH_CNT_DCB, and ICPU_CFG::FDMA_CH_INJ_TOKEN_CNT.CH_INJ_TOKEN_CNT. There is one bit per channel. This field is set when a channel is a) disabled or b) active and scheduled for disabling.	0xF

5.7.8.8 ICPU_CFG:FDMA:FDMA_CH_ACTIVATE

Parent: [ICPU_CFG:FDMA](#)

Instances: 1

TABLE 5-347: FIELDS IN FDMA_CH_ACTIVATE

Field Name	Bit	Access	Description	Default
CH_ACTIVATE	3:0	One-shot	Enables specific FDMA channels, there is one bit per channel. Setting a bit in this field will clear a corresponding pending ICPU_CFG::FDMA_CH_DISABLE.CH_DISABLE request. Bits in this field are cleared immediately when set.	0x0

5.7.8.9 ICPU_CFG:FDMA:FDMA_CH_DISABLE

Parent: [ICPU_CFG:FDMA](#)

Instances: 1

TABLE 5-348: FIELDS IN FDMA_CH_DISABLE

Field Name	Bit	Access	Description	Default
CH_DISABLE	3:0	One-shot	Schedules specific FDMA channels to be disabled, there is one bit per channel. The channel will finish the current DCB and then disable (after writing the DCB status word). Bits in this field is cleared either when the channel disables or by writing ICPU_CFG::FDMA_CH_ACTIVATE.CH_ACTIVATE).	0x0

5.7.8.10 ICPU_CFG:FDMA:FDMA_CH_FORCEDIS

Parent: [ICPU_CFG:FDMA](#)

Instances: 1

TABLE 5-349: FIELDS IN FDMA_CH_FORCEDIS

Field Name	Bit	Access	Description	Default
CH_FORCEDIS	3:0	One-shot	Immediately disable specific FDMA channels, there is one bit per channel. Unlike ICPU_CFG::FDMA_CH_DISABLE using CH_FORCEDIS will not take the state of the channel into account, if the channel is actively extracting or injecting from/to QS there is no guarantee that it will be functional after disabling the channel.	0x0

5.7.8.11 ICPU_CFG:FDMA:FDMA_CH_CNT

Parent: [ICPU_CFG:FDMA](#)

Instances: 4

TABLE 5-350: FIELDS IN FDMA_CH_CNT

Field Name	Bit	Access	Description	Default
CH_CNT_FRM	31:16	R/W	This field is incremented every time the channel saves status for a DCB that has EOF. This counter can only be modified safely when the corresponding channel is disabled (see ICPU_CFG:FDMA_CH_STAT.CH_STAT for more information).	0x0000
CH_CNT_DCB	15:8	R/W	This field is incremented every time the channel loads a DCB. This counter can be modified safely while the corresponding channel is safe (see ICPU_CFG::FDMA_CH_SAFE.CH_SAFE for more information).	0x00
CH_CNT_SIG	7:0	R/W	This field is incremented every time the channel loads a DCB that has the DATAL.SIG field set. The FDMA can generate interrupt whenever this counter is incremented (see FDMA_INTR_SIG:INTR_SIG for more information). This counter can be modified safely while the corresponding channel is safe (see ICPU_CFG::FDMA_CH_SAFE.CH_SAFE for more information).	0x00

5.7.8.12 ICPU_CFG:FDMA:FDMA_CH_INJ_TOKEN_CNT

Parent: [ICPU_CFG:FDMA](#)

Instances: 2

TABLE 5-351: FIELDS IN FDMA_CH_INJ_TOKEN_CNT

Field Name	Bit	Access	Description	Default
CH_INJ_TOKEN_CNT	7:0	R/W	Every time a channel activates with a DCB that has the TOKEN field set this counter is decremented by one. Channels that loads a DCB with the TOKEN field set cannot activate unless this counter is different from zero. This counter can be written by software, or incremented automatically by using the token tick counter (see ICPU_CFG::FDMA_CH_TOKEN_TICK_CNT for more information). This counter can be modified safely when automatic incrementing is not enabled and the corresponding injection channel is in safe mode (see ICPU_CFG:FDMA_CH_INJ_TOKEN_CNT and ICPU_CFG::FDMA_CH_SAFE.CH_SAFE for more information).	0x00

5.7.8.13 ICPU_CFG:FDMA:FDMA_CH_INJ_TOKEN_TICK_RLD

Parent: [ICPU_CFG:FDMA](#)

Instances: 2

TABLE 5-352: FIELDS IN FDMA_CH_INJ_TOKEN_TICK_RLD

Field Name	Bit	Access	Description	Default
CH_INJ_TOKEN_TICK_RLD	31:0	R/W	Automatic incrementing of the token counter is enabled by setting this field different from 0. This field holds the reload value for the ICPU_CFG::FDMA_CH_INJ_TOKEN_TICK_CNT. Note: When changing the value of this field the same value should also be written to the ICPU_CH_INJ_TOKEN_TICK_CNT field, this is needed for speeding up token counter increments when changing from a high reload value to a low reload value. 0: Token tick counter is disabled n: Add one token every $n * 198.4\text{ns}$ clock cycles	0x00000000

5.7.8.14 ICPU_CFG:FDMA:FDMA_CH_INJ_TOKEN_TICK_CNT

Parent: [ICPU_CFG:FDMA](#)

Instances: 2

TABLE 5-353: FIELDS IN FDMA_CH_INJ_TOKEN_TICK_CNT

Field Name	Bit	Access	Description	Default
CH_INJ_TOKEN_TICK_CNT	31:0	R/W	Down-counter, when enabled by ICPU_CFG::FDMA_CH_INJ_TOKEN_TICK_RLD this field is decremented by one every 198.4ns. When zero is reached one token will be added to ICPU_CFG::FDMA_CH_INJ_TOKEN_CNT and this counted will load the value from ICPU_CFG::FDMA_CH_INJ_TOKEN_TICK_RLD (subtract one and continue decrementing from that value).	0x00000000

5.7.8.15 ICPU_CFG:FDMA:FDMA_EVT_ERR

Parent: [ICPU_CFG:FDMA](#)

Instances: 1

TABLE 5-354: FIELDS IN FDMA_EVT_ERR

Field Name	Bit	Access	Description	Default
EVT_ERR	3:0	Sticky	Shows if an Error-event has occurred, there is one bit per channel. See ICPU_CFG::FDMA_EVT_ERR_CODE.EVT_ERR_CODE for description of errors for which the FDMA implements run-time checks.	0x0

5.7.8.16 ICPu_CFG:FDMA:FDMA_EVT_ERR_CODE

Parent: [ICPu_CFG:FDMA](#)

Instances: 1

TABLE 5-355: FIELDS IN FDMA_EVT_ERR_CODE

Field Name	Bit	Access	Description	Default
EVT_ERR_CODE	3:0	R/O	This field shows information about Error-events that has been recorded by the FDMA, this can be used for software development and debugging. If multiple errors happen in succession, only the newest of the err-codes is shown. 0:Default (no error has occurred) 1:CH_ACTIVATE set for channel w. DCB_LLQ==NULL 2:Got DCB w. DATAP==NULL 3:Got extraction DCB w. DATAL==0 4:Got extraction DCB w. DATAL<=BLOCKO 5:Got injection DCB w. BLOCKL==0 6:Got injection DCB w. SOF for already active channel 7:Activate attempted for channel w. error indication. 8:Activate attempted for channel enabled for manual mode. 9:Manual mode enabled for channel in active FDMA mode.	0x0

5.7.8.17 ICPu_CFG:FDMA:FDMA_INTR_LLQ

Parent: [ICPu_CFG:FDMA](#)

Instances: 1

TABLE 5-356: FIELDS IN FDMA_INTR_LLQ

Field Name	Bit	Access	Description	Default
INTR_LLQ	3:0	Sticky	Shows if an LLP-event has occurred, there is one bit per channel. See the data sheet for information on when this event can occur.	0x0

5.7.8.18 ICPu_CFG:FDMA:FDMA_INTR_LLQ_ENA

Parent: [ICPu_CFG:FDMA](#)

Instances: 1

TABLE 5-357: FIELDS IN FDMA_INTR_LLQ_ENA

Field Name	Bit	Access	Description	Default
INTR_LLQ_ENA	3:0	R/W	Enables LLP-event to be propagated as interrupt, there is one bit per channel. See ICPu_CFG::FDMA_INTR_LLQ.INTR_LLQ for additional information.	0x0

5.7.8.19 ICPU_CFG:FDMA:FDMA_INTR_FRM

Parent: [ICPU_CFG:FDMA](#)

Instances: 1

TABLE 5-358: FIELDS IN FDMA_INTR_FRM

Field Name	Bit	Access	Description	Default
INTR_FRM	3:0	Sticky	Shows if a FRM-event has occurred, there is one bit per channel. See the data sheet for information on when this event can occur.	0x0

5.7.8.20 ICPU_CFG:FDMA:FDMA_INTR_FRM_ENA

Parent: [ICPU_CFG:FDMA](#)

Instances: 1

TABLE 5-359: FIELDS IN FDMA_INTR_FRM_ENA

Field Name	Bit	Access	Description	Default
INTR_FRM_ENA	3:0	R/W	Enables FRM-event to be propagated as interrupt, there is one bit per channel. See ICPU_CFG::FDMA_INTR_FRM.INTR_FRM for additional information.	0x0

5.7.8.21 ICPU_CFG:FDMA:FDMA_INTR_SIG

Parent: [ICPU_CFG:FDMA](#)

Instances: 1

TABLE 5-360: FIELDS IN FDMA_INTR_SIG

Field Name	Bit	Access	Description	Default
INTR_SIG	3:0	Sticky	Shows if a SIG-event has occurred, there is one bit per channel. See the data sheet for information on when this event can occur.	0x0

5.7.8.22 ICPU_CFG:FDMA:FDMA_INTR_SIG_ENA

Parent: [ICPU_CFG:FDMA](#)

Instances: 1

TABLE 5-361: FIELDS IN FDMA_INTR_SIG_ENA

Field Name	Bit	Access	Description	Default
INTR_SIG_ENA	3:0	R/W	Enables SIG-event to be propagated as interrupt, there is one bit per channel. See ICPU_CFG::FDMA_INTR_SIG.INTR_SIG for additional information.	0x0

5.7.8.23 ICPU_CFG:FDMA:FDMA_INTR_ENA

Parent: [ICPU_CFG:FDMA](#)

Instances: 1

TABLE 5-362: FIELDS IN FDMA_INTR_ENA

Field Name	Bit	Access	Description	Default
INTR_ENA	3:0	R/W	Enables propagation of enabled channel LLP-event, FRM-event and SIG-event as interrupt, there is one bit per channel. ERR-events are always propagated when interrupt is enabled for a channel.	0xF

5.7.8.24 ICPU_CFG:FDMA:FDMA_INTR_IDENT

Parent: [ICPU_CFG:FDMA](#)

Instances: 1

TABLE 5-363: FIELDS IN FDMA_INTR_IDENT

Field Name	Bit	Access	Description	Default
INTR_IDENT	3:0	R/O	Shows currently interrupting channels, there is one bit per channel.	0x0

5.7.8.25 ICPU_CFG:FDMA:FDMA_CH_CFG

Parent: [ICPU_CFG:FDMA](#)

Instances: 4

TABLE 5-364: FIELDS IN FDMA_CH_CFG

Field Name	Bit	Access	Description	Default
STAT_IN_DATA_ENA	4	R/W	When this field is set the FDMA will save the STAT-word to the DATAP-address instead of the DCB's STAT-word position. The DCB's DATAP field will be incremented by 4 (bytes) when the DCB is loaded (the FDMA will continue as if the DATAP field was DATAP+4). This feature is meant to be used for channels that extract or inject from PCIe mapped memory. The ICPU_CFG::FDMA_DCB_LL-P_PREV field is repurposed as DATAP pointer and can no longer be used when this field is set.	0x0

TABLE 5-364: FIELDS IN FDMA_CH_CFG (CONTINUED)

Field Name	Bit	Access	Description	Default
CH_PRIO	3:2	R/W	The FDMA implements a strict priority scheme between all channels - both injection and extraction. Observe: The FDMA does not directly control the order in which ports are serviced in the Queuing System. In order to adjust for this (and to avoid head of line blocking) all extraction channels are automatically assigned the highest priority of the extraction channels with available data. If multiple channels are configured with equal priorities then the following strict scheme is in place: Higher channel number takes priority over lower channel number. This implies that injection takes priority over extraction.	0x0
DONE_STOP_ENA	1	R/W	Set this field to automatically disable the channel after completing any DCB. The channel will disable after saving DCB status. An LLP-event will be generated at the same time as the channel is disabled. Be careful when using this feature, extraction channels may head-of-line block other extraction channels if not immediately re-activated.	0x0
DONEEOF_STOP_ENA	0	R/W	Set this field to automatically disable the channel after completing any DCB with EOF indication. The channel will disable after saving DCB status. An LLP-event will be generated at the same time as the channel is disabled. Be careful when using this feature, extraction channels may head-of-line block other extraction channels if not immediately re-activated.	0x0

5.7.8.26 ICPU_CFG:FDMA:FDMA_GCFG

Parent: [ICPU_CFG:FDMA](#)

Instances: 1

TABLE 5-365: FIELDS IN FDMA_GCFG

Field Name	Bit	Access	Description	Default
INJ_RF_WM	11:7	R/W	Injection resync FIFO fill-level watermark, when exceeded backpressure will be asserted towards SBA. The maximum fill-level for the FIFO is reported via ICPU_CFG::FDMA_G-STAT.INJ_RF_HIGH. n: backpressure when n+1 or more words in buffer.	0x0E
RESERVED	6:3	R/W	Must be set to its default.	0x9

TABLE 5-365: FIELDS IN FDMA_GCFG (CONTINUED)

Field Name	Bit	Access	Description	Default
PD_IGNORE	0	R/W	Set this field to make the FDMA ignore the value of the STAT.PD field when injecting frames. By default the FDMA will treat the PD field in the same way as STAT.ABORT.	0x0

5.7.8.27 ICPU_CFG:FDMA:FDMA_GSTAT

Parent: [ICPU_CFG:FDMA](#)

Instances: 1

TABLE 5-366: FIELDS IN FDMA_GSTAT

Field Name	Bit	Access	Description	Default
INJ_RF_HIGH	10:5	R/O	This field shows the highest fill level that the injection resync FIFO has experienced since reset of the injection logic. The depth of the FIFO is 32 words, reaching a fill-level of 33 (or more) means that overflow has occurred.	0x00

5.7.8.28 ICPU_CFG:FDMA:FDMA_IDLECNT

Parent: [ICPU_CFG:FDMA](#)

Instances: 1

TABLE 5-367: FIELDS IN FDMA_IDLECNT

Field Name	Bit	Access	Description	Default
IDLECNT	23:0	R/O	The counter is reset whenever a channel is enabled and when FDMA moves frame data to or from the queuing system. When the FDMA is idle this counter is incremented once every 198.4ns. The counter saturates at maximum value (approx 3.3 seconds of idle time).	0x000000

5.7.8.29 ICPU_CFG:FDMA:FDMA_CONST

Parent: [ICPU_CFG:FDMA](#)

Instances: 1

TABLE 5-368: FIELDS IN FDMA_CONST

Field Name	Bit	Access	Description	Default
CH_INJ_CNT	15:8	R/O	The number of injection channels.	0x02
CH_XTR_CNT	7:0	R/O	The number of extraction channels.	0x02

5.7.9 ICPU_CFG:PCIE

Parent: [ICPU_CFG](#)

Instances: 1

These are VCore based registers for configuration of PCIe endpoint support logic. PCIe configuration space registers are memory mapped via the PCIE region of the Chip Registers.

TABLE 5-369: REGISTERS IN PCIE

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
PCIE_CFG	0x00000004	1	PCle endpoint configuration	Page 484
PCIE_STAT	0x00000008	1	PCle endpoint status	Page 484
PCIE_AUX_CFG	0x0000000C	1	Auxiliary power configuration	Page 485
PCIESLV_FDMA	0x00000028	1	FDMA access into PCle address space	Page 485
PCIESLV_SBA	0x0000002C	1	SBA access into PCle address space	Page 485
PCIEPCS_CFG	0x00000030	1	PCle PCS configuration	Page 486
PCIE_INTR	0x00000038	1	PCle events	Page 487
PCIE_INTR_ENA	0x0000003C	1	PCle interrupt enable	Page 487
PCIE_INTR_IDENT	0x00000040	1	Currently active PCle interrupts	Page 487
PCIE_INTR_COMMON_CFG	0x00000044	1	PCle outbound interrupt configuration	Page 487
PCIE_INTR_CFG	0x00000048	2 0x00000004	PCle outbound MSI interrupt configuration	Page 488

5.7.9.1 ICPU_CFG:PCle:PCIE_CFG

Parent: [ICPU_CFG:PCle](#)

Instances: 1

TABLE 5-370: FIELDS IN PCIE_CFG

Field Name	Bit	Access	Description	Default
PCIE_BAR_WR_ENA	2	R/W	Set this field to enable write of BAR1 and BAR2 masks via PCIE register target. Only registers PCIE::CONF_SPACE[5], PCIE::CONF_SPACE[6], and PCIE::CONF_SPACE[7] (in the PCIE target) may be written while this field is set. The minimum size for Memory and IO BARs are 64K (mask 0xFFFF). Note: The low 4 bits of all BARs can be written though the PCle target when this field is not set.	0x0
LTSSM_DIS	1	R/W	Set this field to disable initialization of the PCle link. By default the PCle core will start up and try to achieve link when the SERDES is started, by setting this field before starting the SERDES it is possible to make changes to the PCle configuration/registers prior to linking with the root complex.	0x0
MEM_RING_CORE_ENA	0	R/W	Set to add the PCle core memories to the RAM integrity ring.	0x0

5.7.9.2 ICPU_CFG:PCle:PCIE_STAT

Parent: [ICPU_CFG:PCle](#)

Instances: 1

TABLE 5-371: FIELDS IN PCIE_STAT

Field Name	Bit	Access	Description	Default
PM_STATE	2:0	R/O	The current power management state of the PCIe core. 0: D0 1: D1 2: D2 3: D3 4: D0-Uninitialized	0x0

5.7.9.3 ICPU_CFG:PCIE:PCIE_AUX_CFG

Parent: [ICPU_CFG:PCIE](#)

Instances: 1

TABLE 5-372: FIELDS IN PCIE_AUX_CFG

Field Name	Bit	Access	Description	Default
AUX_POWER_FORCE	1	R/W	Set to force "detection" of PCIe auxiliary power to the value of ICPU_CFG::PCIE_AUX_CFG.AUX_POWER_VAL.	0x0
AUX_POWER_VAL	0	R/W	See ICPU_CFG::PCIE_AUX_CFG.AUX_POWER_FORCE for more information.	0x0

5.7.9.4 ICPU_CFG:PCIE:PCIESLV_FDMA

Parent: [ICPU_CFG:PCIE](#)

Instances: 1

TABLE 5-373: FIELDS IN PCIESLV_FDMA

Field Name	Bit	Access	Description	Default
FDMA_OFFSET	1:0	R/W	The FDMA has access to one 1GByte region (0xC0000000 though 0xFFFFFFFF) that maps accesses to PCIe interface. The value of this field is used for address-bits [31:30] towards the PCIe endpoint. Set this field to 1.	0x0

5.7.9.5 ICPU_CFG:PCIE:PCIESLV_SBA

Parent: [ICPU_CFG:PCIE](#)

Instances: 1

TABLE 5-374: FIELDS IN PCIESLV_SBA

Field Name	Bit	Access	Description	Default
SBA_BE	27:24	R/W	This field allows configuration of outbound PCIe-transaction Byte-Enable field. This is applied to all SBA (non-FDMA) initiated outbound PCIe accesses. This field is not used for TYPE=MRr/MWr/MRdLk accesses. Byte-enables are needed in order to support Zero-byte and non-contiguous byte IO and CFG transfers and Zero-byte Messages.	0x0
SBA_MSG_CODE	22:15	R/W	This field allows configuration of outbound PCIe-transaction MSG field. This is applied to all SBA (non-FDMA) initiated outbound PCIe accesses.	0x00
SBA_TC	14:12	R/W	This field allows configuration of outbound PCIe-transaction TC field. This is applied to all SBA (non-FDMA) initiated outbound PCIe accesses.	0x0
SBA_EP	8	R/W	This field allows configuration of outbound PCIe-transaction EP field. This is applied to all SBA (non-FDMA) initiated outbound PCIe accesses.	0x0
SBA_OFFSET	1:0	R/W	SBA masters (non-FDMA) has access to one 1GByte region (0xC0000000 though 0xFFFFFFFF) that maps accesses to PCIe interface. The value of this field is used for address-bits [31:30] towards the PCIe endpoint. Set this field to 0.	0x0

5.7.9.6 ICPU_CFG:PCle:PCIEPCS_CFG

Parent: [ICPU_CFG:PCle](#)

Instances: 1

TABLE 5-375: FIELDS IN PCIEPCS_CFG

Field Name	Bit	Access	Description	Default
WAKE_POL	2	R/W	Polarity of the PCIe WAKE output, WAKE is typically an active low output - but if an amplifier is needed for driving a large WAKE net then polarity may need to be changed. 0: Active low 1: Active high	0x0
WAKE_OE	1	R/W	Set to permanently drive PCIe WAKE output, by default the WAKE output is only driven when active and thusly allowing pull-resistor network. 0: Only drive output when active. 1: Always drive output.	0x0

TABLE 5-375: FIELDS IN PCIEPCS_CFG (CONTINUED)

Field Name	Bit	Access	Description	Default
BEACON_DIS	0	R/W	Set this field to disable outband PCIe beacon signalling when attempting to wake from D3. When beacon is disabled the WAKE# signal (available as alternate GPIO function) must be used instead.	0x0

5.7.9.7 ICPU_CFG:PCIE:PCIE_INTR

Parent: [ICPU_CFG:PCIE](#)

Instances: 1

TABLE 5-376: FIELDS IN PCIE_INTR

Field Name	Bit	Access	Description	Default
INTR_PM_STATE	0	Sticky	This event is set whenever the ICPU_CFG::PCIE_STAT.PM_STATE field is changed.	0x0

5.7.9.8 ICPU_CFG:PCIE:PCIE_INTR_ENA

Parent: [ICPU_CFG:PCIE](#)

Instances: 1

TABLE 5-377: FIELDS IN PCIE_INTR_ENA

Field Name	Bit	Access	Description	Default
INTR_PM_STATE_ENA	0	R/W	Set to enable propagation of the PM_STATE interrupt.	0x0

5.7.9.9 ICPU_CFG:PCIE:PCIE_INTR_IDENT

Parent: [ICPU_CFG:PCIE](#)

Instances: 1

TABLE 5-378: FIELDS IN PCIE_INTR_IDENT

Field Name	Bit	Access	Description	Default
INTR_PM_STATE_IDENT	0	R/O	Set if the PM_STATE interrupt is currently active.	0x0

5.7.9.10 ICPU_CFG:PCIE:PCIE_INTR_COMMON_CFG

Parent: [ICPU_CFG:PCIE](#)

Instances: 1

TABLE 5-379: FIELDS IN PCIE_INTR_COMMON_CFG

Field Name	Bit	Access	Description	Default
WAKEUP_ON_INTR_DIS	2	R/W	Set to disable wake-up on interrupt. By default the PCIe endpoint will attempt to wake up from powerdown when a change in interrupt state is detected.	0x0

TABLE 5-379: FIELDS IN PCIE_INTR_COMMON_CFG (CONTINUED)

Field Name	Bit	Access	Description	Default
LEGACY_MODE_INTR_SEL	1	R/W	Select the external interrupt from the VCore interrupt controller that must be used to generate PCIe legacy interrupt. 0: Use EXT_DST0 1: Use EXT_DST1	0x1
PCIE_INTR_ENA	0	R/W	Set to enable PCIe interrupts. The PCIe endpoint's MSI Capability Register Set must have been configured before enabling interrupts.	0x0

5.7.9.11 ICPU_CFG:PCIE:PCIE_INTR_CFG

Parent: [ICPU_CFG:PCIE](#)

Instances: 2

Replicated per EXT_DST interrupt.

TABLE 5-380: FIELDS IN PCIE_INTR_CFG

Field Name	Bit	Access	Description	Default
TRAFFIC_CLASS	14:12	R/W	Configure MSI interrupt traffic class for corresponding EXT_DST interrupt.	0x0
FALLING_VECTOR_VAL	11:7	R/W	Configure MSI interrupt vector for falling edge of corresponding EXT_DST interrupt.	0x00
RISING_VECTOR_VAL	6:2	R/W	Configure MSI interrupt vector for rising edge of corresponding EXT_DST interrupt.	0x00
INTR_FALLING_ENA	1	R/W	Set to enable MSI interrupt on falling edge of corresponding EXT_DST interrupt.	0x0
INTR_RISING_ENA	0	R/W	Set to enable MSI interrupt on rising edge of corresponding EXT_DST interrupt.	0x0

5.7.10 ICPU_CFG:MANUAL_XTRINJ

Parent: [ICPU_CFG](#)

Instances: 1

This group contains replicated register array for doing manual Extraction and Injection by use of the FDMA engine.

TABLE 5-381: REGISTERS IN MANUAL_XTRINJ

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
MANUAL_XTR	0x00000000	4096 0x00000004	Manual extraction replicated register-array	Page 489
MANUAL_INJ	0x00004000	4096 0x00000004	Manual injection replicated register-array	Page 489
MANUAL_CFG	0x00008000	1	Manual extraction and injection configuration	Page 489
MANUAL_INTR	0x00008004	1	Manual extraction and injection interrupt indications	Page 489
MANUAL_INTR_ENA	0x00008008	1	Manual extraction and injection interrupt enables	Page 490

5.7.10.1 ICPU_CFG:MANUAL_XTRINJ:MANUAL_XTR

Parent: [ICPU_CFG:MANUAL_XTRINJ](#)

Instances: 4096

TABLE 5-382: FIELDS IN MANUAL_XTR

Field Name	Bit	Access	Description	Default
XTR	31:0	R/W	Manual extraction is done by reading from this block of registers. The manual extraction status word is accessed by reading the last word-address in this block. Manual extraction has to be enabled via ICPU_CFG::MANUAL_CFG.XTR_ENA.	0x00000000

5.7.10.2 ICPU_CFG:MANUAL_XTRINJ:MANUAL_INJ

Parent: [ICPU_CFG:MANUAL_XTRINJ](#)

Instances: 4096

TABLE 5-383: FIELDS IN MANUAL_INJ

Field Name	Bit	Access	Description	Default
INJ	31:0	R/W	Manual injection is done by writing to this block of registers. The manual injection status word is located at the first word-address in this block. Manual injection has to be enabled via ICPU_CFG::MANUAL_CFG.INJ_ENA.	0x00000000

5.7.10.3 ICPU_CFG:MANUAL_XTRINJ:MANUAL_CFG

Parent: [ICPU_CFG:MANUAL_XTRINJ](#)

Instances: 1

TABLE 5-384: FIELDS IN MANUAL_CFG

Field Name	Bit	Access	Description	Default
INJ_ENA	1	R/W	Set to enable manual injection by using FDMA channel number 3. When manual injection is enabled; the FDMA cannot be used for regular FDMA injection operations (on any injection channel).	0x0
XTR_ENA	0	R/W	Set to enable manual extraction by using FDMA channel number 1. When manual extraction is enabled; the FDMA cannot be used for regular FDMA extraction operations (on any extraction channel).	0x0

5.7.10.4 ICPU_CFG:MANUAL_XTRINJ:MANUAL_INTR

Parent: [ICPU_CFG:MANUAL_XTRINJ](#)

Instances: 1

TABLE 5-385: FIELDS IN MANUAL_INTR

Field Name	Bit	Access	Description	Default
INTR_INJ_RDY	2	R/O	Set when there is room for more injection data-words in injection fifo.	0x0
INTR_XTR_ANY_RDY	1	R/O	Set when any extraction word is ready for extraction.	0x0
INTR_XTR_SOF_RDY	0	R/O	Set when there is an extraction word containing SOF ready for extraction.	0x0

5.7.10.5 ICPU_CFG:MANUAL_XTRINJ:MANUAL_INTR_ENA

Parent: [ICPU_CFG:MANUAL_XTRINJ](#)

Instances: 1

TABLE 5-386: FIELDS IN MANUAL_INTR_ENA

Field Name	Bit	Access	Description	Default
INTR_INJ_RDY_ENA	2	R/W	Set to enable FDMA interrupt while there is room for more injection data. This interrupt is asserted for as long as there is free space in the injection buffers.	0x0
INTR_XTR_ANY_RDY_ENA	1	R/W	Set to enable FDMA interrupt while any data is ready for manual extraction. This interrupt is asserted for as long as there is data ready in the extraction buffer.	0x0
INTR_XTR_SOF_RDY_ENA	0	R/W	Set to enable FDMA interrupt when a new frame is waiting to be extracted. This event is asserted when a frame-word with sof set is waiting to be extracted. If a previous frame is only partially extracted then no interrupt will be generated until the previous frame is completely extracted.	0x0

5.8 VCAP_CORE

TABLE 5-387: REGISTER GROUPS IN VCAP_CORE

Register Group Name	Offset within Target	Instances and Address Spacing	Description	Details
VCAP_CORE_CFG	0x00000000	1		Page 490
VCAP_CORE_CACHE	0x00000008	1		Page 494
VCAP_CORE_STICKY	0x0000020C	1		Page 496
VCAP_CONST	0x00000210	1		Page 496
TCAM_BIST	0x0000022C	1	Build in test for TCAM	Page 498

5.8.1 VCAP_CORE:VCAP_CORE_CFG

Parent: [VCAP_CORE](#)

Instances: 1

TABLE 5-388: REGISTERS IN VCAP_CORE_CFG

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
VCAP_UPDATE_CTRL	0x00000000	1		Page 491
VCAP_MV_CFG	0x00000004	1		Page 493

5.8.1.1 VCAP_CORE:VCAP_CORE_CFG:VCAP_UPDATE_CTRL

Parent: [VCAP_CORE:VCAP_CORE_CFG](#)

Instances: 1

TABLE 5-389: FIELDS IN VCAP_UPDATE_CTRL

Field Name	Bit	Access	Description	Default
UPDATE_CMD	24:22	R/W	Specifies the operation to be carried out when the vcap_update_shot field is set. The COPY operations will read or write the content of the VCAP_CORE_CACHE to the TCAM/RAMs at the address specified in UPDATE_ADDR. When writing default actions UPDATE_ENTRY_DIS must be set to avoid overwriting entries at the low addresses in the TCAM. The MOVE operations will move one or more rows up or down starting from the address in UPDATE_ADDR. The number of rows to move is specified in VCAP_MV_CFG.MV_NUM_POS. Moving a row up will decrease its address by MV_NUM_POS, moving a row down will increase its address by MV_NUM_POS. The number of rows to include in the move is specified in VCAP_MV_CFG.MV_SIZE. If a row is moved to an destination address that is less than zero, i.e. if $UPDATE\_ADDR - MV\_NUM\_POS < 0$, the row will be invalidated in the TCAM and the action and counter RAM will be set to zero. The INIT operations will set the range of rows specified by UPDATE_ADDR and VCAP_MV_CFG.MV_SIZE to the value of cache. Note: init starts from the address specified in UPDATE_ADDR. 000: Copy entry and/or action from cache to TCAM/RAM 001: Copy entry and/or action from TCAM/RAM to cache 010: Move entry and/or action up (decreasing addresses) 011: Move entry and/or action down (increasing addresses) 100: Initialize all entries and/or actions with the value in the cache.	0x0
UPDATE_ENTRY_DIS	21	R/W	Specifies whether the operation specified in vcap_update_cmd is applied to entries. 0: Entries are copied/moved/init. 1: Entries are not copied/moved/init.	0x0

TABLE 5-389: FIELDS IN VCAP_UPDATE_CTRL (CONTINUED)

Field Name	Bit	Access	Description	Default
UPDATE_ACTION_DIS	20	R/W	Specifies whether the operation specified in vcap_update_cmd is applied to actions. 0: Actions are copied/moved/init. 1: Actions are not copied/moved/init.	0x0
UPDATE_CNT_DIS	19	R/W	Specifies whether the operation specified in vcap_update_cmd is applied to cnts. 0: Counters are copied/moved/init. 1: Counters are not copied/moved/init.	0x0
UPDATE_ADDR	18:3	R/W	The entry/action address (row-wise) used for copy/move operations. When accessing the default actions the offset specified in VCAP_CONST:ENTRY_CNT should be added to the default action addr, i.e. Default action 0: set UPDATE_ADDR to VCAP_CONST:ENTRY_CNT + 0. Default action 1: set UPDATE_ADDR to VCAP_CONST:ENTRY_CNT + 1. Default action n: set UPDATE_ADDR to VCAP_CONST:ENTRY_CNT + n.	0x0000
UPDATE_SHOT	2	One-shot	Initiate the operation specified in vcap_update_cmd. The bit is cleared by hw when operation has finished.	0x0
CLEAR_CACHE	1	One-shot	Reset all registers in VCAP_CORE_CACHE. The register is cleared by hw when the operation has finished.	0x0
MV_TRAFFIC_IGN	0	R/W	Ignore interrupting traffic during move operation. If a lookup is performed during a move operation the move is finished from where it was interrupted. When this field is set counters are not guaranteed to be up-to-date after the move.	0x0

5.8.1.2 VCAP_CORE:VCAP_CORE_CFG:VCAP_MV_CFG

Parent: [VCAP_CORE:VCAP_CORE_CFG](#)

Instances: 1

TABLE 5-390: FIELDS IN VCAP_MV_CFG

Field Name	Bit	Access	Description	Default
MV_NUM_POS	31:16	R/W	Specifies the number of positions the row must be moved up or down during a move operation. 0x0 The row is moved one position up or down. 0x1: The row is moved two positions up or down. 0xn: The row is moved n+1 positions up or down.	0x0000

TABLE 5-390: FIELDS IN VCAP_MV_CFG (CONTINUED)

Field Name	Bit	Access	Description	Default
MV_SIZE	15:0	R/W	Specifies the number of rows that must be moved up or down during a move operation. This field is also used to define the range that is initialized using the init feature. 0x0: The row at address UPDATE_ADDR is moved up or down. 0x1: The row at address UPDATE_ADDR through UPDATE_ADDR+1 are moved up or down. 0x2: The row at address UPDATE_ADDR through UPDATE_ADDR+2 are moved up or down. 0xn: The row at address UPDATE_ADDR through UPDATE_ADDR+n are moved up or down.	0x0000

5.8.2 VCAP_CORE:VCAP_CORE_CACHE

Parent: [VCAP_CORE](#)

Instances: 1

TABLE 5-391: REGISTERS IN VCAP_CORE_CACHE

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
VCAP_ENTRY_DAT	0x00000000	32 0x00000004		Page 494
VCAP_MASK_DAT	0x00000080	32 0x00000004		Page 495
VCAP_ACTION_DAT	0x00000100	32 0x00000004		Page 495
VCAP_CNT_DAT	0x00000180	32 0x00000004		Page 495
VCAP_TG_DAT	0x00000200	1		Page 496

5.8.2.1 VCAP_CORE:VCAP_CORE_CACHE:VCAP_ENTRY_DAT

Parent: [VCAP_CORE:VCAP_CORE_CACHE](#)

Instances: 32

TABLE 5-392: FIELDS IN VCAP_ENTRY_DAT

Field Name	Bit	Access	Description	Default
ENTRY_DAT	31:0	R/W	The cache register that holds a single TCAM entry data row. The register is replicated 32 times where replication index 0 is the 32 LSBs of the cache. Note that the physical width of the entry cache is specified in VCAP_CONST.ENTRY_WIDTH and that there are only instantiated flops in the CSR target for this width.	0x00000000

5.8.2.2 VCAP_CORE:VCAP_CORE_CACHE:VCAP_MASK_DAT

Parent: [VCAP_CORE:VCAP_CORE_CACHE](#)

Instances: 32

TABLE 5-393: FIELDS IN VCAP_MASK_DAT

Field Name	Bit	Access	Description	Default
MASK_DAT	31:0	R/W	The cache register that holds a single TCAM entry mask row. The register is replicated 32 times where replication index 0 is the 32 LSBs of the cache. Note that the physical width of the mask cache is specified in VCAP_CONST.ENTRY_WIDTH and that there are only instantiated flops in the CSR target for this width.	0x00000000

5.8.2.3 VCAP_CORE:VCAP_CORE_CACHE:VCAP_ACTION_DAT

Parent: [VCAP_CORE:VCAP_CORE_CACHE](#)

Instances: 32

TABLE 5-394: FIELDS IN VCAP_ACTION_DAT

Field Name	Bit	Access	Description	Default
ACTION_DAT	31:0	R/W	The cache register that holds a single action ram data row. The register is replicated 32 times where replication index 0 is the 32 LSBs of the cache. Note that the physical width of the entry cache is specified in VCAP_CONST.ACTION_WIDTH and that there are only instantiated flops in the CSR target for this width.	0x00000000

5.8.2.4 VCAP_CORE:VCAP_CORE_CACHE:VCAP_CNT_DAT

Parent: [VCAP_CORE:VCAP_CORE_CACHE](#)

Instances: 32

TABLE 5-395: FIELDS IN VCAP_CNT_DAT

Field Name	Bit	Access	Description	Default
CNT_DAT	31:0	R/W	The cache register that holds a single counter ram data row. The register is replicated 32 times where replication index 0 is the 32 LSBs of the cache. Note that the physical width of the entry cache is specified in VCAP_CONST.CNT_WIDTH and that there are only instantiated flops in the CSR target for this width.	0x00000000

5.8.2.5 VCAP_CORE:VCAP_CORE_CACHE:VCAP_TG_DAT

Parent: [VCAP_CORE:VCAP_CORE_CACHE](#)

Instances: 1

TABLE 5-396: FIELDS IN VCAP_TG_DAT

Field Name	Bit	Access	Description	Default
TG_DAT	31:0	R/W	This cache register holds the TypeGroup id for each subword in the TCAM. If the VCAP supports multiple subwords, i.e. VCAP_CONST.ENTRY_SWCNT > 1, the TypeGroup ids are place back to back with subword 0 at the LSBs.	0x00000000

5.8.3 VCAP_CORE:VCAP_CORE_STICKY

Parent: [VCAP_CORE](#)

Instances: 1

TABLE 5-397: REGISTERS IN VCAP_CORE_STICKY

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
VCAP_STICKY	0x00000000	1		Page 496

5.8.3.1 VCAP_CORE:VCAP_CORE_STICKY:VCAP_STICKY

Parent: [VCAP_CORE:VCAP_CORE_STICKY](#)

Instances: 1

TABLE 5-398: FIELDS IN VCAP_STICKY

Field Name	Bit	Access	Description	Default
VCAP_ROW_DELETED_STICKY	0	Sticky	A move operation has resulted in one or more rows have been deleted.	0x0

5.8.4 VCAP_CORE:VCAP_CONST

Parent: [VCAP_CORE](#)

Instances: 1

TABLE 5-399: REGISTERS IN VCAP_CONST

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
ENTRY_WIDTH	0x00000000	1		Page 497
ENTRY_CNT	0x00000004	1		Page 497
ENTRY_SWCNT	0x00000008	1		Page 497
ENTRY_TG_WIDTH	0x0000000C	1		Page 497
ACTION_DEF_CNT	0x00000010	1		Page 498
ACTION_WIDTH	0x00000014	1		Page 498
CNT_WIDTH	0x00000018	1		Page 498

5.8.4.1 VCAP_CORE:VCAP_CONST:ENTRY_WIDTH

Parent: [VCAP_CORE:VCAP_CONST](#)

Instances: 1

TABLE 5-400: FIELDS IN ENTRY_WIDTH

Field Name	Bit	Access	Description	Default
ENTRY_WIDTH	9:0	R/O	The width of a TCAM entry including Type-Group ids.	0x000

5.8.4.2 VCAP_CORE:VCAP_CONST:ENTRY_CNT

Parent: [VCAP_CORE:VCAP_CONST](#)

Instances: 1

TABLE 5-401: FIELDS IN ENTRY_CNT

Field Name	Bit	Access	Description	Default
ENTRY_CNT	15:0	R/O	The number of entries in the TCAM.	0x0000

5.8.4.3 VCAP_CORE:VCAP_CONST:ENTRY_SWCNT

Parent: [VCAP_CORE:VCAP_CONST](#)

Instances: 1

TABLE 5-402: FIELDS IN ENTRY_SWCNT

Field Name	Bit	Access	Description	Default
ENTRY_SWCNT	5:0	R/O	The number of supported subwords in the TCAM.	0x00

5.8.4.4 VCAP_CORE:VCAP_CONST:ENTRY_TG_WIDTH

Parent: [VCAP_CORE:VCAP_CONST](#)

Instances: 1

TABLE 5-403: FIELDS IN ENTRY_TG_WIDTH

Field Name	Bit	Access	Description	Default
ENTRY_TG_WIDTH	5:0	R/O	The width of a single TypeGroup id.	0x00

5.8.4.5 VCAP_CORE:VCAP_CONST:ACTION_DEF_CNT

Parent: [VCAP_CORE:VCAP_CONST](#)

Instances: 1

TABLE 5-404: FIELDS IN ACTION_DEF_CNT

Field Name	Bit	Access	Description	Default
ACTION_DEF_CNT	9:0	R/O	The number of default actions.	0x000

5.8.4.6 VCAP_CORE:VCAP_CONST:ACTION_WIDTH

Parent: [VCAP_CORE:VCAP_CONST](#)

Instances: 1

TABLE 5-405: FIELDS IN ACTION_WIDTH

Field Name	Bit	Access	Description	Default
ACTION_WIDTH	9:0	R/O	The width of the action RAM.	0x000

5.8.4.7 VCAP_CORE:VCAP_CONST:CNT_WIDTH

Parent: [VCAP_CORE:VCAP_CONST](#)

Instances: 1

TABLE 5-406: FIELDS IN CNT_WIDTH

Field Name	Bit	Access	Description	Default
CNT_WIDTH	9:0	R/O	The width of the counter RAM.	0x000

5.8.5 VCAP_CORE:TCAM_BIST

Parent: [VCAP_CORE](#)

Instances: 1

TABLE 5-407: REGISTERS IN TCAM_BIST

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
TCAM_CTRL	0x00000000	1	Control of the TCAM	Page 498
TCAM_STAT	0x00000008	1	Status for the TCAM	Page 499

5.8.5.1 VCAP_CORE:TCAM_BIST:TCAM_CTRL

Parent: [VCAP_CORE:TCAM_BIST](#)

Instances: 1

TABLE 5-408: FIELDS IN TCAM_CTRL

Field Name	Bit	Access	Description	Default
TCAM_BIST	1	One-shot	Set this field to start manual BIST of the TCAM. This field will be cleared once BIST is complete. The BIST procedure requires that the TCAM is initialized before start, setting TCAM_INIT at the same time as setting this field will first initialize the TCAM and then run BIST.	0x0
TCAM_INIT	0	One-shot	Set this field to start manual initialization of the TCAM. This field is cleared once initialization is complete. The TCAM has random contents after reset and must be initialized prior to usage.	0x0

5.8.5.2 VCAP_CORE:TCAM_BIST:TCAM_STAT

Parent: [VCAP_CORE:TCAM_BIST](#)

Instances: 1

TABLE 5-409: FIELDS IN TCAM_STAT

Field Name	Bit	Access	Description	Default
BIST_ERR	2	R/O	Set if BIST failed.	0x0
BIST_BUSY	1	R/O	Set when the BIST is still running. When checking the BIST result this field must be cleared.	0x0
TCAM_RDY	0	R/O	Indicates the current operational state of the TCAM. '0': Busy with initialization. '1': Ready to be used.	0x0

5.9 OAM_MEP

The Versatile OAM MEP Processor (VOP) block implements the HW support for implementing OAM MEP.

The VOP consists of 75 Versatile OAM Endpoints (VOE) that each implements an OAM MEP.

The VOP includes the following:

- 64 service/path VOEs
- 11 port VOEs

TABLE 5-410: REGISTER GROUPS IN OAM_MEP

Register Group Name	Offset within Target	Instances and Address Spacing	Description	Details
COMMON	0x00006000	1	Miscellaneous configuration for OAM_MEP	Page 500
VOE	0x00000000	75 0x00000100	Configuration per Versatile OAM MEP Endpoints (VOE)	Page 514

TABLE 5-410: REGISTER GROUPS IN OAM_MEP (CONTINUED)

Register Group Name	Offset within Target	Instances and Address Spacing	Description	Details
PORT_PM	0x00004C00	88 0x00000008	Performance counters per port VOE per priority	Page 543
RX_VOE_PM	0x00005000	512 0x00000004	RX performance counters per service VOE per priority	Page 544
TX_VOE_PM	0x00005800	512 0x00000004	Tx performance counters per service VOE per priority	Page 545

5.9.1 OAM_MEP:COMMON

Parent: [OAM_MEP](#)

Instances: 1

This register group is mainly used to configure properties which are global across all VOEs.

This includes the following functions:

- VOP enable (MEP_CTRL)
- CPU extraction destination for different OAM PDU types (CPU_CFG, CPU_CFG_1)
- LOC scan configuration (CCM_PERIOD_CFG, CCM_CTRL)
- Generic OpCodes (OAM_GENERIC_CFG)
- Interrupt (MASTER_INT_CTRL)
- VOE Multicast Address (COMMON_MEP_MC_MAC_LSB, COMMON_MEP_MC_MAC_MSB)

The following registers are used for VOE specific configuration / status:

- VOE interrupt (INTR)
- Mapping of path VOEs (VOE_MAP_CTRL)
- VOE LM counter options (VOE_CNT_CTRL)

TABLE 5-411: REGISTERS IN COMMON

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
MEP_CTRL	0x00000000	1	Miscellaneous MEP controls	Page 501
CPU_CFG	0x00000004	1	Configuring destination for CPU copied frames.	Page 502
CPU_CFG_1	0x00000008	1	Configuring destination for CPU copied frames.	Page 504
OAM_SDLB_CPU_COPY	0x0000000C	1	Service policer selection	Page 506
OAM_GENERIC_CFG	0x00000010	8 0x00000004	Generic OAM PDU OpCodes configuration	Page 506
CCM_PERIOD_CFG	0x00000030	7 0x00000004	Configures time the the different CCM periods	Page 507
CCM_CTRL	0x0000004C	1	CCM scan configuration	Page 508
CCM_SCAN_STICKY	0x00000050	1	CCM scan diagnostic	Page 509
MASTER_INTR_CTRL	0x00000054	1	Interrupt control	Page 510

TABLE 5-411: REGISTERS IN COMMON (CONTINUED)

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
INTR	0x00000058	3 0x00000004	VOE interrupt	Page 510
VOE_CFG	0x00000064	64 0x00000004	Down-MEP or Up-MEP configuration	Page 510
UPMEP_LM_CNT_VLD	0x00000164	64 0x00000004	Up-MEP counter configuration	Page 510
VOE_MAP_CTRL	0x00000264	64 0x00000004	Path MEP configuration	Page 512
VOE_CNT_CTRL	0x00000364	75 0x00000004	MEP performance counter configuration	Page 513
COMMON_MEP_MC_MAC_LSB	0x00000490	1	OAM MEP multicast MAC address configuration (LSB)	Page 513
COMMON_MEP_MC_MAC_MSB	0x00000494	1	OAM MEP multicast MAC address configuration (MSB)	Page 514

5.9.1.1 OAM_MEP:COMMON:MEP_CTRL

Parent: [OAM_MEP:COMMON](#)

Instances: 1

This register contains various settings that are global to all VOEs.

TABLE 5-412: FIELDS IN MEP_CTRL

Field Name	Bit	Access	Description	Default
CCM_LM_UPD_RSV_ENA	14	R/W	The CCM PDU does not have a dedicated register for the RxFCf counter. However there is a reserved field which can be used to hold this value. Asserting this field enables updating reserved field for RX CCM LM PDU with Rx FCf.	0x1
LMR_UPD_RxFcL_ENA	13	R/W	Down-MEP only (No effect for Up-MEP): The LMR does not have any fields reserved to hold the RxFCf counter upon reception of an LMR frame. The VOE supports adding the RxFCf counter following the other three counters in the LMR PDU. This requires extending the first TLV from 12 --> 16 bytes. This solution is a proprietary solution, which allows forwarding the RxFCf value to an external CPU. Asserting this fields enables updating TLV offset to 16 and setting inserting RxFcLocal after TxFcb.	0x1

TABLE 5-412: FIELDS IN MEP_CTRL (CONTINUED)

Field Name	Bit	Access	Description	Default
CCM_SCAN_ENA	1	R/W	Enables or disables automated CCM scan to be used for CCM Loss of Continuity (LOC) handling.	0x0
MEP_ENA	0	R/W	Enables or disables the functionality of the VOEs in the VOP. The single VOEs can still be configured even though the VOP is disabled.	0x0
EXT_CPU_PORTMASK	12:2	R/W	Configures the front ports to be used for external CPU or FPGA handling. Ports used for OAM PDU injection or OAM PDU extraction must be asserted in this port-mask. 0: port is a front port 1: Port is a CPU extract port.	0x000

5.9.1.2 OAM_MEP:COMMON:CPU_CFG

Parent: [OAM_MEP:COMMON](#)

Instances: 1

For frames which are copied to the CPU, the actual destination to where such frames are forwarded is specified in this register.

Note:

If the destination of the OAM frame is an external port (*_EXT_PORT_ENA = 1), this port must be configured as an OAM extract port. This is done by asserting the OAM extract port in the port mask MEP_CTRL.EXT_CPU_PORTMASK.

OAM frames may be forwarded to the CPU using register: OAM_CPU_COPY_CTRL or because certain OAM PDUs are forwarded unconditionally to the CPU.

Most HW supported OAM PDUs have their own configuration, while some related PDUs share a single configuration.

The configuration for each OAM PDU consists of two bit fields:

1) *_EXT_PORT_ENA

Depending on the value of this bit field, the CPU copied frame is forwarded to an external port (=1) or an internal CPU extraction queue (=0).

2) *_CPU_QU

In case the frame is forwarded to an internal CPU extraction queue, this bit field indicates the target CPU extraction queue.

In case the frame is forwarded to an external, this bit field indicates the number of the physical port to which the frame is forwarded.

OAM PDU types that do not have a specific configuration use the default configuration DEF_COPY_CPU_QU.

The details are described for each bit field.

TABLE 5-413: FIELDS IN CPU_CFG

Field Name	Bit	Access	Description	Default
DEF_COPY_QU	28:25	R/W	This bit field is a default forwarding configuration which determines where to forward OAM PDU, which do not have a specific forwarding queue. Also frames extracted based on the following configuration: VOE: BASIC_CTRL.CPU_HITME_ONCE_CCM_TLV will be extracted to this queue. Note, this queue cannot be located on an external front port.	0x0
CPU_ERR_QU	23:20	R/W	Any errored frame which is extracted to the CPU is forwarded to the destination indicated by this bit field. Applicable frames are frames destined for the CPU extraction queue CPU_ERR_QU due to the following configurations: - CPU_HITME_ONCE_CCM_MEL_TOO_LOW - CPU_COPY_CCM_MEL_TOO_LOW - CPU_HITME_ONCE_MEL_TOO_LOW - CPU_COPY_MEL_TOO_LOW - MAC_ADDR_ERR_REDIR_ENA Further this queue is used if LBR or LTR PDUs are extracted for a MIP. Note, this queue cannot be located on an external front port.	0x0
DMR_EXT_PORT_ENA	19	R/W	This bit field configures the destination for the following OAM PDUs in case they are extracted to the CPU: 1) DMR 2) 1DM 0: Frames are sent to CPU extraction queue configured by DMR_CPU_QU 1: Frames are sent to front port configured by DMR_CPU_QU.	0x0
DMR_CPU_QU	18:15	R/W	Configures CPU extraction queue or external CPU port	0x0

TABLE 5-413: FIELDS IN CPU_CFG (CONTINUED)

Field Name	Bit	Access	Description	Default
LMR_EXT_PORT_ENA	14	R/W	This bit field configures the destination for the following OAM PDUs in case they are extracted to the CPU: 1) LMR 0: Frames are sent to CPU extraction queue configured by LMR_CPU_QU 1: Frames are sent to front port configured by LMR_CPU_QU.	0x0
LMR_CPU_QU	13:10	R/W	Configures CPU extraction queue or external CPU port	0x0
LBR_EXT_PORT_ENA	9	R/W	This bit field configures the destination for the following OAM PDUs in case they are extracted to the CPU: 1) LBR Also used for extracting RX TST frames (If TST frame handling is enabled using VOE:OAM_FWD_CTRL.GENERIC_FWD_MASK(6)) 0: Frames are sent to CPU extraction queue configured by LBR_CPU_QU 1: Frames are sent to front port configured by LBR_CPU_QU.	0x0
LBR_CPU_QU	8:5	R/W	Configures CPU extraction queue or external CPU port Also used for extracting RX TST frames (If TST frame handling is enabled using VOE:OAM_FWD_CTRL.GENERIC_FWD_MASK(6))	0x0
CCM_LM_EXT_PORT_ENA	4	R/W	This bit field configures the destination for the following OAM PDUs in case they are extracted to the CPU: 1) CCM-LM (I.e. CCM frame with counter values != 0) 0: Frames are sent to CPU extraction queue configured by CCM_LM_CPU_QU 1: Frames are sent to front port configured by CCM_LM_CPU_QU.	0x0
CCM_LM_CPU_QU	3:0	R/W	Configures CPU extraction queue or external CPU port to which CCM-LM frames are extracted.	0x0

5.9.1.3 OAM_MEP:COMMON:CPU_CFG_1

Parent: [OAM_MEP:COMMON](#)

Instances: 1

See description for CPU_CFG

TABLE 5-414: FIELDS IN CPU_CFG_1

Field Name	Bit	Access	Description	Default
LT_EXT_PORT_ENA	24	R/W	This bit field configures the destination for the following OAM PDUs in case they are extracted to the CPU: 1) LTR 2) LTM 0: Frames are sent to CPU extraction queue configured by LT_CPU_QU 1: Frames are sent to front port configured by LT_CPU_QU.	0x0
LT_CPU_QU	23:20	R/W	Configures CPU extraction queue or external CPU port	0x0
DMM_EXT_PORT_ENA	19	R/W	This bit field configures the destination for the following OAM PDUs in case they are extracted to the CPU: 1) DMM 0: Frames are sent to CPU extraction queue configured by DMM_CPU_QU 1: Frames are sent to front port configured by DMM_CPU_QU.	0x0
DMM_CPU_QU	18:15	R/W	Configures CPU extraction queue or external CPU port	0x0
LMM_EXT_PORT_ENA	14	R/W	This bit field configures the destination for the following OAM PDUs in case they are extracted to the CPU: 1) LMM 0: Frames are sent to CPU extraction queue configured by LMM_CPU_QU 1: Frames are sent to front port configured by LMM_CPU_QU.	0x0
LMM_CPU_QU	13:10	R/W	Configures CPU extraction queue or external CPU port	0x0
LBM_EXT_PORT_ENA	9	R/W	This bit field configures the destination for the following OAM PDUs in case they are extracted to the CPU: 1) LBM 0: Frames are sent to CPU extraction queue configured by LBM_CPU_QU 1: Frames are sent to front port configured by LBM_CPU_QU.	0x0
LBM_CPU_QU	8:5	R/W	Configures CPU extraction queue or external CPU port	0x0
CCM_CPU_QU	3:0	R/W	Configures CPU extraction queue or external CPU port to which CCM frames are extracted. Note, this queue cannot be located on an external front port.	0x0

5.9.1.4 OAM_MEP:COMMON:OAM_SDLB_CPU_COPY

Parent: [OAM_MEP:COMMON](#)

Instances: 1

TABLE 5-415: FIELDS IN OAM_SDLB_CPU_COPY

Field Name	Bit	Access	Description	Default
OAM_SDLB_CPU_IDX	9:0	R/W	Selected OAM frames which are forwarded to the CPU pass through this register. The OAM frames for which this S-DLB is applied depend on whether the VOE is configured for Down-MEP or Up-MEP. Further it depends on the loopback mode configured for the VOE. Description is to be updated.	0x000

5.9.1.5 OAM_MEP:COMMON:OAM_GENERIC_CFG

Parent: [OAM_MEP:COMMON](#)

Instances: 8

Down-MEP:

The VOE supports eight generic OpCodes, which can be extracted to the CPU or forwarded independently of other OpCodes. An OpCode which is not explicitly supported and which is not configured as a generic OpCode is treated as an unknown OpCode.

The configuration of Generic OpCodes is global across all the VOEs.

The extraction and forwarding of generic OpCodes can be configured individually per VOE.

Note that VOE.OAM_FWD_CTRL.OAM_FWD_MASK(6) is reserved for enabling TST HW processing in the VOE. Hence if forwarding of Generic OpCode(6) is to be used, there must not be TST frames in the frame flow to be handled by the VOE.

Up-MEP:

GenericOpcode(8 - 1) is reserved for special use and cannot be used by the user. Other generic OpCodes can be used to select OAM PDUs which should be handled as data for SLA testing. This is intended to be used for TST frames.

TABLE 5-416: FIELDS IN OAM_GENERIC_CFG

Field Name	Bit	Access	Description	Default
GENERIC_OPCODE_EX-T_PORT_ENA	16	R/W	Configures if GENERIC_OPCODE_CPU_QU is a CPU extraction queue or a front port 0: Frames are sent to CPU extraction queue configured by GENERIC_OPCODE_CPU_QU 1: Frames are sent to front port configured by GENERIC_OPCODE_CPU_QU.	0x0

TABLE 5-416: FIELDS IN OAM_GENERIC_CFG (CONTINUED)

Field Name	Bit	Access	Description	Default
GENERIC_OPCODE_CPU_QU	11:8	R/W	Configures CPU extraction queue or external CPU port, according to the configuration of bit field: GENERIC_OPCODE_EXT_PORT_ENA	0x0
GENERIC_OPCODE_VAL	7:0	R/W	Each OpCode value is handled according to the used replication index. Handling is enabled by the corresponding VOE configuration: VOE:OAM_CPU_COPY_CTRL:GENERIC_COPY_MASK VOE:OAM_FWD_CTRL:GENERIC_FWD_MASK VOE:OAM_CNT_OAM_CTRL:GENERIC_OAM_CNT_MASK VOE:OAM_CNT_DATA_CTRL:GENERIC_DATA_CNT_MASK	0x00
GENERIC_DMAC_CHK_DIS	17	R/W	If this bit is asserted, the DMAC check is disabled for this generic OpCode, regardless of the configuration of the bit field: VOE:BASIC_CTRL.RX_DMAC_CHK_SEL 0: Perform DMAC check according to the value configured in: VOE:BASIC_CTRL.RX_DMAC_CHK_SEL 1: No DMAC for this Generic OpCode.	0x0

5.9.1.6 OAM_MEP:COMMON:CCM_PERIOD_CFG**Parent:** [OAM_MEP:COMMON](#)**Instances:** 7

7 Independent timers are implemented to be used with the Loss of Continuity (LOC) scan.

This register implements the timeout period for every one of the 7 LOC counters.

The timeout period is specified in the number of CCM base ticks between every LOC timeout.

For configuration of the CCM base tick, see bit field:

- CCM_CTRL.CCM_BASE_TICK_CNT

The default value for the CCM base tick is 198,4 ns, which means that the LOC timer counters are incremented every 198,4 ns.

A VOE can be configured for LOC checking based on one of the 7 timeout counters.

When the timeout counter configured for a VOE has timed out 3 times without a CCM frame being received, the LOC interrupt is asserted.

TABLE 5-417: FIELDS IN CCM_PERIOD_CFG

Field Name	Bit	Access	Description	Default
CCM_PERIOD_VAL	31:0	R/W	Configures Period timers in unit of 198.2 ns Default: 3.3ms. 51: 10us 505: 100us 5045: 1 ms 16650: 3.3ms 50454: 10ms 504541: 100ms 5045409: 1s	0x0000410A

5.9.1.7 OAM_MEP:COMMON:CCM_CTRL

Parent: [OAM_MEP:COMMON](#)

Instances: 1

Configures CCM hardware scan

TABLE 5-418: FIELDS IN CCM_CTRL

Field Name	Bit	Access	Description	Default
CCM_BASE_TICK_CNT	31:24	R/W	Specifies the number of system clock cycles for each CCM base time tick. The current system clock of Serval is 6.4 ns (156 MHz) With the current default setting (=31), the CCM base time tick occurs every 198.4 ns. For every CCM base time tick, the 7 LOC counters, are incremented. The Period of the 7 LOC counters are specified in register: - COMMON:CCM_PERIOD_VAL 0: Disable automatic HW scan 1: One clock between interval increment ... n: n clock between interval increment	0x1F
CCM_SPACE_BETWEEN_ENTRY_SCAN	23:16	R/W	Specifies the number of clk cycle before another scan entry can be attempted. This can be used to space the LOC scanning of the VOE's	0x14

TABLE 5-418: FIELDS IN CCM_CTRL (CONTINUED)

Field Name	Bit	Access	Description	Default
CCM_FORCE_HW_SCAN_ENA	14:8	R/W	Specifies an instant CCM miss count increment for this type (once current scan completes, which can be ensured using CCM_FORCE_HW_SCAN_STOP_ENA). The bit is cleared by HW when increment is complete 0: No force x1x: Force one CCM miss count increment for that type	0x00
CCM_FORCE_HW_SCAN_STOP_ENA	0	R/W	Forces stop of current CCM miss count scan. 0: No force 1: Force stop of current CCM miss count scan The bit is cleared by HW when scan stop is complete	0x0

5.9.1.8 OAM_MEP:COMMON:CCM_SCAN_STICKY

Parent: [OAM_MEP:COMMON](#)

Instances: 1

Sticky bits indicate the current status of the LOC scanning.

TABLE 5-419: FIELDS IN CCM_SCAN_STICKY

Field Name	Bit	Access	Description	Default
CCM_SCAN_ONGOING_STATUS	30:24	R/O	High when CCM miss count scan is ongoing. 0: No event has occurred 1: CCM Miss count scan ongoing.	0x00
CCM_SCAN_COMPLETED_STICKY	2	Sticky	Set when CCM_SCAN completes. 0: No event has occurred 1: CCM miss count scan completed Bit is cleared by writing a 1 to this position.	0x0
CCM_SCAN_STARTED_STICKY	1	Sticky	Set when CCM_SCAN starts. 0: No event has occurred 1: CCM miss count scan started Bit is cleared by writing a 1 to this position.	0x0
CCM_SCAN_START_DELAYED_STICKY	0	Sticky	Set when a CCM scan could not start because a scan is already ongoing. The configured Period time in OMA_MEP:CCM_PERIOD must be incremented. 0: No event has occurred 1: Scan could not start in time Bit is cleared by writing a 1 to this position.	0x0

VSC7416-01

5.9.1.9 OAM_MEP:COMMON:MASTER_INTR_CTRL

Parent: [OAM_MEP:COMMON](#)

Instances: 1

This is the combined interrupt output from the VOP.

To determine the VOE source of the interrupt, read register OAM_MEP::INTR.

TABLE 5-420: FIELDS IN MASTER_INTR_CTRL

Field Name	Bit	Access	Description	Default
OAM_MEP_INTR	8	R/O	Status of the combined interrupt in the VOP.	0x0
OAM_MEP_INTR_ENA	25	R/W	Configures if OAM_MEP_INTR interrupts the CPU. 0: Disable interrupt 1: Enable interrupt	0x0

5.9.1.10 OAM_MEP:COMMON:INTR

Parent: [OAM_MEP:COMMON](#)

Instances: 3

Interrupt per VOE. To enable interrupts out of the VOP, see MASTER_INTR_CTRL.

TABLE 5-421: FIELDS IN INTR

Field Name	Bit	Access	Description	Default
CCM_INTR	31:0	R/O	Status of interrupts per VOE (1 bit per VOE). The interrupts are cleared by clearing the sticky bit set in OAM_MEP:VOE:STICKY or by disabling the interrupt source in OAM_MEP:VOE:INTR_ENA. 0: No interrupt 1: Interrupt asserted	0x00000000

5.9.1.11 OAM_MEP:COMMON:VOE_CFG

Parent: [OAM_MEP:COMMON](#)

Instances: 64

TABLE 5-422: FIELDS IN VOE_CFG

Field Name	Bit	Access	Description	Default
UPMEP_VOE	0	R/W	Configure whether the VOE is implementing an Up-MEP or Down-MEP. 0: VOE implements Down-MEP 1: VOE implements Up-MEP	0x0

5.9.1.12 OAM_MEP:COMMON:UPMEP_LM_CNT_VLD

Parent: [OAM_MEP:COMMON](#)

Instances: 64

For Up-MEP, the LM counter values cannot be stored into the OAM frame IFH and forwarded to the CPU. Instead the counter values are sampled when the Rx LM OAM PDU is received by the VOE and they must be read by SW. For each of the LM OAM PDUs supported, the corresponding fields in this register indicate whether there is an LM counter value ready for reading.

TABLE 5-423: FIELDS IN UPMEP_LM_CNT_VLD

Field Name	Bit	Access	Description	Default
CCM_LM_RX_CNT_VLD	2	R/O	This field is asserted by the VOE HW when an Up-MEP has sampled the RX LM counter and saved the value into the following register: VOE:CCM_CAP_RX_FCL.CCM_CAP_RX_FCL This is done in response to receiving an CCM-LM frame. The field is cleared when SW reads the counter value in the above register. If the CPU has not cleared this bit before the next CCM-LM frame is received, the VOE asserts the following sticky bit to indicate that an RX_LM sample was lost: VOE:UPMEP_LM_CNT_STICKY.CCM_LM_RX_CNT_STICKY	0x0
LMM_RX_CNT_VLD	1	R/O	This field is asserted by the VOE HW when an Up-MEP has sampled the RX LM counter and saved the value into the following register: VOE:UPMEP_LMR_RX_LM_CNT.UPMEP_LMR_RX_LM_CNT This is done in response to receiving an LMM frame. The field is cleared when CPU reads the counter value in the above register. If the CPU has not cleared this bit before the next LMM frame is received, the VOE asserts the following sticky bit to indicate that an RX_LM sample was lost. VOE:UPMEP_LM_CNT_STICKY.LMM_RX_CNT_STICKY	0x0

TABLE 5-423: FIELDS IN UPMEP_LM_CNT_VLD (CONTINUED)

Field Name	Bit	Access	Description	Default
LMR_RX_CNT_VLD	0	R/O	This field is asserted by the VOE HW when an Up-MEP has sampled the RX LM counter and saved the value into the following register: VOE:CCM_CAP_TX_FCF.CCM_CAP_TX_FCF This is done in response to receiving an LMR frame. The field is cleared when CPU reads the counter value in the above register. If the CPU has not cleared this bit before the next LMR frame is received, the VOE asserts the following sticky bit to indicate that an RX_LM sample was lost: VOE:UPMEP_LM_CNT_STICKY.LMR_RX_CNT_STICKY	0x0

5.9.1.13 OAM_MEP:COMMON:VOE_MAP_CTRL

Parent: [OAM_MEP:COMMON](#)

Instances: 64

Configures a service VOE to be part of a path VOE data flow.

When a path VOE is assigned to a service VOE, it causes all data on the service VOE (including data and OAM PDUs) to be counted as data in the path VOE's LM counters.

TABLE 5-424: FIELDS IN VOE_MAP_CTRL

Field Name	Bit	Access	Description	Default
PATH_VOEID	12:7	R/W	Configures the associated working path VOE Must be enabled by PATH_VOEID_ENA.	0x00
PATH_VOEID_P	6:1	R/W	Configures the associated PROTECT path VOE when a service is protected. This is used if the associated service is protected via the IPT table (ANA:IPT:IPT_CFG) Must be enabled by PATH_VOEID_ENA.	0x00
PATH_VOEID_ENA	0	R/W	Configures if a service MEP is part of a path MEP data. 0: No path VOE is assigned to the service VOE 1: A path VOE is assigned to the service VOE.	0x0

5.9.1.14 OAM_MEP:COMMON:VOE_CNT_CTRL

Parent: [OAM_MEP:COMMON](#)

Instances: 75

Configures the LM counters per VOE.

TABLE 5-425: FIELDS IN VOE_CNT_CTRL

Field Name	Bit	Access	Description	Default
CNT_YELLOW_ENA	8	R/W	Configures if yellow frames are included in the LM counters.	0x1
PRIO_CNT_MASK	7:0	R/W	Configures if priorities are counted separately or as priority 7. Each bit in the bit field represents a priority. If a priority bit is set in the MASK, frames with the corresponding priority are counted in a separate priority counter. If priority bit is not set in MASK the priority will be counted as priority 7. The priority mask is respected by a Service VOE when updating any configured Path VOE and the Port VOE on frame reception / transmission. Hence the Service VOE will look at the PRI-O_CNT_MASK for any configured Path VOE and the Port VOE and update their LM counters according to the priority configuration. 0: Priority is counted in priority 7 of the performance counters 1: Priority is counted separately in the performance counters.	0xFF

5.9.1.15 OAM_MEP:COMMON:COMMON_MEP_MC_MAC_LSB

Parent: [OAM_MEP:COMMON](#)

Instances: 1

Each VOE can be addressed using either a common multicast MAC address or a VOE specific unicast MAC address.

This register configures the multicast address common to all the VOE's.

The full MAC address is a concatenation of the following registers:

- 1) COMMON:COMMON_MEP_MC_MAC_LSB
- 2) COMMON:COMMON_MEP_MC_MAC_MSB

The DMAC check to be performed for each VOE is configured in VOE:BASIC_CTRL.RX_DMACHK_SEL.

TABLE 5-426: FIELDS IN COMMON_MEP_MC_MAC_LSB

Field Name	Bit	Access	Description	Default
MEP_MC_MAC_LSB	30:3	R/W	Common VOE multicast Address - LSB	0xC200003

5.9.1.16 OAM_MEP:COMMON:COMMON_MEP_MC_MAC_MSB

Parent: [OAM_MEP:COMMON](#)

Instances: 1

Common VOE multicast Address - MSB

TABLE 5-427: FIELDS IN COMMON_MEP_MC_MAC_MSB

Field Name	Bit	Access	Description	Default
MEP_MC_MAC_MSB	15:0	R/W	Common VOE multicast Address - MSB	0x0180

5.9.2 OAM_MEP:VOE

Parent: [OAM_MEP](#)

Instances: 75

This register group contains registers specific to each VOE.

Service/path VOEs are indexed: 0-63

Port VOEs are indexed: 64 (port 0) - 74 (port 10)

TABLE 5-428: REGISTERS IN VOE

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
BASIC_CTRL	0x00000000	1	Configures miscellaneous parameters	Page 515
MEL_CTRL	0x00000004	1	Configures handles in relation to MEL	Page 517
OAM_CPU_COPY_CTRL	0x00000008	1	Configuration of CPU copy for the supported OAM PDU OpCodes	Page 518
OAM_FWD_CTRL	0x0000000C	1	Configuration forwarding for the supported OAM PDU OpCodes	Page 519
OAM_SDLB_CTRL	0x00000010	1	Override S-DLB for OAM frames	Page 521
OAM_CNT_OAM_CTRL	0x00000014	1	Configuration which OAM PDU are to be included in selected counter	Page 521
OAM_CNT_DATA_CTRL	0x00000018	1	Configuration of which OAM PDUs should be counted as data	Page 523
MEP_UC_MAC_LSB	0x0000001C	1	OAM MEP MAC unicast address configuration (LSB)	Page 524

TABLE 5-428: REGISTERS IN VOE (CONTINUED)

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
MEP_UC_MAC_MSB	0x00000020	1	OAM MEP MAC unicast address configuration (MSB)	Page 525
OAM_HW_CTRL	0x00000024	1	OAM handling configuration	Page 525
LOOPBACK_CTRL	0x00000028	1	OAM loopback configuration	Page 527
LBR_RX_FRM_CNT	0x0000002C	1	Count number of received LBR frames	Page 529
LBR_RX_TRANSID_CFG	0x00000030	1	Configuration of expected transaction ID	Page 529
LBM_TX_TRANSID_UPDATE	0x00000034	1	Configuration of Tx transaction ID update	Page 530
LBM_TX_TRANSID_CFG	0x00000038	1	Configuration of transmitted transaction ID	Page 530
CCM_CFG	0x0000003C	1	Miscellaneous CCM configuration	Page 531
CCM_RX_VLD_FC_CNT	0x00000040	1	Count number of received good CCM frames	Page 533
CCM_RX_INVLD_FC_CNT	0x00000044	1	Count number of received invalid CCM frames	Page 534
CCM_CAP_TX_FCf	0x00000048	1	Latest received CCM_TX_FCf value	Page 534
CCM_CAP_RX_FCI	0x0000004C	1	Latest RX_FCI counter value	Page 534
UPMEP_LMR_RX_LM_CNT	0x00000050	1	Up-MEP Rx loss measurement frame counter	Page 535
CCM_TX_SEQ_CFG	0x00000054	1	Configuration of CCM Sequence number	Page 536
CCM_RX_SEQ_CFG	0x00000058	1	Configuration of expected CCM Sequence number	Page 536
CCM_MEPID_CFG	0x0000005C	1	Configuration of CCM MEPID	Page 537
CCM_MEGID_CFG	0x00000060	12 0x00000004	Configuration of CCM MEGID	Page 537
OAM_RX_STICKY	0x00000090	1	Indicate reception of OpCode	Page 538
STICKY	0x00000094	1	CCM sticky status indications	Page 539
UPMEP_LM_CNT_STICKY	0x00000098	1	Up-MEP loss measurement sticky indications	Page 540
INTR_ENA	0x0000009C	1	Configures interrupt sources	Page 541
RX_SEL_OAM_CNT	0x000000A0	1	Count selected OAM PDU received by VOE	Page 542
RX_OAM_FRM_CNT	0x000000A4	1	Count non-selected OAM PDU received by VOE	Page 542
TX_SEL_OAM_CNT	0x000000A8	1	Count selected OAM PDU transmitted by VOE	Page 543
TX_OAM_FRM_CNT	0x000000AC	1	Count non-selected OAM PDU transmitted by VOE	Page 543

5.9.2.1 OAM_MEP:VOE:BASIC_CTRL

Parent: [OAM_MEP:VOE](#)

Instances: 1

Basic VOE configuration.

TABLE 5-429: FIELDS IN BASIC_CTRL

Field Name	Bit	Access	Description	Default
MAC_ADDR_ERR_REDIR_ENA	10	R/W	Configures if frames failing the DMAC check should be redirected to the CPU extraction queue OAM_MEP::CPU_CFG.CPU_ERR_QU.	0x0
RX_DMACHK_SEL	9:8	R/W	Configures the DA check performed when received from the peer MEP. This check can be disabled for Generic OpCodes. See COMMON:OAM_GENERIC_CFG.GENERIC_DMACHK_DIS 0: No DMAC check 1: Check DMAC against MEP_UC_MAC_MSB and MEP_UC_MAC_LSB 2: Check DMAC against MEP_MC_MAC_MSB and MEP_MC_MAC_LSB 3: Check DMAC against either MEP_UC_MAC_MSB and MEP_UC_MAC_LSB or MEP_MC_MAC_MSB and MEP_MC_MAC_LSB	0x0
CPU_HITME_ONCE_CCM_TLV	7	R/W	802.1ag specifies that CCM frames may have more than one TLV. This bit field allows the next CCM frame with Second TLV to CPU. (IEEE802.1AG relevant only). Frames extracted to the CPU based on this setting are forwarded to CPU extraction queue COMMON:CPU_CFG.DEF_COPY_QU. 0: Disable (no CPU hitme once copy) 1: Enable the next CCM_PDU with a second TLV to be sent to CPU. The bit is cleared by HW when a CPU copy has been made.	0x0
CCM_ERR_CPU_HITME_ONCE	6	R/W	Send next CCM frame that fails the required check to CPU configured in OAM_MEP::CPU_CFG.CPU_ERR_QU. 0: Disable (no CPU hitme once copy) 1: Enable the next frame failing the CCM PDU check to be sent to CPU. The bit is cleared by HW when a CPU copy has been made.	0x0

TABLE 5-429: FIELDS IN BASIC_CTRL (CONTINUED)

Field Name	Bit	Access	Description	Default
CCM_ERR_CPU_COPY_ENA	5	R/W	Configures if CCM frame failing the required check is sent to CPU configured in OAM_MEP::CPU_CFG.CPU_ERR_QU.	0x0
CCM_NXT_CPU_HITME_ONCE	4	R/W	Extract next CCM or CCM-LM OAM PDU frame to CPU according to configuration in OAM_MEP::CCM_CFG_1.CCM_CPU_QU 0: Disable (no CPU hitme once copy) 1: Enable the next frame to be sent to CPU. The bit is cleared by HW when a CPU copy has been made.	0x0
MIP_UNICAST_ONLY_COPY_ENA	3	R/W	Configures the VOE only make CPU copy of unicast DMAC addresses. 0: Allows unicast and multicast copy handling 1: Allows only unicast copy handling (MIP)	0x0
VOE_AS_MIP_ENA	2	R/W	Down-MEP only (Must be = 0 for Up-MEP) Configures the VOE to be a MIP. 0: Allows MEP handling 1: Enable MIP handling	0x0
VOE_ENA	1	R/W	Configures if MEP is active or not. 0: Disable MEP handling 1: Enable MEP handling	0x1
UPMEP_KILL_ALL_RX	0	R/W	Up-MEP only (Has not effect on Down-MEP) Asserting this bit causes all DATA RX frames received at the Up-MEP to be discarded. OAM frames are unaffected by this setting and they may be looped or extracted as is standard for every PDU type.	0x0

5.9.2.2 OAM_MEP:VOE:MEL_CTRL

Parent: OAM_MEP:VOE

Instances: 1

VOE configuration related to MEL checking.

TABLE 5-430: FIELDS IN MEL_CTRL

Field Name	Bit	Access	Description	Default
MEL_VAL	19:17	R/W	This register configures the MEL value for the VOE. OAM PDUs with MEL above this MEL_VAL are not effected. OAM PDUs with MEL equal to this MEL_VAL are optionally handled by HW or copied to CPU.	0x0

TABLE 5-430: FIELDS IN MEL_CTRL (CONTINUED)

Field Name	Bit	Access	Description	Default
UPMEP_LB	16	R/W	When VOE is configured as an Up-MEP, it can run in two modes: 1) In-service - in this configuration there is customer traffic on the service. This configuration does implement HW support for loopback of OAM frames. 2) Loopback - in this configuration there are no costumer frames in the service. This configuration implements HW loopback of OAM frames. Note: This bit is not used when the VOE is configured for Down-MEP 0: In-service (No HW loopback) 1: Loopback (HW loopback)	0x0
MEP_PORTMASK	15:5	R/W	For VOE configured for Up-MEP operation, this register points to the front port where the Up-MEP is located. For VOE configured for Down-MEP, this register is not used.	0x7FF
CPU_HITME_ONCE_MEL_TOO_LOW	3	R/W	Configures if the first OAM PDU (excluding CCM) discarded due to a MEL too low is copied to the CPU extraction queue configured in OAM_MEP::CPU_CFG.CPU_ERR_QU. 0: Disable (no CPU hit me once copy) 1: Enable the next frame failing the MEL check to be sent to CPU The bit is cleared by HW when a CPU copy has been made.	0x0
CPU_COPY_MEL_TOO_LOW	2	R/W	Configures if OAM PDUs (excluding CCM) failing the MEL check (MEL too low) are copied to the CPU extraction queue configured in OAM_MEP::CPU_CFG.CPU_ERR_QU.	0x0
CPU_HITME_ONCE_CCM_MEL_TOO_LOW	1	R/W	Configures if the first OAM CCM PDU discarded due to a MEL too low is copied to the CPU extraction queue configured in OAM_MEP::CPU_CFG.CPU_ERR_QU. 0: Disable (no CPU hit me once copy) 1: Enable the next frame failing the MEL check to be sent to CPU The bit is cleared by HW when a CPU copy has been made.	0x0
CPU_COPY_CCM_MEL_TOO_LOW	0	R/W	Configures if OAM CCM PDUs failing the MEL check (MEL too low) are copied to the CPU extraction queue configured in OAM_MEP::CPU_CFG.CPU_ERR_QU.	0x0

5.9.2.3 OAM_MEP:VOE:OAM_CPU_COPY_CTRL

Parent: [OAM_MEP:VOE](#)

Instances: 1

Configures CPU copy for the supported OAM PDU OpCodes. This CPU copy functionality operates independently of OAM_FWD_CTRL

TABLE 5-431: FIELDS IN OAM_CPU_COPY_CTRL

Field Name	Bit	Access	Description	Default
GENERIC_COPY_MASK	19:12	R/W	Enable / disable CPU copy of Generic OpCode configured in OAM_MEP::OAM_GENERIC_CFG 0: No CPU copy 1: Copy to CPU	0x00
UNK_OPCODE_CPU_COPY_ENA	11	R/W	Configures CPU copy of all non specific and generic OpCodes 0: No CPU copy 1: Copy to CPU	0x0
LTM_CPU_COPY_ENA	10	R/W	CPU copy of OAM LTM PDU. 0: No CPU copy 1: Copy to CPU	0x0
LTR_CPU_COPY_ENA	9	R/W	CPU copy of OAM LTR PDU. 0: No CPU copy 1: Copy to CPU	0x0
LMM_CPU_COPY_ENA	8	R/W	CPU copy of OAM LMM PDU. 0: No CPU copy 1: Copy to CPU	0x0
LMR_CPU_COPY_ENA	7	R/W	CPU copy of OAM LMR PDU. 0: No CPU copy 1: Copy to CPU	0x0
LBM_CPU_COPY_ENA	6	R/W	CPU copy of OAM LBM PDU. 0: No CPU copy 1: Copy to CPU	0x0
LBR_CPU_COPY_ENA	5	R/W	CPU copy of OAM LBR PDU. 0: No CPU copy 1: Copy to CPU	0x0
DMM_CPU_COPY_ENA	4	R/W	CPU copy of OAM DMM PDU. 0: No CPU copy 1: Copy to CPU	0x0
DMR_CPU_COPY_ENA	3	R/W	CPU copy of OAM DMR PDU. 0: No CPU copy 1: Copy to CPU	0x0
SM_CPU_COPY_ENA	2	R/W	CPU copy of OAM 1DM PDU. 0: No CPU copy 1: Copy to CPU	0x0
CCM_CPU_COPY_ENA	1	R/W	CPU copy of OAM CCM PDU. 0: No CPU copy 1: Copy to CPU	0x0
CCM_LM_CPU_COPY_ENA	0	R/W	CPU copy of OAM CCM-LM PDU.	0x0

5.9.2.4 OAM_MEP:VOE:OAM_FWD_CTRL

Parent: [OAM_MEP:VOE](#)

Instances: 1

Configures CPU forwarding for the supported OAM PDU OpCodes. This forwarding functionality operates independently of OAM_CPU_COPY_CTRL.

VSC7416-01

This can be used to allow forwarding of OAM PDUs which would normally be stopped by the Rx or Tx MEL check.

The only currently know use for this is Ring PDU (OpCode: 40)

TABLE 5-432: FIELDS IN OAM_FWD_CTRL

Field Name	Bit	Access	Description	Default
GENERIC_FWD_MASK	18:11	R/W	Enable / disable forwarding of Generic OpCode configured in OAM_MEP::OAM_GENERIC_CFG Bit (6) is reserved for special purpose. This bit is used to enable TST frame processing by the VOE. Hence generic OpCode(6) should not be used for other purposes. 0x:x No forwarding of frames with MEL matching MEL_CTRL.MEL_VAL x1x: Allow forwarding	0x00
LTM_FWD_ENA	10	R/W	Forwarding of OAM LTM PDU. 0: No forwarding of frames with MEL matching MEL_CTRL.MEL_VAL 1: Allow forwarding	0x0
LTR_FWD_ENA	9	R/W	Forwarding of OAM LTR PDU. 0: No forwarding of frames with MEL matching MEL_CTRL.MEL_VAL 1: Allow forwarding	0x0
LMM_FWD_ENA	8	R/W	Forwarding of OAM LMM PDU. 0: No forwarding of frames with MEL matching MEL_CTRL.MEL_VAL 1: Allow forwarding	0x0
LMR_FWD_ENA	7	R/W	Forwarding of OAM LMR PDU. 0: No forwarding of frames with MEL matching MEL_CTRL.MEL_VAL 1: Allow forwarding	0x0
LBM_FWD_ENA	6	R/W	Forwarding of OAM LBM PDU. 0: No forwarding of frames with MEL matching MEL_CTRL.MEL_VAL 1: Allow forwarding	0x0
LBR_FWD_ENA	5	R/W	Forwarding of OAM LBR PDU. 0: No forwarding of frames with MEL matching MEL_CTRL.MEL_VAL 1: Allow forwarding	0x0
DMM_FWD_ENA	4	R/W	Forwarding of OAM DMM PDU. 0: No forwarding of frames with MEL matching MEL_CTRL.MEL_VAL 1: Allow forwarding	0x0
DMR_FWD_ENA	3	R/W	Forwarding of OAM DMR PDU. 0: No forwarding of frames with MEL matching MEL_CTRL.MEL_VAL 1: Allow forwarding	0x0
SM_FWD_ENA	2	R/W	Forwarding of OAM 1DM PDU. 0: No forwarding of frames with MEL matching MEL_CTRL.MEL_VAL 1: Allow forwarding	0x0

TABLE 5-432: FIELDS IN OAM_FWD_CTRL (CONTINUED)

Field Name	Bit	Access	Description	Default
CCM_FWD_ENA	1	R/W	Forwarding of OAM CCM PDU. 0: No forwarding of frames with MEL matching MEL_CTRL.MEL_VAL 1: Allow forwarding	0x0
CCM_LM_FWD_ENA	0	R/W	Forwarding of OAM CCM-LM PDU. 0: No forwarding of frames with MEL matching MEL_CTRL.MEL_VAL 1: Allow forwarding	0x0

5.9.2.5 OAM_MEP:VOE:OAM_SDLB_CTRL**Parent:** [OAM_MEP:VOE](#)**Instances:** 1

It is possible to override the SDLB for the OAM frames not counted as data frames.

I.e. if service DLB overriding is enabled, the SDLB for OAM frames, not selected to be counted as data frames by register OAM_CNT_DATA_CTRL, is overwritten by the S-DLB value configured in OAM_SDLB_IDX.

TABLE 5-433: FIELDS IN OAM_SDLB_CTRL

Field Name	Bit	Access	Description	Default
OAM_SDLB_ENA	10	R/W	Enable overriding the S-DLB for OAM frames not counted as data frames.	0x0
OAM_SDLB_IDX	9:0	R/W	If enabled by OAM_SDLB_ENA, the S-DLB for OAM frames are overwritten by this value.	0x000

5.9.2.6 OAM_MEP:VOE:OAM_CNT_OAM_CTRL**Parent:** [OAM_MEP:VOE](#)**Instances:** 1

The OAM frames handled by the VOE can be counted separately in Rx and Tx direction.

In each direction there are two counters:

1) Default OAM counter

This counter counts all the PDU types which are not selected using the OAM_CNT_OAM_CTRL register:

- RX_OAM_FRM_CNT
- TX_OAM_FRM_CNT

2) Selected OAM counter:

This counter counts all the PDU types selected for counting using the OAM_CNT_OAM_CTRL register:

- RX_SEL_OAM_CNT
- TX_SEL_OAM_CNT

Any OAM PDU is counted in exactly one of the above registers.

VSC7416-01

I.e. as default all OAM PDUs are not selected, and they are all counted in the default OAM counters: RX / TX _ OAM _ FRM _ CNT.

Using this register (OAM_CNT_OAM_CTRL), PDUs can be moved to the selected counters:
RX / TX SEL_OAM_CNT.

The selection of OAM PDUs for the selected counter is done commonly for the Tx and Rx direction.

TABLE 5-434: FIELDS IN OAM_CNT_OAM_CTRL

Field Name	Bit	Access	Description	Default
GENERIC_OAM_CNT_MASK	19:12	R/W	Enable / disable count as selected OAM of Generic OpCodes per MEP per generic OpCode. x0x: Count as other OAM x1x: Count as selected OAM	0x00
UNK_OPCODE_OAM_CNT_ENA	11	R/W	Enable / disable count as selected OAM for all non specific and generic OpCodes 0: Count as other OAM 1: Count as selected OAM	0x0
LTM_OAM_CNT_ENA	10	R/W	Enable / disable count of OAM PDU LTM as selected OAM. 0: Count as other OAM 1: Count as selected OAM	0x0
LTR_OAM_CNT_ENA	9	R/W	Enable / disable count of OAM PDU LTR as selected OAM. 0: Count as other OAM 1: Count as selected OAM	0x0
LMM_OAM_CNT_ENA	8	R/W	Enable / disable count of OAM PDU LMM as selected OAM. 0: Count as other OAM 1: Count as selected OAM	0x0
LMR_OAM_CNT_ENA	7	R/W	Enable / disable count of OAM PDU LMR as selected OAM. 0: Count as other OAM 1: Count as selected OAM	0x0
LBM_OAM_CNT_ENA	6	R/W	Enable / disable count of OAM PDU LBM as selected OAM. 0: Count as other OAM 1: Count as selected OAM	0x0
LBR_OAM_CNT_ENA	5	R/W	Enable / disable count of OAM PDU LBR as selected OAM. 0: Count as other OAM 1: Count as selected OAM	0x0
DMM_OAM_CNT_ENA	4	R/W	Enable / disable count of OAM PDU DMM as selected OAM. 0: Count as other OAM 1: Count as selected OAM	0x0
DMR_OAM_CNT_ENA	3	R/W	Enable / disable count of OAM PDU DMR as selected OAM. 0: Count as other OAM 1: Count as selected OAM	0x0

TABLE 5-434: FIELDS IN OAM_CNT_OAM_CTRL (CONTINUED)

Field Name	Bit	Access	Description	Default
SM_OAM_CNT_ENA	2	R/W	Enable / disable count of OAM PDU 1DM as selected OAM. 0: Count as other OAM 1: Count as selected OAM	0x0
CCM_OAM_CNT_ENA	1	R/W	Enable / disable count of OAM PDU CCM as selected OAM. 0: Count as other OAM 1: Count as selected OAM	0x0
CCM_LM_OAM_CNT_ENA	0	R/W	Enable / disable count of OAM PDU CCM with LM content as selected OAM.	0x0

5.9.2.7 OAM_MEP:VOE:OAM_CNT_DATA_CTRL**Parent:** [OAM_MEP:VOE](#)**Instances:** 1

Default behavior is that all OAM PDUs handled by a VOE (i.e. MEG levels between OAM PDU and VOE must match) are not counted as data. This is according to Y.1731

Using this register (OAM_CNT_DATA_CTRL) it is possible to configure the OAM PDUs separately to be counted as data.

Frames are counted separately in the Rx and Tx direction.

The data counters are located:

Service VOE:

RX_VOE_PM.MEP_RX_FRM_CNT

TX_VOE_PM.MEP_TX_FRM_CNT

Port VOE:

PORT_PM.RX_FRM_CNT

PORT_PM.TX_FRM_CNT

TABLE 5-435: FIELDS IN OAM_CNT_DATA_CTRL

Field Name	Bit	Access	Description	Default
GENERIC_DATA_CNT_MASK	18:11	R/W	Enable / disable data count of Generic OpCodes per MEP per generic OpCode.	0x00
UNK_OPCODE_DATA_CNT_ENA	10	R/W	Configures data count of all non specific and generic OpCodes 0: Do not count as data 1: Count as data	0x0
LTM_DATA_CNT_ENA	9	R/W	Enable / disable count of OAM PDU LTM as data in performance counters. 0: Do not count as data 1: Count as data	0x0

TABLE 5-435: FIELDS IN OAM_CNT_DATA_CTRL (CONTINUED)

Field Name	Bit	Access	Description	Default
LTR_DATA_CNT_ENA	8	R/W	Enable / disable count of OAM PDU LTR as data in performance counters. 0: Do not count as data 1: Count as data	0x0
LMM_DATA_CNT_ENA	7	R/W	Enable / disable count of OAM PDU LMM as data in performance counters. 0: Do not count as data 1: Count as data	0x0
LMR_DATA_CNT_ENA	6	R/W	Enable / disable count of OAM PDU LMR as data in performance counters. 0: Do not count as data 1: Count as data	0x0
LBM_DATA_CNT_ENA	5	R/W	Enable / disable count of OAM PDU LBM as data in performance counters. 0: Do not count as data 1: Count as data	0x0
LBR_DATA_CNT_ENA	4	R/W	Enable / disable count of OAM PDU LBR as data in performance counters. 0: Do not count as data 1: Count as data	0x0
DMM_DATA_CNT_ENA	3	R/W	Enable / disable count of OAM PDU DMM as data in performance counters. 0: Do not count as data 1: Count as data	0x0
DMR_DATA_CNT_ENA	2	R/W	Enable / disable count of OAM PDU DMR as data in performance counters. 0: Do not count as data 1: Count as data	0x0
SM_DATA_CNT_ENA	1	R/W	Enable / disable count of OAM PDU 1DM as data in performance counters. 0: Do not count as data 1: Count as data	0x0
CCM_DATA_CNT_ENA	0	R/W	Enable / disable count of OAM PDU CCM as data in performance counters. 0: Do not count as data 1: Count as data	0x0

5.9.2.8 OAM_MEP:VOE:MEP_UC_MAC_LSB

Parent: [OAM_MEP:VOE](#)

Instances: 1

Configures the VOE Unicast MAC address (LSB).

This address can be checked when frames arrive, depending on the configuration of:

VOE:BASIC_CTRL.RX_DMACHK_SEL

TABLE 5-436: FIELDS IN MEP_UC_MAC_LSB

Field Name	Bit	Access	Description	Default
MEP_UC_MAC_LSB	31:0	R/W	VOE Unicast MAC - LSB	0x00000000

5.9.2.9 OAM_MEP:VOE:MEP_UC_MAC_MSB

Parent: [OAM_MEP:VOE](#)

Instances: 1

Configures the VOE Unicast MAC address (MSB).

This address can be checked when frames arrive, depending on the configuration of:

VOE:BASIC_CTRL.RX_DMACHK_SEL

TABLE 5-437: FIELDS IN MEP_UC_MAC_MSB

Field Name	Bit	Access	Description	Default
MEP_UC_MAC_MSB	15:0	R/W	VOE Unicast MAC - LSB	0x0000

5.9.2.10 OAM_MEP:VOE:OAM_HW_CTRL

Parent: [OAM_MEP:VOE](#)

Instances: 1

Configures per OAM OpCode if it is handled by hardware.

Only OAM PDU types enabled in this register are subject to HW support in the VOE.

TABLE 5-438: FIELDS IN OAM_HW_CTRL

Field Name	Bit	Access	Description	Default
LMM_ENA	8	R/W	Configures if LMM is handled by the MEP. If enabled LMM frames applicable for the MEP are either looped or CPU redirected (based on OAM_MEP:VOE:LOOPBACK_CTRL.LB_LMM_ENA). Looped LMM frames are updated as: RxFCf, TxFCb updated with the associated Performance counters DMAC updated to received SMAC SMAC set to VOE MAC address (OAM_MEP:VOE:MEP_UC_MAC_MSB and OAM_MEP:VOE:MEP_UC_MAC_LSB) OpCode changed to LMR. CPU redirected frames are forwarded to CPU extraction queue or external port configured OAM_MEP::CPU_CFG_1.LMM_CPU_QU and get RxFCf updated.	0x0

TABLE 5-438: FIELDS IN OAM_HW_CTRL (CONTINUED)

Field Name	Bit	Access	Description	Default
LMR_ENA	7	R/W	Configures if LMR is handled by the MEP. If enabled LMR frames applicable for the MEP are CPU redirected to CPU extraction queue or external port configured in OAM_MEP::CPU_CFG.LMR_CPU_QU. Handled frames get RxFCI written into REW_VAL in the CPU extraction header and optionally also extending the PDU with RxFCI after TxFCb based on OAM_MEP::MEP_CTRL.LMR_UPD_RxFCI_ENA.	0x0
LBM_ENA	6	R/W	Configures if LBM is handled by the MEP. If enabled LVM frames applicable for the MEP are looped. Looped LBM frames are updated as: DMAC updated to received SMAC SMAC set to VOE MAC address (OAM_MEP:VOE:MEP_UC_MAC_MSB and OAM_MEP:VOE:MEP_UC_MAC_LSB) OpCode changed to LBR. Also used to enable HW processing of Tx TST frames (If TST frame handling is enabled using VOE:OAM_FWD_CTRL.GENERIC_FWD_MASK(6)).	0x0
LBR_ENA	5	R/W	Configures if LBR is handled by the MEP. If enabled LBR frames applicable for the MEP checking received PDU transaction ID. Also used to enable HW processing of RX TST frames (If TST frame handling is enabled using VOE:OAM_FWD_CTRL.GENERIC_FWD_MASK(6)).	0x0
DMM_ENA	4	R/W	Configures if DMM is handled by the MEP. If enabled DMM frames applicable for the MEP are either looped or CPU redirected (based on OAM_MEP:VOE:LOOPBACK_CTRL.LB_DMM_ENA). Looped DMM frames are updated as: RxTimeStampf, TxTimeStampb updated with the associated Performance counters DMAC updated to received SMAC SMAC set to VOE MAC address (OAM_MEP:VOE:MEP_UC_MAC_MSB and OAM_MEP:VOE:MEP_UC_MAC_LSB) OpCode changed to DMR. CPU redirected frames are forwarded to CPU extraction queue or external port configured OAM_MEP::CPU_CFG_1.DMM_CPU_QU and get RxTimeStampf updated.	0x0

TABLE 5-438: FIELDS IN OAM_HW_CTRL (CONTINUED)

Field Name	Bit	Access	Description	Default
DMR_ENA	3	R/W	Configures if DMR is handled by the MEP. If enabled DMR frames applicable for the MEP are CPU redirected to CPU extraction queue or external port configured in OAM_MEP::CPU_CFG.DMR_CPU_QU. Handled frames get PDU RxTimeStampb updated..	0x0
SDM_ENA	2	R/W	Configures if 1DM is handled by the MEP. If enabled 1DM frames applicable for the MEP are CPU redirected to CPU extraction queue or external port configured in OAM_MEP::CPU_CFG.DMR_CPU_QU. Handled frames get PDU RxTimeStampb updated..	0x0
CCM_ENA	1	R/W	Configures if CCM is handled by the MEP. If enabled, CCM frames applicable for the MEP are checked before being used for clearing CCM_MISS_CNT.	0x0
CCM_LM_ENA	0	R/W	Configures if CCM with LM information is handled by the MEP. Requires that CCM_ENA is also set. Handled frames get RxFCI in written into REW_VAL in the CPU extraction header and optionally also the reserved field in the PDU with RxFCI based on OAM_MEP::MEP_CTRL.CCM_LM_UPD_RSV_ENA. Handled frames get RxFCI and TxFCf stored to be used for next transmitted CCM frame with LM information.	0x0

5.9.2.11 OAM_MEP:VOE:LOOPBACK_CTRL

Parent: [OAM_MEP:VOE](#)

Instances: 1

Configures Down-MEP loopback of OAM functions.

Loopback VOE functionality is not supported for Up-MEP configurations.

For Up-MEP the following bit fields have a special meaning:

- LB_ISDX[4:0]: See description of bit field.
- LB_ISDX_P[4:0]: See description of bit field.

TABLE 5-439: FIELDS IN LOOPBACK_CTRL

Field Name	Bit	Access	Description	Default
LB_LMM_ENA	25	R/W	Configures if LMM frames are looped or sent to CPU. Note: This must not be enabled for Up-MEP configurations. 0: OAM LMM PDU is sent to CPU queue: CPU_CFG_1:LMM_CPU_QU 1: OAM LMM PDU is looped to LB_PORT-MASK and OpCode and MAC addresses are updated). Configuration is only used is OAM_HW_CTRL.LMM_ENA is set.	0x1
LB_DMM_ENA	24	R/W	Configures if DMM frames are looped or sent to CPU. Note: This must not be enabled for Up-MEP configurations. 0: OAM DMM PDU is sent to CPU queue: CPU_CFG_1:LMM_CPU_QU 1: OAM DMM PDU is looped to LB_PORT-MASK and OpCode and MAC addresses are updated). Configuration is only used is OAM_HW_CTRL.DMM_ENA is set.	0x1
LB_ISDX	23:14	R/W	Down-MEP: Configures ISDX for loopback traffic. Up-MEP: LB_ISDX[4:0]: Contains the current value of the LM Rx sequence number. This 5 bit value is sent to the CPU in the CPU extraction header when a valid OAM PDU with LM content is received in in-service mode. After the value is sent to the CPU, it is incremented.	0x000
LB_ISDX_P	13:4	R/W	Down-MEP: Configures protected ISDX for loopback traffic. Up-MEP: LB_ISDX_P[4:0]: Contains the current value of the DM Rx sequence number. This 5 bit value is sent to the CPU in the CPU extraction header when a valid OAM PDU with DM content is received in in-service mode. After the value is sent to the CPU, it is incremented.	0x000

TABLE 5-439: FIELDS IN LOOPBACK_CTRL (CONTINUED)

Field Name	Bit	Access	Description	Default
LB_PORTNUM_P	3:0	R/W	This register is used for Down-MEP implementations only (has no effect for Up-MEP). For a VOE implementing an unprotected Service MEP, this register must be configured to the VOE PORT (The port on which the VOE is located) For a VOE implementing a protected service MEP, this register must be programmed with the port number of the protect port.	0x0

5.9.2.12 OAM_MEP:VOE:LBR_RX_FRM_CNT**Parent:** [OAM_MEP:VOE](#)**Instances:** 1

Counts number of received OAM LBR PDU.

TABLE 5-440: FIELDS IN LBR_RX_FRM_CNT

Field Name	Bit	Access	Description	Default
LBR_RX_FRM_CNT	31:0	R/W	Number of RX LBR frames. Note: Also used to count number of RX TST frames (If TST frames are enabled using VOE:OAM_FWD_CTRL.GENERIC_FWD_MASK(6)).	0x00000000

5.9.2.13 OAM_MEP:VOE:LBR_RX_TRANSID_CFG**Parent:** [OAM_MEP:VOE](#)**Instances:** 1

Configures expected transaction ID of the next RX LBR frame.

This value is checked against the value in the next RX LBR frame, if LBR function is enabled for the VOE.

In case of an error, the following sticky bit is asserted:

VOE:OAM_RX_STICKY.LBR_TRANSID_ERR_STICKY.

TABLE 5-441: FIELDS IN LBR_RX_TRANSID_CFG

Field Name	Bit	Access	Description	Default
LBR_RX_TRANSID	31:0	R/W	Configures one less than next expected Transaction ID. Used for checking the LBR transaction ID. Updated by Hardware to received OAM PDU transaction ID. Note: Also updated for RX TST frames (If TST frames are enabled using VOE:OAM_FWD_CTRL.GENERIC_FWD_MASK(6)).	0x00000000

5.9.2.14 OAM_MEP:VOE:LBM_TX_TRANSID_UPDATE

Parent: [OAM_MEP:VOE](#)

Instances: 1

If set, the transaction id is updated for LBM frames transmitted from this VOE.

This can be used to avoid overwriting the Tx ID for externally generated LBM frames.

TABLE 5-442: FIELDS IN LBM_TX_TRANSID_UPDATE

Field Name	Bit	Access	Description	Default
LBM_TRANSID_UPDATE_ENA	0	R/W	If set, the transaction id is updated for LBM frames transmitted from this VOE. This can be used to avoid overwriting the Tx ID for externally generated LBM frames. Note: Also used to update TX TST frames (If TST frames are enabled using VOE:OAM_FWD_CTRL.GENERIC_FWD_MASK(6)).	0x1

5.9.2.15 OAM_MEP:VOE:LBM_TX_TRANSID_CFG

Parent: [OAM_MEP:VOE](#)

Instances: 1

Configures transmitted transaction ID for LBM.

LBM Tx frames are updated with this value if LBM_TRANSID_UPDATE_ENA is set.

TABLE 5-443: FIELDS IN LBM_TX_TRANSID_CFG

Field Name	Bit	Access	Description	Default
LBM_TX_TRANSID	31:0	R/W	Configures one less than transmitted transaction ID for LBM. Auto incremented by hardware upon usage. Note: Also used to update Tx TST frames (If TST frames are enabled using VOE:OAM_F-WD_CTRL.GENERIC_FWD_MASK(6)).	0x00000000

5.9.2.16 OAM_MEP:VOE:CCM_CFG

Parent: [OAM_MEP:VOE](#)

Instances: 1

Misc configuration for CCM frame handling.

TABLE 5-444: FIELDS IN CCM_CFG

Field Name	Bit	Access	Description	Default
CCM_LM_PERIOD	26:24	R/W	Down-MEP only (Not supported for Up-MEP VOE can insert LM information into CCM frames injected by internal or external CPU. The insertion of LM information into CCM frames occurs based on a LOC timeout counter. This bit field configures which of the 7 LOC timeout counters to use for LM insertion. This configuration does not affect the CCM PDU period field. When the selected LOC timeout counter times out, LM information is inserted into the next CCM frame transmitted by the VOE. Note that the rate at which LM information is inserted is twice the rate indicated by the PERIOD_VAL of the selected timeout counter. I.e. if the CCM_PERIOD_CFG.CCM_PERIOD_VAL of the selected timeout counter is set to 10 ms, LM information is inserted every 5 ms. 0: Disable automatic insertion of loss measurements in OAM CCM PDU n: Automatic insertion of loss measurements in next OAM CCM PDU when the corresponding OAM_MEP::CCM_PERIOD_CFG occurs.	0x0

TABLE 5-444: FIELDS IN CCM_CFG (CONTINUED)

Field Name	Bit	Access	Description	Default
CCM_LM_INSERT_NXT	23	R/W	Down-MEP only (Not supported for Up-MEP) Indicate that CCM_LM info is inserted into next CCM Tx frame. Automatically set by hardware every time CCM_LM_PERIOD has occurred. Can also be forced by SW writing a 1 to this register.	0x0
CCM_SEQ_INCR_ENA	22	R/W	Configures auto increment of CCM SEQ (upon use). This is only valid if CCM_SEQ_UPD_ENA = 1	0x0
CCM_SEQ_UPD_ENA	21	R/W	If asserted (=1), the sequence number is updated for CCM frames transmitted from this VOE. This can be used to avoid overwriting the Tx ID for externally generated LBM frames.	0x1
CCM_RX_SEQ_CHK_ENA	20	R/W	Enables checking the incoming CCM sequence number.	0x0
CCM_TX_RDI	19	R/W	Specifies the RDI to be transmitted in the next CCM frame.	0x0
CCM_MISS_CNT	18:16	R/W	Specifies the number of CCM half periods has passed without CCM messages received from the peer MEP. Cleared by HW when receiving a CCM message from the peer associated with MEP. Can be configured to periodically increment CCM miss count via HW configured by VOP::CCM_CTRL. n<7: No Loss of continuity n==7: Loss of continuity	0x0
CCM_PRIO	15:13	R/W	Configures expected CCM priority.	0x0
CCM_PERIOD	12:10	R/W	Configures expected CCM period. The PERIOD fields of the incoming CCM frames are checked against this value. Further the configuration of this bit field specifies the LOC timeout counter to be used for LOC scan for this VOE.	0x0
CCM_MEGID_CHK_ENA	9	R/W	Configures CCM MEG ID check. 0: No MEGID check 1: Check MEG ID	0x0

TABLE 5-444: FIELDS IN CCM_CFG (CONTINUED)

Field Name	Bit	Access	Description	Default
CCM_MEPID_CHK_ENA	8	R/W	If this bit is asserted, the value of the MEP ID in incoming CCM frames are verified against the value programmed in: VOE:CCM_MEPID_CFG.CCM_MEPID Errors in the incoming MEP ID are reported in the following error bit: VOE:CCM_CFG.CCM_MEPID_ERR	0x1
CCM_PERIOD_ERR	7	R/W	Status of period check for the latest received CCM frame. 0: No non zero period err 1: Non zero period err detected	0x0
CCM_PRIO_ERR	6	R/W	Status of latest priority check done for CCM frames 0: No priority err 1: Priority err detected	0x0
CCM_ZERO_PERIOD_ERR	5	R/W	Status of latest Zero period check 0: No Zero period err 1: Zero period err detected	0x0
CCM_RX_RDI	4	R/W	Specifies the RDI received in the latest CCM messages received from the peer MEP.	0x0
CCM_MEPID_ERR	2	R/W	Status of latest MEPID check, checking the MEPID of the incoming CCM frame against the value programmed in: VOE:CCM_MEPID_CFG.CCM_MEPID MEPID check is enabled by setting: VOE:CCM_CFG.CCM_MEPID_CHK_ENA= 1 0: No MEPID err 1: MEPID err detected	0x0
CCM_MEGID_ERR	1	R/W	Status of latest MEGID check, validating incoming CCM MEGID against the value configured in: VOE:CCM_MEGID_CFG.CCM_MEGID. MEGID checking is enabled by: VOE:CCM_CFG.CCM_MEGID_CHK_ENA= 1 0: No MEGID err 1: MEGID err detected	0x0
OAM_MEL_ERR	0	R/W	Status of latest MEL check 0: No MEL err 1: MEL err detected	0x0

5.9.2.17 OAM_MEP:VOE:CCM_RX_VLD_FC_CNT

Parent: OAM_MEP:VOE

Instances: 1

VSC7416-01

Counts number of received OAM CCM PDU that passes required ingress check and therefore clear the CCM miss counter.

TABLE 5-445: FIELDS IN CCM_RX_VLD_FC_CNT

Field Name	Bit	Access	Description	Default
CCM_RX_VLD_FC_CNT	31:0	R/W	See Register description.	0x00000000

5.9.2.18 OAM_MEP:VOE:CCM_RX_INVLD_FC_CNT

Parent: [OAM_MEP:VOE](#)

Instances: 1

Counts number of received OAM CCM PDU that fails required ingress check and therefore doesn't clear the CCM miss counter.

TABLE 5-446: FIELDS IN CCM_RX_INVLD_FC_CNT

Field Name	Bit	Access	Description	Default
CCM_RX_INVLD_FC_CNT	31:0	R/W	See Register description	0x00000000

5.9.2.19 OAM_MEP:VOE:CCM_CAP_TX_FCf

Parent: [OAM_MEP:VOE](#)

Instances: 1

Register use differs when used for Up-MEP and Down-MEP:

Down-MEP:

Latest received CCM_TX_FCf value.

Value is used by VOE for writing outgoing CCM-LM frames (TxFCf)

Up-MEP:

Value of RX_LM frame counter sampled when the latest LMM frame was received.

Value must be read by the CPU.

When the register value is written, the VOE asserts the valid bit:

UPMEP_LM_CNT_VLD.LMM_RX_CNT_VLD

Reading the register from CPU clears the valid bit.

If the register has not been read by the CPU, when a new sample arrives, the following sticky bit is asserted:

UPMEP_LM_CNT_STICKY.CCM_LMM_RX_CNT_STICKY

TABLE 5-447: FIELDS IN CCM_CAP_TX_FCF

Field Name	Bit	Access	Description	Default
CCM_CAP_TX_FCF	31:0	R/W	See Register description.	0x00000000

5.9.2.20 OAM_MEP:VOE:CCM_CAP_RX_FCI

Parent: [OAM_MEP:VOE](#)

Instances: 1

Register use differs when used for Up-MEP and Down-MEP:

Down-MEP:

Value of RX frame counter sampled when the latest CCM-LM frame was received.

Value is used by VOE for writing outgoing CCM-LM frames (RxFcf)

Up-MEP:

Value of RX_LM frame counter sampled when the latest CCM-LM frame was received.

Value must be read by the CPU.

When the register value is written, the VOE asserts the valid bit:

UPMEP_LM_CNT_VLD.CCM_LM_RX_CNT_VLD

Reading the register from CPU clears the valid bit.

If the register has not been read by the CPU, when a new sample arrives, the following sticky bit is asserted:

UPMEP_LM_CNT_STICKY.CCM_LM_RX_CNT_STICKY

TABLE 5-448: FIELDS IN CCM_CAP_RX_FCL

Field Name	Bit	Access	Description	Default
CCM_CAP_RX_FCI	31:0	R/W	See Register description	0x00000000

5.9.2.21 OAM_MEP:VOE:UPMEP_LMR_RX_LM_CNT

Parent: [OAM_MEP:VOE](#)

Instances: 1

Register used for Up-MEP only:

Value of RX_LM frame counter sampled when the latest LMR frame was received.

Value must be read by the CPU.

When the register value is written, the VOE asserts the valid bit:

UPMEP_LM_CNT_VLD.LMR_RX_CNT_VLD

Reading the register from CPU clears the valid bit.

If the register has not been read by the CPU, when a new sample arrives, the following sticky bit is asserted:

UPMEP_LM_CNT_STICKY.CCM_LMR_RX_CNT_STICKY

TABLE 5-449: FIELDS IN UPMEP_LMR_RX_LM_CNT

Field Name	Bit	Access	Description	Default
UPMEP_LMR_RX_LM_CNT	31:0	R/W	See register description.	0x00000000

5.9.2.22 OAM_MEP:VOE:CCM_TX_SEQ_CFG

Parent: [OAM_MEP:VOE](#)

Instances: 1

Configures CCM Sequence number to be inserted into the next CCM or CCM-LM frame.

TABLE 5-450: FIELDS IN CCM_TX_SEQ_CFG

Field Name	Bit	Access	Description	Default
CCM_TX_SEQ	31:0	R/W	See Register Description	0x00000000

5.9.2.23 OAM_MEP:VOE:CCM_RX_SEQ_CFG

Parent: [OAM_MEP:VOE](#)

Instances: 1

Down-MEP:

Configures the expected value of the Sequence number for the next CCM frame.

The value is updated every time a CCM frame is received.

The value is updated with the CCM sequence number plus 1.

The value is verified if the RX CCM Sequence number check is enabled:

VOE:CCM_CFG.CCM_RX_SEQ_CHK_ENA = 1

If the received Sequence Number is not as expected, the following sticky bit is asserted:

CCM_RX_SEQ_ERR_STICKY

Up-MEP:

For Up-MEP, only the lower 16 bit of the RX CCM SEQ is verified.

The remaining bits contain the RX_LM_SN for the latest received LM frame as follows:

[20:16] contains RX_LM_SN for latest CCM-LM received.

Note: Corresponds to the LM sample in: CCM_CAP_RX_FCL

[25:21] contains RX_LM_SN for latest: LMR received (only in in-service mode).

Note: Corresponds to the LM sample in: UPMEP_LMR_RX_LM_CNT

[30:26] contains RX_LM_SN for latest: LMM received

Note: Corresponds to the LM sample in: CCM_CAP_TX_FCF

TABLE 5-451: FIELDS IN CCM_RX_SEQ_CFG

Field Name	Bit	Access	Description	Default
CCM_RX_SEQ	31:0	R/W	Down-MEP: Configures CCM RX SEQ. Up-MEP: Holds the values of the RX_LM_SN for OAM LM frames.	0x00000000

5.9.2.24 OAM_MEP:VOE:CCM_MEPID_CFG

Parent: [OAM_MEP:VOE](#)

Instances: 1

Configures MEP ID value to be verified against incoming CCM frames.

TABLE 5-452: FIELDS IN CCM_MEPID_CFG

Field Name	Bit	Access	Description	Default
CCM_MEPID	15:0	R/W	Configures CCM MEP ID value for this VOE. If CCM MEP ID checking is enabled (CCM_MEPID_CHK_ENA = 1), the MEP ID of incoming CCM frames are verified against the value configured in CCM_MEPID. If the received value differs from the value configured in MEP ID, the incoming CCM frame is discarded (No LOC update) and the error is signalled using the following error bit: VOE:CCM_CFG.CCM_MEPID_ERR	0x0000

5.9.2.25 OAM_MEP:VOE:CCM_MEGID_CFG

Parent: [OAM_MEP:VOE](#)

Instances: 12

Configures 48 byte MEG ID (lowest replication index correspond to MSB) to verified in incoming CCM frames.

TABLE 5-453: FIELDS IN CCM_MEGID_CFG

Field Name	Bit	Access	Description	Default
CCM_MEGID	31:0	R/W	Configures the MEG ID to be used for CCM checking. The checking is enabled by setting: VOE:CCM_CFG.CCM_MEGID_CHK_ENA=1 Any errors in the incoming MEG ID are reported in error bit: VOE.CCM_CFG.CCM_MEGID_ERR	0x00000000

5.9.2.26 OAM_MEP:VOE:OAM_RX_STICKY

Parent: [OAM_MEP:VOE](#)

Instances: 1

Sticky bits are asserted when OAM PDU of specified type is received at the correct MEL level and when the DMAC is valid for this VOE.

TABLE 5-454: FIELDS IN OAM_RX_STICKY

Field Name	Bit	Access	Description	Default
GENERIC_RX_STICKY_MASK	24:17	R/W	Rx reception Generic Opcodes per MEP per generic OpCode.	0x00
UNK_OPCODE_RX_STICKY	16	R/W	Asserted when receiving an OAM PDU with an OpCode which is a dedicated or configured as a generic Opcodes 0: No PDU received 1: PDU with specified OpCode received	0x0
LTM_RX_STICKY	15	R/W	Asserted when receiving an LTM PDU 0: No PDU received 1: PDU with specified OpCode received	0x0
LTR_RX_STICKY	14	R/W	Asserted when receiving an LTR PDU 0: No PDU received 1: PDU with specified OpCode received	0x0
LMM_RX_STICKY	13	R/W	Asserted when receiving an LMM PDU 0: No PDU received 1: PDU with specified OpCode received	0x0
LMR_RX_STICKY	12	R/W	Asserted when receiving an LMR PDU 0: No PDU received 1: PDU with specified OpCode received	0x0
LBM_RX_STICKY	11	R/W	Asserted when receiving an LBM PDU 0: No PDU received 1: PDU with specified OpCode received	0x0
LBR_RX_STICKY	10	R/W	Asserted when receiving an LBR PDU 0: No PDU received 1: PDU with specified OpCode received	0x0

TABLE 5-454: FIELDS IN OAM_RX_STICKY (CONTINUED)

Field Name	Bit	Access	Description	Default
DMM_RX_STICKY	9	R/W	Asserted when receiving an DMM PDU 0: No PDU received 1: PDU with specified OpCode received	0x0
DMR_RX_STICKY	8	R/W	Asserted when receiving an DMR PDU 0: No PDU received 1: PDU with specified OpCode received	0x0
SM_RX_STICKY	7	R/W	Asserted when receiving an SM PDU 0: No PDU received 1: PDU with specified OpCode received	0x0
CCM_RX_STICKY	6	R/W	Asserted when receiving an CCM or CCM-LM PDU 0: No PDU received 1: PDU with specified OpCode received	0x0
CCM_LM_RX_STICKY	5	R/W	Asserted when receiving an CCM-LM PDU 0: No PDU received 1: PDU with specified OpCode received	0x0
LBR_TRANSID_ERR_STICKY	4	R/W	LBR TRANSACTION ID error check detected	0x0
CCM_RX_SEQ_ERR_STICKY	3	R/W	CCM PDU received sequence error detected	0x0
OAM_AS_DATA_STICKY	2	R/W	OAM traffic received with MEL larger than this MEP MEL	0x0
MEP_EGR_BLOCK_STICKY	1	R/W	OAM_MEP is discarding frame in egress lookup due to OAM MEL <= VOE_MEL	0x0
MAC_ADDR_ERR_STICKY	0	R/W	Indicate incorrect MAC address received. 0: No event 1: MAC addr err occurred	0x0

5.9.2.27 OAM_MEP:VOE:STICKY**Parent:** [OAM_MEP:VOE](#)**Instances:** 1

Group of sticky bits asserted based in checking different CCM field values against the configured values.

Sticky bits are updated when the state of the check changes.

I.e. the sticky bits are asserted both when an error is asserted and when it is deasserted.

TABLE 5-455: FIELDS IN STICKY

Field Name	Bit	Access	Description	Default
CCM_PERIOD_STICKY	7	R/W	Indicate change in the status of the PERIOD check of incoming CCM frames against the value programmed in CCM_PERIOD. 0: No event 1: Changed non zero Period occurred	0x0
CCM_PRIO_STICKY	6	R/W	Indicate change in the status of the PRIORITY check of incoming CCM frames against the value programmed in CCM_PRIO. 0: No event 1: Changed CCM Priority occurred	0x0

TABLE 5-455: FIELDS IN STICKY (CONTINUED)

Field Name	Bit	Access	Description	Default
CCM_ZERO_PERIOD_STICKY	5	R/W	Indicate change in reception of frames with zero period. (According to G.8081 CCM CCM frames with PERIOD = 0 is an illegal value.) 0: No event 1: Changed Zero Period occurred	0x0
CCM_RX_RDI_STICKY	4	R/W	Indicate RDI changed value in latest CCM frame Clear this bit to receive new RX RDI up/down interrupts for this MEP. 0: No event 1: Changed RDI occurred	0x0
CCM_LOC_STICKY	3	R/W	Indicate changed Loss of Continuity (LOC) state Clear this bit to receive new up/down interrupts for this MEP. 0: LOC has not changed state 1: Changed LOC state (LOC is indicated based on whether CCM_MISS_CNT is less than 7)	0x0
CCM_MEPID_STICKY	2	R/W	Indicate a change in the status of the CCM MEPID checking of incoming CCM frames against the value programmed in: VOE:CCM_MEPID 0: No event 1: PDU with changed MEPID seen	0x0
CCM_MEGID_STICKY	1	R/W	Indicate a change in the status of the CCM MEGID checking of incoming CCM frames against the value programmed in: VOE:CCM_MEGID_CFG.CCM_MEGID MEG ID check is enabled by setting: VOE:CCM_CFG.CCM_MEGID_CHK_ENA=1 0: No event 1: PDU with changed MEGID seen	0x0
OAM_MEL_STICKY	0	R/W	Indicate unexpected MEL. 0: No event 1: PDU with unexpected MEL seen	0x0

5.9.2.28 OAM_MEP:VOE:UPMEP_LM_CNT_STICKY

Parent: [OAM_MEP:VOE](#)

Instances: 1

UPMEP only:

If the Up-MEP RX counter samples, based on LM frame reception has not been read by the CPU prior to reception of the next LM frame, a sticky bit is asserted in this register.

TABLE 5-456: FIELDS IN UPMEP_LM_CNT_STICKY

Field Name	Bit	Access	Description	Default
CCM_LM_RX_CNT_STICKY	2	Sticky	Up-MEP only. Asserted if a CCM-LM frame is received and the corresponding sample valid flag, has not been cleared by the CPU. Sample valid: CCM_LM_RX_CNT_VLD Assertion of this sticky bit indicates that an RX_LM sample was overwritten by the VOE before it was read by the CPU. The valid flag is cleared by reading the RX LM sample value in: CCM_CAP_RX_FCI	0x0
LMM_RX_CNT_STICKY	1	Sticky	Up-MEP only. Asserted if a LMM frame is received and the corresponding sample valid flag, has not been cleared by the CPU. Sample valid: LMM_RX_CNT_VLD Assertion of this sticky bit indicates that an RX_LM sample was overwritten by the VOE before it was read by the CPU. The valid flag is cleared by reading the RX LM sample value in: CCM_CAP_TX_FCI	0x0
LMR_RX_CNT_STICKY	0	Sticky	Up-MEP only. Asserted if a LMM frame is received and the corresponding sample valid flag, has not been cleared by the CPU. Sample valid: LMR_RX_CNT_VLD Assertion of this sticky bit indicates that an RX_LM sample was overwritten by the VOE before it was read by the CPU. The valid flag is cleared by reading the RX LM sample value in: UPMEP_LMR_RX-_LM_CNT	0x0

5.9.2.29 OAM_MEP:VOE:INTR_ENA

Parent: [OAM_MEP:VOE](#)

Instances: 1

Configures which sources should cause interrupts.

Status of interrupt sources can be found in OAM_MEP:VOE:STICKY.

Interrupting VOE can be found in OAM_MEP::INTR.

TABLE 5-457: FIELDS IN INTR_ENA

Field Name	Bit	Access	Description	Default
CCM_PERIOD_INTR_ENA	7	R/W	Configures interrupt based on CCM_PERIOD_STICKY 0: Disable interrupt 1: Enable interrupt	0x0
CCM_PRIO_INTR_ENA	6	R/W	Configures interrupt based on CCM_PRIO_STICKY 0: Disable interrupt 1: Enable interrupt	0x0
CCM_ZERO_PERIOD_INTR_ENA	5	R/W	Configures interrupt based on CCM_ZERO_PERIOD_STICKY 0: Disable interrupt 1: Enable interrupt	0x0
CCM_RX_RDI_INTR_ENA	4	R/W	Configures interrupt based on CCM_RX_RDI_STICKY 0: Disable interrupt 1: Enable interrupt	0x0
CCM_LOC_INTR_ENA	3	R/W	Configures interrupt based on CCM_LOC_STICKY 0: Disable interrupt 1: Enable interrupt	0x0
CCM_MEPID_INTR_ENA	2	R/W	Configures interrupt based on CCM_MEPID_STICKY 0: Disable interrupt 1: Enable interrupt	0x0
CCM_MEGID_INTR_ENA	1	R/W	Configures interrupt based on CCM_MEGID_STICKY 0: Disable interrupt 1: Enable interrupt	0x0
OAM_MEL_INTR_ENA	0	R/W	Configures interrupt based on OAM_MEL_STICKY 0: Disable interrupt 1: Enable interrupt	0x0

5.9.2.30 OAM_MEP:VOE:RX_SEL_OAM_CNT

Parent: [OAM_MEP:VOE](#)

Instances: 1

Count selected OAM PDU received MEP.

Selected OAM PDUs are configured in VOE:OAM_CNT_OAM_CTRL.*

TABLE 5-458: FIELDS IN RX_SEL_OAM_CNT

Field Name	Bit	Access	Description	Default
RX_SEL_OAM_FRM_CNT	31:0	R/W	See Register Description	0x00000000

5.9.2.31 OAM_MEP:VOE:RX_OAM_FRM_CNT

Parent: [OAM_MEP:VOE](#)

Instances: 1

Count non-selected OAM frames received by VOE

Selected OAM PDUs are configured in VOE:OAM_CNT_OAM_CTRL.*

TABLE 5-459: FIELDS IN RX_OAM_FRM_CNT

Field Name	Bit	Access	Description	Default
RX_OAM_FRM_CNT	31:0	R/W	See Register description	0x00000000

5.9.2.32 OAM_MEP:VOE:TX_SEL_OAM_CNT

Parent: [OAM_MEP:VOE](#)

Instances: 1

Count selected OAM PDU transmitted VOE.

Selected OAM PDUs are configured in VOE:OAM_CNT_OAM_CTRL.*

TABLE 5-460: FIELDS IN TX_SEL_OAM_CNT

Field Name	Bit	Access	Description	Default
TX_SEL_OAM_FRM_CNT	31:0	R/W	See Register Description	0x00000000

5.9.2.33 OAM_MEP:VOE:TX_OAM_FRM_CNT

Parent: [OAM_MEP:VOE](#)

Instances: 1

Count non-selected OAM frames transmitted by VOE.

Selected OAM PDUs are configured in VOE:OAM_CNT_OAM_CTRL.*

TABLE 5-461: FIELDS IN TX_OAM_FRM_CNT

Field Name	Bit	Access	Description	Default
TX_OAM_FRM_CNT	31:0	R/W	See Register Description	0x00000000

5.9.3 OAM_MEP:PORT_PM

Parent: [OAM_MEP](#)

Instances: 88

Contains the LM frame counters for all port VOEs.

Each port counts frames on eight priorities.

To read the counters from port X priority Y, the following index must be used.

$\text{Cnt\_idx} = X * 8 \text{ (number of priorities)} + Y$

TABLE 5-462: REGISTERS IN PORT_PM

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
PORT_RX_FRM_CNT	0x00000000	1	Rx port MEP frame counters	Page 544
PORT_TX_FRM_CNT	0x00000004	1	Tx port MEP frame counters	Page 544

5.9.3.1 OAM_MEP:PORT_PM:PORT_RX_FRM_CNT

Parent: [OAM_MEP:PORT_PM](#)

Instances: 1

TABLE 5-463: FIELDS IN PORT_RX_FRM_CNT

Field Name	Bit	Access	Description	Default
RX_FRM_CNT	31:0	R/W	Rx LM frame counters per port VOE per priority.	0x00000000

5.9.3.2 OAM_MEP:PORT_PM:PORT_TX_FRM_CNT

Parent: [OAM_MEP:PORT_PM](#)

Instances: 1

TABLE 5-464: FIELDS IN PORT_TX_FRM_CNT

Field Name	Bit	Access	Description	Default
TX_FRM_CNT	31:0	R/W	Tx LM frame counters per port VOE per priority.	0x00000000

5.9.4 OAM_MEP:RX_VOE_PM

Parent: [OAM_MEP](#)

Instances: 512

Contains the RX LM frame counters for all service/path VOEs.

Each service VOE counts frames on eight priorities.

To read the counters from service VOE= X priority Y, the following index must be used.

$\text{Cnt\_idx} = X * 8 \text{ (number of priorities)} + Y$

TABLE 5-465: REGISTERS IN RX_VOE_PM

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
MEP_RX_FRM_CNT	0x00000000	1	Rx service MEP frame counters	Page 545

5.9.4.1 OAM_MEP:RX_VOE_PM:MEP_RX_FRM_CNT

Parent: [OAM_MEP:RX_VOE_PM](#)

Instances: 1

TABLE 5-466: FIELDS IN MEP_RX_FRM_CNT

Field Name	Bit	Access	Description	Default
RX_FRM_CNT	31:0	R/W	RX LM frame counters by VOE.	0x00000000

5.9.5 OAM_MEP:TX_VOE_PM

Parent: [OAM_MEP](#)

Instances: 512

Contains the Tx LM frame counters for all service/path VOEs.

Each service VOE counts frames on eight priorities.

To read the counters from service VOE= X priority Y, the following index must be used.

Cnt_idx = X * 8 (number of priorities) + Y

TABLE 5-467: REGISTERS IN TX_VOE_PM

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
MEP_TX_FRM_CNT	0x00000000	1	Tx service MEP frame counters	Page 545

5.9.5.1 OAM_MEP:TX_VOE_PM:MEP_TX_FRM_CNT

Parent: [OAM_MEP:TX_VOE_PM](#)

Instances: 1

TABLE 5-468: FIELDS IN MEP_TX_FRM_CNT

Field Name	Bit	Access	Description	Default
TX_FRM_CNT	31:0	R/W	Tx LM frame counters.	0x00000000

5.10 PCIE**TABLE 5-469: REGISTER GROUPS IN PCIE**

Register Group Name	Offset within Target	Instances and Address Spacing	Description	Details
PCIE	0x00000000	1	PCIe endpoint configuration space	Page 545

5.10.1 PCIE:PCIE

Parent: [PCIE](#)

Instances: 1

TABLE 5-470: REGISTERS IN PCIE

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
CONF_SPACE	0x00000000	448 0x00000004	PCI configuration space	Page 546
ATU_REGION	0x00000900	1	Address translation region	Page 549
ATU_CFG1	0x00000904	1	Address translation configuration register 1	Page 549
ATU_CFG2	0x00000908	1	Address translation configuration register 2	Page 549
ATU_BASE_ADDR_LOW	0x0000090C	1	Address translation lower base address	Page 550
ATU_LIMIT_ADDR	0x00000914	1	Address translation limit address	Page 550
ATU_TGT_ADDR_LOW	0x00000918	1	Address translation lower target address	Page 550
ATU_TGT_ADDR_HIGH	0x0000091C	1	Address translation upper target address	Page 550

5.10.1.1 PCIE:PCIE:CONF_SPACE

Parent: [PCIE:PCIE](#)

Instances: 448

This is not the complete configuration + extended configuration space, only up to byte address 6FF.

TABLE 5-471: FIELDS IN CONF_SPACE

Field Name	Bit	Access	Description	Default
CONF_SPACE	31:0	R/W	The contents of the configuration (and extended) region is: 00: 0xB002101B Header Type 0 01: 0x00100000 02: 0x02800000 03: 0x00000000 04: 0x00000000 05: 0x00000000 06: 0x0000000C 07: 0x00000000 08: 0x00000000 09: 0x00000000 10: 0x00000000 11: 0x0000101B 12: 0x00000000 13: 0x00000040 14: 0x00000000 15: 0x000001FF 16: 0x5A635001 PM Capability Register Set 17: 0x00000000 ... 20: 0x00847005 MSI Capability Register Set 21: 0x00000000 22: 0x00000000 23: 0x00000000 ... 28: 0x00020010 PCI Express Capability Register Set 29: 0x00008740 30: 0x00002010 31: 0x00003C11 32: 0x00110000 ... 64: 0x14010001 Advanced Error Reporting Capability Register Set 65: 0x00000000 66: 0x00000000 67: 0x00462030 68: 0x00000000 69: 0x00002000 70: 0x000000A0 71: 0x00000000 72: 0x00000000 73: 0x00000000 74: 0x00000000 ... 80: 0x00010002 Virtual Channel Capability Register Set 81: 0x00000000 82: 0x00000000 83: 0x00000000 84: 0x00000000 85: 0x800000FF 86: 0x00000000 ...	0x00000000

5.10.1.2 PCIE:PCIE:ATU_REGION

Parent: [PCIE:PCIE](#)

Instances: 1

The address translation unit supports 2 outbound regions. The registers PCIE::ATU_CFG1, PCIE::ATU_CFG2, PCIE::ATU_BASE_ADDR_LOW, PCIE::ATU_BASE_ADDR_HIGH, PCIE::ATU_LIMIT_ADDR, PCIE::ATU_TGT_ADDR_LOW, and PCIE::ATU_TGT_ADDR_HIGH all map to the currently configured region (as configured in this register).

TABLE 5-472: FIELDS IN ATU_REGION

Field Name	Bit	Access	Description	Default
ATU_IDX	0	R/W	Selects region index, set to 0 or 1.	0x0

5.10.1.3 PCIE:PCIE:ATU_CFG1

Parent: [PCIE:PCIE](#)

Instances: 1

TABLE 5-473: FIELDS IN ATU_CFG1

Field Name	Bit	Access	Description	Default
ATU_ATTR	10:9	R/W	When the address of an outbound TLP is matched to this region, then the ATTR field of the TLP is changed to the value in this register.	0x0
ATU_TD	8	R/W	When the address of an outbound TLP is matched to this region, then the TD field of the TLP is changed to the value in this register.	0x0
ATU_TC	7:5	R/W	When the address of an outbound TLP is matched to this region, then the TC field of the TLP is changed to the value in this register.	0x0
ATU_TYPE	4:0	R/W	When the address of an outbound TLP is matched to this region, then the TYPE field of the TLP is changed to the value in this register. 0: MRd/MWr 1: MRdLk 2: IORd/IOWr 4: CfgRd0/CfgWr0 5: CfgRd1/CfgWr1 16-23: Msg/MsgD	0x00

5.10.1.4 PCIE:PCIE:ATU_CFG2

Parent: [PCIE:PCIE](#)

Instances: 1

TABLE 5-474: FIELDS IN ATU_CFG2

Field Name	Bit	Access	Description	Default
ATU_REGION_ENA	31	R/W	This bit must be set for address translation to take place.	0x0
ATU_MSG_CODE	7:0	R/W	When the address of an outbound TLP is matched to this region, and the translated TLP TYPE field is Msg or MsgD; then the Message field of the TLP is changed to the value in this register.	0x00

5.10.1.5 PCIE:PCIE:ATU_BASE_ADDR_LOW

Parent: [PCIE:PCIE](#)

Instances: 1

TABLE 5-475: FIELDS IN ATU_BASE_ADDR_LOW

Field Name	Bit	Access	Description	Default
ATU_BASE_ADDR_LOW	31:16	R/W	Bits 31:16 of the starting address of the address region to be translated.	0x0000

5.10.1.6 PCIE:PCIE:ATU_LIMIT_ADDR

Parent: [PCIE:PCIE](#)

Instances: 1

This register configures the window which is translated.

For example: If PCIE::ATU_BASE_ADDR_LOW is set to 0x00010000, PCIE::ATU_LIMIT_ADDR is set to 0x0005FFFF, PCIE::ATU_TGT_ADDR_LOW is set to 0x20000000, and PCIE::ATU_TGT_ADDR_HIGH is set to 0x00007000. Then outbound TLPs with PCIe addresses in the range 0x00010000 through 0x0005FFFF is mapped to 64bit address 0x0000700020000000 through 0x000070002004FFFF. Header fields of the mapped TLPs are configured via PCIE::ATU_CFG1 and PCIE::ATU_CFG2.

TABLE 5-476: FIELDS IN ATU_LIMIT_ADDR

Field Name	Bit	Access	Description	Default
ATU_LIMIT_ADDR	31:16	R/W	Bits 31:16 of the ending address of the address region to be translated.	0x0000
RESERVED	15:0	R/O	Must be set to its default.	0xFFFF

5.10.1.7 PCIE:PCIE:ATU_TGT_ADDR_LOW

Parent: [PCIE:PCIE](#)

Instances: 1

TABLE 5-477: FIELDS IN ATU_TGT_ADDR_LOW

Field Name	Bit	Access	Description	Default
ATU_TGT_ADDR_LOW	31:16	R/W	Bits 31:16 of the new address of the translated region.	0x0000

5.10.1.8 PCIE:PCIE:ATU_TGT_ADDR_HIGH

Parent: [PCIE:PCIE](#)

Instances: 1

TABLE 5-478: FIELDS IN ATU_TGT_ADDR_HIGH

Field Name	Bit	Access	Description	Default
ATU_TGT_ADDR_HIGH	31:0	R/W	Bits 63:32 of the new address of the translated region. Set to 0 to force new address into 32bit PCIe memory space.	0x00000000

5.11 QSYS

TABLE 5-479: REGISTER GROUPS IN QSYS

Register Group Name	Offset within Target	Instances and Address Spacing	Description	Details
SYSTEM	0x00015A00	1	Switch configuration	Page 551
QMAP	0x00015ADC	13 0x00000004	Mapping frames to scheduling elements	Page 555
SGRP	0x00018000	1024 0x00000004		Page 557
TIMED_FRAME_CFG	0x00015B10	1		Page 557
RES_QOS_ADV	0x00015B38	1		Page 560
RES_CTRL	0x00016000	1024 0x00000008	Watermarks and status for egress queue system	Page 561
DROP_CFG	0x00015B7C	1	Watermarks for egress queue system	Page 563
MMGT	0x00015B80	1	Memory manager status	Page 563
HSCH	0x00000000	692 0x00000080	Configuration of scheduling system and shapers	Page 563
HSCH_MISC	0x00015B88	1	Miscellaneous scheduler configuration	Page 568

5.11.1 QSYS:SYSTEM

Parent: [QSYS](#)

Instances: 1

TABLE 5-480: REGISTERS IN SYSTEM

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
PORT_MODE	0x00000000	13 0x00000004	Per device port configuration	Page 552
SWITCH_PORT_MODE	0x00000034	12 0x00000004	Various switch port mode settings	Page 552
STAT_CNT_CFG	0x00000064	1	Statistics configuration	Page 553
EEE_CFG	0x00000068	11 0x00000004	Control Energy Efficient Ethernet operation per front port.	Page 553

TABLE 5-480: REGISTERS IN SYSTEM (CONTINUED)

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
EEE_THRES	0x00000094	1	Thresholds for delayed EEE queues	Page 553
IGR_NO_SHARING	0x00000098	1	Control shared memory users	Page 554
EGR_NO_SHARING	0x0000009C	1	Control shared memory users	Page 554
SW_STATUS	0x000000A0	12 0x00000004	Various status info per switch port	Page 554
EXT_CPU_CFG	0x000000D0	1	External CPU port configuration	Page 554
CPU_GROUP_MAP	0x000000D8	1	Map CPU extraction queues to CPU ports	Page 555

5.11.1.1 QSYS:SYSTEM:PORT_MODE

Parent: [QSYS:SYSTEM](#)

Instances: 13

These configurations exists per frontport and for each of the two CPU ports (11+12).

TABLE 5-481: FIELDS IN PORT_MODE

Field Name	Bit	Access	Description	Default
DEQUEUE_DIS	1	R/W	Disable dequeuing from the egress queues. Frames are not discarded, but may become aged when dequeuing is reenabled.	0x0
DEQUEUE_LATE	0	R/W	Delay dequeuing from the egress queues. This might be necessary for minimising flow control tails.	0x0

5.11.1.2 QSYS:SYSTEM:SWITCH_PORT_MODE

Parent: [QSYS:SYSTEM](#)

Instances: 12

TABLE 5-482: FIELDS IN SWITCH_PORT_MODE

Field Name	Bit	Access	Description	Default
PORT_ENA	13	R/W	Enable port for any frame transfer. Frames to or from a port with PORT_ENA cleared are discarded.	0x0
RESERVED	12	R/W	Must be set to its default.	0x1
RESERVED	11	R/W	Must be set to its default.	0x1
INGRESS_DROP_MODE	9	R/W	When enabled for a port, frames -from- that port are discarded when the controlling watermarks are hit. If disabled - the frame will stay in memory until resources are available. If INGRESS_DROP_MODE or EGRESS_DROP_MODE applies for a frame copy, it will be discarded.	0x1

TABLE 5-482: FIELDS IN SWITCH_PORT_MODE (CONTINUED)

Field Name	Bit	Access	Description	Default
TX_PFC_ENA	8:1	R/W	When set the MAC sends PRIO pause control frames in the Tx direction when congested.	0x00
TX_PFC_MODE	0	R/W	When set, a congested priority request pause of all lower priorities as well.	0x0

5.11.1.3 QSYS:SYSTEM:STAT_CNT_CFG

Parent: [QSYS:SYSTEM](#)

Instances: 1

TABLE 5-483: FIELDS IN STAT_CNT_CFG

Field Name	Bit	Access	Description	Default
TX_GREEN_CNT_MODE	5	R/W	Counter mode for the Tx priority counters for green frames (C_TX_GREEN_PRIO_x) 0: Count octets 1: Count frames	0x1
TX_YELLOW_CNT_MODE	4	R/W	Counter mode for the Tx priority counters for green frames (C_TX_YELLOW_PRIO_x) 0: Count octets 1: Count frames	0x1
DROP_GREEN_CNT_MODE	3	R/W	Counter mode for the drop counters for green frames (C_DR_GREEN_PRIO_x) 0: Count octets 1: Count frames	0x1
DROP_YELLOW_CNT_MODE	2	R/W	Counter mode for the drop counters for green frames (C_DR_YELLOW_PRIO_x) 0: Count octets 1: Count frames	0x1
DROP_COUNT_EGRESS	0	R/W	When set, a frame discarded due to lack of resources is counted on the egress port instead of the ingress. Side effect is a slower processing of multiple drops on the same frame, causing potential head-of-line blocking.	0x0

5.11.1.4 QSYS:SYSTEM:EEE_CFG

Parent: [QSYS:SYSTEM](#)

Instances: 11

TABLE 5-484: FIELDS IN EEE_CFG

Field Name	Bit	Access	Description	Default
EEE_FAST_QUEUES	7:0	R/W	Queues set in this mask activate the egress port immediately when any of the queues have data available.	0x00

5.11.1.5 QSYS:SYSTEM:EEE_THRES

Parent: [QSYS:SYSTEM](#)

Instances: 1

TABLE 5-485: FIELDS IN EEE_THRES

Field Name	Bit	Access	Description	Default
EEE_HIGH_BYTES	15:8	R/W	Maximum number of bytes in a queue before egress port is activated. Unit is 60 bytes.	0x00
EEE_HIGH_FRAMES	7:0	R/W	Maximum number of frames in a queue before the egress port is activated. Unit is 1 frame.	0x00

5.11.1.6 QSYS:SYSTEM:IGR_NO_SHARING

Parent: [QSYS:SYSTEM](#)

Instances: 1

TABLE 5-486: FIELDS IN IGR_NO_SHARING

Field Name	Bit	Access	Description	Default
IGR_NO_SHARING	11:0	R/W	Control whether frames received on the port may use shared resources. If ingress port or queue has reserved memory left to use, frame enqueueing is always allowed. 0: Use shared memory as well 1: Do not use shared memory	0x000

5.11.1.7 QSYS:SYSTEM:EGR_NO_SHARING

Parent: [QSYS:SYSTEM](#)

Instances: 1

TABLE 5-487: FIELDS IN EGR_NO_SHARING

Field Name	Bit	Access	Description	Default
EGR_NO_SHARING	11:0	R/W	Control whether frames forwarded to the port may use shared resources. If egress port or queue has reserved memory left to use, frame enqueueing is always allowed. 0: Use shared memory as well 1: Do not use shared memory	0x000

5.11.1.8 QSYS:SYSTEM:SW_STATUS

Parent: [QSYS:SYSTEM](#)

Instances: 12

TABLE 5-488: FIELDS IN SW_STATUS

Field Name	Bit	Access	Description	Default
EQ_AVAIL	7:0	R/O	Status bit per egress queue indicating whether data is ready for transmission.	0x00

5.11.1.9 QSYS:SYSTEM:EXT_CPU_CFG

Parent: [QSYS:SYSTEM](#)

Instances: 1

TABLE 5-489: FIELDS IN EXT_CPU_CFG

Field Name	Bit	Access	Description	Default
EXT_CPU_PORT	12:8	R/W	Select the port to use as the external CPU port.	0x0C
EXT_CPUQ_MSK	7:0	R/W	Frames destined for a CPU extraction queue set in this mask are sent to the external CPU defined by EXT_CPU_PORT instead of the internal CPU.	0x00

5.11.1.10 QSYS:SYSTEM:CPU_GROUP_MAP

Parent: [QSYS:SYSTEM](#)

Instances: 1

TABLE 5-490: FIELDS IN CPU_GROUP_MAP

Field Name	Bit	Access	Description	Default
CPU_GROUP_MAP	7:0	R/W	Map the 8 CPU extraction queues to the two CPU ports. Bit <n> set to 1 directs CPU extraction queue <n> to CPU port 12. Bit <n> set to 0 directs CPU extraction queue <n> to CPU port 11.	0x00

5.11.2 QSYS:QMAP

Parent: [QSYS](#)

Instances: 13

TABLE 5-491: REGISTERS IN QMAP

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
QMAP	0x00000000	1	Mapping frames to scheduling elements per egress port	Page 555

5.11.2.1 QSYS:QMAP:QMAP

Parent: [QSYS:QMAP](#)

Instances: 1

TABLE 5-492: FIELDS IN QMAP

Field Name	Bit	Access	Description	Default
SE_BASE	12:5	R/W	Base index of scheduling elements associated with the egress port. Scheduling element number = SE_BASE + SE_IDX_SEL. Input port number on scheduling element = SE_INP_SEL. Example: SE_BASE = 0, SE_IDX_SEL = 0 (ingress port), and SE_INP_SEL = 1 (QoS class). Scheduling elements 0, 1, ..., 11 are assigned to the egress port. For a frame received on port 5 and classified to QoS class 2, input port 2 on scheduling element 5 is selected for the frame. Note that the highest indexable scheduling element number is 217 and the highest indexable input port number is 11. Default: When SE_IDX_SEL=0, a default configuration overrules other configured values in QMAP. In this configuration, 1st level scheduling towards the egress queues operates at the QoS class, and 2nd level towards the egress port operates at the source port. SE_BASE is automatically set to 8 times the egress port number.	0x00
SE_IDX_SEL	4:2	R/W	Scheduling element selector for the egress port. The egress port use the scheduling elements selected by SE_BASE + SE_IDX_SEL. See SE_BASE. 000: Default configuration for port scheduling 001: One level queueing. Only one scheduling element in use. 010: Reserved 011: Reserved 100: Add the frame's ingress port to SE_BASE to select scheduling element. 101: Add the frame's QoS class to SE_BASE to select the scheduling element. 110: Add the frame's service group to SE_BASE to select the scheduling element. 111: Add the frame's service group modulus 12 to SE_BASE to select the scheduling element.	0x0

TABLE 5-492: FIELDS IN QMAP (CONTINUED)

Field Name	Bit	Access	Description	Default
SE_INP_SEL	1:0	R/W	Select the input port on the scheduling element selected by SE_BASE and SE_IDX_SEL. See SE_BASE. 00: Use the frame's ingress port as input port number on the scheduling element. 01: Use the frame's QoS class as input port number on the scheduling element. 10: Use the frame's service group as input port number on the scheduling element. 11: Use the frame's service group modulus 12 as input port number on the scheduling element.	0x0

5.11.3 QSYS:SGRP

Parent: [QSYS](#)

Instances: 1024

TABLE 5-493: REGISTERS IN SGRP

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
ISDX_SGRP	0x00000000	1	Configures service group per ISDX	Page 557

5.11.3.1 QSYS:SGRP:ISDX_SGRP

Parent: [QSYS:SGRP](#)

Instances: 1

TABLE 5-494: FIELDS IN ISDX_SGRP

Field Name	Bit	Access	Description	Default
ISDX_SGRP	6:0	R/W	Service group value used in the egress scheduler. The service group enables per service egress scheduling as the service group selects which set of egress queues to use. See QSYS::QMAP for more details.	0x00

5.11.4 QSYS:TIMED_FRAME_CFG

Parent: [QSYS](#)

Instances: 1

TABLE 5-495: REGISTERS IN TIMED_FRAME_CFG

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
TFRM_MISC	0x00000000	1	Timed frame configuration	Page 558
TFRM_PORT_DLY	0x00000004	1	Interleaved transmission setup	Page 558

TABLE 5-495: REGISTERS IN TIMED_FRAME_CFG (CONTINUED)

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
TFRM_TIMER_CFG_1	0x00000008	1	Configures timers for timed transmissions	Page 558
TFRM_TIMER_CFG_2	0x0000000C	1	Configures timers for timed transmissions	Page 559
TFRM_TIMER_CFG_3	0x00000010	1	Configures timers for timed transmissions	Page 559
TFRM_TIMER_CFG_4	0x00000014	1	Configures timers for timed transmissions	Page 559
TFRM_TIMER_CFG_5	0x00000018	1	Configures timers for timed transmissions	Page 559
TFRM_TIMER_CFG_6	0x0000001C	1	Configures timers for timed transmissions	Page 559
TFRM_TIMER_CFG_7	0x00000020	1	Configures timers for timed transmissions	Page 560
TFRM_TIMER_CFG_8	0x00000024	1	Configures timers for timed transmissions	Page 560

5.11.4.1 QSYS:TIMED_FRAME_CFG:TFRM_MISC

Parent: [QSYS:TIMED_FRAME_CFG](#)

Instances: 1

TABLE 5-496: FIELDS IN TFRM_MISC

Field Name	Bit	Access	Description	Default
TIMED_CANCEL_SLOT	18:9	R/W	Specify slot of timed frame to cancel	0x000
TIMED_CANCEL_1SHOT	8	One-shot	Set this field to cancel a timed transmission. Auto cleared when complete.	0x0
RESERVED	6:0	R/W	Must be set to its default.	0x02

5.11.4.2 QSYS:TIMED_FRAME_CFG:TFRM_PORT_DLY

Parent: [QSYS:TIMED_FRAME_CFG](#)

Instances: 1

TABLE 5-497: FIELDS IN TFRM_PORT_DLY

Field Name	Bit	Access	Description	Default
TFRM_PORT_DLY_ENA	12:0	R/W	When multiple timed frames are ready for transmission at the same time, they are by default transmitted as fast as they can, resulting in a burst of frames. By enabling the interleave mode, timer 8 must run out between each frame being transmitted on a port. 0: Send frames in burst 1: Send frames interleaved	0x0000

5.11.4.3 QSYS:TIMED_FRAME_CFG:TFRM_TIMER_CFG_1

Parent: [QSYS:TIMED_FRAME_CFG](#)

Instances: 1

TABLE 5-498: FIELDS IN TFRM_TIMER_CFG_1

Field Name	Bit	Access	Description	Default
TFRM_TIMER_CFG_1	31:0	R/W	Configures timer 1. Default 3.3 ms. Unit: 198.2 ns	0x0000410A

5.11.4.4 QSYS:TIMED_FRAME_CFG:TFRM_TIMER_CFG_2

Parent: [QSYS:TIMED_FRAME_CFG](#)

Instances: 1

TABLE 5-499: FIELDS IN TFRM_TIMER_CFG_2

Field Name	Bit	Access	Description	Default
TFRM_TIMER_CFG_2	31:0	R/W	Configures timer 2. Default 10 us. Unit: 198.2 ns	0x00000033

5.11.4.5 QSYS:TIMED_FRAME_CFG:TFRM_TIMER_CFG_3

Parent: [QSYS:TIMED_FRAME_CFG](#)

Instances: 1

TABLE 5-500: FIELDS IN TFRM_TIMER_CFG_3

Field Name	Bit	Access	Description	Default
TFRM_TIMER_CFG_3	31:0	R/W	Configures timer 3. Default 100 us. Unit: 198.2 ns	0x000001F9

5.11.4.6 QSYS:TIMED_FRAME_CFG:TFRM_TIMER_CFG_4

Parent: [QSYS:TIMED_FRAME_CFG](#)

Instances: 1

TABLE 5-501: FIELDS IN TFRM_TIMER_CFG_4

Field Name	Bit	Access	Description	Default
TFRM_TIMER_CFG_4	31:0	R/W	Configures timer 4. Default 1 ms. Unit: 198.2 ns	0x000013B5

5.11.4.7 QSYS:TIMED_FRAME_CFG:TFRM_TIMER_CFG_5

Parent: [QSYS:TIMED_FRAME_CFG](#)

Instances: 1

TABLE 5-502: FIELDS IN TFRM_TIMER_CFG_5

Field Name	Bit	Access	Description	Default
TFRM_TIMER_CFG_5	31:0	R/W	Configures timer 5. Default 10 ms. Unit: 198.2 ns	0x0000C516

5.11.4.8 QSYS:TIMED_FRAME_CFG:TFRM_TIMER_CFG_6

Parent: [QSYS:TIMED_FRAME_CFG](#)

VSC7416-01

Instances: 1

TABLE 5-503: FIELDS IN TFRM_TIMER_CFG_6

Field Name	Bit	Access	Description	Default
TFRM_TIMER_CFG_6	31:0	R/W	Configures timer 6. Default 100 ms. Unit: 198.2 ns	0x0007B2DD

5.11.4.9 QSYS:TIMED_FRAME_CFG:TFRM_TIMER_CFG_7

Parent: [QSYS:TIMED_FRAME_CFG](#)

Instances: 1

TABLE 5-504: FIELDS IN TFRM_TIMER_CFG_7

Field Name	Bit	Access	Description	Default
TFRM_TIMER_CFG_7	31:0	R/W	Configures timer 7. Default 1sec. Unit: 198.2 ns	0x004CFCA1

5.11.4.10 QSYS:TIMED_FRAME_CFG:TFRM_TIMER_CFG_8

Parent: [QSYS:TIMED_FRAME_CFG](#)

Instances: 1

TABLE 5-505: FIELDS IN TFRM_TIMER_CFG_8

Field Name	Bit	Access	Description	Default
TFRM_TIMER_CFG_8	31:0	R/W	Configures timer 8. Default 198.2 ns. Unit: 198.2 ns	0x00000001

5.11.5 QSYS:RES_QOS_ADV

Parent: [QSYS](#)

Instances: 1

TABLE 5-506: REGISTERS IN RES_QOS_ADV

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
RED_PROFILE	0x00000000	16 0x00000004	Random Early Discard (RED) configuration	Page 560
RES_QOS_MODE	0x00000040	1	Shared QOS resource mode	Page 561

5.11.5.1 QSYS:RES_QOS_ADV:RED_PROFILE

Parent: [QSYS:RES_QOS_ADV](#)

Instances: 16

Configuration of Weighted Random Early Discard (WRED) profile per QoS class per drop precedence level. Profiles 0-7 are for frames with QoS class 0-7 and drop precedence level 0. Profiles 8-15 are for frames with QoS class 0-7 and drop precedence level 1.

TABLE 5-507: FIELDS IN RED_PROFILE

Field Name	Bit	Access	Description	Default
WM_RED_LOW	21:11	R/W	When enqueueing a frame, WRED is active if the ingress memory consumption by the frame's QoS class is above WM_RED_LEVEL. The probability of random early discarding is calculated as: (Memory consumption by the frame's QoS class - WM_RED_LOW)/(WM_RED_HIGH - WM_RED_LOW). Unit is 960 bytes.	0x7FF
WM_RED_HIGH	10:0	R/W	See WM_RED_LOW. Unit is 960 bytes.	0x7FF

5.11.5.2 QSYS:RES_QOS_ADV:RES_QOS_MODE

Parent: [QSYS:RES_QOS_ADV](#)

Instances: 1

TABLE 5-508: FIELDS IN RES_QOS_MODE

Field Name	Bit	Access	Description	Default
RES_QOS_RSRVD	7:0	R/W	When a qos class is enabled in this mask, the class will have guaranteed shared space. The watermarks found in RES_CFG are used for setting the amount of space set aside.	0x00

5.11.6 QSYS:RES_CTRL

Parent: [QSYS](#)

Instances: 1024

TABLE 5-509: REGISTERS IN RES_CTRL

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
RES_CFG	0x00000000	1	Watermark configuration	Page 561
RES_STAT	0x00000004	1	Resource status	Page 562

5.11.6.1 QSYS:RES_CTRL:RES_CFG

Parent: [QSYS:RES_CTRL](#)

Instances: 1

The queue system tracks four resource consumptions:

Resource 0: Memory tracked per source

Resource 1: Frame references tracked per source

Resource 2: Memory tracked per destination

Resource 3: Frame references tracked per destination

Before a frame is added to the queue system, some conditions must be met:

VSC7416-01

- Reserved memory for the specific (SRC, PRIO) or for the specific SRC is available

OR

- Reserved memory for the specific (DST,PRIO) or for the specific DST is available

OR

- Shared memory is available

The frame reference resources are checked for availability like the memory resources. Enqueuing of a frame is allowed if both the memory resource check and the frame reference resource check succeed.

The extra resources consumed when enqueueing a frame are first taken from the reserved (SRC,PRIO), next from the reserved SRC, and last from the shared memory area. The same is done for DST. Both memory consumptions and frame reference consumptions are updated.

The register is laid out in the following way for source memory (resource 0):

Index 0-95: Q_RSRV - Reserved amount for (SRC,PRIO) at index $8 \times \text{SRC} + \text{PRIO}$

Index 216-223: PRIO_SHR - Maximum allowed use of the shared buffer for PRIO at index $\text{PRIO} + 216$

Index 224-235: P_RSRV - Reserved amount for SRC at index $\text{SRC} + 224$.

Index 254: COL_SHR - Maximum allowed use of the shared buffer for frames with DP=1.

Index 255: COL_SHR Maximum allowed use of the shared buffer for source.

Note The default values depend on the index used.

The layout is similar for resources 1, 2, and 3. Resource 1 uses indices 256-511, resource 2 uses indices 512-767, and resource 3 uses indices 768-1023.

The allocation unit for memory tracking is 60 bytes, and the allocation unit for reference tracking is 1 frame reference. All frames are prepended with a 16-byte header.

TABLE 5-510: FIELDS IN RES_CFG

Field Name	Bit	Access	Description	Default
WM_MULTIPLIER	11	R/W	Selects multiplier unit for watermark. 0: Multiply watermark value with 1. 1: Multiply watermark value with 16.	0x0
WM_HIGH	10:0	R/W	Watermark value, multiplied with WM_MUL- TIPLIER.	0x000

5.11.6.2 QSYS:RES_CTRL:RES_STAT

Parent: [QSYS:RES_CTRL](#)

Instances: 1

TABLE 5-511: FIELDS IN RES_STAT

Field Name	Bit	Access	Description	Default
INUSE	29:15	R/W	Current consumption for corresponding watermark in RES_CFG.	0x0000
MAXUSE	14:0	R/W	Maximum consumption since last read for corresponding watermark in RES_CFG.	0x0000

5.11.7 QSYS:DROP_CFG

Parent: [QSYS](#)

Instances: 1

TABLE 5-512: REGISTERS IN DROP_CFG

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
EGR_DROP_MODE	0x00000000	1	Configures egress ports for flowcontrol	Page 563

5.11.7.1 QSYS:DROP_CFG:EGR_DROP_MODE

Parent: [QSYS:DROP_CFG](#)

Instances: 1

TABLE 5-513: FIELDS IN EGR_DROP_MODE

Field Name	Bit	Access	Description	Default
EGRESS_DROP_MODE	11:0	R/W	When enabled for a port, frames -to- that port are discarded when the controlling watermarks are hit. If disabled - the frame will stay in memory until resources are available. If INGRESS_DROP_MODE or EGRESS_DROP_MODE applies for a frame copy, it will be discarded.	0x000

5.11.8 QSYS:MMGT

Parent: [QSYS](#)

Instances: 1

TABLE 5-514: REGISTERS IN MMGT

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
EQ_CTRL	0x00000000	1	Egress queue status	Page 563

5.11.8.1 QSYS:MMGT:EQ_CTRL

Parent: [QSYS:MMGT](#)

Instances: 1

TABLE 5-515: FIELDS IN EQ_CTRL

Field Name	Bit	Access	Description	Default
FP_FREE_CNT	15:0	R/O	Number of free frame references.	0x0000

5.11.9 QSYS:HSCH

Parent: [QSYS](#)

Instances: 692

VSC7416-01

General purpose scheduling elements are used to form a hierarchy of queue scheduling. There are 256 elements.

The first 218 elements connect directly to frame queues. These elements have extra shapers on inputs 6 and 7, which can be used for strict/DWRR mixed priority scheduling. These shapers are configurable through this set of registers, at indexes $256+2*SE\_NO$ (shaper 6) and indexes $256+2*SE\_NO+1$ (shaper 7). Only the CIR_CFG and CIR_STATE registers are valid for these indexes.

Example: To set the priority 7 shaper on scheduling element 67, access HSCH:391:CIR_CFG.

TABLE 5-516: REGISTERS IN HSCH

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
CIR_CFG	0x00000000	1	Shaping configuration of the SE	Page 564
EIR_CFG	0x00000004	1	Excess rate configuration	Page 564
SE_CFG	0x00000008	1	Configuration of shaper and scheduling algorithm	Page 565
SE_DWRR_CFG	0x0000000C	12 0x00000004	Configuration of the DWRR costs	Page 565
SE_CONNECT	0x0000003C	1	Configuration of the connections between SEs	Page 565
SE_DLB_SENSE	0x00000040	1	Configuration of which fill levels in the queue system that the DLB shapers use to trigger excess information rate	Page 566
CIR_STATE	0x00000044	1	CIR status	Page 567
EIR_STATE	0x00000048	1	EIR status	Page 568

5.11.9.1 QSYS:HSCH:CIR_CFG

Parent: [QSYS:HSCH](#)

Instances: 1

TABLE 5-517: FIELDS IN CIR_CFG

Field Name	Bit	Access	Description	Default
CIR_RATE	20:6	R/W	Leak rate for this shaper. Unit is 100 kbps.	0x0000
CIR_BURST	5:0	R/W	Burst capacity of this shaper. Unit is 4 kilobytes. The shaper is disabled when CIR_BURST=0.	0x00

5.11.9.2 QSYS:HSCH:EIR_CFG

Parent: [QSYS:HSCH](#)

Instances: 1

TABLE 5-518: FIELDS IN EIR_CFG

Field Name	Bit	Access	Description	Default
EIR_RATE	21:7	R/W	Leak rate for this shaper. Unit is 100 kbps.	0x0000

TABLE 5-518: FIELDS IN EIR_CFG (CONTINUED)

Field Name	Bit	Access	Description	Default
EIR_BURST	6:1	R/W	Burst capacity of this shaper. Unit is 4 kilobytes. The dual leaky bucket shaper operates as a single leaky bucket shaper when EIR_BURST=0.	0x00
EIR_MARK_ENA	0	R/W	If a frame is scheduled with excess information rate, the frame color can be changed to yellow.	0x0

5.11.9.3 QSYS:HSCH:SE_CFG

Parent: [QSYS:HSCH](#)

Instances: 1

TABLE 5-519: FIELDS IN SE_CFG

Field Name	Bit	Access	Description	Default
SE_DWRR_CNT	9:6	R/W	Number of inputs running with DWRR algorithm, otherwise strict. Strict inputs always have the highest input index.	0x0
SE_RR_ENA	5	R/W	The DWRR algorithm is replaced with frame-based round robin.	0x0
SE_AVB_ENA	4	R/W	Enable AVB mode for this shaper, creating burst capacity only when data is available.	0x0
SE_FRM_MODE	3:2	R/W	Accounting mode for this shaper. 0: Line rate. Shape bytes including HSCH_MISC::FRM_ADJ. 1: Data rate. Shape bytes excluding IPG. 2: Frame rate. Shape frames with rate unit = 100 fps and burst unit = 32.8 frames. 3: Frame rate. Shape framed with rate unit = 1 fps and burst unit = 0.3 frames.	0x0
SE_EXC_ENA	1	R/W	Allow bandwidth out of this element to exceed configured levels. Still - arbitrations always prefers inputs not exceeding configured levels.	0x0

5.11.9.4 QSYS:HSCH:SE_DWRR_CFG

Parent: [QSYS:HSCH](#)

Instances: 12

TABLE 5-520: FIELDS IN SE_DWRR_CFG

Field Name	Bit	Access	Description	Default
DWRR_COST	4:0	R/W	DWRR cost configuration	0x00

5.11.9.5 QSYS:HSCH:SE_CONNECT

Parent: [QSYS:HSCH](#)

Instances: 1

TABLE 5-521: FIELDS IN SE_CONNECT

Field Name	Bit	Access	Description	Default
SE_OUTP_IDX	24:17	R/W	The output port on this SE is connected to the SE selected by SE_OUTP_IDX. SE_OUTP_CON selects the input port number on the SE. If the SE_TERMINAL is set, the output port on this SE is connected to the egress port selected by SE_OUTP_IDX instead. The connections are automatically configured after reset to a source port fair strict priority scheme: SEs 0-7 connect to SE 218: Priority scheduler for port 0 SEs 8-15 connect to SE 219: Priority scheduler for port 1 ... SEs 96-103 connect to SE 230: Priority scheduler for CPU port 12 SEs 218+N connect to egress port N with SE_TERMINAL set.	0x00
SE_INP_IDX	16:9	R/W	The inputs on this SE are the output ports of the consecutive range of SEs starting at index SE_INP_IDX and ending at SE_INP_IDX+SE_INP_CNT-1. The 2616 queues connect 12 at a time to the first 218 SEs, and the SE_INP_IDX does not apply for those.	0x00
SE_OUTP_CON	8:5	R/W	The input port on the next SE, which this SE connects to. The next SE is selected by SE_OUTP_IDX.	0x0
SE_INP_CNT	4:1	R/W	Number of input ports on to this SE. Default values matches the default scheduling hierarchy described in the datasheet.	0xC
SE_TERMINAL	0	R/W	If set, this SE is connected to egress port SE_OUTP_IDX	0x0

5.11.9.6 QSYS:HSCH:SE_DLB_SENSE

Parent: [QSYS:HSCH](#)

Instances: 1

TABLE 5-522: FIELDS IN SE_DLB_SENSE

Field Name	Bit	Access	Description	Default
SE_DLB_PRIO	13:11	R/W	QoS class used when SE_DLB_PRIO_ENA is set.	0x0
SE_DLB_SPORT	10:7	R/W	Source port used when SE_DLB_SPORT_ENA is set.	0x0
SE_DLB_DPORT	6:3	R/W	Destination port used when SE_DLB_DPORT_ENA is set.	0x0

TABLE 5-522: FIELDS IN SE_DLB_SENSE (CONTINUED)

Field Name	Bit	Access	Description	Default
SE_DLB_PRIO_ENA	2	R/W	Enable priority awareness by the DLB shaper. If set, the DLB shaper allows excess information rate when the egress buffer use for QoS class SE_DLB_PRIO has reached threshold BUF_PRIO_SHR_E[SE_DLB_PRIO] minus 3000 bytes. If both SE_DLB_PRIO_ENA and SE_DLB_SPORT_ENA are set, the excess information rate is allowed when the ingress buffer use for QoS class per source port has reached threshold BUF_Q_RSRV_I[SE_DLB_PRIO, SE_DLB_SPORT] minus 3000 bytes. If both SE_DLB_PRIO_ENA and SE_DLB_DPORT_ENA are set, the excess information rate is allowed when the egress buffer use for QoS class per destination port has reached threshold BUF_Q_RSRV_E[SE_DLB_PRIO, SE_DLB_DPORT] minus 3000 bytes. If multiple awareness functions (SE_DLB_PRIO_ENA, SE_DLB_SPORT_ENA, SE_DLB_DPORT_ENA) are enabled, all relevant thresholds must be exceeded before excess information rate is allowed.	0x0
SE_DLB_SPORT_ENA	1	R/W	Enable source port awareness by the DLB shaper. If set, the DLB shaper allows excess information rate when the ingress buffer use for source port SE_DLB_SPORT has reached threshold BUF_P_RSRV_I[SE_DLB_SPORT] minus 3000 bytes. If multiple awareness functions (SE_DLB_PRIO_ENA, SE_DLB_SPORT_ENA, SE_DLB_DPORT_ENA) are enabled, all relevant thresholds must be exceeded before excess information rate is allowed.	0x0
SE_DLB_DPORT_ENA	0	R/W	Enable destination port awareness by the DLB shaper. If set, the DLB shaper allows excess information rate when the egress buffer use for destination port SE_DLB_SPORT has reached threshold BUF_P_RSRV_E[SE_DLB_DPORT] minus 3000 bytes. If multiple awareness functions (SE_DLB_PRIO_ENA, SE_DLB_SPORT_ENA, SE_DLB_DPORT_ENA) are enabled, all relevant thresholds must be exceeded before excess information rate is allowed.	0x0

5.11.9.7 QSYS:HSCH:CIR_STATE

Parent: [QSYS:HSCH](#)

VSC7416-01

Instances: 1

TABLE 5-523: FIELDS IN CIR_STATE

Field Name	Bit	Access	Description	Default
CIR_LVL	25:4	R/W	Current fill level. Unit is 0.5 bits.	0x000000

5.11.9.8 QSYS:HSCH:EIR_STATE

Parent: [QSYS:HSCH](#)

Instances: 1

TABLE 5-524: FIELDS IN EIR_STATE

Field Name	Bit	Access	Description	Default
EIR_LVL	21:0	R/W	Current fill level. Unit is 0.5 bits.	0x000000

5.11.10 QSYS:HSCH_MISC

Parent: [QSYS](#)

Instances: 1

TABLE 5-525: REGISTERS IN HSCH_MISC

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
HSCH_MISC_CFG	0x00000000	1	Common config for HSCH and policer module	Page 568

5.11.10.1 QSYS:HSCH_MISC:HSCH_MISC_CFG

Parent: [QSYS:HSCH_MISC](#)

Instances: 1

TABLE 5-526: FIELDS IN HSCH_MISC_CFG

Field Name	Bit	Access	Description	Default
SE_CONNECT_VLD	7	R/W	Set to one when the SE connection configuration is valid. Set to zero when reconfiguring SE inter connections	0x1
FRM_ADJ	6:2	R/W	Values to add each frame when frame length adjustment is in use.	0x14
QSHP_EXC_ENA	0	R/W	When set closed queue shapers will be scheduled as excessed, and can be scheduled if no other inputs are available.	0x0

5.12 REW

TABLE 5-527: REGISTER GROUPS IN REW

Register Group Name	Offset within Target	Instances and Address Spacing	Description	Details
PORT	0x00000000	13 0x00000080	Per port configurations for Rewriter	Page 569
COMMON	0x00000690	1	Common configurations for Rewriter	Page 572
PPT	0x00000680	1	Path group protection table per path	Page 573

5.12.1 REW:PORT

Parent: [REW](#)

Instances: 13

TABLE 5-528: REGISTERS IN PORT

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
PORT_VLAN_CFG	0x00000000	1	Port VLAN configuration	Page 569
TAG_CFG	0x00000004	1	Tagging configuration	Page 569
PORT_CFG	0x00000008	1	Special port configuration	Page 570
DSCP_CFG	0x0000000C	1	DSCP updates	Page 571
PCP_DEI_QOS_MAP_CFG	0x00000010	16 0x00000004	Mapping of DP level and QoS class to PCP and DEI values	Page 571
PTP_CFG	0x00000050	1	Precision time protocol configuration	Page 571
PTP_DLY1_CFG	0x00000054	1	Asymmetric delay	Page 572

5.12.1.1 REW:PORT:PORT_VLAN_CFG

Parent: [REW:PORT](#)

Instances: 1

TABLE 5-529: FIELDS IN PORT_VLAN_CFG

Field Name	Bit	Access	Description	Default
PORT_TPID	31:16	R/W	Tag Protocol Identifier for port.	0x0000
PORT_DEI	15	R/W	DEI value for port.	0x0
PORT_PCP	14:12	R/W	PCP value for port.	0x0
PORT_VID	11:0	R/W	VID value for port.	0x001

5.12.1.2 REW:PORT:TAG_CFG

Parent: [REW:PORT](#)

Instances: 1

TABLE 5-530: FIELDS IN TAG_CFG

Field Name	Bit	Access	Description	Default
TAG_CFG	8:7	R/W	Enable VLAN port tagging. 0: Port tagging disabled. 1: Tag all frames, except when VID=PORT_VLAN_CFG.PORT_VID or VID=0. 2: Tag all frames, except when VID=0. 3: Tag all frames.	0x0
TAG_TPID_CFG	6:5	R/W	Select TPID EtherType in port tag. 0: Use 0x8100. 1: Use 0x88A8. 2: Use custom value from PORT_VLAN_CFG.PORT_TPID. 3: Use PORT_VLAN_CFG.PORT_TPID, unless ingress tag was a C-tag (EtherType = 0x8100)	0x0
TAG_VID_CFG	4	R/W	Select VID in port tag. 0: Use classified VID. 1: Use PORT_VLAN_CFG.PORT_VID.	0x0
TAG_PCP_CFG	3:2	R/W	Select PCP in port tag. 0: Classified PCP 1: PORT_PCP 2: DP and QoS mapped to PCP (PCP_DEI_QOS_MAP_CFG) 3: QoS class	0x0
TAG_DEI_CFG	1:0	R/W	Select DEI in port tag. 0: Classified DEI 1: PORT_DEI 2: DP and QoS mapped to DEI (PCP_DEI_QOS_MAP_CFG) 3: DP level	0x0

5.12.1.3 REW:PORT:PORT_CFG

Parent: [REW:PORT](#)

Instances: 1

TABLE 5-531: FIELDS IN PORT_CFG

Field Name	Bit	Access	Description	Default
ES0_ENA	5	R/W	Enable ES0 lookup.	0x0
FCS_UPDATE_NONCPU_CFG	4:3	R/W	FCS update mode for frames not received on the CPU port. 0: Update FCS if frame data has changed 1: Never update FCS 2: Always update FCS	0x0
FCS_UPDATE_CPU_ENA	2	R/W	If set, update FCS for all frames injected by the CPU. If cleared, never update the FCS.	0x1
FLUSH_ENA	1	R/W	If set, all frames destined for the egress port are discarded.	0x0

TABLE 5-531: FIELDS IN PORT_CFG (CONTINUED)

Field Name	Bit	Access	Description	Default
AGE_DIS	0	R/W	Disable frame ageing for this egress port.	0x0

5.12.1.4 REW:PORT:DSCP_CFG

Parent: [REW:PORT](#)

Instances: 1

TABLE 5-532: FIELDS IN DSCP_CFG

Field Name	Bit	Access	Description	Default
DSCP_REWR_CFG	1:0	R/W	Egress DSCP rewrite. 0: No update of DSCP value in frame. 1: Update with DSCP value from analyzer. 2: Update with DSCP value from analyzer remapped through DSCP_REMAP_CFG. 3: Update with DSCP value from analyzer remapped based on drop precedence level through DSCP_REMAP_CFG or DSCP_REMAP_DP1_CFGI.	0x0

5.12.1.5 REW:PORT:PCP_DEI_QOS_MAP_CFG

Parent: [REW:PORT](#)

Instances: 16

TABLE 5-533: FIELDS IN PCP_DEI_QOS_MAP_CFG

Field Name	Bit	Access	Description	Default
DEI_QOS_VAL	3	R/W	Map the frame's DP level and QoS class to a DEI value. DEI = PCP_DEI_QOS_MAP_CFG[8*DP level + QoS class].DEI_QOS_VAL. This must be enabled in TAG_CFG.TAG_DEI_CFG.	0x0
PCP_QOS_VAL	2:0	R/W	Map the frame's DP level and QoS class to a PCP value. PCP = PCP_DEI_QOS_MAP_CFG[8*DP level + QoS class].PCP_QOS_VAL. This must be enabled in TAG_CFG.TAG_PCP_CFG.	0x0

5.12.1.6 REW:PORT:PTP_CFG

Parent: [REW:PORT](#)

Instances: 1

TABLE 5-534: FIELDS IN PTP_CFG

Field Name	Bit	Access	Description	Default
PTP_BACKPLANE_MODE	7	R/W	Backplane mode for PTP. If set, the frame's receive timestamp is written into the PTP header (4 reserved bytes at offset 16).	0x0

TABLE 5-534: FIELDS IN PTP_CFG (CONTINUED)

Field Name	Bit	Access	Description	Default
PTP_1STEP_DIS	2	R/W	Overrules the one-step and the origin PTP commands from IS2. This disables updating the correction field or the origin timestamp field in the PTP frame.	0x0
PTP_2STEP_DIS	1	R/W	Overrules the two-step PTP command from IS2. This effectively disables saving the frame's departure time into the timestamp FIFO queue.	0x0

5.12.1.7 REW:PORT:PTP_DLY1_CFG

Parent: [REW:PORT](#)

Instances: 1

TABLE 5-535: FIELDS IN PTP_DLY1_CFG

Field Name	Bit	Access	Description	Default
PTP_DLY1_VAL	31:0	R/W	Delay to add to frame's residence time. This is applicable to one-step PTP and it is enabled by setting the IS2 action REW_OP[5] to 1. Unit is ns. Value is signed.	0x00000000

5.12.2 REW:COMMON

Parent: [REW](#)

Instances: 1

TABLE 5-536: REGISTERS IN COMMON

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
DSCP_REMAP_DP1_CFG	0x00000000	64 0x00000004	Remap table of DSCP values for frames with drop precedence set	Page 572
DSCP_REMAP_CFG	0x00000100	64 0x00000004	Remap table of DSCP values	Page 573
STAT_CFG	0x00000200	1	Egress statistics mode	Page 573

5.12.2.1 REW:COMMON:DSCP_REMAP_DP1_CFG

Parent: [REW:COMMON](#)

Instances: 64

TABLE 5-537: FIELDS IN DSCP_REMAP_DP1_CFG

Field Name	Bit	Access	Description	Default
DSCP_REMAP_DP1_VAL	5:0	R/W	One to one DSCP remapping table common for all ports. This table is used if DSCP_CFG.DSCP_REWR_ENA=3 and DP=1.	0x00

5.12.2.2 REW:COMMON:DSCP_REMAP_CFG

Parent: [REW:COMMON](#)

Instances: 64

TABLE 5-538: FIELDS IN DSCP_REMAP_CFG

Field Name	Bit	Access	Description	Default
DSCP_REMAP_VAL	5:0	R/W	One to one DSCP remapping table common for all ports. This table is used if DSCP_CFG.DSCP_REWR_ENA=2 or if DSCP_CFG.DSCP_REWR_ENA=3 and DP=0.	0x00

5.12.2.3 REW:COMMON:STAT_CFG

Parent: [REW:COMMON](#)

Instances: 1

TABLE 5-539: FIELDS IN STAT_CFG

Field Name	Bit	Access	Description	Default
STAT_MODE	1:0	R/W	Configures the egress service statistics mode. The following options for looking up the egress service counter set are available: 00: Use ESDX from ES0 if hit, otherwise ISDX 01: Use ESDX from ES0 if hit, otherwise no counting. 10: Use ISDX. 11: Use 8*SRC_PORT + QOS_CLASS.	0x0

5.12.3 REW:PPT

Parent: [REW](#)

Instances: 1

TABLE 5-540: REGISTERS IN PPT

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
PPT	0x00000000	3 0x00000004	Path protection table	Page 573

5.12.3.1 REW:PPT:PPT

Parent: [REW:PPT](#)

Instances: 3

First register in this group maps to path protection group indexes 31:0.

Second register in this group maps to path protection group indexes 63:32.

Third register in this group maps to path protection group indexes 75:64.

TABLE 5-541: FIELDS IN PPT

Field Name	Bit	Access	Description	Default
PATH_STAT	31:0	R/W	Path protection status with one entry per path protection group. The path protection group is selected by ES0_action.PPT_IDX from the ES0 lookup. Bit n in PPT[0].PATH_CFG returns the status for path protection group n, n <= 31. Bit n in PPT[1].PATH_CFG returns the status for path protection group n+32, 32 <= n <= 63. Bit n in PPT[2].PATH_CFG returns the status for path protection group n+64, 64 <= n <= 75. Each bit is encoded the following way: 0: Path protection group uses working path. 1: Path protection group uses protected path.	0x00000000

5.13 SBA

TABLE 5-542: REGISTER GROUPS IN SBA

Register Group Name	Offset within Target	Instances and Address Spacing	Description	Details
SBA	0x00000000	1	Shared bus arbiter registers	Page 574

5.13.1 SBA:SBA

Parent: [SBA](#)

Instances: 1

Configurations for the shared bus of the VCore System.

TABLE 5-543: REGISTERS IN SBA

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
PL_CPU	0x00000000	1	Master 1 arbitration priority	Page 575
PL_PCIE	0x00000004	1	Master 2 arbitration priority	Page 575
PL_CSR	0x00000008	1	Master 3 arbitration priority	Page 575
WT_EN	0x0000004C	1	Weighted-token arbitration scheme enable	Page 575
WT_TCL	0x00000050	1	Clock tokens refresh period	Page 575
WT_CPU	0x00000054	1	Master 1 clock tokens	Page 576
WT_PCIE	0x00000058	1	Master 2 clock tokens	Page 576
WT_CSR	0x0000005C	1	Master 3 clock tokens	Page 576

5.13.1.1 SBA:SBA:PL_CPU

Parent: [SBA:SBA](#)

Instances: 1

TABLE 5-544: FIELDS IN PL_CPU

Field Name	Bit	Access	Description	Default
PL1	3:0	R/W	Arbitration priority for the master. When multiple masters request the bus at the same time, the one with the highest priority is granted bus access. Values 0x1 through 0xF, higher values are prioritized over lower values.	0xE

5.13.1.2 SBA:SBA:PL_PCIE

Parent: [SBA:SBA](#)

Instances: 1

TABLE 5-545: FIELDS IN PL_PCIE

Field Name	Bit	Access	Description	Default
PL2	3:0	R/W	See SBA::PL1 for description.	0xD

5.13.1.3 SBA:SBA:PL_CSR

Parent: [SBA:SBA](#)

Instances: 1

TABLE 5-546: FIELDS IN PL_CSR

Field Name	Bit	Access	Description	Default
PL3	3:0	R/W	See SBA::PL1 for description.	0xC

5.13.1.4 SBA:SBA:WT_EN

Parent: [SBA:SBA](#)

Instances: 1

When weighted token arbitration is enabled, each master on the shared bus is granted a configurable number of tokens at the start of each refresh period. The length of each refresh period is configurable. In each clock-cycle that a master uses the bus, the token counter for that master decreases. Once all tokens are spent, the master is forced to a low priority. A master with tokens remaining, always takes priority over masters with no tokens remaining.

TABLE 5-547: FIELDS IN WT_EN

Field Name	Bit	Access	Description	Default
WT_EN	0	R/W	Set this field to enable weighted-token arbitration scheme.	0x0

5.13.1.5 SBA:SBA:WT_TCL

Parent: [SBA:SBA](#)

Instances: 1

TABLE 5-548: FIELDS IN WT_TCL

Field Name	Bit	Access	Description	Default
WT_TCL	15:0	R/W	Refresh period length for the weighted-token arbitration scheme.	0xFFFF

5.13.1.6 SBA:SBA:WT_CPU

Parent: [SBA:SBA](#)

Instances: 1

TABLE 5-549: FIELDS IN WT_CPU

Field Name	Bit	Access	Description	Default
WT_CL1	15:0	R/W	Number of tokens to grant the master at the start of each refresh period for weighted-token arbitration scheme. If configured with a value of zero, the master is considered to have infinite tokens.	0x0FFF

5.13.1.7 SBA:SBA:WT_PCIE

Parent: [SBA:SBA](#)

Instances: 1

TABLE 5-550: FIELDS IN WT_PCIE

Field Name	Bit	Access	Description	Default
WT_CL2	15:0	R/W	See SBA::WT_CL1 for description.	0x0FFF

5.13.1.8 SBA:SBA:WT_CSR

Parent: [SBA:SBA](#)

Instances: 1

TABLE 5-551: FIELDS IN WT_CSR

Field Name	Bit	Access	Description	Default
WT_CL3	15:0	R/W	See SBA::WT_CL1 for description.	0x0FFF

5.14 SYS

TABLE 5-552: REGISTER GROUPS IN SYS

Register Group Name	Offset within Target	Instances and Address Spacing	Description	Details
SYSTEM	0x00000518	1	Switch configuration	Page 577
ENCAPSULATIONS	0x00000630	1	MPLS encapsulation table	Page 582
PAUSE_CFG	0x0000065C	1	Watermarks for egress queue system	Page 583
MMGT	0x000006F0	1	Memory manager status	Page 585

TABLE 5-552: REGISTER GROUPS IN SYS (CONTINUED)

Register Group Name	Offset within Target	Instances and Address Spacing	Description	Details
STAT	0x00000000	320 0x00000004	Frame statistics	Page 586
PTP	0x0000070C	1	Precision Time Protocol configuration	Page 587
PTP_TOD	0x00000500	1	Precision Time Protocol time of day	Page 589

5.14.1 SYS:SYSTEM

Parent: [SYS](#)

Instances: 1

TABLE 5-553: REGISTERS IN SYSTEM

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
RESET_CFG	0x00000000	1	Core reset control	Page 577
VLAN_ETYPE_CFG	0x00000008	1	S-tag Ethernet Type	Page 578
PORT_MODE	0x0000000C	13 0x00000004	Per device port configuration	Page 578
FRONT_PORT_MODE	0x00000040	11 0x00000004	Various Ethernet port configurations per front port	Page 580
FRM_AGING	0x0000006C	1	Frame aging configuration	Page 580
STAT_CFG	0x00000070	1	Statistics configuration	Page 580
MPLS_QOS_MAP_CFG	0x000000A4	16 0x00000004	Tables for translating QoS class and DP level to TC values	Page 581
IO_PATH_DELAY	0x000000E4	12 0x00000004	Local path delay compensation per front port	Page 581

5.14.1.1 SYS:SYSTEM:RESET_CFG

Parent: [SYS:SYSTEM](#)

Instances: 1

Controls reset and initialization of the switching core. Proper startup sequence is:

- Enable memories
- Initialize memories
- Enable core

TABLE 5-554: FIELDS IN RESET_CFG

Field Name	Bit	Access	Description	Default
CORE_ENA	7	R/W	Switch core is enabled when this field is set.	0x0
MEM_ENA	6	R/W	Core memory controllers are enabled when this field is set.	0x0

TABLE 5-554: FIELDS IN RESET_CFG (CONTINUED)

Field Name	Bit	Access	Description	Default
MEM_INIT	5	One-shot	Initialize core memories. Field is automatically cleared when operation is complete (approx. 40 us).	0x0
ENCAP_CNT	4:0	R/W	Number of encapsulation records taken out from main packet memory. This field must be set before the switching core is enabled. 0: No records 1: 64 records 2: 96 records 3:... (unit is 32)	0x00

5.14.1.2 SYS:SYSTEM:VLAN_ETYPE_CFG

Parent: [SYS:SYSTEM](#)

Instances: 1

TABLE 5-555: FIELDS IN VLAN_ETYPE_CFG

Field Name	Bit	Access	Description	Default
VLAN_S_TAG_ETYPE_VAL	15:0	R/W	Custom Ethernet Type for S-tags. Tags with TPID = 0x88A8 are always recognized as S-tags.	0x88A8

5.14.1.3 SYS:SYSTEM:PORT_MODE

Parent: [SYS:SYSTEM](#)

Instances: 13

These configurations exists per frontport and for each of the two CPU ports (11+12).

TABLE 5-556: FIELDS IN PORT_MODE

Field Name	Bit	Access	Description	Default
INCL_INJ_HDR	5:4	R/W	Set the mode for the formatting of incoming frames. If INCL_INJ_HDR>0, incoming frames are expected to contain the selected prefix followed by the 16-byte CPU injection header as the first part of the frame. Frames are forwarded based on the contents in the CPU injection header instead of normal forwarding. Three different prefixes are supported: - No prefix. - Short prefix: 0x88800005. - Long prefix: any DMAC, any SMAC, Ether-Type=0x8880, payload=0x0005. In modes 2 and 3, if the incoming frame's format does not comply with the prefix, then sticky bit INJ_HDR_PREFIX_ERR is set. A non-complying frame is discarded. 0: No CPU injection header (normal frames) 1: CPU injection header without prefix 2: CPU injection header with short prefix 3: CPU injection header with long prefix	0x0
INCL_XTR_HDR	3:2	R/W	Set the mode for the formatting of outgoing frames. If INCL_XTR_HDR>0, outgoing frames are prepended the selected prefix followed by the 16-byte CPU extraction header as the first part of the frame. Inserting the CPU extraction header prevents other rewriter operations on the same frame. Three different prefixes are supported: - No prefix. - Short prefix: 0x88800005. - Long prefix: DMAC=0xFFFFFFFFFFFF, SMAC=0xFEFFFFFFF, Ether-Type=0x8880, payload=0x0005. 0: No CPU extraction header (normal frames) 1: CPU extraction header without prefix 2: CPU extraction header with short prefix 3: CPU extraction header with long prefix	0x0
MPLS_ENA	1	R/W	Enable processing of MPLS encapsulations on the port. If set, the MPLS link layer header is processed by the MPLS classifier. Otherwise, the MPLS-encapsulated frames are processed by the basic classifier as normal frames.	0x0

TABLE 5-556: FIELDS IN PORT_MODE (CONTINUED)

Field Name	Bit	Access	Description	Default
INJ_HDR_ERR	0	Sticky	If set, a frame has been received with prefix not complying with the setting in INCL_INJ_HDR.	0x0

5.14.1.4 SYS:SYSTEM:FRONT_PORT_MODE

Parent: [SYS:SYSTEM](#)

Instances: 11

TABLE 5-557: FIELDS IN FRONT_PORT_MODE

Field Name	Bit	Access	Description	Default
HDX_MODE	0	R/W	Enables the queue system to support the half duplex mode. Must be set for a port when enabled for half-duplex mode (MAC_MODE_ENA.FDX_ENA cleared).	0x0

5.14.1.5 SYS:SYSTEM:FRM_AGING

Parent: [SYS:SYSTEM](#)

Instances: 1

TABLE 5-558: FIELDS IN FRM_AGING

Field Name	Bit	Access	Description	Default
AGE_TX_ENA	20	R/W	If set, frames exceeding the frame aging time while waiting head-of-line in a port due to a link partner signaling pause are allowed to be aged. Note that frames aged in this case are not counted by the normal frame aging counter C_TX_AGED.	0x1
MAX_AGE	19:0	R/W	Frames are aged and removed from the queue system when the frame's age timer becomes two. The frame age timer is increased for all frames whenever the configured time, MAX_AGE, has passed. The unit is 6.5us. Effectively, this means that a frame is aged when the frame has waited in the queue system between one or two times the period specified by MAX_AGE. A value of zero disables frame aging.	0x00000

5.14.1.6 SYS:SYSTEM:STAT_CFG

Parent: [SYS:SYSTEM](#)

Instances: 1

TABLE 5-559: FIELDS IN STAT_CFG

Field Name	Bit	Access	Description	Default
STAT_CLEAR_SHOT	14:10	One-shot	Set STAT_CLEAR_SHOT to clear counters in the counter group for the port or service index selected by STAT_VIEW. Auto-cleared when complete (1us). Multiple counter groups can be cleared at the same time by setting multiple bits in STAT_CLEAR_SHOT. Bit 0: Group 0 - Rx counters. Bit 1: Group 1 - Tx counters. Bit 2: Group 2 - Drop counters. Bit 3: Group 3 - Ingress service counters. Bit 4: Group 4 - Egress service counters.	0x00
STAT_VIEW	9:0	R/W	Selects the subset of counters accessible through the SYS:STAT:CNT register. For groups 0-2, STAT_VIEW defines the port to read counters for. For groups 3-4, STAT_VIEW defines the service index to read counters for.	0x000

5.14.1.7 SYS:SYSTEM:MPLS_QOS_MAP_CFG

Parent: [SYS:SYSTEM](#)

Instances: 16

The table to use is given as 8 x DP + TC_SEL. The TC_SEL field comes from the MPLS remarking configuration (SYS::ENCAP_DATA). The tables are looked up for each MPLS label being pushed.

TABLE 5-560: FIELDS IN MPLS_QOS_MAP_CFG

Field Name	Bit	Access	Description	Default
MPLS_TC_VAL	23:0	R/W	The frame's classified QoS class selects the TC bits to use: Bits 2:0: TC bits for QoS class 0 Bits 5:3: TC bits for QoS class 1 Bits 8:6: TC bits for QoS class 2 Bits 11:9: TC bits for QoS class 3 Bits 14:12: TC bits for QoS class 4 Bits 17:15: TC bits for QoS class 5 Bits 20:18: TC bits for QoS class 6 Bits 23:21: TC bits for QoS class 7	0x000000

5.14.1.8 SYS:SYSTEM:IO_PATH_DELAY

Parent: [SYS:SYSTEM](#)

Instances: 12

TABLE 5-561: FIELDS IN IO_PATH_DELAY

Field Name	Bit	Access	Description	Default
RX_PATH_DELAY	23:12	R/W	Delay through the local system receive path, which is added to the frame's receive timestamp. The value is signed. Unit: Nanoseconds	0xF41
TX_PATH_DELAY	11:0	R/W	Delay through the local system transmit path, which is added to the frame's transmit timestamp. The value is signed. Unit: Nanoseconds	0x0B9

5.14.2 SYS:ENCAPSULATIONS

Parent: [SYS](#)

Instances: 1

TABLE 5-562: REGISTERS IN ENCAPSULATIONS

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
ENCAP_CTRL	0x00000000	1	Controls access to the MPLS encapsulation table	Page 582
ENCAP_DATA	0x00000004	10 0x00000004	MPLS encapsulation entry	Page 582

5.14.2.1 SYS:ENCAPSULATIONS:ENCAP_CTRL

Parent: [SYS:ENCAPSULATIONS](#)

Instances: 1

TABLE 5-563: FIELDS IN ENCAP_CTRL

Field Name	Bit	Access	Description	Default
ENCAP_ID	10:1	R/W	Index used to select the MPLS encapsulation entry, which is written when setting ENCAP_WR.	0x000
ENCAP_WR	0	One-shot	If set, the MPLS encapsulation entry in ENCAP_DATA is written to the MPLS encapsulation table. Cleared when operation is complete.	0x0

5.14.2.2 SYS:ENCAPSULATIONS:ENCAP_DATA

Parent: [SYS:ENCAPSULATIONS](#)

Instances: 10

An MPLS encapsulation entry consists of 320 bits with the most significant bits at index 0 and the least significant bits at index 9. The 320 bits are encoded as follows:

Bits 319:272: 48-bit MPLS remarking configuration word.

Bits 271:0: MPLS encapsulation data (MPLS link layer header and MPLS label stack). Data is always pushed starting from the most significant bits of this field. ES0_ACTION.ENCAP_LEN controls how much of the data is pushed.

The MPLS remarking configuration word is encoded as follows:

Bit 298: USE_IVID - If set, the frame's classified VID is used in the VLAN tag

Bit 297: USE_DP - If set, the frame's DP level is used in the VLAN tag as DEI. This overrides USE_IDEI.

Bit 296: USE_IPCP - If set, the frame's classified PCP is used in the VLAN tag

Bit 295: USE_IDEI - If set, the frame's classified DEI is used in the VLAN tag

Bit 294: TAGGED - Set if the MPLS encapsulation data contains a VLAN tag after the SMAC

Bit 293: USE_ITTL_2 - If set, the incoming frame's TTL is used in label 2

Bit 292: USE_ITTL_1 - If set, the incoming frame's TTL is used in label 1

Bit 291: USE_ITTL_0 - If set, the incoming frame's TTL is used in label 0

Bit 290: USE_ISBIT_2 - If set, the incoming frame's S-bit is used in label 2

Bit 289: USE_ISBIT_1 - If set, the incoming frame's S-bit is used in label 1

Bit 288: USE_ISBIT_0 - If set, the incoming frame's S-bit is used in label 0

Bit 283: TC_MAP_2 - If set, the TC bits for label 2 are mapped

Bit 282:280: TC_SEL_2 - Used when mapping label 2's TC bits in MPLS_QOS_MAP_CFG

Bit 279: TC_MAP_1 - If set, the TC bits for label 2 are mapped

Bit 278:276: TC_SEL_1 - Used when mapping label 1's TC bits in MPLS_QOS_MAP_CFG

Bit 275: TC_MAP_0 - If set, the TC bits for label 2 are mapped

Bit 274:272: TC_SEL_0 - Used when mapping label 0's TC bits in MPLS_QOS_MAP_CFG

TABLE 5-564: FIELDS IN ENCAP_DATA

Field Name	Bit	Access	Description	Default
ENCAP_DATA	31:0	R/W	32 bits of encapsulation entry data.	0x00000000

5.14.3 SYS:PAUSE_CFG

Parent: [SYS](#)

Instances: 1

TABLE 5-565: REGISTERS IN PAUSE_CFG

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
PAUSE_CFG	0x00000000	12 0x00000004	Watermarks for flow control condition per switch port.	Page 584
PAUSE_TOT_CFG	0x00000030	1	Configure total memory pause condition	Page 584
ATOP	0x00000034	12 0x00000004	Tail dropping level	Page 584
ATOP_TOT_CFG	0x00000064	1	Total raw memory use before tail dropping is activated	Page 584
MAC_FC_CFG	0x00000068	11 0x00000004	MAC Flow Control Configuration Register	Page 585

VSC7416-01

5.14.3.1 SYS:PAUSE_CFG:PAUSE_CFG

Parent: [SYS:PAUSE_CFG](#)

Instances: 12

TABLE 5-566: FIELDS IN PAUSE_CFG

Field Name	Bit	Access	Description	Default
PAUSE_START	24:13	R/W	Start pausing ingress stream when the amount of memory consumed by the port exceeds this watermark. The TOTPAUSE condition must also be met. See RES_CFG	0xFFF
PAUSE_STOP	12:1	R/W	Stop pausing ingress stream when the amount of memory consumed by the port is below this watermark. See RES_CFG.	0xFFF
PAUSE_ENA	0	R/W	Enable pause feedback to the MAC, allowing transmission of pause frames or HDX collisions to limit ingress data rate.	0x0

5.14.3.2 SYS:PAUSE_CFG:PAUSE_TOT_CFG

Parent: [SYS:PAUSE_CFG](#)

Instances: 1

TABLE 5-567: FIELDS IN PAUSE_TOT_CFG

Field Name	Bit	Access	Description	Default
PAUSE_TOT_START	23:12	R/W	Assert TOTPAUSE condition when total memory allocation is above this watermark. See RES_CFG	0x000
PAUSE_TOT_STOP	11:0	R/W	Deassert TOTPAUSE condition when total memory allocation is below this watermark. See RES_CFG	0x000

5.14.3.3 SYS:PAUSE_CFG:ATOP

Parent: [SYS:PAUSE_CFG](#)

Instances: 12

TABLE 5-568: FIELDS IN ATOP

Field Name	Bit	Access	Description	Default
ATOP	11:0	R/W	When a source port consumes more than this level in the packet memory, frames are tail dropped, unconditionally of destination. See RES_CFG	0xFFF

5.14.3.4 SYS:PAUSE_CFG:ATOP_TOT_CFG

Parent: [SYS:PAUSE_CFG](#)

Instances: 1

TABLE 5-569: FIELDS IN ATOP_TOT_CFG

Field Name	Bit	Access	Description	Default
ATOP_TOT	11:0	R/W	Tail dropping is activate on a port when the port use has exceeded the ATOP watermark for the port, and the total memory use has exceeded this watermark. See RES_CFG	0xFFFF

5.14.3.5 SYS:PAUSE_CFG:MAC_FC_CFG

Parent: [SYS:PAUSE_CFG](#)

Instances: 11

TABLE 5-570: FIELDS IN MAC_FC_CFG

Field Name	Bit	Access	Description	Default
FC_LINK_SPEED	27:26	R/W	Configures the link speed. This is used to evaluate the time specification in incoming pause frames. 0: 2500 Mbps 1: 1000 Mbps 2: 100 Mbps 3: 10 Mbps	0x1
FC_LATENCY_CFG	25:20	R/W	Accepted reaction time for link partner after the port has transmitted a pause frame. Frames starting after this latency are aborted. Unit is 64 byte times.	0x07
ZERO_PAUSE_ENA	18	R/W	If set, a zero-delay pause frame is transmitted when flow control is deasserted. '0': Don't send zero pause frame. '1': Send zero pause frame.	0x0
TX_FC_ENA	17	R/W	When set the MAC will send pause control frames in the Tx direction. '0': Don't send pause control frames '1': Send pause control frames	0x0
RX_FC_ENA	16	R/W	When set the MAC obeys received pause control frames '0': Don't obey received pause control frames '1': Obey received pause control frames.	0x0
PAUSE_VAL_CFG	15:0	R/W	Pause timer value inserted in generated pause frames. 0: Insert timer value 0 in TX pause frame. ... N: Insert timer value N in TX pause frame.	0x0000

5.14.4 SYS:MMGT

Parent: [SYS](#)

Instances: 1

TABLE 5-571: REGISTERS IN MMGT

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
MMGT	0x00000000	1	Packet memory status	Page 586

5.14.4.1 SYS:MMGT:MMGT

Parent: [SYS:MMGT](#)

Instances: 1

TABLE 5-572: FIELDS IN MMGT

Field Name	Bit	Access	Description	Default
FREECNT	15:0	R/O	Number of 192-byte free memory words.	0x0000

5.14.5 SYS:STAT

Parent: [SYS](#)

Instances: 320

TABLE 5-573: REGISTERS IN STAT

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
CNT	0x00000000	1	Counter values	Page 586

5.14.5.1 SYS:STAT:CNT

Parent: [SYS:STAT](#)

Instances: 1

TABLE 5-574: FIELDS IN CNT

Field Name	Bit	Access	Description	Default
CNT	31:0	R/W	Counter values. The counters are layed in five counter groups: Group 0: 0-63: Receive counters for port specified in STAT_CFG.STAT_VIEW. Group 1: 64-127: Transmit counters for port specified in STAT_CFG.STAT_VIEW. Group 2: 128-191: Drop counters for port specified in STAT_CFG.STAT_VIEW. Group 3: 192-255: Ingress Service counters for ingress service point specified in STAT_CFG.STAT_VIEW. Group 4: 256-319: Egress Service counters for egress service point specified in STAT_CFG.STAT_VIEW. QSYS::STAT_CNT_CFG and ANA::AGENC-TRL control whether bytes or frames are counted for specific counters. Counters are cleared through STAT_CFG.STAT_CLEAR_SHOT.	0x00000000

5.14.6 SYS:PTP

Parent: [SYS](#)

Instances: 1

TABLE 5-575: REGISTERS IN PTP

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
PTP_STATUS	0x00000000	1	Stored timestamp and time-stamp queue status	Page 587
PTP_TXSTAMP	0x00000004	1	Timestamp value	Page 588
PTP_NXT	0x00000008	1	Advancing the timestamp queue	Page 588
PTP_CFG	0x0000000C	1	System configuration of PTP	Page 588

5.14.6.1 SYS:PTP:PTP_STATUS

Parent: [SYS:PTP](#)

Instances: 1

TABLE 5-576: FIELDS IN PTP_STATUS

Field Name	Bit	Access	Description	Default
PTP_TXSTAMP_OAM	29	R/O	If set, the timestamp is generated on request by the Versatile OAM engine.	0x0

TABLE 5-576: FIELDS IN PTP_STATUS (CONTINUED)

Field Name	Bit	Access	Description	Default
PTP_OVFL	28	R/O	If set, the timestamp queue has overflowed implying a timestamp entry could not be enqueued. The PTP_OVFL bit is not cleared until the timestamp queue is completely empty.	0x0
PTP_MESS_VLD	27	R/O	A timestamp entry is ready for reading. PTP_MESS_ID, PTP_MESS_TXPORT, and PTP_DELAY contain the data of the timestamp entry.	0x0
PTP_MESS_ID	26:21	R/O	Timestamp identifier for head-of-line timestamp entry.	0x00
PTP_MESS_TXPORT	20:16	R/O	The transmit port for the head-of-line timestamp entry.	0x00
PTP_MESS_SEQ_ID	15:0	R/O	Timestamp sequence identifier for head-of-line timestamp entry.	0x0000

5.14.6.2 SYS:PTP:PTP_TXSTAMP

Parent: [SYS:PTP](#)

Instances: 1

TABLE 5-577: FIELDS IN PTP_TXSTAMP

Field Name	Bit	Access	Description	Default
PTP_TXSTAMP	29:0	R/O	The timestamp value for the head-of-line timestamp entry. The timestamp value is the frame's departure time plus IO_PATH_DELAY.TX_PATH_DELAY. Unit is ns.	0x00000000
PTP_TXSTAMP_SEC	31	R/O	The timestamp value for the head-of-line timestamp entry. The timestamp value is the frame's departure time plus IO_PATH_DELAY.TX_PATH_DELAY. This is the LSB of the TOD_SEC value at time of event.	0x0

5.14.6.3 SYS:PTP:PTP_NXT

Parent: [SYS:PTP](#)

Instances: 1

TABLE 5-578: FIELDS IN PTP_NXT

Field Name	Bit	Access	Description	Default
PTP_NXT	0	One-shot	Advance to the next timestamp entry. Registers PTP_STATUS and PTP_DELAY points to the next entry.	0x0

5.14.6.4 SYS:PTP:PTP_CFG

Parent: [SYS:PTP](#)

Instances: 1

TABLE 5-579: FIELDS IN PTP_CFG

Field Name	Bit	Access	Description	Default
PTP_STAMP_WID	7:2	R/W	Determines how many LSbits from the time-stamp value are used by the timestamp transfer function.	0x20
PTP_CF_ROLL_MODE	1:0	R/W	Configures how rollover protection is done when doing add/subtract transfer. 0: The four LSbits of CF sub-nano is filled out with the MSbits of the timestamp. 1: The LSbit of the CF sub-nano field is set to the MSbit of the timestamp. 2: The CF bit 62 is set to the MSbit of the timestamp. 3: Reserved.	0x0

5.14.7 SYS:PTP_TOD

Parent: [SYS](#)

Instances: 1

TABLE 5-580: REGISTERS IN PTP_TOD

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
PTP_TOD_MSB	0x00000000	1	Most significant bits of TOD seconds	Page 589
PTP_TOD_LSB	0x00000004	1	Least significant bits of TOD seconds	Page 589
PTP_TOD_NSEC	0x00000008	1	TOD nanoseconds	Page 590

5.14.7.1 SYS:PTP_TOD:PTP_TOD_MSB

Parent: [SYS:PTP_TOD](#)

Instances: 1

TABLE 5-581: FIELDS IN PTP_TOD_MSB

Field Name	Bit	Access	Description	Default
PTP_TOD_MSB	15:0	R/W	Bits 47:32 in 48-bit time of day. Unit is seconds.	0x0000

5.14.7.2 SYS:PTP_TOD:PTP_TOD_LSB

Parent: [SYS:PTP_TOD](#)

Instances: 1

TABLE 5-582: FIELDS IN PTP_TOD_LSB

Field Name	Bit	Access	Description	Default
PTP_TOD_LSB	31:0	R/W	Bits 31:0 in 48-bit time of day. Unit is seconds.	0x00000000

5.14.7.3 SYS:PTP_TOD:PTP_TOD_NSEC

Parent: [SYS:PTP_TOD](#)

Instances: 1

TABLE 5-583: FIELDS IN PTP_TOD_NSEC

Field Name	Bit	Access	Description	Default
PTP_TOD_NSEC	29:0	R/O	Nanoseconds controlled by the PTP one-second timer.	0x00000000

5.15 TWI

TABLE 5-584: REGISTER GROUPS IN TWI

Register Group Name	Offset within Target	Instances and Address Spacing	Description	Details
TWI	0x00000000	1	Two-Wire Interface controller	Page 590

5.15.1 TWI:TWI

Parent: [TWI](#)

Instances: 1

TABLE 5-585: REGISTERS IN TWI

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
CFG	0x00000000	1	TWI configuration	Page 591
TAR	0x00000004	1	Target address	Page 592
SAR	0x00000008	1	Slave address	Page 593
DATA_CMD	0x00000010	1	Rx/Tx data buffer and command	Page 593
SS_SCL_HCNT	0x00000014	1	Standard speed TWI clock SCL high count	Page 594
SS_SCL_LCNT	0x00000018	1	Standard speed TWI clock SCL low count	Page 595
FS_SCL_HCNT	0x0000001C	1	Fast speed TWI clock SCL high count	Page 595
FS_SCL_LCNT	0x00000020	1	Fast speed TWI clock SCL low count	Page 595
INTR_STAT	0x0000002C	1	Interrupt status	Page 596
INTR_MASK	0x00000030	1	Interrupt mask	Page 596
RAW_INTR_STAT	0x00000034	1	Raw interrupt status	Page 597
RX_TL	0x00000038	1	Receive FIFO threshold	Page 599
TX_TL	0x0000003C	1	Transmit FIFO threshold	Page 600
CLR_INTR	0x00000040	1	Clear combined and individual interrupt	Page 600
CLR_RX_UNDER	0x00000044	1	Clear RX_UNDER interrupt	Page 600

TABLE 5-585: REGISTERS IN TWI (CONTINUED)

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
CLR_RX_OVER	0x00000048	1	Clear RX_OVER interrupt	Page 601
CLR_TX_OVER	0x0000004C	1	Clear TX_OVER interrupt	Page 601
CLR_RD_REQ	0x00000050	1	Clear RD_REQ interrupt	Page 601
CLR_TX_ABRT	0x00000054	1	Clear TX_ABRT interrupt	Page 601
CLR_RX_DONE	0x00000058	1	Clear RX_DONE interrupt	Page 601
CLR_ACTIVITY	0x0000005C	1	Clear ACTIVITY interrupt	Page 602
CLR_STOP_DET	0x00000060	1	Clear STOP_DET interrupt	Page 602
CLR_START_DET	0x00000064	1	Clear START_DET interrupt	Page 602
CLR_GEN_CALL	0x00000068	1	Clear GEN_CALL interrupt	Page 602
CTRL	0x0000006C	1	TWI control	Page 603
STAT	0x00000070	1	TWI status	Page 603
TXFLR	0x00000074	1	Transmit FIFO level	Page 604
RXFLR	0x00000078	1	Receive FIFO level	Page 604
TX_ABRT_SOURCE	0x00000080	1	Transmit abort source	Page 605
SDA_SETUP	0x00000094	1	SDA setup	Page 606
ACK_GEN_CALL	0x00000098	1	Acknowledge general call	Page 607
ENABLE_STATUS	0x0000009C	1	Enable status	Page 607

5.15.1.1 TWI:TWI:CFG

Parent: [TWI:TWI](#)

Instances: 1

TABLE 5-586: FIELDS IN CFG

Field Name	Bit	Access	Description	Default
SLAVE_DIS	6	R/W	This bit controls whether the TWI controller has its slave disabled. If this bit is set (slave is disabled), the controller functions only as a master and does not perform any action that requires a slave. '0': slave is enabled '1': slave is disabled	0x1

TABLE 5-586: FIELDS IN CFG (CONTINUED)

Field Name	Bit	Access	Description	Default
RESTART_ENA	5	R/W	Determines whether RESTART conditions may be sent when acting as a master. Some older slaves do not support handling RESTART conditions; however, RESTART conditions are used in several operations. When RESTART is disabled, the master is prohibited from performing the following functions: * Change direction within a transfer (split) * Send a START BYTE * Combined format transfers in 7-bit addressing modes * Read operation with a 10-bit address * Send multiple bytes per transfer By replacing RESTART condition followed by a STOP and a subsequent START condition, split operations are broken down into multiple transfers. If the above operations are performed, it will result in setting RAW_INTR_STAT.TX_ABRT. '0': disable '1': enable	0x1
MASTER_10BITADDR	4	R/W	Controls whether transfers starts in 7- or 10-bit addressing mode when acting as a master. '0': 7-bit addressing '1': 10-bit addressing	0x0
SLAVE_10BITADDR	3	R/W	Controls whether the the TWI controller responds to 7- or 10-bit addresses in slave mode. In 7-bit mode; transactions that involve 10-bit addressing are ignored and only the lower 7 bits of the SAR register are compared. '0': 7-bit addressing. '1': 10-bit addressing.	0x0
SPEED	2:1	R/W	These bits control at which speed the TWI controller operates; its setting is relevant only in master mode. Hardware protects against illegal values being programmed by software. '1': standard mode (100 kbit/s) '2': fast mode (400 kbit/s)	0x2
MASTER_ENA	0	R/W	This bit controls whether the TWI master is enabled. '0': master disabled '1': master enabled	0x1

5.15.1.2 TWI:TWI:TAR

Parent: [TWI:TWI](#)

Instances: 1

TABLE 5-587: FIELDS IN TAR

Field Name	Bit	Access	Description	Default
GC_OR_START_ENA	11	R/W	This bit indicates whether software performs a General Call or START BYTE command. '0': ignore bit 10 GC_OR_START and use TAR normally '1': perform special TWI command as specified in GC_OR_START bit	0x0
GC_OR_START	10	R/W	If TAR.SPECIAL is set to 1, then this bit indicates whether a General Call or START byte command is to be performed. '0': General Call Address - after issuing a General Call, only writes may be performed. Attempting to issue a read command results in setting RAW_INTR_STAT.TX_ABRT. The TWI controller remains in General Call mode until the TAR.SPECIAL field is cleared. '1': START BYTE	0x0
TAR	9:0	R/W	This is the target address for any master transaction. When transmitting a General Call, these bits are ignored. To generate a START BYTE, the CPU needs to write only once into these bits. If the TAR and SAR are the same, loopback exists but the FIFOs are shared between master and slave, so full loopback is not feasible. Only one direction loopback mode is supported (simplex), not duplex. A master cannot transmit to itself; it can transmit to only a slave.	0x055

5.15.1.3 TWI:TWI:SAR

Parent: [TWI:TWI](#)

Instances: 1

TABLE 5-588: FIELDS IN SAR

Field Name	Bit	Access	Description	Default
SAR	9:0	R/W	The SAR holds the slave address when the TWI is operating as a slave. For 7-bit addressing, only SAR[6:0] is used. This register can be written only when the TWI interface is disabled (ENABLE = 0).	0x055

5.15.1.4 TWI:TWI:DATA_CMD

Parent: [TWI:TWI](#)

Instances: 1

TABLE 5-589: FIELDS IN DATA_CMD

Field Name	Bit	Access	Description	Default
CMD	8	R/W	This bit controls whether a read or a write is performed. This bit does not control the direction when the TWI acts as a slave. It controls only the direction when it acts as a master. When a command is entered in the TX FIFO, this bit distinguishes the write and read commands. In slave-receiver mode, this bit is a "don't care" because writes to this register are not required. In slave-transmitter mode, a "0" indicates that CPU data is to be transmitted and as DATA. When programming this bit, remember the following: attempting to perform a read operation after a General Call command has been sent results in a TX_ABRT interrupt (RAW_INTR_STAT.R_TX_ABRT), unless TAR.SPECIAL has been cleared. If a "1" is written to this bit after receiving a RD_REQ interrupt, then a TX_ABRT interrupt occurs. NOTE: It is possible that while attempting a master TWI read transfer, a RD_REQ interrupt may have occurred simultaneously due to a remote TWI master addressing this controller. In this type of scenario, the TWI controller ignores the DATA_CMD write, generates a TX_ABRT interrupt, and waits to service the RD_REQ interrupt. '1' = Read '0' = Write	0x0
DATA	7:0	R/W	This register contains the data to be transmitted or received on the TWI bus. If you are writing to this register and want to perform a read, this field is ignored by the controller. However, when you read this register, these bits return the value of data received on the TWI interface.	0x00

5.15.1.5 TWI:TWI:SS_SCL_HCNT

Parent: [TWI:TWI](#)

Instances: 1

The clock for the TWI controller is the VCore system clock. This field must be set accordingly to the VCore system frequency; value = (4us / VCore clock period) - 8.

Example: a 178.6MHz clock correspond to a period of 5.6ns, for this frequency this field must not be set lower than (round up): $707 = (4\mu s / 5.6ns) - 8$.

TABLE 5-590: FIELDS IN SS_SCL_HCNT

Field Name	Bit	Access	Description	Default
SS_SCL_HCNT	15:0	R/W	This register sets the SCL clock divider for the high-period in standard speed. This value must result in a high period of no less than 4us.	0x033A

5.15.1.6 TWI:TWI:SS_SCL_LCNT

Parent: [TWI:TWI](#)

Instances: 1

The clock for the TWI controller is the VCore system clock. This field must be set accordingly to the VCore system frequency; value = $(4.7\mu\text{s} / \text{VCore clock period}) - 1$.

Example: a 178.6MHz clock correspond to a period of 5.6ns, for this frequency this field must not be set lower than (round up): $839 = (4.7\mu\text{s} / 5.6\text{ns}) - 1$.

TABLE 5-591: FIELDS IN SS_SCL_LCNT

Field Name	Bit	Access	Description	Default
SS_SCL_LCNT	15:0	R/W	This register sets the SCL clock divider for the low-period in standard speed. This value must result in a value no less than 4.7us.	0x03D3

5.15.1.7 TWI:TWI:FS_SCL_HCNT

Parent: [TWI:TWI](#)

Instances: 1

The clock for the TWI controller is the VCore system clock. This field must be set accordingly to the VCore system frequency; value = $(0.6\mu\text{s} / \text{VCore clock period}) - 8$.

Example: a 178.6MHz clock correspond to a period of 5.6ns, for this frequency this field must not be set lower than (round up): $100 = (0.6\mu\text{s} / 5.6\text{ns}) - 8$.

TABLE 5-592: FIELDS IN FS_SCL_HCNT

Field Name	Bit	Access	Description	Default
FS_SCL_HCNT	15:0	R/W	This register sets the SCL clock divider for the high-period in fast speed. This value must result in a value no less than 0.6us.	0x0075

5.15.1.8 TWI:TWI:FS_SCL_LCNT

Parent: [TWI:TWI](#)

Instances: 1

The clock for the TWI controller is the VCore system clock. This field must be set accordingly to the VCore system frequency; value = $(1.3\mu\text{s} / \text{VCore clock period}) - 1$.

Example: a 178.6MHz clock correspond to a period of 5.6ns, for this frequency this field must not be set lower than (round up): $232 = (1.3\mu\text{s} / 5.6\text{ns}) - 1$.

TABLE 5-593: FIELDS IN FS_SCL_LCNT

Field Name	Bit	Access	Description	Default
FS_SCL_LCNT	15:0	R/W	This register sets the SCL clock divider for the low-period in fast speed. This value must result in a value no less than 1.3us.	0x010E

5.15.1.9 TWI:TWI:INTR_STAT

Parent: [TWI:TWI](#)

Instances: 1

Each field in this register has a corresponding mask field in the INTR_MASK register. These fields are cleared by reading the matching interrupt clear register. The unmasked raw versions of these fields are available in the RAW_INTR_STAT register.

See RAW_INTR_STAT for a description of these fields

TABLE 5-594: FIELDS IN INTR_STAT

Field Name	Bit	Access	Description	Default
GEN_CALL	11	R/O		0x0
START_DET	10	R/O		0x0
STOP_DET	9	R/O		0x0
ACTIVITY	8	R/O		0x0
RX_DONE	7	R/O		0x0
TX_ABRT	6	R/O		0x0
RD_REQ	5	R/O		0x0
TX_EMPTY	4	R/O		0x0
TX_OVER	3	R/O		0x0
RX_FULL	2	R/O		0x0
RX_OVER	1	R/O		0x0
RX_UNDER	0	R/O		0x0

5.15.1.10 TWI:TWI:INTR_MASK

Parent: [TWI:TWI](#)

Instances: 1

These fields mask the corresponding interrupt status fields (RAW_INTR_STAT). They are active high; a value of 0 prevents the corresponding field in RAW_INTR_STAT from generating an interrupt.

TABLE 5-595: FIELDS IN INTR_MASK

Field Name	Bit	Access	Description	Default
M_GEN_CALL	11	R/W		0x1
M_START_DET	10	R/W		0x0
M_STOP_DET	9	R/W		0x0
M_ACTIVITY	8	R/W		0x0
M_RX_DONE	7	R/W		0x1
M_TX_ABRT	6	R/W		0x1
M_RD_REQ	5	R/W		0x1

TABLE 5-595: FIELDS IN INTR_MASK (CONTINUED)

Field Name	Bit	Access	Description	Default
M_TX_EMPTY	4	R/W		0x1
M_TX_OVER	3	R/W		0x1
M_RX_FULL	2	R/W		0x1
M_RX_OVER	1	R/W		0x1
M_RX_UNDER	0	R/W		0x1

5.15.1.11 TWI:TWI:RAW_INTR_STAT

Parent: [TWI:TWI](#)

Instances: 1

Unlike the INTR_STAT register, these fields are not masked so they always show the true status of the TWI controller.

TABLE 5-596: FIELDS IN RAW_INTR_STAT

Field Name	Bit	Access	Description	Default
R_GEN_CALL	11	R/O	Set only when a General Call address is received and it is acknowledged. It stays set until it is cleared either by disabling TWI controller or when the CPU reads bit 0 of the CLR_GEN_CALL register. The TWI controller stores the received data in the Rx buffer.	0x0
R_START_DET	10	R/O	Indicates whether a START or RESTART condition has occurred on the TWI regardless of whether the TWI controller is operating in slave or master mode.	0x0
R_STOP_DET	9	R/O	Indicates whether a STOP condition has occurred on the TWI controller regardless of whether the TWI controller is operating in slave or master mode.	0x0
R_ACTIVITY	8	R/O	This bit captures TWI activity and stays set until it is cleared. There are four ways to clear it: * Disabling the TWI controller * Reading the CLR_ACTIVITY register * Reading the CLR_INTR register * VCore system reset Once this bit is set, it stays set unless one of the four methods is used to clear it. Even if the TWI controller module is idle, this bit remains set until cleared, indicating that there was activity on the bus.	0x0
R_RX_DONE	7	R/O	When the TWI controller is acting as a slave-transmitter, this bit is set to 1 if the master does not acknowledge a transmitted byte. This occurs on the last byte of the transmission, indicating that the transmission is done.	0x0

TABLE 5-596: FIELDS IN RAW_INTR_STAT (CONTINUED)

Field Name	Bit	Access	Description	Default
R_TX_ABRT	6	R/O	This bit is set to 1 when the TWI controller is acting as a master is unable to complete a command that the processor has sent. The conditions that set this field are: * No slave acknowledges the address byte. * The addressed slave receiver does not acknowledge a byte of data. * Attempting to send a master command when configured only to be a slave. * When CFG.RESTART_ENA is set to 0 (RESTART condition disabled), and the processor attempts to issue a TWI function that is impossible to perform without using RESTART conditions. * High-speed master code is acknowledged (this controller does not support high-speed). * START BYTE is acknowledged. * General Call address is not acknowledged. * When a read request interrupt occurs and the processor has previously placed data in the Tx buffer that has not been transmitted yet. This data could have been intended to service a multi-byte RD_REQ that ended up having fewer numbers of bytes requested. * The TWI controller loses arbitration of the bus between transfers and is then accessed as a slave-transmitter. * If a read command is issued after a General Call command has been issued. Disabling the TWI reverts it back to normal operation. * If the CPU attempts to issue read command before a RD_REQ is serviced. Anytime this bit is set, the contents of the transmit and receive buffers are flushed.	0x0
R_RD_REQ	5	R/O	This bit is set to 1 when the TWI controller acts as a slave and another TWI master is attempting to read data from this controller. The TWI controller holds the TWI bus in a wait state (SCL=0) until this interrupt is serviced, which means that the slave has been addressed by a remote master that is asking for data to be transferred. The processor must respond to this interrupt and then write the requested data to the DATA_CMD register. This bit is set to 0 just after the required data is written to the DATA_CMD register.	0x0

TABLE 5-596: FIELDS IN RAW_INTR_STAT (CONTINUED)

Field Name	Bit	Access	Description	Default
R_TX_EMPTY	4	R/O	This bit is set to 1 when the transmit buffer is at or below the threshold value set in the TX_TL register. It is automatically cleared by hardware when the buffer level goes above the threshold. When ENABLE is 0, the TX FIFO is flushed and held in reset. There the TX FIFO looks like it has no data within it, so this bit is set to 1, provided there is activity in the master or slave state machines. When there is no longer activity, then with ENABLE_STATUS.BUSY=0, this bit is set to 0.	0x0
R_TX_OVER	3	R/O	Set during transmit if the transmit buffer is filled to TX_BUFFER_DEPTH and the processor attempts to issue another TWI command by writing to the DATA_CMD register. When the module is disabled, this bit keeps its level until the master or slave state machines go into idle, and when ENABLE_STATUS.BUSY goes to 0, this interrupt is cleared.	0x0
R_RX_FULL	2	R/O	Set when the receive buffer reaches or goes above the RX_TL threshold in the RX_TL register. It is automatically cleared by hardware when buffer level goes below the threshold. If the module is disabled (ENABLE=0), the RX FIFO is flushed and held in reset; therefore the RX FIFO is not full. So this bit is cleared once the ENABLE field is programmed with a 0, regardless of the activity that continues.	0x0
R_RX_OVER	1	R/O	Set if the receive buffer is completely filled to RX_BUFFER_DEPTH and an additional byte is received from an external TWI device. The TWI controller acknowledges this, but any data bytes received after the FIFO is full are lost. If the module is disabled (ENABLE=0), this bit keeps its level until the master or slave state machines go into idle, and when ENABLE_STATUS.BUSY goes to 0, this interrupt is cleared.	0x0
R_RX_UNDER	0	R/O	Set if the processor attempts to read the receive buffer when it is empty by reading from the DATA_CMD register. If the module is disabled (ENABLE=0), this bit keeps its level until the master or slave state machines go into idle, and when ENABLE_STATUS.BUSY goes to 0, this interrupt is cleared.	0x0

5.15.1.12 TWI:TWI:RX_TL

Parent: [TWI:TWI](#)

Instances: 1

TABLE 5-597: FIELDS IN RX_TL

Field Name	Bit	Access	Description	Default
RX_TL	2:0	R/W	Controls the level of entries (or above) that triggers the RX_FULL interrupt (bit 2 in RAW_INTR_STAT register). The valid range is 0-7. A value of 0 sets the threshold for 1 entry, and a value of 7 sets the threshold for 8 entries.	0x0

5.15.1.13 TWI:TWI:TX_TL

Parent: [TWI:TWI](#)

Instances: 1

TABLE 5-598: FIELDS IN TX_TL

Field Name	Bit	Access	Description	Default
TX_TL	2:0	R/W	Controls the level of entries (or below) that trigger the TX_EMPTY interrupt (bit 4 in RAW_INTR_STAT register). The valid range is 0-7. A value of 0 sets the threshold for 0 entries, and a value of 7 sets the threshold for 7 entries.	0x0

5.15.1.14 TWI:TWI:CLR_INTR

Parent: [TWI:TWI](#)

Instances: 1

TABLE 5-599: FIELDS IN CLR_INTR

Field Name	Bit	Access	Description	Default
CLR_INTR	0	R/O	Read this register to clear the combined interrupt, all individual interrupts, and the TX_ABRT_SOURCE register. This bit does not clear hardware clearable interrupts but software clearable interrupts. Refer to Bit 9 of the TX_ABRT_SOURCE register for an exception to clearing TX_ABRT_SOURCE.	0x0

5.15.1.15 TWI:TWI:CLR_RX_UNDER

Parent: [TWI:TWI](#)

Instances: 1

TABLE 5-600: FIELDS IN CLR_RX_UNDER

Field Name	Bit	Access	Description	Default
CLR_RX_UNDER	0	R/O	Read this register to clear the R_RX_UNDER interrupt (bit 0) of the RAW_INTR_STAT register.	0x0

5.15.1.16 TWI:TWI:CLR_RX_OVER

Parent: [TWI:TWI](#)

Instances: 1

TABLE 5-601: FIELDS IN CLR_RX_OVER

Field Name	Bit	Access	Description	Default
CLR_RX_OVER	0	R/O	Read this register to clear the R_RX_OVER interrupt (bit 1) of the RAW_INTR_STAT register.	0x0

5.15.1.17 TWI:TWI:CLR_TX_OVER

Parent: [TWI:TWI](#)

Instances: 1

TABLE 5-602: FIELDS IN CLR_TX_OVER

Field Name	Bit	Access	Description	Default
CLR_TX_OVER	0	R/O	Read this register to clear the R_TX_OVER interrupt (bit 3) of the RAW_INTR_STAT register.	0x0

5.15.1.18 TWI:TWI:CLR_RD_REQ

Parent: [TWI:TWI](#)

Instances: 1

TABLE 5-603: FIELDS IN CLR_RD_REQ

Field Name	Bit	Access	Description	Default
CLR_RD_REQ	0	R/O	Read this register to clear the R_RD_REQ interrupt (bit 5) of the RAW_INTR_STAT register.	0x0

5.15.1.19 TWI:TWI:CLR_TX_ABRT

Parent: [TWI:TWI](#)

Instances: 1

TABLE 5-604: FIELDS IN CLR_TX_ABRT

Field Name	Bit	Access	Description	Default
CLR_TX_ABRT	0	R/O	Read this register to clear the R_TX_ABRT interrupt (bit 6) of the RAW_INTR_STAT register, and the TX_ABRT_SOURCE register. Refer to Bit 9 of the TX_ABRT_SOURCE register for an exception to clearing TX_ABRT_SOURCE.	0x0

5.15.1.20 TWI:TWI:CLR_RX_DONE

Parent: [TWI:TWI](#)

Instances: 1

TABLE 5-605: FIELDS IN CLR_RX_DONE

Field Name	Bit	Access	Description	Default
CLR_RX_DONE	0	R/O	Read this register to clear the R_RX_DONE interrupt (bit 7) of the RAW_INTR_STAT register.	0x0

5.15.1.21 TWI:TWI:CLR_ACTIVITY

Parent: [TWI:TWI](#)

Instances: 1

TABLE 5-606: FIELDS IN CLR_ACTIVITY

Field Name	Bit	Access	Description	Default
CLR_ACTIVITY	0	R/O	Reading this register clears the ACTIVITY interrupt if the TWI controller is not active anymore. If the TWI controller is still active on the bus, the ACTIVITY interrupt bit continues to be set. It is automatically cleared by hardware if the module is disabled and if there is no further activity on the bus. The value read from this register to get status of the R_ACTIVITY interrupt (bit 8) of the RAW_INTR_STAT register.	0x0

5.15.1.22 TWI:TWI:CLR_STOP_DET

Parent: [TWI:TWI](#)

Instances: 1

TABLE 5-607: FIELDS IN CLR_STOP_DET

Field Name	Bit	Access	Description	Default
CLR_STOP_DET	0	R/O	Read this register to clear the R_STOP_DET interrupt (bit 9) of the RAW_INTR_STAT register.	0x0

5.15.1.23 TWI:TWI:CLR_START_DET

Parent: [TWI:TWI](#)

Instances: 1

TABLE 5-608: FIELDS IN CLR_START_DET

Field Name	Bit	Access	Description	Default
CLR_START_DET	0	R/O	Read this register to clear the R_START_DET interrupt (bit 10) of the RAW_INTR_STAT register.	0x0

5.15.1.24 TWI:TWI:CLR_GEN_CALL

Parent: [TWI:TWI](#)

Instances: 1

TABLE 5-609: FIELDS IN CLR_GEN_CALL

Field Name	Bit	Access	Description	Default
CLR_GEN_CALL	0	R/O	Read this register to clear the R_GEN_CALL interrupt (bit 11) of RAW_INTR_STAT register.	0x0

5.15.1.25 TWI:TWI:CTRL

Parent: [TWI:TWI](#)

Instances: 1

TABLE 5-610: FIELDS IN CTRL

Field Name	Bit	Access	Description	Default
ENABLE	0	R/W	Controls whether the TWI controller is enabled. Software can disable the controller while it is active. However, it is important that care be taken to ensure that the controller is disabled properly. When TWI controller is disabled, the following occurs: * The TX FIFO and RX FIFO get flushed. * The interrupt bits in the RAW_INTR_STAT register are cleared. * Status bits in the INTR_STAT register are still active until the TWI controller goes into IDLE state. If the module is transmitting, it stops as well as deletes the contents of the transmit buffer after the current transfer is complete. If the module is receiving, the controller stops the current transfer at the end of the current byte and does not acknowledge the transfer. '0': Disables TWI controller '1': Enables TWI controller	0x0

5.15.1.26 TWI:TWI:STAT

Parent: [TWI:TWI](#)

Instances: 1

TABLE 5-611: FIELDS IN STAT

Field Name	Bit	Access	Description	Default
SLV_ACTIVITY	6	R/O	Slave FSM Activity Status. When the Slave Finite State Machine (FSM) is not in the IDLE state, this bit is set. '0': Slave FSM is in IDLE state so the Slave part of the controller is not Active '1': Slave FSM is not in IDLE state so the Slave part of the controller is Active	0x0

TABLE 5-611: FIELDS IN STAT (CONTINUED)

Field Name	Bit	Access	Description	Default
MST_ACTIVITY	5	R/O	Master FSM Activity Status. When the Master Finite State Machine (FSM) is not in the IDLE state, this bit is set. '0': Master FSM is in IDLE state so the Master part of the controller is not Active '1': Master FSM is not in IDLE state so the Master part of the controller is Active	0x0
RFF	4	R/O	Receive FIFO Completely Full. When the receive FIFO is completely full, this bit is set. When the receive FIFO contains one or more empty location, this bit is cleared. '0': Receive FIFO is not full '1': Receive FIFO is full	0x0
RFNE	3	R/O	Receive FIFO Not Empty. Set when the receive FIFO contains one or more entries and is cleared when the receive FIFO is empty. This bit can be polled by software to completely empty the receive FIFO. '0': Receive FIFO is empty '1': Receive FIFO is not empty	0x0
TFE	2	R/O	Transmit FIFO Completely Empty. When the transmit FIFO is completely empty, this bit is set. When it contains one or more valid entries, this bit is cleared. This bit field does not request an interrupt. '0': Transmit FIFO is not empty '1': Transmit FIFO is empty	0x1
TFNF	1	R/O	Transmit FIFO Not Full. Set when the transmit FIFO contains one or more empty locations, and is cleared when the FIFO is full. '0': Transmit FIFO is full '1': Transmit FIFO is not full	0x1
BUS_ACTIVITY	0	R/O	TWI Activity Status.	0x0

5.15.1.27 TWI:TWI:TXFLR

Parent: [TWI:TWI](#)

Instances: 1

TABLE 5-612: FIELDS IN TXFLR

Field Name	Bit	Access	Description	Default
TXFLR	2:0	R/O	Transmit FIFO Level. Contains the number of valid data entries in the transmit FIFO.	0x0

5.15.1.28 TWI:TWI:RXFLR

Parent: [TWI:TWI](#)

Instances: 1

TABLE 5-613: FIELDS IN RXFLR

Field Name	Bit	Access	Description	Default
RXFLR	2:0	R/O	Receive FIFO Level. Contains the number of valid data entries in the receive FIFO.	0x0

5.15.1.29 TWI:TWI:TX_ABRT_SOURCE

Parent: [TWI:TWI](#)

Instances: 1

TABLE 5-614: FIELDS IN TX_ABRT_SOURCE

Field Name	Bit	Access	Description	Default
ABRT_SLVRD_INTX	15	R/W	When the processor side responds to a slave mode request for data to be transmitted to a remote master and user writes a 1 to DATA_CMD.CMD.	0x0
ABRT_SLV_ARBLOST	14	R/W	Slave lost the bus while transmitting data to a remote master. TX_ABRT_SOURCE[12] is set at the same time. Note: Even though the slave never "owns" the bus, something could go wrong on the bus. This is a fail safe check. For instance, during a data transmission at the low-to-high transition of SCL, if what is on the data bus is not what is supposed to be transmitted, then the TWI controller no longer own the bus.	0x0
ABRT_SLVFLUSH_TXFIFO	13	R/W	Slave has received a read command and some data exists in the TX FIFO so the slave issues a TX_ABRT interrupt to flush old data in TX FIFO.	0x0
ARB_LOST	12	R/W	Master has lost arbitration, or if TX_ABRT_SOURCE[14] is also set, then the slave transmitter has lost arbitration. Note: the TWI controller can be both master and slave at the same time.	0x0
ABRT_MASTER_DIS	11	R/W	User tries to initiate a Master operation with the Master mode disabled.	0x0
ABRT_10B_RD_NORSTR	10	R/W	The restart is disabled (RESTART_ENA bit (CFG[5]) = 0) and the master sends a read command in 10-bit addressing mode.	0x0

TABLE 5-614: FIELDS IN TX_ABRT_SOURCE (CONTINUED)

Field Name	Bit	Access	Description	Default
ABRT_SBYTE_NORSTRT	9	R/W	To clear Bit 9, the source of the ABRT_SBYTE_NORSTRT must be fixed first; restart must be enabled (CFG[5]=1), the SPECIAL bit must be cleared (TAR[11]), or the GC_OR_START bit must be cleared (TAR[10]). Once the source of the ABRT_SBYTE_NORSTRT is fixed, then this bit can be cleared in the same manner as other bits in this register. If the source of the ABRT_SBYTE_NORSTRT is not fixed before attempting to clear this bit, bit 9 clears for one cycle and then gets re-asserted. '1': The restart is disabled (RESTART_ENA bit (CFG[5]) = 0) and the user is trying to send a START Byte.	0x0
ABRT_SBYTE_ACKDET	7	R/W	Master has sent a START Byte and the START Byte was acknowledged (wrong behavior).	0x0
ABRT_GCALL_READ	5	R/W	TWI controller in master mode sent a General Call but the user programmed the byte following the General Call to be a read from the bus (DATA_CMD[9] is set to 1).	0x0
ABRT_GCALL_NOACK	4	R/W	TWI controller in master mode sent a General Call and no slave on the bus acknowledged the General Call.	0x0
ABRT_TXDATA_NOACK	3	R/W	This is a master-mode only bit. Master has received an acknowledgement for the address, but when it sent data byte(s) following the address, it did not receive an acknowledge from the remote slave(s).	0x0
ABRT_10ADDR2_NOACK	2	R/W	Master is in 10-bit address mode and the second address byte of the 10-bit address was not acknowledged by any slave.	0x0
ABRT_10ADDR1_NOACK	1	R/W	Master is in 10-bit address mode and the first 10-bit address byte was not acknowledged by any slave.	0x0
ABRT_7B_ADDR_NOACK	0	R/W	Master is in 7-bit addressing mode and the address sent was not acknowledged by any slave.	0x0

5.15.1.30 TWI:TWI:SDA_SETUP

Parent: [TWI:TWI](#)

Instances: 1

This field must be set accordingly to the VCore system frequency; value = 100ns / VCore clock period.

Example: a 178.6MHz clock correspond to a period of 5.6ns, for this frequency and fast TWI speed this field must not be set lower than (round up): $18 = 100\text{ns} / 5.6\text{ns}$. For normal TWI speed this field must not be set lower than (round up): $45 = 250\text{ns} / 5.6\text{ns}$.

TABLE 5-615: FIELDS IN SDA_SETUP

Field Name	Bit	Access	Description	Default
SDA_SETUP	7:0	R/W	This register controls the amount of time delay (in terms of number of VCore clock periods) introduced in the rising edge of SCL, relative to SDA changing, when the TWI controller services a read request in a slave-receiver operation. The minimum for fast mode is 100ns, for normal mode the minimum is 250ns.	0x15

5.15.1.31 TWI:TWI:ACK_GEN_CALL

Parent: [TWI:TWI](#)

Instances: 1

TABLE 5-616: FIELDS IN ACK_GEN_CALL

Field Name	Bit	Access	Description	Default
ACK_GEN_CALL	0	R/W	ACK General Call. When set to 1, the TWI controller responds with a ACK when it receives a General Call. Otherwise, the controller responds with a NACK.	0x1

5.15.1.32 TWI:TWI:ENABLE_STATUS

Parent: [TWI:TWI](#)

Instances: 1

TABLE 5-617: FIELDS IN ENABLE_STATUS

Field Name	Bit	Access	Description	Default
SLV_FIFO_FILLED_AND_FLUSHED	2	R/O	Slave FIFO Filled and Flushed. This bit indicates if a Slave-Receiver operation has been aborted with at least 1 data byte received from a TWI transfer due to the setting of ENABLE from 1 to 0. When read as 1, the TWI controller is deemed to have been actively engaged in an aborted TWI transfer (with matching address) and the data phase of the TWI transfer has been entered, even though the data byte has been responded with a NACK. When read as 0, the TWI controller is deemed to have been disabled when the TWI bus is idle.	0x0

TABLE 5-617: FIELDS IN ENABLE_STATUS (CONTINUED)

Field Name	Bit	Access	Description	Default
SLV_RX_ABORTED	1	R/O	Slave-Receiver Operation Aborted. This bit indicates if a Slave-Receiver operation has been aborted due to the setting of the ENABLE register from 1 to 0. When read as 1, the TWI controller is deemed to have forced a NACK during any part of a TWI transfer, irrespective of whether the TWI address matches the slave address set in the TWI controller (SAR register). When read as 0, the TWI controller is deemed to have been disabled when the TWI bus is idle.	0x0
BUSY	0	R/O	When read as 1, the TWI controller is deemed to be actively involved in an TWI transfer, irrespective of whether being in an address or data phase for all master or slave modes. When read as 0, the TWI controller is deemed completely inactive.	0x0

5.16 UART

TABLE 5-618: REGISTER GROUPS IN UART

Register Group Name	Offset within Target	Instances and Address Spacing	Description	Details
UART	0x00000000	1	UART	Page 608

5.16.1 UART:UART

Parent: [UART](#)

Instances: 1

TABLE 5-619: REGISTERS IN UART

Register Name	Offset within Register Group	Instances and Address Spacing	Description	Details
RBR_THR	0x00000000	1	Receive buffer / transmit holding / divisor (low)	Page 609
IER	0x00000004	1	Interrupt enable / divisor (high)	Page 609
IIR_FCR	0x00000008	1	Interrupt identification / FIFO control register	Page 610
LCR	0x0000000C	1	Line control	Page 611
MCR	0x00000010	1	Modem control	Page 612
LSR	0x00000014	1	Line status	Page 613
MSR	0x00000018	1	Modem status	Page 615
SCR	0x0000001C	1	Scratchpad	Page 616
USR	0x0000007C	1	UART status	Page 616

5.16.1.1 UART:UART:RBR_THR

Parent: [UART:UART](#)

Instances: 1

When the LCR.DLAB is set, this register is the lower 8 bits of the 16-bit Divisor register that contains the baud rate divisor for the UART.

The output baud rate is equal to the VCore system clock frequency divided by sixteen times the value of the baud rate divisor, as follows: $\text{baud rate} = (\text{VCore clock freq}) / (16 * \text{divisor})$. Note that with the Divisor set to zero, the baud clock is disabled and no serial communications occur. In addition, once this register is set, wait at least 0.1us before transmitting or receiving data.

TABLE 5-620: FIELDS IN RBR_THR

Field Name	Bit	Access	Description	Default
RBR_THR	7:0	R/W	Use this register to access the Rx and Tx FIFOs. When reading: The data in this register is valid only if LSR.DR is set. If FIFOs are disabled (IIR_FCR.FIFOE), the data in this register must be read before the next data arrives, otherwise it is overwritten, resulting in an overrun error. When FIFOs are enabled (IIR_FCR.FIFOE), this register accesses the head of the receive FIFO. If the receive FIFO is full and this register is not read before the next data character arrives, then the data already in the FIFO is preserved, but any incoming data is lost and an overrun error occurs. When writing: Data should only be written to this register when the LSR.THRE indicates that there is room in the FIFO. If FIFOs are disabled (IIR_FCR.FIFOE), writes to this register while LSR.THRE is zero, causes the register to be overwritten. When FIFOs are enabled (IIR_FCR.FIFOE) and LSR.THRE is set, 16 characters may be written to this register before the FIFO is full. Any attempt to write data when the FIFO is full results in the write data being lost.	0x00

5.16.1.2 UART:UART:IER

Parent: [UART:UART](#)

Instances: 1

When the LCR.DLAB is set, this register is the upper 8 bits of the 16-bit Divisor register that contains the baud rate divisor for the UART. For more information and a description of how to calculate the baud rate, see RBR_THR.

TABLE 5-621: FIELDS IN IER

Field Name	Bit	Access	Description	Default
PTIME	7	R/W	Programmable THRE interrupt mode enable. This is used to enable or disable the generation of THRE interrupt. 0: Disabled 1: Enabled	0x0

TABLE 5-621: FIELDS IN IER (CONTINUED)

Field Name	Bit	Access	Description	Default
EDSSI	3	R/W	Enable modem status interrupt. This is used to enable or disable the generation of Modem Status interrupt. This is the fourth highest priority interrupt. 0: Disabled 1: Enabled	0x0
ELSI	2	R/W	Enable receiver line status interrupt. This is used to enable or disable the generation of Receiver Line Status interrupt. This is the highest priority interrupt. 0: Disabled 1: Enabled	0x0
ETBEI	1	R/W	Enable transmit holding register empty interrupt. This is used to enable or disable the generation of Transmitter Holding Register Empty interrupt. This is the third highest priority interrupt. 0: Disabled 1: Enabled	0x0
ERBFI	0	R/W	Enable received data available interrupt. This is used to enable or disable the generation of Received Data Available interrupt and the Character Timeout interrupt (if FIFOs are enabled). These are the second highest priority interrupts. 0: Disabled 1: Enabled	0x0

5.16.1.3 UART:UART:IIR_FCR

Parent: [UART:UART](#)

Instances: 1

This register has special meaning when reading, here the lowest 4 bits indicate interrupting sources. The encoding is as follows:

0110; type: Receiver line status, priority: Highest. Overrun/parity/ framing errors or break interrupt. Cleared by reading LSR.

0100; type: Received data available, priority: Second. RCVR FIFO trigger level reached. Cleared when FIFO drops below the trigger level.

1100; type: Character timeout indication, priority: Second. No characters in or out of the RCVR FIFO during the last four character times and there is at least 1 character in it during this time. Cleared by reading the receiver buffer register.

0010; type: Transmit holding register empty, priority: Third. Transmitter holding register empty (Prog. THRE Mode disabled) or XMIT FIFO at or below threshold (Prog. THRE Mode enabled). Cleared by reading the IIR register (if source of interrupt); or, writing into THR (THRE Mode disabled) or XMIT FIFO above threshold (THRE Mode enabled).

0000; type: Modem status, priority: Fourth. Clear to send. Note that if auto flow control mode is enabled, a change in CTS (that is, DCTS set) does not cause an interrupt. Cleared by reading the Modem status register.

0111; type: Busy detect indication, priority: Fifth. Master has tried to write to the Line Control register while the UART is busy (USR[0] is set to one). Cleared by reading the UART status register.

0001: No interrupting sources.

TABLE 5-622: FIELDS IN IIR_FCR

Field Name	Bit	Access	Description	Default
FIFOSE_RT	7:6	R/W	When reading this field, the current status of the FIFO is returned; 00 for disabled or 11 for enabled. Writing this field selects the trigger level in the receive FIFO at which the Received Data Available interrupt is generated (see encoding.) In auto flow control mode, it is used to determine when to generate back-pressure using the RTS signal. 00: 1 character in the Rx FIFO 01: Rx FIFO 1/4 full 10: Rx FIFO 1/2 full 11: Rx FIFO 2 less than full	0x1
TET	5:4	R/W	Tx empty trigger. When the THRE mode is enabled (IER.PTIME), this field selects the empty threshold level at which the THRE Interrupts are generated. 00: Tx FIFO empty 01: 2 characters in the Tx FIFO 10: Tx FIFO 1/4 full 11: Tx FIFO 1/2 full	0x0
XFIFOR	2	R/W	This description is valid for writes only. Reading this field has special meaning; for more information, see the general register description. Tx FIFO Reset. This resets the control portion of the transmit FIFO and treats the FIFO as empty. Note that this bit is self-clearing. It is not necessary to clear this bit.	0x0
RFIFOR	1	R/W	This description is valid for writes only. Reading this field has special meaning; for more information, see the general register description. Rx FIFO Reset. This resets the control portion of the receive FIFO and treats the FIFO as empty. Note that this bit is self-clearing. It is not necessary to clear this bit.	0x0
FIFOE	0	R/W	This description is valid for writes only. Reading this field has special meaning; for more information, see the general register description. FIFO Enable. This enables or disables the transmit (XMIT) and receive (RCVR) FIFOs. Whenever the value of this bit is changed, both the XMIT and RCVR controller portion of FIFOs are reset.	0x0

5.16.1.4 UART:UART:LCR

Parent: [UART:UART](#)

Instances: 1

Writes can be made to this register, with the exception of the BC field, only when UART is not busy, that is, when USR.BUSY is zero. This register can always be read.

TABLE 5-623: FIELDS IN LCR

Field Name	Bit	Access	Description	Default
DLAB	7	R/W	Divisor latch access bit. This bit is used to enable reading and writing of the Divisor registers (RBR_THR and IER) to set the baud rate of the UART. To access other registers, this bit must be cleared after initial baud rate setup.	0x0
BC	6	R/W	Break control bit. This bit is used to cause a break condition to be transmitted to the receiving device. If set to one, the serial output is forced to the spacing (logic 0) state. When not in Loopback Mode, as determined by MCR[4], the serial line is forced low until the Break bit is cleared.	0x0
EPS	4	R/W	Even parity select. This bit is used to select between even and odd parity, when parity is enabled (PEN set to one). If set to one, an even number of logic 1s is transmitted or checked. If set to zero, an odd number of logic 1s is transmitted or checked.	0x0
PEN	3	R/W	Parity enable. This bit is used to enable or disable parity generation and detection in both transmitted and received serial characters. 0: Parity disabled 1: Parity enabled	0x0
STOP	2	R/W	Number of stop bits. This is used to select the number of stop bits per character that the peripheral transmits and receives. If set to zero, one stop bit is transmitted in the serial data. If set to one and the data bits are set to 5 (LCR.DLS), one and a half stop bits are transmitted. Otherwise, two stop bits are transmitted. Note that regardless of the number of stop bits selected, the receiver checks only the first stop bit. 0: 1 stop bit 1: 1.5 stop bits when LCR.DLS is zero, otherwise, 2 stop bits	0x0
DLS	1:0	R/W	Data length select. This is used to select the number of data bits per character that the peripheral transmits and receives. The following settings specify the number of bits that may be selected. 00: 5 bits 01: 6 bits 10: 7 bits 11: 8 bits	0x0

5.16.1.5 UART:UART:MCR

Parent: [UART:UART](#)

Instances: 1

TABLE 5-624: FIELDS IN MCR

Field Name	Bit	Access	Description	Default
AFCE	5	R/W	Auto flow control enable. This mode requires that FIFOs are enabled and that MCR.RTS is set. 0: Auto flow control mode disabled 1: Auto flow control mode enabled	0x0
LB	4	R/W	Loopback Bit. This is used to put the UART into a diagnostic mode for test purposes. The transmit line is held high, while serial transmit data is looped back to the receive line internally. In this mode, all the interrupts are fully functional. In addition, in loopback mode, the modem control input CTS is disconnected, and the modem control output RTS is looped back to the input internally.	0x0
RTS	1	R/W	Request to send. This is used to directly control the Request to Send (RTS) output. The RTS output is used to inform the partner that the UART is ready to exchange data. The RTS is still controlled from this field when Auto RTS Flow Control is enabled (MCR.AFCE), but the output can be forced high by the flow control mechanism. If this field is cleared, the UART permanently indicates backpressure to the partner. 0: RTS is set high 1: RTS is set low	0x0

5.16.1.6 UART:UART:LSR

Parent: [UART:UART](#)

Instances: 1

TABLE 5-625: FIELDS IN LSR

Field Name	Bit	Access	Description	Default
RFE	7	R/W	Receiver FIFO error bit. This bit is only valid when FIFOs are enabled. This is used to indicate whether there is at least one parity error, framing error, or break indication in the FIFO. This bit is cleared when the LSR is read, the character with the error is at the top of the receiver FIFO, and there are no subsequent errors in the FIFO. 0: No error in Rx FIFO 1: Error in Rx FIFO	0x0
TEMT	6	R/W	Transmitter empty bit. If FIFOs are enabled, this bit is set whenever the Transmitter Shift Register and the FIFO are both empty.	0x1

TABLE 5-625: FIELDS IN LSR (CONTINUED)

Field Name	Bit	Access	Description	Default
THRE	5	R/W	If FIFO (IIR_FCR.FIFOE) and THRE mode are enabled (IER.PTIME), this bit indicates that the Tx FIFO is full. Otherwise, this bit indicates that the Tx FIFO is empty.	0x1
BI	4	R/W	Break interrupt bit. This is used to indicate the detection of a break sequence on the serial input data. It is set whenever the serial input is held in a logic 0 state for longer than the sum of start time + data bits + parity + stop bits. A break condition on serial input causes one and only one character, consisting of all-zeros, to be received by the UART. In the FIFO mode, the character associated with the break condition is carried through the FIFO and is revealed when the character is at the top of the FIFO. Reading the LSR clears the BI bit. In the non-FIFO mode, the BI indication occurs immediately and persists until the LSR is read.	0x0
FE	3	R/W	Framing error bit. This is used to indicate the a framing error in the receiver. A framing error occurs when the receiver does not detect a valid STOP bit in the received data. A framing error is associated with a received character. Therefore, in FIFO mode, an error is revealed when the character with the framing error is at the top of the FIFO. When a framing error occurs, the UART tries to resynchronize. It does this by assuming that the error was due to the start bit of the next character and then continues to receive the other bit, that is, data and/or parity, and then stops. Note that this field is set if a break interrupt has occurred, as indicated by Break Interrupt (LSR.BI). This field is cleared on read. 0: No framing error 1: Framing error	0x0
PE	2	R/W	Parity error bit. This is used to indicate the occurrence of a parity error in the receiver if the Parity Enable bit (LCR.PEN) is set. A parity error is associated with a received character. Therefore, in FIFO mode, an error is revealed when the character with the parity error arrives at the top of the FIFO. Note that this field is set if a break interrupt has occurred, as indicated by Break Interrupt (LSR.BI). This field is cleared on read. 0: No parity error 1: Parity error	0x0

TABLE 5-625: FIELDS IN LSR (CONTINUED)

Field Name	Bit	Access	Description	Default
OE	1	R/W	Overrun error bit. This is used to indicate the occurrence of an overrun error. This occurs if a new data character was received before the previous data was read. In non-FIFO mode, the OE bit is set when a new character arrives before the previous character was read. When this happens, the data in the RBR is overwritten. In FIFO mode, an overrun error occurs when the FIFO is full and a new character arrives at the receiver. The data in the FIFO is retained and the data in the receive shift register is lost. This field is cleared on read. 0: No overrun error 1: Overrun error	0x0
DR	0	R/W	Data ready. This is used to indicate that the receiver contains at least one character in the receiver FIFO. This bit is cleared when the RX FIFO is empty. 0: No data ready 1: Data ready	0x0

5.16.1.7 UART:UART:MSR

Parent: [UART:UART](#)

Instances: 1

TABLE 5-626: FIELDS IN MSR

Field Name	Bit	Access	Description	Default
CTS	4	R/O	Clear to send. This field indicates the current state of the modem control line, CTS. When the Clear to Send input (CTS) is asserted, it is an indication that the partner is ready to exchange data with the UART. 0: CTS input is deasserted (logic 0) 1: CTS input is asserted (logic 1)	0x0
DCTS	0	R/O	Delta clear to send. This is used to indicate that the modem control line, CTS, has changed since the last time the MSR was read. Reading the MSR clears the DCTS bit. Note: If the DCTS bit is not set, the CTS signal is asserted, and a reset occurs (software or otherwise), then the DCTS bit is set when the reset is removed, if the CTS signal remains asserted. A read of the MSR after reset can be performed to prevent unwanted interrupts. 0: No change on CTS since the last read of the MSR 1: Change on CTS since the last read of the MSR	0x0

5.16.1.8 UART:UART:SCR

Parent: [UART:UART](#)

Instances: 1

TABLE 5-627: FIELDS IN SCR

Field Name	Bit	Access	Description	Default
SCR	7:0	R/W	This register is for programmers to use as a temporary storage space. It has no functional purpose for the UART.	0x00

5.16.1.9 UART:UART:USR

Parent: [UART:UART](#)

Instances: 1

TABLE 5-628: FIELDS IN USR

Field Name	Bit	Access	Description	Default
BUSY	0	R/O	UART busy. 0: UART is idle or inactive 1: UART is busy (actively transferring data)	0x0

6.0 ELECTRICAL SPECIFICATIONS

This section provides the DC characteristics, AC characteristics, recommended operating conditions, and stress ratings for the VSC7416-01 device.

6.1 DC Characteristics

This section contains the DC specifications for the VSC7416-01 device.

6.1.1 INTERNAL PULL-UP OR PULL-DOWN RESISTORS

Internal pull-up or pull-down resistors are specified in the following table. For more information about signals with internal pull-up or pull-down resistors, see [Section 7.2, Pins by Function](#).

All internal pull-up resistors are connected to their respective I/O supply.

TABLE 6-1: INTERNAL PULL-UP OR PULL-DOWN RESISTORS

Parameter	Symbol	Minimum	Typical	Maximum	Unit
Internal pull-up resistor, GPIO and SI	R_{PU}	33	53	90	k Ω
Internal pull-up resistor, all other pins	R_{PU}	96	120	144	k Ω
Internal pull-down resistor	R_{PD}	96	120	150	k Ω

6.1.2 REFERENCE CLOCKS

The following table lists the DC specifications for the differential reference clock signals. Differential and single-ended modes are supported. For more information about single-ended mode operation, see [Section 9.3.1, Single-Ended Ref-Clk Input](#).

TABLE 6-2: REFCLK INPUT DC SPECIFICATIONS

Parameter	Symbol	Minimum	Maximum	Unit
Input voltage range	V_{IP}, V_{IN}	-25	1260	mV
Input differential voltage, peak-to-peak	$ V_{ID} $	150 ¹	1000	mV
Input common-mode voltage	V_{CM}	0	1200 ²	mV
Differential input impedance	R_I	80	120	Ω

1. To meet jitter specifications, the minimum input differential voltage must be 400 mV. When using a single-ended clock input, the RefClk_P low voltage level must be lower than $VDD\_A - 200$ mV, and the high voltage level must be higher than $VDD\_A + 200$ mV.
2. The maximum common-mode voltage is given without any differential signal, because the input is internally AC-coupled. The common-mode voltage is only limited by the maximum and minimum input voltage range and the input signal's differential amplitude.

6.1.2.1 Reference Clock Output

The following table lists the DC specifications for the ClkOut2 driver.

TABLE 6-3: CLKOUT2 DRIVER DC SPECIFICATIONS

Parameter	Symbol	Minimum	Maximum	Unit	Condition
Output high voltage, V_{OA} or V_{OB}	V_{OH}		1050	mV	$R_L = 100 \Omega \pm 1\%$
Output low voltage, V_{OA} or V_{OB}	V_{OL}	0		mV	$R_L = 100 \Omega \pm 1\%$
Output differential peak voltage	$ V_{OD} $	150	400	mV	$R_L = 100 \Omega \pm 1\%$
Output offset voltage	$ V_{OS} $	420	580	mV	$R_L = 100 \Omega \pm 1\%$
DC output impedance, single-ended	R_O	40	140	Ω	
R_O mismatch between A and B	ΔR_O		10	%	
Change in $ V_{OD} $ between 0 and 1	$\Delta V_{OD} $		25	mV	$R_L = 100 \Omega \pm 1\%$

TABLE 6-3: CLKOUT2 DRIVER DC SPECIFICATIONS (CONTINUED)

Parameter	Symbol	Minimum	Maximum	Unit	Condition
Change in V_{OS} between 0 and 1	$\Delta V_{OS} $		25	mV	$R_L = 100\ \Omega \pm 1\%$.

6.1.3 DDR2/DDR3 SDRAM INTERFACE

The DDR2 SDRAM interface supports the requirements of SDRAM devices as described in the JEDEC DDR2 specifications. The SDRAM interface signals are compatible with JESD79-2E (DDR2 SDRAM Specification, April 2008) and the JESD8-15A (Stub Series Terminated Logic for 1.8V (SSTL₁₈), September 2003). The SSTL I/O buffers have programmable on-die termination (ODT).

The following table lists the DC specifications for DDR2 SDRAM interface signals.

TABLE 6-4: DDR2 SDRAM SIGNAL DC SPECIFICATIONS

Parameter	Symbol	Minimum	Maximum	Unit	Condition
Input reference voltage ⁽¹⁾	DDR_VREF	49% V_{DD_IODDR}	51% V_{DD_IODDR}	V	$V_{DD\_IODDR} = 1.8\text{ V}$
Input voltage high	$V_{IH(DC)}$	DDR_VREF + 0.125	$V_{DD\_IODDR} + 0.3$	V	
Input voltage low	$V_{IL(DC)}$	-0.3	DDR_VREF - 0.125	V	
Input leakage current	$ I_L $		58	μA	$0\text{ V} \leq V_I \leq V_{DD\_IODDR}$.
Output source DC current ⁽²⁾	I_{OH}	-6		mA	External 50 Ω termination to $V_{DD\_IODDR}/2$.
Output sink DC current ⁽²⁾	I_{OL}	6		mA	External 50 Ω termination to $V_{DD\_IODDR}/2$.

1. DDR_VREF is expected to track variations in V_{DD_IODDR} . Peak-to-peak AC noise on DDR_VREF must not exceed $\pm 2\%$ of DDR_VREF.
2. With 40 Ω output driver impedance.

The DDR3 SDRAM interface supports the requirements of SDRAM devices as described in the JEDEC DDR3 specifications. The SDRAM interface signals are compatible with JESD79-3E (DDR3 SDRAM Specification, July 2010). The SSTL I/O buffers have programmable on-die termination (ODT).

The following table lists the DC specifications for DDR3 SDRAM interface signals.

TABLE 6-5: DDR3 SDRAM SIGNAL DC SPECIFICATIONS

Parameter	Symbol	Minimum	Maximum	Unit	Condition
Input reference voltage ⁽¹⁾	DDR_VREF	49% V_{DD_IODDR}	51% V_{DD_IODDR}	V	$V_{DD\_IODDR} = 1.5\text{ V}$
Input voltage high	$V_{IH(DC)}$	DDR_VREF + 0.100	V_{DD_IODDR}	V	
Input voltage low	$V_{IL(DC)}$	-0.3	DDR_VREF - 0.100	V	
Input leakage current	$ I_L $		12	μA	$0\text{ V} \leq V_I \leq V_{DD\_IODDR}$.
Output source DC current ⁽²⁾	I_{OH}	-6		mA	External 40 Ω termination to $V_{DD\_IODDR}/2$.
Output sink DC current ⁽²⁾	I_{OL}	6		mA	External 40 Ω termination to $V_{DD\_IODDR}/2$.

1. DDR_VREF is expected to track variations in V_{DD_IODDR} . Peak-to-peak AC noise on DDR_VREF must not exceed $\pm 2\%$ of DDR_VREF.
2. With 50 Ω output driver impedance.

6.1.4 SGMII DC DEFINITIONS AND TEST CIRCUITS

This section provides information about the definitions and test circuits that apply to certain parameters for the Enhanced SerDes and SerDes interfaces. The following illustrations show the DC definitions for the SGMII inputs and outputs.

FIGURE 6-1: SGMII DC INPUT DEFINITIONS




FIGURE 6-2: SGMII DC TRANSMIT TEST CIRCUIT





FIGURE 6-3: SGMII DC DEFINITIONS




$$\Delta|V_{OD}| = | |V_{OAH} - V_{OBL}| - |V_{OBH} - V_{OAL}| |$$

$$\Delta V_{OS} = | \frac{1}{2}(V_{OAH} + V_{OBL}) - \frac{1}{2}(V_{OAL} + V_{OBH}) |$$

The following illustrations show the SGMII DC driver output impedance test circuit and the DC input definitions.

FIGURE 6-4: SGMII DC DRIVER OUTPUT IMPEDANCE TEST CIRCUIT



6.1.5 ENHANCED SERDES INTERFACE

All DC specifications for the Enhanced SerDes interface are compliant with the OIF-CEI version 2.0 requirements where applicable.

The Enhanced SerDes interface supports three major modes: SGMII, 2.5G, and SFP. The values in the following table apply to modes specified.

TABLE 6-6: ENHANCED SERDES DRIVER DC SPECIFICATIONS

Parameter	Symbol	Minimum	Maximum	Unit	Condition
Output differential peak voltage ⁽¹⁾ , 1.0 V	$ V_{ODp} $	150	400	mV	$V_{DD\_VS} = 1.0\text{ V}$. $R_L = 100\ \Omega \pm 1\%$.
Output differential peak voltage ⁽¹⁾ , 1.2 V	$ V_{ODp} $	150	600	mV	$V_{DD\_VS} = 1.2\text{ V}$. $R_L = 100\ \Omega \pm 1\%$.
DC output impedance, single-ended, SGMII mode	$ R_O $	40	140	Ω	$R_L = 100\ \Omega \pm 1\%$. $V_C = 1.0\text{ V}$ and 1.2 V .
R_O mismatch between A and B ⁽²⁾ , SGMII mode	$\Delta R_O $		10	%	$V_C = 1.0\text{ V}$ and 1.2 V
Change in $ V_{OD} $ between 0 and 1, SGMII mode	$\Delta V_{OD} $		25	mV	$R_L = 100\ \Omega \pm 1\%$
Change in V_{OS} between 0 and 1, SGMII mode	$\Delta V_{OS} $		25	mV	$R_L = 100\ \Omega \pm 1\%$
Output current, drivers shorted to GND, SGMII	$ I_{OSA} $, $ I_{OSB} $		40	mA	
Output current, drivers shorted together, SGMII	$ I_{OSAB} $		12	mA	

1. Voltage is adjustable in 64 steps. For more information about setting the adjustable voltages, see OB_LEV in the HSIO:SERDES6G_ANA_CFG:SERDES6G_OB_CFG1 register. To meet the return loss specification, the maximum swing is 600 mV peak-to-peak for $V_{DD\_VS} = 1.0\text{ V}$ and 950 mV peak-to-peak for $V_{DD\_VS} = 1.2\text{ V}$.
2. Matching of reflection coefficients. For more information about test methods, see IEEE 1596.3-1996.

The following table lists the DC specifications for the Enhanced SerDes receivers. AC-coupling is required for most applications. For more information, see [Section 9.4.4, Enhanced SerDes Interface](#).

TABLE 6-7: ENHANCED SERDES RECEIVER DC SPECIFICATIONS

Parameter	Symbol	Minimum	Typical	Maximum	Unit
Input voltage range, V_{IA} or V_{IB}	V_I	-0.25		1.2	V
Input differential voltage ⁽¹⁾ , peak-to-peak	$ V_{ID} $	50		800	mV
Receiver differential input impedance	R_I	80	100	120	Ω

1. Range specified is for optimal operation.

6.1.6 SERDES (SGMII) INTERFACE

The SerDes output drivers are designed to operate in an SGMII/LVDS mode and in a high-drive/PECL mode (SFP and 1000BASE-KX modes). The SGMII/LVDS mode meets or exceeds the DC requirements of the Serial-GMII Specification version 1.9, unless otherwise noted.

The following table lists the DC specifications for the SGMII driver. The values are valid for all configurations, unless stated otherwise.

TABLE 6-8: SERDES DRIVER DC SPECIFICATIONS

Parameter	Symbol	Minimum	Maximum	Unit	Condition
Output high voltage, V_{OA} or V_{OB}	V_{OH}		1250	mV	$R_L = 100\ \Omega \pm 1\%$.
Output low voltage, V_{OA} or V_{OB}	V_{OL}	0		mV	$R_L = 100\ \Omega \pm 1\%$.
Output differential peak voltage ⁽¹⁾ , 1.0 V	$ V_{OD} $	150	400	mV	$V_{DD\_VS} = 1.0\text{ V}$, $R_L = 100\ \Omega \pm 1\%$.
Output differential peak voltage ⁽¹⁾ , 1.2 V	$ V_{OD} $	150	600	mV	$V_{DD\_VS} = 1.2\text{ V}$, $R_L = 100\ \Omega \pm 1\%$.
Output offset voltage ⁽²⁾ , 1.0 V	$ V_{OS} $	420	580	mV	$V_{DD\_VS} = 1.0\text{ V}$, $R_L = 100\ \Omega \pm 1\%$.
Output offset voltage ⁽²⁾ , 1.2 V		445	605	mV	$V_{DD\_VS} = 1.2\text{ V}$, $R_L = 100\ \Omega \pm 1\%$.
DC output impedance, single-ended, SGMII mode	R_O	40	140	Ω	$V_C = 1.0\text{ V}$ and 1.2 V.
R_O mismatch between A and B ⁽³⁾ , SGMII mode	ΔR_O		10	%	$V_C = 1.0\text{ V}$ and 1.2 V.
Change in $ V_{OD} $ between 0 and 1, SGMII mode	$\Delta V_{OD} $		25	mV	$R_L = 100\ \Omega \pm 1\%$.
Change in V_{OS} between 0 and 1, SGMII mode	$\Delta V_{OS} $		25	mV	$R_L = 100\ \Omega \pm 1\%$.
Output current, drivers shorted to GND, SGMI mode	$ I_{OSA} $, $ I_{OSB} $		40	mA	
Output current, drivers shorted together, SGMII mode	$ I_{OSAB} $		12	mA	

1. Voltage is adjustable in 14 steps. For more information about setting the adjustable voltages, see the OB_AMP_CTRL bit. To meet the return loss specification, the maximum amplitude is 600 mV peak-to-peak for $V_{DD\_VS} = 1.0\text{ V}$ and 950 mV peak-to-peak for $V_{DD\_VS} = 1.2\text{ V}$.
2. Requires AC-coupling for SGMII compliance.
3. Matching of reflection coefficients. For more information about test methods, see IEEE 1596.3-1996.

The following table lists the DC specifications for the SGMII receivers.

TABLE 6-9: SERDES RECEIVER DC SPECIFICATIONS

Parameter	Symbol	Minimum	Maximum	Unit	Condition
Input voltage range, V_{IA} or V_{IB}	V_I	-25	1250	mV	
Input differential peak voltage	$ V_{ID} $	50	1000	mV	
Input common-mode voltage ⁽¹⁾	V_{IC}	0	V_{DD_A} ⁽²⁾	mV	Without any differential signal (internally AC-coupled)
Receiver differential input impedance	R_I	80	120	Ω	
Input differential hysteresis	V_{HYST}	25		mV	

1. SGMII compliancy requires external AC-coupling. When interfacing with specific Microchip devices, DC-coupling is possible. For more information, contact your Microchip sales representative.

- The maximum common-mode voltage is given without any differential signal, because the input is internally AC-coupled. The common-mode voltage is only limited by the maximum and minimum input voltage range and the input signal's differential amplitude.

6.1.7 PCI EXPRESS INTERFACE DC SPECIFICATIONS

All DC specifications for the PCI express interface are compliant with PCI Express Base Specification Revision 1.0 where applicable.

The values in the following table apply to modes specified in the condition column.

TABLE 6-10: PCI EXPRESS DRIVER DC SPECIFICATIONS

Parameter	Symbol	Minimum	Maximum	Unit	Condition
Output differential peak voltage ⁽¹⁾ , 1.0 V,	$ V_{ODp} $	250	400	mV	$V_{DD\_VS} = 1.0 \text{ V}$, $R_L = 100 \Omega \pm 1\%$.
Output differential peak voltage ⁽¹⁾ , 1.2 V,	V_{ODp}	360	600	mV	$V_{DD\_VS} = 1.2 \text{ V}$, $R_L = 100 \Omega \pm 1\%$
DC output impedance, differential	R_O	40	140	Ω	$R_L = 100 \Omega \pm 1\%$, $V_C = 1.0 \text{ V}$ and 1.2 V .
R_O mismatch between A and B ⁽²⁾	ΔR_O		10	%	$V_C = 1.0 \text{ V}$ and 1.2 V .
Change in $ V_{OD} $ between 0 and 1	$\Delta V_{OD} $		25	mV	$R_L = 100 \Omega \pm 1\%$.
Change in V_{OS} between 0 and 1	ΔV_{OS}		25	mV	$R_L = 100 \Omega \pm 1\%$.
Output current, driver shorted to GND	$ I_{OSA} $, $ I_{OSB} $		40	mA	
Output current, drivers shorted together	$ I_{OSAB} $		12	mA	

- Voltage is adjustable in 64 steps. For more information about setting the adjustable voltages, see OB_LEV in the HSIO:SERDES6G_ANA_CFG:SERDES6G_OB_CFG1 register. For $V_{DD\_VS} = 1.0 \text{ V}$, the output voltage does not meet the full power defined in *PCI Express 3.0 Base Specification Revision 3.0*; however, it does meet the low power specification and can be used for short links. To meet the return loss specification, the maximum amplitude is 600 mV peak-to-peak for $V_{DD\_VS} = 1.0 \text{ V}$ and 950 mV peak-to-peak for $V_{DD\_VS} = 1.2 \text{ V}$.
- Matching of reflection coefficients. For more information about test methods, see *PCI Express 3.0 Base Specification Revision 3.0*.

The following table lists the DC specifications for the PCI Express receivers. In most applications, AC-coupling is required. For more information, see [Section 9.4.5, PCI Express Interface](#).

TABLE 6-11: PCI EXPRESS RECEIVER DC SPECIFICATIONS

Parameter	Symbol	Minimum	Typical	Maximum	Unit
Input voltage range, V_{IA} or V_{IB}	V_I	-0.25		1.2	V
Input differential voltage	$ V_{ID} $	50		800	mV
Input common-mode voltage	V_{IC}	0		1200	mV
Receiver differential input impedance	R_I	80	100	120	Ω

6.1.8 MIIM, GPIO, SI, JTAG, AND MISCELLANEOUS SIGNALS

This section provides the DC specifications for the MII Management (MIIM), GPIO, SI, JTAG, and miscellaneous signals. The following I/O signals comply with the specifications provided in this section.

TABLE 6-12:

GPIO[31:0]	JTAG_TDO	PLL2_Enable
SI_Clk	JTAG_TCK	JTAG_ICE_nEN
SI_DI	JTAG_TDI	RCVRD_CLK[9:0]
SI_DO	nReset	Test_Mode
SI_nEn	Reserved	
JTAG_nTRST	RefClk2_Sel[2:0]	
JTAG_TMS	VCore_CFG[2:0]	

The outputs and inputs meet or exceed the requirements of the LVTTTL and LVCMOS standard, JEDEC JESD8-B (September 1999) standard, unless otherwise stated. The inputs are Schmitt-trigger for noise immunity.

TABLE 6-13: MIIM, GPIO, SI, JTAG, AND MISCELLANEOUS SIGNALS DC SPECIFICATIONS

Parameter	Symbol	Minimum	Typical	Maximum	Unit	Condition
Output high voltage ⁽¹⁾ , $I_{OH} = -2 \text{ mA}$	V_{OH}	2.1	2.35		V	
Output high voltage ⁽¹⁾ , $I_{OH} = -12 \text{ mA}$	V_{OH}	1.7	2.0		V	
Output low voltage, $I_{OL} = 2 \text{ mA}$	V_{OL}			0.4		
Output low voltage, $I_{OL} = 12 \text{ mA}$	V_{OL}			0.7	V	
Input high voltage	V_{IH}	1.85		3.6	V	
Input low voltage	V_{IL}	-0.3		0.8	V	
Input high current ⁽²⁾	I_{IH}			10	μA	$V_I = V_{DD\_IO}$
Input low current ⁽²⁾	I_{IL}	-10			μA	$V_I = 0 \text{ V}$
Input capacitance	C_I			10	pF	

1. $V_{DD\_IO} = 2.38 \text{ V}$ for minimum, and $V_{DD\_IO} = 2.50 \text{ V}$ for typical.

2. Input high current and input low current equals the maximum leakage current, excluding the current in the built-in pull resistors.

6.1.9 THERMAL DIODE

The VSC7416-01 device include an on-die diode and internal circuitry for monitoring die temperature (junction temperature). The operation and accuracy of the diode is not guaranteed and should only be used as a reference.

A thermal sensor, located on the board or in a stand-alone measurement kit, can monitor and display the die temperature of the switch for thermal management or instrumentation purposes.

Temperature measurement using a thermal diode is very sensitive to noise.

The following table provides the diode parameter and interface specifications. Note that the THERMDC_VSS pin is connected to VSS internally in the devices.

TABLE 6-14: THERMAL DIODE PARAMETERS

Parameter	Symbol	Typical	Maximum	Unit
Forward bias current	IFW		1	mA
Diode ideality factor	n	1.008		

Notes Microchip does not support or recommend operation of the thermal diode under reverse bias.

The ideality factor, n , represents the deviation from ideal diode behavior as exemplified by the diode equation:

$$I_{FW} = I_S \times \left(e^{V_d \times \frac{q}{nkT}} - 1 \right)$$

where, I_S = saturation current, q = electronic charge, V_d = voltage across the diode, k = Boltzmann Constant, and T = absolute temperature (Kelvin).

6.2 AC Characteristics

This section provides the AC specifications for the VSC7416-01 device.

6.2.1 REFERENCE CLOCKS

The signal applied to the RefClk differential inputs must comply with the requirements listed in the following table at the pin of the device. Use of a 25 MHz single-ended reference clock is not recommended. However, to implement an SGMII chip interconnect using a 25 MHz single-ended reference clock and achieve error-free data transfer on that interface, use an Ethernet PHY with higher jitter tolerance than specified in the standard, such as Microchip's VSC8512-01 or VSC8522-01. For more information about SGMII interoperability when using a 25 MHz single-ended reference clock, contact your Microchip representative.

TABLE 6-15: REFCLK AC SPECIFICATIONS

Parameter	Symbol	Minimum	Typical	Maximum	Unit	Condition
RefClk frequency RefClk_Sel = 000	f	-100 ppm	125	100 ppm	MHz	
RefClk frequency RefClk_Sel = 001	f	-100 ppm	156.25	100 ppm	MHz	
RefClk frequency, RefClk_Sel = 100	f	-100 ppm	25	100 ppm	MHz	
RefClk frequency RefClk_Sel = 010	f	-100 ppm	250	100 ppm	MHz	
Clock duty cycle		40		60	%	Measured at 50% threshold
Rise time and fall time	t_R, t_F			1.5	ns	20% to 80% threshold
RefClk input RMS jitter tolerance, bandwidth from 12 kHz to 500 kHz				20	ps	
RefClk input RMS jitter tolerance, bandwidth from 500 kHz to 15 MHz				4	ps	
RefClk input RMS jitter tolerance, bandwidth from 15 MHz to 40 MHz				20	ps	
RefClk input RMS jitter tolerance, bandwidth from 40 MHz to 80 MHz				100	ps	
Jitter gain from RefClk to SerDes output, bandwidth from 0 Hz to 100 kHz				0.3	dB	
Jitter gain from RefClk to SerDes output, bandwidth from 100 kHz to 7 MHz				3	dB	
Jitter gain from RefClk to SerDes output, bandwidth greater than 7 MHz				$3 - 20 \times \log(f/7 \text{ MHz})$	dB	

6.2.1.1 Reference Clock Output

The following table lists the AC specifications for the ClkOut2 driver.


TABLE 6-16: CLKOUT2 DRIVER AC SPECIFICATION

Parameter	Symbol	Minimum	Maximum	Unit	Condition
Clock duty cycle	t_C	45	55		Measures at 50% threshold
V_{OD} rise time and fall time	t_R, t_F	100	300	ps	20% to 80% of V_S , $R_L = 100 \Omega \pm 1$
Intrapair skew	t_{SKEW}		100	ps	
RMS jitter, bandwidth from 12 kHz to 500 kHz			20	ps	
RMS jitter, bandwidth from 500 kHz to 15 MHz			4	ps	
RMS jitter, bandwidth from 15 MHz to 40 MHz			20	ps	
RMS jitter, bandwidth from 40 kHz to 80 kHz			100	ps	

6.2.2 RESET TIMING

The nReset signal waveform and the required measurement points for the timing specification are shown in the following illustration.

FIGURE 6-5: NRESET SIGNAL TIMING SPECIFICATIONS



The signal applied to the nReset input must comply with the specifications listed in the following table at the reset pin of the device.

TABLE 6-17: NRESET TIMING SPECIFICATIONS


Parameter	Symbol	Minimum	Maximum	Unit
nReset assertion time after power supplies and clock stabilize	t_W	2		ms
Recovery time from reset inactive to device fully active	t_{REC}		50	ms
nReset pulse width	$t_{W(RL)}$	100		ns
Hold time for GPIO mapped strapping pins relative to nReset	t_{HCFG}	50		ns

6.2.3 DDR2 SDRAM INTERFACE

This section provides the AC characteristics for the DDR2/DDR3 SDRAM interface in DDR2 mode.

The following illustration shows the DDR2 SDRAM input timing diagram.

FIGURE 6-6: DDR2 SDRAM INPUT TIMING DIAGRAM




The following table lists the AC specifications for the DDR2 SDRAM input signals.

TABLE 6-18: DDR2 SDRAM INPUT SIGNAL AC CHARACTERISTICS

Parameter	Symbol	Minimum	Maximum	Unit	Condition
Input voltage high	$V_{\text{IH(AC)}}$	$\text{DDR\_V}_{\text{REF}} + 0.20$	$V_{\text{DD\_IODDR}} + 0.3$	V	
Input voltage low	$V_{\text{IL(AC)}}$	-0.3	$\text{DDR\_V}_{\text{REF}} - 0.20$	V	
Differential input voltage	$V_{\text{ID(AC)}}$	0.5	$V_{\text{DD\_IODDR}}$	V	
Differential crosspoint voltage	$V_{\text{IX(AC)}}$	$0.5 \times V_{\text{DD\_IODDR}} - 0.175$	$0.5 \times V_{\text{DD\_IODDR}} + 0.175$	V	
DDR_DQ[7:0] input setup time relative to DDR_LDQS	t_{DS}		350	ps	
DDR_DQ[15:8] input setup time relative to DDR_UDQS	t_{DS}		350	ps	
DDR_DQ[7:0] input hold time relative to DDR_LDQS	t_{DH}		250	ps	
DDR_DQ[15:8] input hold time relative to DDR_UDQS	t_{DH}		250	ps	
Quarter period offset from clock edge	t_{90}	$0.25 \times t_{\text{CLK}}$	$0.25 \times t_{\text{CLK}}$		$t_{\text{CLK}} = 3.2 \text{ ns}$

FIGURE 6-7: DDR2 SDRAM OUTPUT TIMING DIAGRAM



The following table lists the AC characteristics for the DDR2 SDRAM output signals.

TABLE 6-19: DDR2 SDRAM OUTPUT SIGNAL AC CHARACTERISTICS


Parameter	Symbol	Minimum	Typical	Maximum	Unit
DDR_Ck cycle time 312.5 MHz (DDR625) ⁽¹⁾	t_{CLK}		3.20		ns

TABLE 6-19: DDR2 SDRAM OUTPUT SIGNAL AC CHARACTERISTICS (CONTINUED)

Parameter	Symbol	Minimum	Typical	Maximum	Unit
DDR_CK/CKn duty cycle		48		52	%
DDR_A, DDR_BA, DDR_CkE, DDR_nRAS, DDR_nCAS and DDR_nWE output setup time relative to DDR_CK/CKn	t_{AS}	700			ps
DDR_A, DDR_BA, DDR_CkE, DDR_nRAS, DDR_nCAS and DDR_nWE output hold time relative to DDR_CK/CKn	t_{AH}	800			ps
DDR_CK/CKn to DDR_DQS skew	t_{CLKDQS}	-400		400	ps
DDR_DQ[7:0]/DDR_LDM output setup time with relative to DDR_LDQS	t_{DS}	400			ps
DDR_DQ[15:8]/DDR_UDM output setup time with relative to DDR_UDQS	t_{DS}	400			ps
DDR_DQ[7:0]/DDR_LDM output hold time relative to DDR_DQS	t_{DH}	500			ps
DDR_DQ[15:8]/DDR_LDM output hold time relative to DDR_LDQS	t_{DH}	500			ps
DDR_DQS preamble start	t_{DQSPRE}	$0.4 \times t_{CLK}$		$0.6 \times t_{CLK}$	ps
DDR_DQS postamble end	t_{DQSPST}	$0.4 \times t_{CLK}$		$0.6 \times t_{CLK}$	ps

1. Timing reference is DDR_Ck/DDR_Ckn crossing ± 0.1 V.


The following illustration shows the test load circuit for the DDR2 outputs.

FIGURE 6-8: TEST LOAD CIRCUIT FOR DDR2 OUTPUTS

6.2.4 DDR3 SDRAM INTERFACE

This section provides the AC characteristics for the DDR2/DDR3 SDRAM interface in DDR3 mode.

The following illustration shows the DDR3 SDRAM input timing diagram.

FIGURE 6-9: DDR3 SDRAM INPUT TIMING DIAGRAM

The following table lists the AC specifications for the DDR3 SDRAM input signals.

TABLE 6-20: DDR3 SDRAM INPUT SIGNAL AC CHARACTERISTICS

Parameter	Symbol	Minimum	Maximum	Unit	Condition
Input voltage high	$V_{IH(AC)}$	$DDR\_V_{REF} + 0.175$	$V_{DD\_IODDR} + 0.3$	V	

1. These requirements are dependent on the operation frequency of the DDR SDRAM interface. Stated limits are for DDR3-800.

[illegible]

PROSTATE 3.101

[illegible]


TABLE 6-21: DDR3 SDRAM OUTPUT SIGNAL AC CHARACTERISTICS (CONTINUED)

Parameter	Symbol	Minimum	Typical	Maximum	Unit
DDR_DQ[7:0]/DDR_LDM output setup time with relative to DDR_LDQS	t_{DS}	400			ps
DDR_DQ[15:8]/DDR_UDM output setup time with relative to DDR_UDQS	t_{DS}	400			ps
DDR_DQ[7:0]/DDR_LDM output hold time relative to DDR_LDQS	t_{DH}	500			ps
DDR_DQ[15:8]/DDR_UDM output hold time relative to DDR_UDQS	t_{DH}	500			ps
DDR_xDQS preamble start	t_{DQSPRE}	$0.4 \times t_{CLK}$		$0.6 \times t_{CLK}$	ps
DDR_xDQS postamble end	t_{DQSPST}	$0.4 \times t_{CLK}$		$0.6 \times t_{CLK}$	ps

1. Timing reference is DDR_C/DDR_Ckn crossing ± 0.1 V.

The following illustration shows the test load circuit for the DDR3 outputs.

FIGURE 6-11: TEST LOAD CIRCUIT FOR DDR3 OUTPUTS




6.2.5 ENHANCED SERDES INTERFACE

All AC specifications for the Enhanced SerDes interface are compliant with the OIF-CEI version 2.0 requirements where applicable.

The Enhanced SerDes interface supports three major modes: SGMII, 2.5G, and SFP. The values in the tables in the following sections apply to modes listed in the condition column and are based on the test circuit shown in [Figure 6-2](#). The transmit and receive eye specifications in the tables relate to the eye diagrams shown in the following illustration, with the compliance load as defined in the test circuit.

FIGURE 6-12: ENHANCED SERDES TRANSIENT PARAMETERS



6.2.5.1 Enhanced SerDes Outputs

The following table provides the AC specifications for the Enhanced SerDes outputs in SGMII mode.

TABLE 6-22: ENHANCED SERDES OUTPUT AC SPECIFICATIONS IN SGMII MODE

Parameter	Symbol	Minimum	Typical	Maximum	Unit	Condition
Data rate		–100 ppm	1.25	100 ppm	Gbps	
VOD ringing compared to VS, SGMII mode	V _{RING}			±10	%	R _L = 100 Ω ± 1%.
VOD rise time and fall time, SGMII mode	t _R , t _F	100		200	ps	20% to 80% of V _S , R _L = 100 Ω ± 1%.
Differential output peak-to-peak voltage	V _{OD}			30	mV	Tx disabled.
Differential output return loss, 1000BASE-KX mode, 50 MHz to 625 MHz	RL _{TX_DIFF}	≥10			dB	R _L = 100 Ω ± 1%.
Differential output return loss, 1000BASE-KX mode, 625 MHz to 1250 MHz		10–10 × log (f/625 MHz)			dB	R _L = 100 Ω ± 1%.
Common-mode return loss, 1000BASE-KX mode	RL _{CM}	6			dB	50 MHz to 625 MHz.
Intrapair skew, SGMII mode	t _{SKEW}			20	ps	

The following table provides the AC specifications for the Enhanced SerDes outputs in 2.5G mode.

TABLE 6-23: ENHANCED SERDES OUTPUT AC SPECIFICATIONS IN 2.5G MODE

Parameter	Symbol	Minimum	Typical	Maximum	Unit	Condition
Data rate		–100 ppm	3.125	100 ppm	Gbps	
V _{OD} rise time and fall time, SGMII mode	t _R , t _F	60		130	ps	20% to 80% of V _S , R _L = 100 Ω ± 1%.
Differential output peak-to-peak voltage	V _{OD}			30	mV	Tx disabled.
Differential output return loss, 100 MHz to 625 MHz	RL _{TX_DIFF}	10			dB	R _L = 100 Ω ± 1%.
Differential output return loss, 625 MHz to 3.125 GHz	RL _{TX_DIFF}	10–10 × log (f/625 MHz)			dB	R _L = 100 Ω ± 1%.
Eye mask (T_X1)				0.175	UI	
Eye mask (T_X2)				0.390	UI	
Eye mask (T_Y1)		200			mV	
Eye mask (T_Y2)				400	mV	

6.2.5.2 Enhanced SerDes Driver Jitter Characteristics

The following table lists the jitter characteristics for the Enhanced SerDes driver in SGMII mode.

TABLE 6-24: ENHANCED SERDES DRIVER JITTER CHARACTERISTICS IN SGMII MODE

Parameter	Symbol	Maximum	Unit	Condition
Total output jitter	$t_{JIT(O)}$	192	ps	Measured according to IEEE 802.3.38.5.
Deterministic output jitter	$t_{JIT(OD)}$	80	ps	Measured according to IEEE 802.3.38.5.

6.2.5.3 Enhanced SerDes Inputs

The following table lists the AC specifications for the Enhanced SerDes inputs in SGMII mode.

TABLE 6-25: ENHANCED SERDES INPUT AC SPECIFICATIONS IN SGMII MODE

Parameter	Minimum	Typical	Maximum	Unit	Condition
Data rate	-100 ppm	1.25	100 ppm	Gbps	
Differential input return loss	10			dB	50 MHz to 625 MHz, $R_L = 100 \Omega \pm 1\%$.
Common-mode input return loss	6			dB	50 MHz to 625 MHz.

The following table lists the AC specifications for the Enhanced SerDes inputs in 2.5G mode.

TABLE 6-26: ENHANCED SERDES INPUT AC SPECIFICATIONS IN 2.5G MODE

Parameter	Symbol	Minimum	Typical	Maximum	Unit	Condition
Data rate		-100 ppm	3.125	100 ppm	Gbps	
Differential input return loss	RL_{RX_DIFF}	10			dB	100 MHz to 2.5 GHz, $R_L = 100 \Omega \pm 1\%$.
Common-mode input return loss		6			dB	100 MHz to 2.5 GHz.
Eye mask (R_X1)				0.275	UI	
Eye mask (R_Y1)		100			mV	
Eye mask (R_Y2)				800	mV	

6.2.5.4 Enhanced SerDes Receiver Jitter Tolerance

The following table lists jitter tolerances for the Enhanced SerDes receiver in SGMII mode.

TABLE 6-27: ENHANCED SERDES RECEIVER JITTER TOLERANCE IN SGMII MODE

Parameter	Symbol	Minimum	Unit	Condition
Total input jitter tolerance, 1000BASE-KX and SFP modes	$t_{JIT(I)}$	600	ps	Above 637 kHz. Measured according to IEEE 802.3 38.6.8.
Deterministic input jitter tolerance, 1000BASE-KX and SFP mode	$t_{JIT(ID)}$	370	ps	Above 637 kHz. Measured according to IEEE 802.3 38.6.8.
Cycle distortion input jitter tolerance, 100BASE-FX mode	D_{CD}	1.4	ns	Measured according to ISO/IEC 9314-3:1990.
Data-dependent input jitter tolerance, 100BASE-FX mode	D_{DJ}	2.2	ns	Measured according to ISO/IEC 9314-3:1990.

TABLE 6-27: ENHANCED SERDES RECEIVER JITTER TOLERANCE IN SGMII MODE

Parameter	Symbol	Minimum	Unit	Condition
Random input jitter tolerance, peak-to-peak, 100BASE-FX mode	R_J	2.27	ns	Measured according to ISO/IEC 9314-3:1990.

6.2.6 SERDES (SGMII) INTERFACE

In SGMII mode, the SGMII interface is compliant with Serial-GMII Specification, version 1.9.

In 1000BASE-KX mode, the SGMII interface is compliant with IEEE 802.3 clause 70.


In SFP mode, the SGMII interface is compliant with the SFP MSA standard.

In 100BASE-FX mode, the SGMII interface is compliant with IEEE 802.3 clause 26.

The rise time and fall time parameters and other transient performance specifications are defined in the following illustration. The definition of V_S is the difference between the steady state high and low voltage of the differential signal.

In addition, the signals are monotonic between 20% and 80% of V_S when loaded with $100\ \Omega \pm 1\%$.

FIGURE 6-13: SGMII TRANSIENT PARAMETERS



All SerDes driver signals comply with the conditions listed in the following table when measured with the test circuit shown in [Figure 6-2](#).

6.2.6.1 SerDes Outputs

The following table lists the AC specifications for the SerDes outputs. The values in the following table are valid for all configurations, unless otherwise specified.

TABLE 6-28: SERDES OUTPUT AC SPECIFICATIONS

Parameter	Symbol	Minimum	Maximum	Unit	Condition
V_{OD} ringing compared to V_S , SGMII mode	V_{RING}		± 10	%	$R_L = 100\ \Omega \pm 1\%$.
V_{OD} rise time and fall time, SGMII mode	t_R, t_F	100	200	ps	20% to 80% of V_S , $R_L = 100\ \Omega \pm 1\%$.
Differential output peak-to-peak voltage	V_{OD}		30	mV	Tx disabled.
Differential output return loss, 1000BASE-KX mode, 50 MHz to 625 MHz	RL_{TX_DIFF}	≥ 10		dB	$R_L = 100\ \Omega \pm 1\%$.
Differential output return loss, 1000BASE-KX mode, 625 MHz to 1250 MHz	RL_{TX_DIFF}	$10 - 10 \times \log(f/625\text{ MHz})$		dB	$R_L = 100\ \Omega \pm 1\%$.
Common-mode return loss, 1000BASE-KX mode	RL_{CM}	6		dB	50 MHz to 625 MHz.
Intrapair skew, SGMII mode	t_{SKEW}		20	ps	

6.2.6.2 SerDes Driver Jitter Characteristics

The following table lists the jitter characteristics for the SerDes driver.

TABLE 6-29: SERDES DRIVER JITTER CHARACTERISTICS

Parameter	Symbol	Maximum	Unit	Condition
Total output jitter	$t_{JIT(O)}$	192	ps	Measured according to IEEE 802.3.38.5.
Deterministic output jitter	$t_{JIT(OD)}$	80	ps	Measured according to IEEE 802.3.38.5.

6.2.6.3 SerDes Inputs

The following table lists the AC specifications for the SerDes inputs.

TABLE 6-30: SERDES INPUT AC SPECIFICATIONS

Parameter	Maximum	Unit	Condition
Differential input return loss, 1000BASE-KX mode, 50 MHz to 625 MHz	≥ 10	dB	$R_L = 100 \Omega \pm 1\%$.
Differential input return loss, 1000BASE-KX mode, 625 MHz to 1250 MHz	$10 - 10 \times \log(f/625 \text{ MHz})$	dB	$R_L = 100 \Omega \pm 1\%$.

6.2.6.4 SerDes Receiver Jitter Tolerance

The following table lists jitter tolerances for the SerDes receiver.

TABLE 6-31: SERDES RECEIVER JITTER TOLERANCE

Parameter	Symbol	Minimum	Unit	Condition
Total input jitter tolerance, 1000BASE-KX and SFP modes	$t_{JIT(I)}$	600	ps	Above 637 kHz. Measured according to IEEE 802.3ap-2007 70.7.2.1 and Annex 69A.
Deterministic input jitter tolerance, 1000BASE-KX and SFP modes	$t_{JIT(ID)}$	370	ps	Above 637 kHz. Measured according to IEEE 802.3ap-2007 70.7.2.1 and Annex 69A.
Cycle distortion input jitter tolerance, 100BASE-FX mode	D_{CD}	1.4	ns	Measured according to ISO/IEC 9314-3:1990.
Data-dependent input jitter tolerance, 100BASE-FX mode	D_{DJ}	2.2	ns	Measured according to ISO/IEC 9314-3:1990.
Random input jitter tolerance, peak-to-peak, 100BASE-FX mode	R_J	2.27	ns	Measured according to ISO/IEC 9314-3:1990.

6.2.7 PCI EXPRESS INTERFACE

The following table provides the AC specifications for the PCI Express output.

TABLE 6-32: PCI EXPRESS OUTPUT AC SPECIFICATION

Parameter	Symbol	Minimum	Typical	Maximum	Unit	Condition
Unit interval, 2.5G	UI					400 ps.
V_{OD} rise time and fall time	t_R, t_F	60		130	ps	20% to 80% of V_S , $R_L =$
Differential output voltage	V_{OD}			30	mVp-p	Tx disabled.

TABLE 6-32: PCI EXPRESS OUTPUT AC SPECIFICATION

Parameter	Symbol	Minimum	Typical	Maximum	Unit	Condition
Differential output return loss, 50 MHz to 1250 MHz	RL_{TX_DIFF}	10			dB	$R_L = 100 \Omega \pm 1\%$.
Intrapair skew	t_{SKEW}			20	ps	
De-emphasized differential output voltage (Ratio)		-3.0	-3.5	-4.0	dB	
Eye mask (T_X1)				0.175	UI	
Eye mask (T_X2)				0.390	UI	
Eye mask (T_Y1)		200			mV	
Eye mask (T_Y2)				600	mV	

The following table lists the AC specifications for the PCI Express input.

TABLE 6-33: PCI EXPRESS INPUT AC SPECIFICATIONS

Parameter	Symbol	Minimum	Maximum	Unit	Condition
Unit interval, 2.5G	UI				400 ps.
Differential input return loss	RL_{RX_DIFF}	10		dB	50 MHz to 1.25 GHz, $R_L = 100 \Omega \pm 1\%$.
Common-mode input return loss		6		dB	50 MHz to 1.25 GHz.
Eye mask (R_X1)			0.3	UI	
Eye mask (R_Y1)		85		mV	
Eye mask (R_Y2)			600	mV	


Note that the device does not accept spread spectrum modulated PCIe input.

6.2.8 MII MANAGEMENT

All AC specifications for the MII Management (MIIM) interface meet or exceed the requirements of IEEE 802.3-2002 (clause 22.2-4).

All MIIM AC timing requirements are specified relative to the input low and input high threshold levels. The following illustration shows the MIIM waveforms and required measurement points for the signals.

FIGURE 6-14: MIIM TIMING DIAGRAM



The setup time of MDIO relative to the rising edge of MDC is defined as the length of time between when the MDIO exits and remains out of the switching region and when MDC enters the switching region. The hold time of MDIO relative to the rising edge of MDC is defined as the length of time between when MDC exits the switching region and when MDIO enters the switching region.

All MIIM signals comply with the specifications in the following table. The MDIO signal requirements are requested at the pin of the device.

TABLE 6-34: MIIM TIMING SPECIFICATIONS


Parameter	Symbol	Minimum	Maximum	Unit	Condition
MDC frequency ⁽¹⁾	f	0.488	20.83	MHz	
MDC cycle time ⁽²⁾	t_C	48	2048	ns	
MDC time high	$t_{W(CH)}$	20		ns	$C_L = 50 \text{ pF}$
MDC time low	$t_{W(CL)}$	20		ns	$C_L = 50 \text{ pF}$
MDC input rise time and fall time (slave mode)	t_R, t_F		10	ns	Between $V_{IL(MAX)}$ and $V_{IH(MIN)}$
MDIO setup time to MDC on write	$t_{SU(W)}$	15		ns	$C_L = 50 \text{ pF}$
MDIO hold time from MDC on write	$t_{H(W)}$	15		ns	$C_L = 50 \text{ pF}$
MDIO setup time to MDC on read	$t_{SU(R)}$	30		ns	$C_L = 50 \text{ pF}$ on MDC
MDIO hold time from MDC on read	$t_{H(R)}$	0		ns	$C_L = 50 \text{ pF}$

1. For the maximum value, the devices support an MDC clock speed of up to 20 MHz for faster communication with the PHYs. If the standard frequency of 2.5 MHz is used, the MIIM interface is designed to meet or exceed the IEEE 802.3 requirements of the minimum MDC high and low times of 160 ns and an MDC cycle time of minimum 400 ns, which is not possible at faster speeds.
2. Calculated as $t_C = 1/f$.

6.2.9 SERIAL CPU INTERFACE (SI) MASTER MODE

All serial CPU interface (SI) timing requirements for master mode are specified relative to the input low and input high threshold levels. The following illustration shows the timing parameters and measurement points.

FIGURE 6-15: SI TIMING DIAGRAM FOR MASTER MODE



All SI signals comply with the specifications shown in the following table. The SI input timing requirements are requested at the pins of the device.

TABLE 6-35: SI TIMING SPECIFICATIONS FOR MASTER MODE

Parameter	Symbol	Minimum	Maximum	Unit	Condition
Clock frequency	f		25 ⁽¹⁾	MHz	
Clock cycle time	t_C	40		ns	
Clock time high	$t_{W(CH)}$	16		ns	
Clock time low	$t_{W(CL)}$	16		ns	

TABLE 6-35: SI TIMING SPECIFICATIONS FOR MASTER MODE (CONTINUED)

Parameter	Symbol	Minimum	Maximum	Unit	Condition
Clock rise time and fall time	t_R, t_F		10	ns	Between $V_{IL(MAX)}$ and $V_{IH(MIN)}$, $C_L = 30$ pF.
DO setup time to clock	$t_{SU(DO)}$	10		ns	
DO hold time from clock	$t_H(DO)$	10		ns	
Enable active before first clock	t_{LEAD}	10		ns	
Enable inactive after clock	t_{LAG}	5		ns	
DI setup time to clock	$t_{SU(DI)}$	22		ns	
DI hold time from clock	$t_H(DI)$	-2		ns	

1. Frequency is programmable. The startup frequency is 7 MHz.

6.2.10 SERIAL CPU INTERFACE (SI) FOR SLAVE MODE

All serial CPU interface (SI) slave mode timing requirements are specified relative to the input low and input high threshold levels. The following illustrations show the timing parameters and measurement points for SI input and output data.

FIGURE 6-16: SI INPUT DATA TIMING DIAGRAM FOR SLAVE MODE





FIGURE 6-17: SI OUTPUT DATA TIMING DIAGRAM FOR SLAVE MODE



All SI signals comply with the specifications shown in the following table. The SI input timing requirements are requested at the pins of the device.

TABLE 6-36: SI TIMING SPECIFICATIONS FOR SLAVE MODE

Parameter	Symbol	Minimum	Maximum	Unit	Condition
Clock frequency	f		25	MHz	
Clock cycle time	t_C	40		ns	


TABLE 6-36: SI TIMING SPECIFICATIONS FOR SLAVE MODE (CONTINUED)

Parameter	Symbol	Minimum	Maximum	Unit	Condition
Clock time high	$t_{W(CH)}$	16		ns	
Clock time low	$t_{W(CL)}$	16		ns	
Clock rise time and fall time	t_R, t_F		10	ns	Between $V_{IL(MAX)}$ and $V_{IH(MIN)}$.
DI setup time to clock	$t_{SU(DI)}$	4		ns	
DI hold time from clock	$t_{H(DI)}$	4		ns	
Enable active before first clock	t_{LEAD}	10		ns	
Enable inactive after clock (input cycle) ⁽¹⁾	t_{LAG1}	25		ns	
Enable inactive after clock (output cycle)	t_{LAG2}	See note ⁽²⁾		ns	
Enable inactive width	$t_{W(EH)}$	20		ns	
DO valid after clock	$t_{V(C)}$		20	ns	$C_L = 30$ pF.
DO hold time from clock	$t_{H(DO)}$	0		ns	$C_L = 0$ pF.
DO disable time ⁽³⁾	t_{DIS}		15	ns	See Table 6-18.

1. t_{LAG1} is defined only for write operations to the device, not for read operations.
2. The last rising edge on the clock is necessary for the external master to read in the data. The lag time depends on the necessary hold time on the external master data input.
3. Pin begins to float when a 300 mV change from the loaded V_{OH} or V_{OL} level occurs.

The following illustration shows the test circuit for the SI_DO disable time.

FIGURE 6-18: TEST CIRCUIT FOR SI_DO DISABLE




6.2.11 JTAG INTERFACE

All AC specifications for the JTAG interface meet or exceed the requirements of IEEE 1149.1-2001.

The following illustration shows the JTAG transmit and receive waveforms and required measurement points for the different signals.

FIGURE 6-19: JTAG INTERFACE TIMING DIAGRAM



All JTAG signals comply with the specifications in the following table. The JTAG receive signal requirements are requested at the pin of the device.


The JTAG_nTRST signal is asynchronous to the clock and does not have a setup or hold time requirement.

TABLE 6-37: JTAG INTERFACE AC SPECIFICATIONS

Parameter	Symbol	Minimum	Maximum	Unit	Condition
TCK frequency	f		10	MHz	
TCK cycle time	t_C	100		ns	
TCK high time	$t_{W(CH)}$	40		ns	
TCK low time	$t_{W(CL)}$	40		ns	
Setup time to TCK rising	t_{SU}	10		ns	
Hold time from TCK rising	t_H	10		ns	
TDO valid after TCK falling	$t_{V(C)}$		28	ns	$C_L = 10 \text{ pF}$
TDO hold time from TCK falling	$t_{H(TDO)}$	0		ns	$C_L = 0 \text{ pF}$
TDO disable time ⁽¹⁾	t_{DIS}		30	ns	
nTRST time low	$t_{W(TL)}$	30		ns	

1. The pin begins to float when a 300 mV change from the actual V_{OH}/V_{OL} level occurs.


The following illustration shows the test circuit for the TDO disable time.

FIGURE 6-20: TEST CIRCUIT FOR TDO DISABLE TIME

6.2.12 SERIAL INPUTS/OUTPUTS

This section provides the AC characteristics for the serial I/O signals: SIO_CLK, SIO_LD, SIO_DO, and SIO_DI. The SI signals are alternate function signals on the GPIO_[0:3] pins. For more information about the GPIO pin mapping, see [Section 7.2.4, General-Purpose Inputs and Outputs](#).

The serial I/O timing diagram is shown in the following illustration.

FIGURE 6-21: SERIAL I/O TIMING DIAGRAM

The following table lists the serial I/O timing specifications.

TABLE 6-38: SERIAL I/O TIMING SPECIFICATIONS

Parameter	Symbol	Minimum	Maximum	Unit
Clock frequency ⁽¹⁾	f		25	MHz
SIO_CLK clock duty cycle	t_C	40	60	%
SIO_DO valid after clock falling	$t_{V(DO)}$		6	ns
SIO_DO hold time from clock falling	$t_{H(DO)}$		6	ns
SIO_LD propagation delay from clock falling	$t_{PD(LD)}$	40		ns
SIO_LD width	$t_{W(LD)}$	10		ns
SIO_DI setup time to clock	$t_{SU(DI)}$	25		ns
SIO_DI hold time from clock	$t_{H(DI)}$	4		ns


1. The SIO clock frequency is programmable.

6.2.13 RECOVERED CLOCK OUTPUTS

This section provides the AC characteristics for the recovered clock output signals: RCVRD_CLK[9:0].

The following illustration shows the test circuit for the recovered clock output signals.

FIGURE 6-22: TEST CIRCUIT FOR RECOVERED CLOCK OUTPUT SIGNALS



The following table lists the AC specifications for the recovered clock outputs.

TABLE 6-39: RECOVERED CLOCK OUTPUT AC SPECIFICATIONS

Parameter	Symbol	Minimum	Maximum	Unit	Condition
RCVRD_CLK[9:0] clock frequency	f		125	MHz	
Clock duty cycle	t_C	40	60	%	Measured at 50% threshold.
RCVRD_CLK[9:0] rise time and fall time	t_R, t_F		1.5	ns	
Squelching delay from SGMII signal to RCVRD_CLK[9:0]			200	ns	Squelch enabled.
Squelching delay from XAUI signal to RCVRD_CLK[9:0]			200	ns	Squelch enabled.
RCVRD_CLK[9:0] peak-to-peak jitter, bandwidth from 12 kHz to 10 MHz. ⁽¹⁾			500	ps	
RCVRD_CLK[9:0] peak-to-peak jitter, bandwidth from 10 MHz to 80 MHz. ⁽¹¹⁾			500	ps	

1. Maximum jitter on the recovered signal.

6.2.14 TWO-WIRE SERIAL INTERFACE

This section provides the AC specifications for the two-wire serial interface signals TWI_SCL and TWI_SDA. The two-wire serial interface signals are alternate function signals on the GPIO_6 and GPIO_7 pins. For more information about the GPIO pin mapping, see [Table 913](#).

The two-wire serial interface signals are compatible with the Philips I²C-BUS specifications, except for the minimum rise time and fall time requirements for fast mode.

FIGURE 6-23: TWO-WIRE SERIAL READ TIMING DIAGRAM





FIGURE 6-24: TWO-WIRE SERIAL WRITE TIMING DIAGRAM

The data in this table assumes that the software-configurable two-wire interface timing parameters are set to valid values for the selected speed. For more information about configuring these registers for the selected speed, see SS_SCL_HCNT, SS_SCL_LCNT, FS_SCL_HCNT, and FS_SCL_LCNT in [Section 6, Registers](#).

TABLE 6-40: TWO-WIRE SERIAL INTERFACE AC SPECIFICATIONS

		Standard Mode		Fast Mode			
Parameter	Symbol	Minimum	Maximum	Minimum	Maximum	Unit	Condition
TWI_SCL clock frequency	f		100		400	kHz	
TWI_SCL low period	t_{LOW}	4.7		1.3		μs	
TWI_SCL high period	t_{HIGH}	4.0		0.6		μs	
TWI_SCL and TWI_SDA rise time			1000		300	ns	
TWI_SCL and TWI_SDA fall time			300		300	ns	
TWI_SDA setup time to TWI_SCL fall	t_{SU_DAT}	250		100	300	ns	
TWI_SDA hold time to TWI_SCL fall ⁽¹⁾	t_{HD_DAT}	300	3450	300	900	ns	300 ns delay enabled in ICPU_CFG::TWI_DELAY register.
Setup time for repeated START condition	t_{SU_STA}	4.7		0.6		μs	
Hold time after repeated START condition	t_{HD_STA}	4.0		0.6		μs	
Bus free time between STOP and START conditions	t_{BUF}	4.7		1.3		μs	
Clock to valid data out ⁽²⁾	t_{VD_DAT}	300		300		ns	
Pulse width of spike suppressed by input filter on TWI_SCL or TWI_SDA		0	5	0	5	ns	

1. An external device must provide a hold time of at least 300 ns for the TWI_SDA signal to bridge the undefined region of the falling edge of the TWI_SCL signal.
2. Some external devices may require more data in hold time (target device's t_{HD_DAT}) than what is provided by t_{VD_DAT} , for example, 300 ns to 900 ns. The minimum value of t_{VD_DAT} is adjustable; the value given represents the recommended minimum value, which is enabled in ICPU_CFG::TWI_DELAY.TWI_DELAY_ENABLE.

6.2.15 1588 PPS TIME TICK OUTPUT

This section provides the AC specifications for the 1588 PPS0 and 1588 PPS1 time tick output signals. The 1588 PPS0 signal is an alternate function signal on the GPIO_31 pin, and 1588 PPS1 is an alternate function signal on the GPIO_16 pin. For more information about the GPIO pin mapping, see [Table 913](#).

TABLE 6-41: 1588 PPS TIME TICK OUTPUT AC SPECIFICATIONS

Parameter	Symbol	Minimum	Maximum	Unit	Condition
1588 PPS frequency ⁽¹⁾	f		25	MHz	
Clock duty cycle ⁽¹⁾		45	55	%	Duty cycle programmed to Measured at 50% threshold.
1588 PPS rise time and fall time	t_R, t_F	1		ns	20% to 80% threshold.
1588 PPS peak-to-peak jitter, bandwidth between 12 kHz and 10 MHz			200	ps	
1588 PPS peak-to-peak jitter, bandwidth between 10 MHz and 80 MHz			200	ps	

1. Frequency and duty cycle are programmable.

6.3 Current and Power Consumption

This section provides the current and power consumption requirements for the VSC7416-01 device.

6.3.1 CURRENT CONSUMPTION

This section provides the operating current consumption parameters for the VSC7416-01 device.

Typical current consumption values are over nominal supply settings at 25 °C case temperature, and maximum traffic load. Maximum current consumption values are over worst-case process, temperature, and supply settings, and maximum traffic load.

The following table lists the typical and maximum operating current consumption values for the VSC7416-01 device.

TABLE 6-42: OPERATING CURRENT

Parameter	Symbol	Typical	Maximum	Unit	Condition
V _{DD} operating current	I _{DD}	0.8	2.2	A	V _{TYP} = 1.0 V.
V _{DD_A} operating current	I _{DD_A}	0.35	0.50	A	V _{TYP} = 1.0 V.
V _{DD_VS} operating current	I _{DD_VS}	0.16	0.20	A	V _{TYP} = 1.0 V or 1.2 V.
V _{DD_IODDR} operating current	I _{DD_IODDR}	0.10	0.165	A	V _{TYP} = 1.5 V or 1.8 V. On-die termination disabled.
V _{DD_IO} and V _{DD_H} operating current ⁽¹⁾	I _{DD_IO}	0.05	0.05	A	With unloaded GPIOs.

1. Typical current for V_{DD_H} is 0.03 A.

6.3.2 POWER CONSUMPTION

This section provides the power consumption parameters for the VSC7416-01 device, based on current consumption and with DDR2 on-die termination disabled.

Typical power consumption values are over nominal supplies and 25 °C case temperature. Maximum power consumption values are over maximum temperature and all supplies at maximum voltages.

The following table lists the typical and maximum power consumption values for the VSC7416-01 device.

TABLE 6-43: POWER CONSUMPTION

Parameter	Typical	Maximum	Unit
Power consumption, SGMII in LVDS mode $V_{DD\_VS} = 1.0\text{ V}$	1.6	3.5	W
Power consumption, SGMII in high-drive mode $V_{DD\_VS} = 1.2\text{ V}$	1.6	3.5	W

6.3.3 POWER SUPPLY SEQUENCING

During power on and off, V_{DD_A} and V_{DD_VS} must never be more than 300 mV above V_{DD} .

V_{DD_VS} must be powered, even if the associated interface is not used. These power supplies must not remain at ground or left floating.

A maximum delay of 100 ms from V_{DD_IODDR} to V_{DD} is recommended. There is no requirement from V_{DD} to V_{DD_IODDR} .

There are no sequencing requirements for V_{DD_IO} .

The nReset and JTAG_nTRST inputs must be held low until all power supply voltages have reached their recommended operating condition values.

6.4 Operating Conditions

The following table lists the recommended operating conditions.

TABLE 6-44: RECOMMENDED OPERATING CONDITIONS

Parameter	Symbol	Minimum	Typical	Maximum	Unit
Power supply voltage for core supply	V_{DD}	0.95	1.00	1.05	V
Power supply voltage for analog circuits	V_{DD_A}	0.95	1.00	1.05	V
Power supply voltage for SerDes, Enhanced SerDes, and PCIe interfaces, 1.0 V	V_{DD_VS}	0.95	1.00	1.05	V
Power supply voltage for SerDes, Enhanced SerDes, and PCIe interfaces, 1.2 V	V_{DD_VS}	1.14	1.20	1.26	V
Power supply voltage for DDR2/3 interface in DDR2 mode	V_{DD_IODDR}	1.70	1.80	1.90	V
Power supply voltage for DDR2/3 interface in DDR3 mode	V_{DD_IODDR}	1.425	1.50	1.575	
Power supply voltage for MIIM, SI, GPIO, and miscellaneous I/O	V_{DD_IO} , V_{DD_H}	2.38	2.50	2.62	V
Operating temperature ⁽¹⁾	T	-40		125	°C

1. Minimum specification is ambient temperature, and the maximum is junction temperature.

6.5 Stress Ratings

Warning Stresses listed in the following table may be applied to devices one at a time without causing permanent damage. Functionality at or exceeding the values listed is not implied. Exposure to these values for extended periods may affect device reliability.

TABLE 6-45: STRESS RATINGS

Parameter	Symbol	Minimum	Maximum	Unit
Power supply voltage for core supply	V_{DD}	-0.3	1.10	V
Power supply voltage for analog circuits	V_{DD_A}	-0.3	1.10	V

TABLE 6-45: STRESS RATINGS (CONTINUED)

Parameter	Symbol	Minimum	Maximum	Unit
Power supply voltage for SerDes, Enhanced SerDes, and PCIe interfaces	V_{DD_VS}	-0.3	1.32	V
Power supply voltage for DDR2/DDR3 interface	V_{DD_IODDR}	-0.3	1.98	V
Power supply voltage for MIIM, SI, GPIO, and miscellaneous I/O	V_{DD_IO} , V_{DD_H}	-0.3	2.75	V
Storage temperature	T_S	-55	125	°C
Electrostatic discharge voltage, charged device model	V_{ESD_CDM}	-250	250	V
Electrostatic discharge voltage, human body model	V_{ESD_HBM}	See note ⁽¹⁾		V

1. This device has completed all required testing as specified in the JEDEC standard JESD22-A114, *Electrostatic Discharge (ESD) Sensitivity Testing Human Body Model (HBM)*, and complies with a Class 2 rating. The definition of Class 2 is any part that passes an ESD pulse of 2000 V, but fails an ESD pulse of 4000 V.

Warning This device can be damaged by electrostatic discharge (ESD) voltage. Microchip recommends that all integrated circuits be handled with appropriate precautions. Failure to observe proper handling and installation procedures may adversely affect reliability of the device.

7.0 PIN DESCRIPTIONS


The VSC7416-01 device has 324 pins, which are described in this section.

The pin information is also provided as an attached Microsoft Excel file, so that you can copy it electronically. In Adobe Reader, double-click the attachment icon.

7.1 Pin Diagram

The following illustration is a representation of the VSC7416-01 device, as seen from the top view looking through the device.

FIGURE 7-1: PACKAGE DRAWING: TOP VIEW



7.2 Pins by Function

This section contains the functional pin descriptions for the VSC7416-01 device. The following table lists the definitions for the pin type symbols.

TABLE 7-1: PIN TYPE SYMBOL DEFINITIONS

Symbol	Pin Type	Description
ABIAS	Analog bias	Analog bias pin.
DIFF	Differential	Differential signal pair.
I	Input	Input signal.
O	Output	Output signal.
I/O	Bidirectional	Bidirectional input or output signal.
A	Analog input	Analog input for sensing variable voltage levels.
PD	Pull-down	On-chip pull-down resistor to VSS.
PU	Pull-up	On-chip pull-up resistor to VDD_IO.
3V		3.3 V-tolerant.
OZ	3-state output	Output.
ST	Schmitt-trigger	Input has Schmitt-trigger circuitry.

TABLE 7-1: PIN TYPE SYMBOL DEFINITIONS (CONTINUED)

Symbol	Pin Type	Description
TD	Termination differential	Internal differential termination.

7.2.1 ANALOG BIAS SIGNALS

The following table lists the pins associated with the analog bias signals.

TABLE 7-2: ANALOG BIAS PINS

Name	Type	Description
SerDes_Rext_[1:0]	A	Analog bias calibration. Connect an external $626\ \Omega \pm 1\%$ resistor between SerDes_Rext_1 and SerDes_Rext_0.

7.2.2 DDR2/DDR3 SDRAM INTERFACE

The following table lists the pins associated with the DDR2 or DDR3 SDRAM interface.

TABLE 7-3: DDR2/DDR3 SDRAM PINS

Name	Type	Description
DDR_A[15:0]	O	SDRAM address outputs. Provide row and column addresses to the SDRAM.
DDR_BA[2:0]	O	SDRAM bank address outputs. DDR_BA[2:0] define to which bank an ACTIVE, READ, WRITE, or PRECHARGE command is applied.
DDR_CK DDR_CKn	0, Diff	SDRAM differential clock. Differential clock to external SDRAM. DDR_CK is the true part of the differential signal. DDR_nCK is the complement part.
DDR_CKE	O	SDRAM clock enable. 0: Disables clock in external SDRAM. 1: Enables clock in external SDRAM.
DDR_DQ[15:0]	I/O	SDRAM data bus.
DDR_LDM DDR_UDM	O	SDRAM data mask outputs. DDR_LDM and DDR_UDM are mask signals for data written to the SDRAM. DDR_LDM corresponds to data on DDR_DQ[7:0], and DDR_UDM corresponds to data on DDR_DQ[15:8].
DDR_nCAS DDR_nRAS DDR_nWE	O	SDRAM command outputs. DDR_nCAS, DDR_nRAS, and DDR_nWE (along with DDR_ODT) define the command being entered.
DDR_nCS[1:0]	O	SDRAM chip select.
DDR_ODT[1:0]	O	Control signals for the attached DDR2/DDR3 SDRAM devices on-die termination.
DDR_Rext	ABIAS	External DDR impedance calibration. Connect the pin through an external $240\ \Omega \pm 1\%$ resistor to ground.
DDR_LDQS/DDR_UDQS DDR_LDQSn/DDR_UDQSn	I/O, Diff	SDRAM differential data strobes. Bidirectional differential signal that follows data direction. Used by SDRAM to capture data on write operations. Edge-aligned with data from SDRAM during read operations and used by the device to capture data.

TABLE 7-3: DDR2/DDR3 SDRAM PINS (CONTINUED)

Name	Type	Description
DDR_Vref[1:0]	ABIAS	Input reference voltage. Provides the input switching reference voltage to the SSTL DDR signals.

7.2.3 ENHANCED SERDES INTERFACE

The following pins are associated with the Enhanced SerDes interface.

TABLE 7-4: ENHANCED SERDES INTERFACE PINS

Name	Type	Description
SerDes_E[1:0]_RxP, N	I, Diff, TD	Differential Enhanced SerDes data inputs.
SerDes_E[1:0]_TxP, N	O, Diff	Differential Enhanced SerDes data outputs.

7.2.4 GENERAL-PURPOSE INPUTS AND OUTPUTS

Most of the GPIO pins have overlaid (alternative) functions that can be enabled through the replicated GPIO_ALT registers. For more information, see DEVCPU_GCB:GPIO:GPIO_ALT in [Section 6, Registers](#).

To enable a particular GPIO pin with the alternate function, set the G_ALT[n] register field in the replicated registers as follows:

- Overlaid mode 1, set GPIO_ALT[1][n], GPIO_ALT[0][n] = 1.
- Overlaid mode 2, set GPIO_ALT[1][n], GPIO_ALT[0][n] = 2.
- Overlaid mode 3, set GPIO_ALT[1][n], GPIO_ALT[0][n] = 3.
- Normal GPIO mode, set GPIO_ALT[1][n], GPIO_ALT[0][n] = 0.

When the MIIM slave mode is enabled through the VCore_CFG strapping pins, specific GPIO pins are overtaken and used for the MIIM slave interface.

During reset, the VCore_CFG interface is sampled and used for VCore configuration. after reset for the device is released these GPIOs can be used for output or inputs. For more information, see [Section 3.1, VCore Configurations](#).

The following table maps the GPIO pins and available overlaid functions.

TABLE 7-5: GPIO PIN MAPPING

Name	Overlaid Function 1	Overlaid Function 2	Overlaid Function 3	Interfaces	Type
GPIO_0	SIO_CLK				I/O, PU, ST, 3V
GPIO_1	SIO_LD				I/O, PU, ST, 3V
GPIO_2	SIO_DO				I/O, PU, ST, 3V
GPIO_3	SIO_DI				I/O, PU, ST, 3V
GPIO_4	TACHO				I/O, PU, ST, 3V
GPIO_5	PWM				I/O, PU, ST, 3V
GPIO_6	TWI_SCL				I/O, PU, ST, 3V
GPIO_7	TWI_SDA				I/O, PU, ST, 3V
GPIO_8	SI_nEN1				I/O, PU, ST, 3V
GPIO_9	SI_nEn2	MDC1			I/O, PU, ST, 3V
GPIO_10	SI_nEn3	MDIO1			I/O, PU, ST, 3V
GPIO_11	SFP0_SD	MDC0	TWI_SCL_M0		I/O, PU, ST, 3V
GPIO_12	SFP1_SD	MDIO1	TWI_SCL_M1		I/O, PU, ST, 3V
GPIO_13	SFP2_SD	UART2_TX	TWI_SCL_M2		I/O, PU, ST, 3V
GPIO_14	SFP3_SD	UART2_RX	TWI_SCL_M3	MIIM_SLV_ADD R	I/O, PU, ST, 3V
GPIO_15	SFP4_SD	1588 LD1/ST1	TWI_SCL_M4	MIIM_SLV_MDC	I/O, PU, ST, 3V

TABLE 7-5: GPIO PIN MAPPING (CONTINUED)

Name	Overlaid Function 1	Overlaid Function 2	Overlaid Function 3	Interfaces	Type
GPIO_16	SFP5_SD	1588 PPS1	TWI_SCL_M5		I/O, PU, ST, 3V
GPIO_17	SFP6_SD	PCIe_WAKE	TWI_SCL_M6		I/O, PU, ST, 3V
GPIO_18	SFP7_SD		TWI_SCL_M7		I/O, PU, ST, 3V
GPIO_19	SFP8_SD		TWI_SCL_M8		I/O, PU, ST, 3V
GPIO_20	SFP9_SD		TWI_SCL_M9		I/O, PU, ST, 3V
GPIO_21	SFP10_SD		TWI_SCL_M10		I/O, PU, ST, 3V
GPIO_22				VCore_CFG0	I/O, PU, ST, 3V
GPIO_23				VCore_CFG1	I/O, PU, ST, 3V
GPIO_24				VCore_CFG2	I/O, PU, ST, 3V
GPIO_25				VCore_CFG3	I/O, PU, ST, 3V
GPIO_26	UART_TX				I/O, PU, ST, 3V
GPIO_27	UART_RX				I/O, PU, ST, 3V
GPIO_28	IRQ0				I/O, PU, ST, 3V
GPIO_29	IRQ1				I/O, PU, ST, 3V
GPIO_30	1588 LD0/ST0				I/O, PU, ST, 3V
GPIO_31	1588 PPS0				I/O, PU, ST, 3V

7.2.5 JTAG INTERFACE

The following table lists the pins associated with the JTAG interface. The JTAG interface can be connected to the boundary scan TAP controller or the internal VCore-III TAP controller for software debug as described under the JTAG_ICE_nEn signal.

The JTAG signals are not 5 V tolerant.

TABLE 7-6: JTAG INTERFACE PINS

Name	Type	Description
JTAG_ICE_nEN	I, PU, ST, 3V	0: Enables the VCore-III JTAG debug interface over the JTAG interface pins. 1: Enables normal JTAG I/O over the JTAG interface.
JTAG_TCK	I, PU, ST, 3V	JTAG clock.
JTAG_TDI	I, PU, ST, 3V	JTAG test data in.
JTAG_TDO	OZ, 3V	JTAG test data out.
JTAG_TMS	I, PU, ST, 3V	JTAG test mode select.
JTAG_TRST	I, PD, ST, 3V	JTAG test reset, active low. For normal device operation, JTAG_TRST should be pulled low.

7.2.6 MISCELLANEOUS SIGNALS

The following table lists the pins associated with a particular interface or facility on the device.

TABLE 7-7: MISCELLANEOUS PINS

Name	Type	Description
nRESET	I, PU, ST, 3V	Global device reset, active low.
Reserved_0	I, PD, 3V	Tie to V _{DD_IO} .
Reserved_1	A	Tie to V _{SS} .

TABLE 7-7: MISCELLANEOUS PINS (CONTINUED)

Name	Type	Description
Reserved_[12:11]	I	Tie to V_{DDA} .
Reserved_[16:15]	I	
Reserved_[20:19]	I	
Reserved_[24:23]	I	
Reserved_[4:2]	O	Leave floating.
Reserved_[8:7]	O	
Reserved_[10:9]	O	
Reserved_[14:13]	O	
Reserved_[18:17]	O	
Reserved_[22:21]	O	
Test_Mode	I, PD, ST, 3V	Tie to V_{SS} .
THERMDA	A	Thermal diode anode (p-junction).
THERMDC_VSS	A	Thermal diode cathode (n-junction). Connected on-die to V_{SS} .

7.2.7 PCI EXPRESS INTERFACE

The following table lists the PCI Express (PCIe) pins.

TABLE 7-8: PCIE INTERFACE PINS

Name	Type	Description
PCIE_RxN PCIE_RxP	I, Diff, TD	Differential PCIe data inputs.
PCIE_TxN PCIE_TxP	O, Diff	Differential PCIe data outputs.

7.2.8 POWER SUPPLIES AND GROUND

The following table lists the power supply and ground pins.

TABLE 7-9: POWER SUPPLY AND GROUND PINS

Name	Type	Description
VDD	Power	1.0 V power supply voltage for core.
VDDA	Power	1.0 V power supply voltage for analog circuits.
VDD_H	Power	2.5 V reference supply for SerDes and PLL.
VDD_IO	Power	2.5 V power supply for GPIOs, RefClk_Sel[2:0], RefClk2_Sel[2:0], PLL2_Enable, RCVRD_CLK[9:0], JTAG interface, serial CPU interface, and miscellaneous I/Os.
VDD_IODDR	Power	1.8 V (DDR2) or 1.5 V (DDR3) power supply for DDR interface.
VDD_VS	Power	1.0 V or 1.2 V power supply for SerDes, Enhanced SerDes, and PCIe interfaces.
VSS	Ground	Ground reference.

7.2.9 SERDES INTERFACE

The following pins are associated with the SerDes (SGMII) interface.

TABLE 7-10: SERDES INTERFACE PINS

Name	Type	Description
SerDes[3:0]_RxN SerDes[3:0]_RxP	I, Diff, TD	Differential SerDes data inputs.
SerDes[3:0]_TxN SerDes[3:0]_TxP	O, Diff	Differential SerDes data outputs.

7.2.10 SERIAL CPU INTERFACE

The serial CPU interface (SI) can be used as a serial slave, master, or boot interface.

As a slave interface, it allows external CPUs to access internal registers. As a master interface, it allows the device to access external devices using a programmable protocol. The serial CPU interface also allows the internal VCore-III CPU system to boot from an attached serial memory device when the VCore_CFG signals are set appropriately.

The following table lists the pins associated with the serial CPU interface (SI).

TABLE 7-11: SERIAL CPU INTERFACE PINS

Name	Type	Description
SI_Clk	I/O, 3V	Slave mode: Input receiving serial interface clock from external master. Master mode: Output controlled directly by software through register bit. Boot mode: Output driven with clock to external serial memory device.
SI_DI	I, 3V	Slave mode: Input receiving serial interface data from external master. Master mode: Input directly read by software from register bit. Boot mode: Input boot data from external serial memory device.
SI_DO	OZ, 3V	Slave mode: Output transmitting serial interface data to external master. Master mode: Output controlled directly by software through register bit. Boot mode: No function.
SI_nEn SI_nEn[3:1] ⁽¹⁾	I/O, 3V	Slave mode: Input used to enable SI slave interface. 0 = Enabled. 1 = Disabled. Master mode: Output controlled directly by software through register bit. Boot mode: Output driven while booting from EEPROM or serial flash to internal VCore-III CPU system. Released when booting is completed.

¹Available as an alternate function on the GPIO_8, GPIO_9, and GPIO_10 pins. For more information about GPIO pin mapping, see [Table 7-5](#).

7.2.11 NETWORK PROCESSOR INTERFACE (NPI)

The following table lists the 1G SGMII network processor interface (NPI) pins.

TABLE 7-12: SGMII NETWORK PROCESSOR INTERFACE PINS

Name	Type	Description
SGMII_NPI_RxN SGMII_NPI_RxP	I, Diff, TD	Differential NPI data inputs.
SGMII_NPI_TxN SGMII_NPI_TxP	O, Diff	Differential NPI data outputs.

7.2.12 SYSTEM CLOCK INTERFACE

The following table lists the pins associated with the system clock interface.

TABLE 7-13: SYSTEM CLOCK INTERFACE PINS

Name	Type	Description
CLKOUT2_N CLKOUT2_P	O, Diff	250 MHz 1588 reference clock output from LCPLL2.

TABLE 7-13: SYSTEM CLOCK INTERFACE PINS (CONTINUED)

Name	Type	Description
PLL2_Enable	I, PD, ST, 3V	Selects the source for the core and IEEE 1588 clocks. 0: RefClk. 1: RefClk2.
RCVRD_CLK[9:0]	OZ, 3V	The output clock frequency can be between 25 MHz and 125 MHz, based on the selected active recovered media programmed for this pin and the divider configuration. These pins are not active when nReset is asserted. Clock outputs can be enabled or disabled from registers. When disabled, the pin is held low.
RefClk_N RefClk_P RefClk2_N RefClk2_P	I, Diff	Reference clock input. The input can be either differential or single-ended. In differential mode, RefClk_P is the true part of the differential signal, and RefClk_N is the complement part of the differential signal. In single-ended mode, RefClk_P is used as single-ended LVTTTL input, and the RefClk_N should be pulled to V _{DD_A} . Required applied frequency depends on RefClk_Sel[2:0] input state. See following description for RefClk_Sel[2:0] pins.
RefClk_Sel[2:0] RefClk2_Sel[2:0]	I, PD, 3V	Reference clock frequency selection. 0: Connect to pull-down or leave floating. 1: Connect to pull-up to V _{DD_IO} . Bit setting: 000: 125 MHz (default). 001: 156.25 MHz. 010: 250 MHz. 011: Reserved. 100: 25 MHz. 101: Reserved. 110: Reserved. 111: Reserved.
1588 LD0/ST0 ⁽¹⁾ 1588 LD1/ST1 ⁽²⁾	I/O, 3V	These pins can be used to load the internal 1588 master timer with a new value or to store the current internal 1588 master timer.
1588 PPS0 ⁽³⁾ 1588 PPS1 ⁽⁴⁾	I/O, 3V	These pins can be programmed independently to either output or input. The pins can be used as either an input pulse for synchronization of the internal 1588 master timer or as programmable divided-frequency outputs from the internal 1588 master timer. The programmable divided frequency is between 25 MHz and 1 pulse per second. The programmed output signals duty cycle depends on the programmed divider factor.

1Available as an alternate function on the GPIO_30 pin.

2Available as an alternate function on the GPIO_15 pin.

3Available as an alternate function on the GPIO_31 pin.

4Available as an alternate function on the GPIO_16 pin.

VSC7416-01

7.3 Pins by Number

This section provides a numeric list of the VSC7416-01 pins.

A1	NC_1	C6	Reserved_12	E11	VSS_12
A2	VSS_1	C7	SGMII_NPI_RxP	E12	VSS_13
A3	Reserved_18	C8	PCIE_RxP	E13	VSS_14
A4	Reserved_8	C9	VDD_IO_3	E14	VDD_IODDR_5
A5	Reserved_14	C10	RCVRD_CLK3	E15	DDR_nCS1
A6	Reserved_10	C11	RCVRD_CLK6	E16	DDR_ODT1
A7	SGMII_NPI_TxP	C12	RCVRD_CLK9	E17	DDR_nWE
A8	PCIE_TxP	C13	RefClk2_P	E18	DDR_CKE
A9	VDD_IO_1	C14	VDD_IODDR_3	F1	VDDA_1
A10	RCVRD_CLK1	C15	DDR_UDQSn	F2	VDDA_2
A11	RCVRD_CLK4	C16	DDR_UDQS	F3	VDDA_3
A12	RCVRD_CLK7	C17	DDR_UDM	F4	VDDA_4
A13	CLKOUT2_P	C18	DDR_DQ14	F5	VDDA_5
A14	VDD_IODDR_1	D1	VDD_VS_1	F6	VDDA_6
A15	Reserved_4	D2	VDD_VS_2	F7	VDDA_7
A16	Reserved_2	D3	Reserved_19	F8	VDDA_8
A17	Reserved_3	D4	RefClk_N	F9	VDDA_9
A18	NC_2	D5	Reserved_15	F10	VDDA_10
B1	VSS_2	D6	Reserved_11	F11	VDDA_11
B2	VSS_3	D7	SGMII_NPI_RxN	F12	VDDA_12
B3	Reserved_17	D8	PCIE_RxN	F13	VSS_15
B4	Reserved_7	D9	RCVRD_CLK0	F14	VDD_IODDR_6
B5	Reserved_13	D10	VDD_IO_4	F15	DDR_BA2
B6	Reserved_9	D11	VDD_IO_5	F16	DDR_BA0
B7	SGMII_NPI_TxN	D12	VDD_IO_6	F17	DDR_BA1
B8	PCIE_TxN	D13	RefClk2_N	F18	DDR_A10
B9	VDD_IO_2	D14	VDD_IODDR_4	G1	Reserved_22
B10	RCVRD_CLK2	D15	DDR_DQ9	G2	Reserved_21
B11	RCVRD_CLK5	D16	DDR_DQ11	G3	Reserved_24
B12	RCVRD_CLK8	D17	DDR_DQ12	G4	Reserved_23
B13	CLKOUT2_N	D18	DDR_Vref	G5	VDDA_13
B14	VDD_IODDR_2	E1	VSS_6	G6	VSS_16
B15	DDR_DQ13	E2	VSS_7	G7	VSS_17
B16	DDR_DQ10	E3	VDD_VS_3	G8	VSS_18
B17	DDR_DQ8	E4	VSS_8	G9	VSS_19
B18	DDR_DQ15	E5	VDD_VS_4	G10	VSS_20
C1	VSS_4	E6	VSS_9	G11	VSS_21
C2	VSS_5	E7	VDD_VS_5	G12	VSS_22
C3	Reserved_20	E8	VSS_10	G13	VSS_23
C4	RefClk_P	E9	VDD_VS_6	G14	VDD_IODDR_7
C5	Reserved_16	E10	VSS_11	G15	DDR_A1

Pins by number (*continued*)

G16	DDR_A3	K7	VSS_35	M16	DDR_DQ5
G17	DDR_A5	K8	VSS_36	M17	DDR_DQ2
G18	DDR_A7	K9	VSS_37	M18	DDR_DQ0
H1	SerDes_E1_TxP	K10	VSS_38	N1	SerDes_E0_TxP
H2	SerDes_E1_TxN	K11	VSS_39	N2	SerDes_E0_TxN
H3	SerDes_E1_RxP	K12	VSS_40	N3	SerDes_E0_RxP
H4	SerDes_E1_RxN	K13	VSS_41	N4	SerDes_E0_RxN
H5	VSS_24	K14	VDD_IODDR_10	N5	VSS_49
H6	VDD_VS_7	K15	DDR_A4	N6	VDD_VS_10
H7	VDD_1	K16	DDR_A2	N7	VSS_50
H8	VDD_2	K17	DDR_A0	N8	VSS_51
H9	VDD_3	K18	DDR_nCAS	N9	VSS_52
H10	VDD_4	L1	SerDes1_TxP	N10	VSS_53
H11	VDD_5	L2	SerDes1_TxN	N11	VSS_54
H12	VDD_6	L3	SerDes1_RxP	N12	VSS_55
H13	VSS_25	L4	SerDes1_RxN	N13	VSS_56
H14	VDD_IODDR_8	L5	VSS_42	N14	VDD_IODDR_13
H15	DDR_A9	L6	VDD_VS_8	N15	DDR_DQ7
H16	DDR_A12	L7	VDD_7	N16	DDR_LDQSn
H17	DDR_CK	L8	VDD_8	N17	DDR_LDQS
H18	DDR_CKn	L9	VDD_9	N18	DDR_LDM
J1	SerDes3_TxP	L10	VDD_10	P1	VDD_VS_11
J2	SerDes3_TxN	L11	VDD_11	P2	VDD_VS_12
J3	SerDes3_RxP	L12	VDD_12	P3	VDD_VS_13
J4	SerDes3_RxN	L13	VSS_43	P4	VDD_VS_14
J5	VDDA_14	L14	VDD_IODDR_11	P5	VDD_VS_15
J6	VSS_26	L15	DDR_nCS0	P6	VSS_57
J7	VSS_27	L16	DDR_ODT0	P7	VDD_IO_7
J8	VSS_28	L17	DDR_nRAS	P8	VDD_IO_8
J9	VSS_29	L18	DDR_A14	P9	VDD_IO_9
J10	VSS_30	M1	SerDes0_TxP	P10	VDD_IO_10
J11	VSS_31	M10	VSS_46	P11	VDD_IO_11
J12	VSS_32	M2	SerDes0_TxN	P12	VDD_IO_12
J13	VSS_33	M3	SerDes0_RxP	P13	VDD_IO_13
J14	VDD_IODDR_9	M4	SerDes0_RxN	P14	VSS_58
J15	DDR_A11	M5	VDDA_16	P15	DDR_DQ6
J16	DDR_A13	M6	VDD_VS_9	P16	DDR_DQ3
J17	DDR_A8	M7	VSS_44	P17	DDR_DQ1
J18	DDR_A6	M8	VDD_13	P18	DDR_DQ4
K1	SerDes2_TxP	M9	VSS_45	R1	SerDes_Rext_0
K2	SerDes2_TxN	M11	VDD_14	R2	SerDes_Rext_1
K3	SerDes2_RxP	M12	VSS_47	R3	VSS_59
K4	SerDes2_RxN	M13	VSS_48	R4	VSS_60
K5	VDDA_15	M14	VDD_IODDR_12	R5	GPIO_8
K6	VSS_34	M15	DDR_A15	R6	GPIO_12

VSC7416-01

R7	GPIO_16
R8	GPIO_20
R9	GPIO_24
R10	GPIO_28
R11	SI_nEn
R12	JTAG_TDO
R13	JTAG_TRST
R14	VDD_IO_14
R15	VSS_61
R16	VSS_62
R17	VSS_63
R18	DDR_Rext
T1	THERMDA
T2	VSS_64
T3	GPIO_2
T4	GPIO_5
T5	GPIO_9
T6	GPIO_13
T7	GPIO_17
T8	GPIO_21
T9	GPIO_25
T10	GPIO_29
T11	SI_DO
T12	JTAG_TDI
T13	RefClk2_Sel2
T14	PLL2_Enable
T15	VDD_IO_15
T16	VDD_IO_16
T17	VDD_IO_17
T18	VDD_IO_18
U1	THERMDC_VSS
U2	GPIO_0
U3	GPIO_3
U4	GPIO_6
U5	GPIO_10
U6	GPIO_14
U7	GPIO_18
U8	GPIO_22
U9	GPIO_26
U10	GPIO_30
U11	SI_DI
U12	JTAG_TCK
U13	RefClk2_Sel1
U14	RefClk_Sel2
U15	RefClk_Sel0

U16	Reserved_0
U17	Test_Mode
U18	VSS_65
V1	NC_3
V2	GPIO_1
V3	GPIO_4
V4	GPIO_7
V5	GPIO_11
V6	GPIO_15
V7	GPIO_19
V8	GPIO_23
V9	GPIO_27
V10	GPIO_31
V11	SI_Clk
V12	JTAG_TMS
V13	RefClk2_Sel0
V14	RefClk_Sel1
V15	nRESET
V16	JTAG_ICE_nEn
V17	Reserved_1
V18	NC_4

VSC7416-01

7.4 Pins by Name

This section provides an alphabetical list of the VSC7416-01 pins.

CLKOUT2_N	B13	DDR_LDQS	N17	GPIO_27	V9
CLKOUT2_P	A13	DDR_LDQSn	N16	GPIO_28	R10
DDR_A0	K17	DDR_nCAS	K18	GPIO_29	T10
DDR_A1	G15	DDR_nCS0	L15	GPIO_30	U10
DDR_A2	K16	DDR_nCS1	E15	GPIO_31	V10
DDR_A3	G16	DDR_nRAS	L17	JTAG_ICE_nEn	V16
DDR_A4	K15	DDR_nWE	E17	JTAG_TCK	U12
DDR_A5	G17	DDR_ODT0	L16	JTAG_TDI	T12
DDR_A6	J18	DDR_ODT1	E16	JTAG_TDO	R12
DDR_A7	G18	DDR_Rext	R18	JTAG_TMS	V12
DDR_A8	J17	DDR_UDM	C17	JTAG_TRST	R13
DDR_A9	H15	DDR_UDQS	C16	NC_1	A1
DDR_A10	F18	DDR_UDQSn	C15	NC_2	A18
DDR_A11	J15	DDR_Vref	D18	NC_3	V1
DDR_A12	H16	GPIO_0	U2	NC_4	V18
DDR_A13	J16	GPIO_1	V2	nRESET	V15
DDR_A14	L18	GPIO_2	T3	PCIE_RxN	D8
DDR_A15	M15	GPIO_3	U3	PCIE_RxP	C8
DDR_BA0	F16	GPIO_4	V3	PCIE_TxN	B8
DDR_BA1	F17	GPIO_5	T4	PCIE_TxP	A8
DDR_BA2	F15	GPIO_6	U4	PLL2_Enable	T14
DDR_CK	H17	GPIO_7	V4	RCVRD_CLK0	D9
DDR_CKE	E18	GPIO_8	R5	RCVRD_CLK1	A10
DDR_CKn	H18	GPIO_9	T5	RCVRD_CLK2	B10
DDR_DQ0	M18	GPIO_10	U5	RCVRD_CLK3	C10
DDR_DQ1	P17	GPIO_11	V5	RCVRD_CLK4	A11
DDR_DQ2	M17	GPIO_12	R6	RCVRD_CLK5	B11
DDR_DQ3	P16	GPIO_13	T6	RCVRD_CLK6	C11
DDR_DQ4	P18	GPIO_14	U6	RCVRD_CLK7	A12
DDR_DQ5	M16	GPIO_15	V6	RCVRD_CLK8	B12
DDR_DQ6	P15	GPIO_16	R7	RCVRD_CLK9	C12
DDR_DQ7	N15	GPIO_17	T7	RefClk_N	D4
DDR_DQ8	B17	GPIO_18	U7	RefClk_P	C4
DDR_DQ9	D15	GPIO_19	V7	RefClk_Sel0	U15
DDR_DQ10	B16	GPIO_20	R8	RefClk_Sel1	V14
DDR_DQ11	D16	GPIO_21	T8	RefClk_Sel2	U14
DDR_DQ12	D17	GPIO_22	U8	RefClk2_N	D13
DDR_DQ13	B15	GPIO_23	V8	RefClk2_P	C13
DDR_DQ14	C18	GPIO_24	R9	RefClk2_Sel0	V13
DDR_DQ15	B18	GPIO_25	T9	RefClk2_Sel1	U13
DDR_LDM	N18	GPIO_26	U9	RefClk2_Sel2	T13

Reserved_0	U16
Reserved_1	V17
Reserved_2	A16
Reserved_3	A17
Reserved_4	A15
Reserved_7	B4
Reserved_8	A4
Reserved_9	B6
Reserved_10	A6
Reserved_11	D6
Reserved_12	C6
Reserved_13	B5
Reserved_14	A5
Reserved_15	D5
Reserved_16	C5
Reserved_17	B3
Reserved_18	A3
Reserved_19	D3
Reserved_20	C3
Reserved_21	G2
Reserved_22	G1
Reserved_23	G4
Reserved_24	G3
SerDes_E0_RxN	N4
SerDes_E0_RxP	N3
SerDes_E0_TxN	N2
SerDes_E0_TxP	N1
SerDes_E1_RxN	H4
SerDes_E1_RxP	H3
SerDes_E1_TxN	H2
SerDes_E1_TxP	H1
SerDes_Rext_0	R1
SerDes_Rext_1	R2
SerDes0_RxN	M4
SerDes0_RxP	M3
SerDes0_TxN	M2
SerDes0_TxP	M1
SerDes1_RxN	L4
SerDes1_RxP	L3
SerDes1_TxN	L2
SerDes1_TxP	L1
SerDes2_RxN	K4
SerDes2_RxP	K3
SerDes2_TxN	K2
SerDes2_TxP	K1

SerDes3_RxN	J4
SerDes3_RxP	J3
SerDes3_TxN	J2
SerDes3_TxP	J1
SGMII_NPI_RxN	D7
SGMII_NPI_RxP	C7
SGMII_NPI_TxN	B7
SGMII_NPI_TxP	A7
SI_Clk	V11
SI_DI	U11
SI_DO	T11
SI_nEn	R11
Test_Mode	U17
THERMDA	T1
THERMDC_VSS	U1
VDD_1	H7
VDD_2	H8
VDD_3	H9
VDD_4	H10
VDD_5	H11
VDD_6	H12
VDD_7	L7
VDD_8	L8
VDD_9	L9
VDD_10	L10
VDD_11	L11
VDD_12	L12
VDD_13	M8
VDD_14	M11
VDD_IO_1	A9
VDD_IO_2	B9
VDD_IO_3	C9
VDD_IO_4	D10
VDD_IO_5	D11
VDD_IO_6	D12
VDD_IO_7	P7
VDD_IO_8	P8
VDD_IO_9	P9
VDD_IO_10	P10
VDD_IO_11	P11
VDD_IO_12	P12
VDD_IO_13	P13
VDD_IO_14	R14
VDD_IO_15	T15
VDD_IO_16	T16

VDD_IO_17	T17
VDD_IO_18	T18
VDD_IODDR_1	A14
VDD_IODDR_2	B14
VDD_IODDR_3	C14
VDD_IODDR_4	D14
VDD_IODDR_5	E14
VDD_IODDR_6	F14
VDD_IODDR_7	G14
VDD_IODDR_8	H14
VDD_IODDR_9	J14
VDD_IODDR_10	K14
VDD_IODDR_11	L14
VDD_IODDR_12	M14
VDD_IODDR_13	N14
VDD_VS_1	D1
VDD_VS_2	D2
VDD_VS_3	E3
VDD_VS_4	E5
VDD_VS_5	E7
VDD_VS_6	E9
VDD_VS_7	H6
VDD_VS_8	L6
VDD_VS_9	M6
VDD_VS_10	N6
VDD_VS_11	P1
VDD_VS_12	P2
VDD_VS_13	P3
VDD_VS_14	P4
VDD_VS_15	P5
VDDA_1	F1
VDDA_2	F2
VDDA_3	F3
VDDA_4	F4
VDDA_5	F5
VDDA_6	F6
VDDA_7	F7
VDDA_8	F8
VDDA_9	F9
VDDA_10	F10
VDDA_11	F11
VDDA_12	F12
VDDA_13	G5
VDDA_14	J5
VDDA_15	K5

Pins by number (*continued*)

VDDA_16	M5	VSS_45	M9
VSS_1	A2	VSS_46	M10
VSS_2	B1	VSS_47	M12
VSS_3	B2	VSS_48	M13
VSS_4	C1	VSS_49	N5
VSS_5	C2	VSS_50	N7
VSS_6	E1	VSS_51	N8
VSS_7	E2	VSS_52	N9
VSS_8	E4	VSS_53	N10
VSS_9	E6	VSS_54	N11
VSS_10	E8	VSS_55	N12
VSS_11	E10	VSS_56	N13
VSS_12	E11	VSS_57	P6
VSS_13	E12	VSS_58	P14
VSS_14	E13	VSS_59	R3
VSS_15	F13	VSS_60	R4
VSS_16	G6	VSS_61	R15
VSS_17	G7	VSS_62	R16
VSS_18	G8	VSS_63	R17
VSS_19	G9	VSS_64	T2
VSS_20	G10	VSS_65	U18
VSS_21	G11		
VSS_22	G12		
VSS_23	G13		
VSS_24	H5		
VSS_25	H13		
VSS_26	J6		
VSS_27	J7		
VSS_28	J8		
VSS_29	J9		
VSS_30	J10		
VSS_31	J11		
VSS_32	J12		
VSS_33	J13		
VSS_34	K6		
VSS_35	K7		
VSS_36	K8		
VSS_37	K9		
VSS_38	K10		
VSS_39	K11		
VSS_40	K12		
VSS_41	K13		
VSS_42	L5		
VSS_43	L13		
VSS_44	M7		

8.0 PACKAGE INFORMATION

The VSC7416XMT-01 package is a lead(Pb)-free, 324-pin, plastic ball grid array (BGA) with a 19 mm × 19 mm body size, 1 mm pin pitch, and 1.8 mm maximum height.


Lead(Pb)-free products from Microchip comply with the temperatures and profiles defined in the joint IPC and JEDEC standard IPC/JEDEC J-STD-020. For more information, see the IPC and JEDEC standard.

This section provides the package drawing, thermal specifications, and moisture sensitivity rating for the VSC7416-01 device.

8.1 Package Drawing

The following illustration shows the package drawing for the VSC7416-01 device. The drawing contains the top view, bottom view, side view, detail views, dimensions, tolerances, and notes.

FIGURE 8-1: PACKAGE DRAWING



	Symbol	Common Dimensions		
		MIN.	NOM.	MAX.
Package :	PBGA			
Body Size:	X		19.000	
	Y		19.000	
Ball Pitch :	e		1.000	
Total Thickness :	A	—	—	1.800
Mold Thickness :	M		0.700	Ref.
Substrate Thickness :	S		0.542	Ref.
Ball Diameter :			0.500	
Stand Off :	A1	0.310	—	0.410
Ball Width :	b	0.440	—	0.640
Package Edge Tolerance :	aaa		0.200	
Mold Parallelism :	bbb		0.350	
Coplanarity:	ddd		0.200	
Ball Offset (Package) :	eee		0.250	
Ball Offset (Ball) :	fff		0.100	
Ball Count :	n		324	
Edge Ball Center to Center :	X		17.000	
	Y		17.000	

8.2 Thermal Specifications

Thermal specifications for these devices are based on the JEDEC JESD51 family of documents. These documents are available on the JEDEC Web site at www.jedec.org. The thermal specifications are modeled using a four-layer test board with two signal layers, a power plane, and a ground plane (2s2p PCB). For more information about the thermal measurement method used for these devices, see the JESD51-1 standard.

TABLE 8-1: THERMAL RESISTANCES

Symbol	°C/W	Parameter
θ_{JCTop}	2.89	Die junction to package case top
θ_{JB}	8.65	Die junction to printed circuit board
θ_{JA}	15.94	Die junction to ambient
θ_{JMA} at 1 m/s	13.3	Die junction to moving air measured at an air speed of 1 m/s
θ_{JMA} at 2 m/s	12.36	Die junction to moving air measured at an air speed of 2 m/s

To achieve results similar to the modeled thermal measurements, the guidelines for board design described in the JESD51 family of publications must be applied. For information about applications using BGA packages with an exposed pad, see the following:

- *JESD51-2A, Integrated Circuits Thermal Test Method Environmental Conditions, Natural Convection (Still Air)*
- *JESD51-6, Integrated Circuit Thermal Test Method Environmental Conditions, Forced Convection (Moving Air)*
- *JESD51-8, Integrated Circuit Thermal Test Method Environmental Conditions, Junction-to-Board*
- *JESD51-9, Test Boards for Area Array Surface Mount Package Thermal Measurements*

8.3 Moisture Sensitivity

This device is rated moisture sensitivity level 4 as specified in the joint IPC and JEDEC standard IPC/JEDEC J-STD-020. For more information, see the IPC and JEDEC standard.

9.0 DESIGN GUIDELINES

This section provides information about design guidelines for the VSC7416-01 device.

9.1 Power Supplies

The following guidelines apply to designing power supplies for use with the VSC7416-01 device:

- Make at least one unbroken ground plane (GND).
- Use the power and ground plane combination as an effective power supply bypass capacitor. The capacitance is proportional to the area of the two planes and inversely proportional to the separation between the planes. Typical values with a 0.25 mm (0.01 inch) separation are 100 pF/in². This capacitance is more effective than a capacitor of equivalent value, because the planes have no inductance or Equivalent Series Resistance (ESR).
- Do not cut up the power or ground planes in an effort to steer current paths. This usually produces more noise, not less. Furthermore, place vias and clearances in a configuration that maintains the integrity of the plane. Groups of vias spaced close together often overlap clearances. This can form a large slot in the plane. As a result, return currents are forced around the slot, which increases the loop area and EMI emissions. Signals should never be placed on a ground plane, because the resulting slot forces return currents around the slot.
- Vias connecting power planes to the supply and ground balls must be at least 0.25 mm (0.010 inch) in diameter, preferably with no thermal relief and plated closed with copper or solder. Use separate (or even multiple) vias for each supply and ground ball.

9.2 Power Supply Decoupling

Each power supply voltage should have both bulk and high-frequency decoupling capacitors. Recommended capacitors are as follows:

- For bulk decoupling, use 10 μ F high capacity and low ESR capacitors or equivalent, distributed across the board.
- For high-frequency decoupling, use 0.1 μ F high frequency (for example, X7R) ceramic capacitors placed on the side of the PCB closest to the plane being decoupled, and as close as possible to the power ball. A larger value in the same housing unit produces even better results.
- Use surface-mounted components for lower lead inductance and pad capacitance. Smaller form factor components are best (that is, 0402 is better than 0603).

9.3 Reference Clock

The device reference clock can be a 25 MHz, 125 MHz, or 156.25 MHz clock signal. It can be either a differential reference clock or a single-ended clock. However, 25 MHz single-ended operation is not recommended to achieve the jitter specification requirements of the SGMII interface. For more information, see [Section 6.1.2, Reference Clocks](#).

9.3.1 SINGLE-ENDED REFCLK INPUT

An external resistor network is required to use a single-ended reference clock. The network limits the amplitude and adjusts the center of the swing.

The following illustrations show configurations for a single-ended reference clock.

FIGURE 9-1: 2.5 V CMOS SINGLE-ENDED REFCLK INPUT RESISTOR NETWORK





FIGURE 9-2: 3.3 V CMOS SINGLE-ENDED REFCLK INPUT RESISTOR NETWORK



9.4 Interfaces

This section provides general recommendations for all interfaces and information related to the specific interfaces on the device.

9.4.1 GENERAL RECOMMENDATIONS

High-speed signals require excellent frequency and phase response up to the third harmonic. The best design would provide excellent frequency and phase response up to the seventh harmonic. The following recommendations can improve signal quality and minimize transmission distances:

Keep traces as short as possible. Initial component placement should be considered very carefully.

- The impedance of the traces must match the impedance of the termination resistors, connectors, and cable. This reduces reflections due to impedance mismatches.
- Differential impedance must be maintained in a 100 Ω differential application. Routing two 50 Ω traces is not adequate. The two traces must be separated by enough distance to maintain differential impedance. When routing differential pairs, keep the two trace lengths identical. Differences in trace lengths translate directly into signal skew. Note that the differential impedance may be affected when separations occur.
- Keep differential pair traces on the same layer of the PCB to minimize impedance discontinuities.
- Do not group all the passive components together. The pads of the components add capacitance to the traces. At the frequencies encountered, this can result in unwanted reductions in impedance. Use surface-mounted 0603 components to reduce this effect.
- Eliminate or reduce stub lengths.
- Reduce, if not eliminate, vias to minimize impedance discontinuities. Remember that vias and their clearance holes in the power/ground planes can cause impedance discontinuities in nearby signals. Keep vias away from other traces.
- Keep signal traces away from other signals that might capacitively couple noise into the signals. A good rule of thumb is to keep the traces apart by ten times the width of the trace.
- Do not route digital signals from other circuits across the area of the high-speed transmitter and receiver signals.
- Using grounded guard traces is typically not effective for improving signal quality. A ground plane is more effective. However, a common use of guard traces is to route them during the layout, but remove them prior to design completion. This has the benefit of enforcing keep-out areas around sensitive high-speed signals so that vias and other traces are not accidentally placed incorrectly.
- When signals in a differential pair are mismatched, the result is a common-mode current. In a well-designed system, common-mode currents should make up less than one percent of the total differential currents. Mode currents represent a primary source of EMI emissions. To reduce common-mode currents, route differential traces so that their lengths are the same. For example, a 5-mm (0.2-inch) length mismatch between differential signals having the rise and fall times of 200 ps results in the common-mode current being up to 18% of the differential current.

Note Care must be taken when choosing proper components (such as the termination resistors) in the designing of the layout of a printed circuit board, because of the high application frequency. The use of surface-mount components is highly recommended to minimize parasitic inductance and lead length of the termination resistor.

Matching the impedance of the PCB traces, connectors, and balanced interconnect media is also highly recommended. Impedance variations along the entire interconnect path must be minimized, because they degrade the signal path and may cause reflections of the signal.

9.4.2 SGMII INTERFACE

The SGMII interface consists of a Tx and Rx differential pair operating at 1250 Mbps.

The SGMII signals can be routed on any PCB trace layer with the following constraints:

- The Tx output signals in a pair should have matched electrical lengths.
- The Rx input signals in a pair should have matched electrical lengths.
- SGMII Tx and Rx pairs must be routed as 100 Ω differential traces with ground plane as reference.
- Keep differential pair traces on the same layer of the PCB to minimize impedance discontinuities.
- AC-coupling of Tx and Rx may be needed, depending on the PHY. If AC-coupled, the inputs are self-biased.
- To reduce the crosstalk between pairs or other PCB lines, it is recommended that the spacing on each side of the pair be larger than four times the track width.

9.4.3 SERIAL INTERFACE

If the serial CPU interface is not used, all input signals can be left floating.

The SI bus consists of the SI_Clk clock signal, the SI_DO and SI_DI data signals, and the SI_nCS0 device select signal.

When routing the SI_Clk signal, be sure to create clean edges. If the SI bus is connected to more than one slave device, route it in a daisy-chain configuration with no stubs. Terminate the SI_Clk signal properly to avoid reflections and double clocking.

If it is not possible (or desirable) to route the bus in a daisy-chain configuration, the SI_Clk signal should be buffered and routed in a star topology from the buffers. Each buffered clock should be terminated at its source.

9.4.4 ENHANCED SERDES INTERFACE

The Enhanced SerDes interface can be used for fiber connections, backplanes, or direct coupler cable connections, such as the CX4 cable.

The Enhanced SerDes interface can operate in several modes. The physical signal bit rate is between 125 Mbps to 3.125 Mbps.

The inputs are self-biased and have internal AC-coupling. In some modes, the interface requires external AC-coupling, because of the input DC voltage limitation. If external AC-coupling capacitors are required, it is recommended to use small form factor components, such as 0603. The small form factor minimizes impedance mismatch by the AC-coupling capacitors, because the size of the form factor approximately matches the trace width commonly used for these signals.

The Enhanced SerDes interface can be directly connected to either 1 Gbps or 2.5 Gbps SFP modules. For these applications, external AC-coupling capacitors are not required, because the SFP module already includes capacitors.

The following table lists the AC-coupling requirements for common Enhanced SerDes connections.

TABLE 9-1: ENHANCED SERDES INTERFACE COUPLING REQUIREMENTS

Enhanced SerDes Connection	Mode	External AC-Coupling Requirement
SFP modules	SFP	Not required
SGMII PHY	SGMII	Required ⁽¹⁾
Enhanced SerDes device	Enhanced SerDes	Required

1. AC-coupling is not required with direct connection to the VSC8512 PHY device.

The Enhanced SerDes interface signals must be routed as a differential pair, with a 100 Ω differential characteristic impedance. The differential intra-pair skew must be below 5 ps in the PCB trace.

To minimize crosstalk between transmitter and receiver, equal drive strength is recommended in both devices in a link.

To minimize crosstalk between differential pairs, the characteristic differential impedance of the signals must be determined only by the distance to the reference plane and the intra-pair coupling and not the distance to the neighboring traces. To separate the transmitter and receiver signals to minimize crosstalk, it is recommended to route the transmitter and receiver signals on as many different PCB layers as feasible.

9.4.5 PCI EXPRESS INTERFACE

The VSC7416-01 device does not accept spread spectrum modulated PCIe input. Although the device only supports PCI Express Base Specification Revision 1.1, the PCIe transmitter and receiver support the PCI Express Base Specification Revision 2.0 Electrical sub-block specifications under the following conditions:

- Only 2.5 gigatransfers per second (GT/s) is supported.

VSC7416-01

- Full swing output signaling is only supported when $V_{DD\_VS} = 1.2\text{ V}$.
- Low swing signaling is supported for $V_{DD\_VS} = 1.2\text{ V}$ and $V_{DD\_VS} = 1.0\text{ V}$.

During PCIe startup (VCORE_CFG = 1001), only low swing signaling with no de-emphasis is supported, regardless of the V_{DD_VS} voltage. After startup, configure the PCIe interface as desired.

9.4.6 TWO-WIRE SERIAL INTERFACE

The two-wire serial interface is capable of suppressing small amplitude glitches less than 5 ns in duration, which is less than the 50 ns duration often quoted for similar interfaces. Because the two-wire serial implementation uses Schmitt-triggered inputs, the VSC7416-01 device has a greater tolerance to low amplitude noise. For glitch-free operation, select the proper pull-up resistor value to ensure that the transition time through the input switching region is less than 5 ns given the line's total capacitive load. For capacitive loads up to 40 pF, a pull-up resistor of 510 Ω or less ensures glitch-free operation for noise levels up to 700 mV peak-to-peak.

9.4.7 DDR2/DDR3 SDRAM INTERFACE

The DDR2/DDR3 SDRAM interface is designed to interface with 8-bit or 16-bit DDR SDRAM devices. The maximum amount of physical memory that can be addressed is one gigabyte. Possible combinations of memory modules are:

- One 8-bit device, connect to CS0 byte lane 0.
- Two 8-bit devices, connect to CS0 and both byte lanes.
- Four 8-bit devices, connect to both chip-selects and both byte lanes.
- One 16-bit device, connect to CS0 and both byte lanes.
- Two 16-bit devices, connect to both chip-selects and both byte lanes.


When using a single 8-bit device the memory controller must be configured for 8-bit mode. All other configurations use 16-bit mode.

All signals on this interface must be connected one-to-one with the corresponding signals on the DDR SDRAM device. When using only one 8-bit device, the DDR_UDQS, DDR_UDQSn, DDR_UDM, and DDR_DQ[15:8] signals must be left unconnected. When using DDR2 SDRAM devices that has only four banks the DDR_BA[2] signal must be left unconnected.

When using four 8-bit devices or two 16-bit devices, the CS1 memory modules must be placed in-between the VSC7416-01 device and the CS0 memory modules so that the DDR-DQS, DDR-DM, and DDR_DQ signals pass the CS1 devices first before reaching the CS0 devices.

The placement of the VSC7416-01 interface signals is optimized for point-to-point routing directly to a single DDR3 16-bit SDRAM.

FIGURE 9-3: 16-BIT DDR3 SDRAM POINT-TO-POINT ROUTING



Because reflections are absorbed by the devices, keep the physical distance of all the SDRAM interface signals as low as possible. Omit external discrete termination on the address, command, control and clock lines.

When routing the DDR interface, attention must be paid to the skew, primary concern is skew within the byte lane between the differential strobe and the single-ended signals. Skew recommendations for the DDR interface are listed in the following table.

TABLE 9-2: RECOMMENDED SKEW BUDGET

Description	Signal	Maximum Skew
Skew within byte lane 0	DDR_LDQS/DDR_LDQSn	50 ps
Skew within byte lane 1	DDR_UDQS/DDR_UDQSn	50 ps

TABLE 9-2: RECOMMENDED SKEW BUDGET

Description	Signal	Maximum Skew
Skew within address, command, and control bus	DDR_CK/DDR_CKn DDR_nRAS DDR_CKE DDR_ODT[1:0] DDR_nCAS DDR_nWE DDR_BA[2:0] DDR_A[15:0]	100 ps
Skew between control bus clock and byte lane 0 clock	DDR_CK/DDR_CKn DDR_LDQS/DDR_LDQSn	1250 ps
Skew between control bus clock and byte lane 1 clock	DDR_CK/DDR_CKn DDR_UDQS/DDR_UDQSn	1250 ps
Control bus differential clock intra-pair skew	DDR_CK/DDR_CKn	5 ps

Power supply recommendations:

- Use a shared voltage reference between the VSC7416-01's device's DDR_Vref supply and the DDR device's reference voltage.
- Generate the DDR_Vref from the V_{DD_IODDR} supply using a resistor divider with value of 1 k Ω and an accuracy of 1% or better.
- Use a decoupling capacitance of at least 0.1 μ F on the supply in a manner similar to V_{DD_IODDR} and V_{SS} to ensure tracking of supply variations; however, the time constant of the resistor divider and decoupling capacitance should not exceed the nReset assertion time after power on.
- V_{DD_IODDR} pins must not share vias. Use at least one via for each V_{DD_IODDR} pin. The extra inductance from sharing vias may cause bit errors in the DDR interface.

Routing recommendations:

- DDR_CK/DDR_CKn must be routed as a differential pair with a 100 Ω differential characteristic impedance.
- DDR_xDQS/DDR_xDQSn must be routed as a differential pair with a 100 Ω differential characteristic impedance.
- To minimize crosstalk, the characteristic impedance of the single-ended signals should be determined predominantly by the distance to the reference plane and not the distance to the neighboring traces.
- The crosstalk should be below -20 dB.
- If the DDR interface is not used, connect V_{DD_IODDR} and DDR_VREF to ground. Leave all other DDR signals unconnected (floating). V_{DD_IODDR} can also be left floating; however, DDR_VREF must also be left floating.

9.4.8 THERMAL DIODE EXTERNAL CONNECTION

The internal on-die thermal diode can be used with an external temperature monitor to easily and accurately measure the junction temperature of the VSC7416-01 device.


The on-die thermal diode has internal connected the diode cathode to V_{SS}, the external temperature sensor must support the thermal diode cathode connected to V_{SS}.

Thermal diode is extremely sensitive to noise. To minimize the temperature measurement errors, follow these guidelines:

- Route the ThermDC and ThermDA signals as a differential pair with a differential impedance less than 100 Ω .
- Place the external temperature monitor as close as is possible to the VSC7416-01 device.
- Add a 47 Ω resistor in series with the external temperature monitor supply to filter noise.
- Place a de-coupling capacitor between the external temperature monitor supply pin and the ThermDC signal. Place the capacitor close to the external temperature sensor, as shown in the following illustration.

Connect the external temperature monitor V_{SS} pin directly to the ThermDC pin, which has the connection to V_{SS}, as shown in the following illustration. Do not connect the external temperature monitor V_{SS} pin to the global V_{SS} plane.

FIGURE 9-4: EXTERNAL TEMPERATURE MONITOR CONNECTION



10.0 ORDERING INFORMATION

The VSC7416XMT-01 package is a lead(Pb)-free, 324-pin, plastic ball grid array (BGA) with a 19 mm × 19 mm body size, 1 mm pin pitch, and 1.8 mm maximum height.

Lead(Pb)-free products from Microchip comply with the temperatures and profiles defined in the joint IPC and JEDEC standard IPC/JEDEC J-STD-020. For more information, see the IPC and JEDEC standard.

The following table lists the ordering information for the VSC7416-01 device.

TABLE 10-1: ORDERING INFORMATION

Part Order Number	Description
VSC7416XMT-01	Lead-free, 324-pin, plastic BGA with a 19 mm × 19 mm body size, 1 mm pin pitch, and 1.8 mm maximum height

11.0 DESIGN CONSIDERATIONS

This section provides information about the design considerations for the VSC7416-01 devices.

TLVs appended to PTP event frames not supported over IPv6

The IEEE 1588-2008 standard supports TLV (Type-Length-Value) extensions inserted after a normal PTP payload. VSC7416-01 does not support TLVs located in PTP event frames (the frames that require timestamping) sent over IPv6/UDP.

The IEEE 1588-2008 standard recommends that TLVs not be added to PTP event frames, however, it is not mandated. If it cannot be guaranteed that PTP event frames are transmitted without TLVs, configure the VSC7416-01 device to drop any PTP/IPv6/UDP event frames that contain TLVs.

This limitation applies only for PTP over IPv6/UDP.

PTP on ports in 10BASE-T mode has large internal asymmetry

When running PTP timestamping, the ingress and egress latency from the device pins to the internal timestamp point must be configured. When running a port in 10 Mb mode, the internal ingress and egress latencies in VSC7416-01 are too large to configure using the internal latency registers. This makes it impossible to compensate for the internal latency asymmetry and will cause a time offset when running PTP.

PTP software may be programmed to compensate for the internal latency when running the device in PTP boundary clock mode, but not in transparent clock mode. As a result, it is not recommended to run PTP over 10 Mb links.

PFC frames with VLAN tag incorrectly accepted

Priority-based flow control (PFC) frames with a VLAN tag are incorrectly accepted as valid pause frames. This issue is not estimated to present any system level impact.

UNH test 32.7.1 Full/Half Duplex Resolution failure

UNH failure is caused by an error in the hardware ANEG state machine, because no ANEG software was running on the device during test.

Contact Microchip for the current API software release.

UNH test 4.4.3 Collisions in Data Outside of slotTime while not Bursting failure

The VSC7416-01 device fails UNH Fast Ethernet Consortium 100 Mbps half duplex test 4.4.3, Collisions in Data Outside of slotTime while not Bursting.

The device is not able to correctly perform a late collision of the last packet byte in a frame. This means that the maximum collision domain is approximately 8 meters shorter than what is physically allowed.

Port statistics incorrect when using Up-MEP on port

Rx port statistics incorrectly include all traffic injected on the port using masquerading, which includes OAM initiated on the port's Up-MEP.

Software must modify port statistics based on injected frame counts. Contact Microchip for the current software release.

Port loss measurement counters incorrect when using Up-MEP on port

Port Down-MEP loss measurements with a link partner is not supported for ports with an Up-MEP, because Rx loss measurement (LM) counters incorrectly include frames injected on the port using masquerading.

ISDX not reliably assigned in IS1 lookups 2 and 3

The ISDX obtained from the first IS1 lookup dictates protection switching operation (IPT/PPT) and VOE selection. IS1 lookups 2 and 3 cannot reliably determine a frame's ISDX if protection switching is enabled.

Disable protection switching using IPT/PPT, and use IS1 and IS2 instead to implement protection switching.

Contact Microchip for the current API software release.

Frames killed by IPT linear protection do not cause VOE lookup

Frames discarded by IPT (inactive Rx segment of 1:1 protection) do not make it to Versatile OAM endpoint (VOE) lookup. Port OAM and Segment/Path OAM on the inactive Segment/Path will not work.

Disable protection switching using IPT/PPT, and use IS1 instead to implement protection switching.

Contact Microchip for the current API software release.

Injection of OAM on one port and extraction on another port not supported

OAM exchange between two Up-MEPs in the same VSC7416 device is not supported. For example, both UNIs of an E-Line residing on different ports on the same VSC7416 device.

False matches possible using IS2 OAM Key

The IS2 OAM Key should only match the EtherTypes of 802.1ag/Y.1731 OAM frames, however, it also incorrectly matches the EtherTypes of EFM OAM and E-LMI frames.

To prevent false matches, use VCAP IS1 keys to match only the 802.1ag/Y.1731 OAM EtherTypes, and signal the result to VCAP IS2 using a Policy Association Group (PAG).

Contact Microchip for the current API software release.

APPENDIX A: REVISION HISTORY

The revision history describes the changes that were implemented in the document. The changes are listed by revision, starting with the latest publication.

Rev. A - 05/2021

The following is a summary of changes in revision A of this document.

- Migrated the document from Vitesse format to Microchip format.
- Updated the Pin Description section by updating the number of pins. For more information, see [Section 7.0, Pin Descriptions](#).
- Updated the Pin diagram section with the top view of the package drawing. For more information, see [Section 7.1, Pin Diagram](#).
- Updated the Package Information section with the BGA package information. For more information, see [Section 8.0, Package Information](#).
- Updated the package diagram. For more information, see [Figure 8-1](#), page 659.
- Updated the Ordering Information section. For more information, see [Section 10.0, Ordering Information](#).

Rev. 4.0 - 12/2012

Revision 4.0 of this datasheet was published in December 2012. This was the first production-level publication of the document.

THE MICROCHIP WEBSITE

Microchip provides online support via our WWW site at www.microchip.com. This website is used as a means to make files and information easily available to customers. Accessible by using your favorite Internet browser, the website contains the following information:

- **Product Support** – Data sheets and errata, application notes and sample programs, design resources, user's guides and hardware support documents, latest software releases and archived software
- **General Technical Support** – Frequently Asked Questions (FAQ), technical support requests, online discussion groups, Microchip consultant program member listing
- **Business of Microchip** – Product selector and ordering guides, latest Microchip press releases, listing of seminars and events, listings of Microchip sales offices, distributors and factory representatives

CUSTOMER CHANGE NOTIFICATION SERVICE

Microchip's customer notification service helps keep customers current on Microchip products. Subscribers will receive e-mail notification whenever there are changes, updates, revisions or errata related to a specified product family or development tool of interest.

To register, access the Microchip website at www.microchip.com. Under "Support", click on "Customer Change Notification" and follow the registration instructions.

CUSTOMER SUPPORT

Users of Microchip products can receive assistance through several channels:

- Distributor or Representative
- Local Sales Office
- Field Application Engineer (FAE)
- Technical Support

Customers should contact their distributor, representative or Field Application Engineer (FAE) for support. Local sales offices are also available to help customers. A listing of sales offices and locations is included in the back of this document.

Technical support is available through the website at: www.microchip.com/support

Note the following details of the code protection feature on Microchip devices:

- Microchip products meet the specification contained in their particular Microchip Data Sheet.
- Microchip believes that its family of products is one of the most secure families of its kind on the market today, when used in the intended manner and under normal conditions.
- There are dishonest and possibly illegal methods used to breach the code protection feature. All of these methods, to our knowledge, require using the Microchip products in a manner outside the operating specifications contained in Microchip's Data Sheets. Most likely, the person doing so is engaged in theft of intellectual property.
- Microchip is willing to work with the customer who is concerned about the integrity of their code.
- Neither Microchip nor any other semiconductor manufacturer can guarantee the security of their code. Code protection does not mean that we are guaranteeing the product as “unbreakable.”

Code protection is constantly evolving. We at Microchip are committed to continuously improving the code protection features of our products. Attempts to break Microchip's code protection feature may be a violation of the Digital Millennium Copyright Act. If such acts allow unauthorized access to your software or other copyrighted work, you may have a right to sue for relief under that Act.

Information contained in this publication regarding device applications and the like is provided only for your convenience and may be superseded by updates. It is your responsibility to ensure that your application meets with your specifications. MICROCHIP MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WHETHER EXPRESS OR IMPLIED, WRITTEN OR ORAL, STATUTORY OR OTHERWISE, RELATED TO THE INFORMATION, INCLUDING BUT NOT LIMITED TO ITS CONDITION, QUALITY, PERFORMANCE, MERCHANTABILITY OR FITNESS FOR PURPOSE. Microchip disclaims all liability arising from this information and its use. Use of Microchip devices in life support and/or safety applications is entirely at the buyer's risk, and the buyer agrees to defend, indemnify and hold harmless Microchip from any and all damages, claims, suits, or expenses resulting from such use. No licenses are conveyed, implicitly or otherwise, under any Microchip intellectual property rights unless otherwise stated.

Trademarks

The Microchip name and logo, the Microchip logo, Adaptec, AnyRate, AVR, AVR logo, AVR Freaks, BesTime, BitCloud, chipKIT, chipKIT logo, CryptoMemory, CryptoRF, dsPIC, FlashFlex, flexPWR, HELDO, IGLOO, JukeBlox, KeeLoq, Klear, LANCheck, LinkMD, maXStylus, maXTouch, MediaLB, megaAVR, Microsemi, Microsemi logo, MOST, MOST logo, MPLAB, OptoLyzer, PackTime, PIC, picoPower, PICSTART, PIC32 logo, PolarFire, Prochip Designer, QTouch, SAM-BA, SenGenuity, SpyNIC, SST, SST Logo, SuperFlash, Symmetricom, SyncServer, Tachyon, TempTrackr, TimeSource, tinyAVR, UNI/O, Vectron, and XMEGA are registered trademarks of Microchip Technology Incorporated in the U.S.A. and other countries.

APT, ClockWorks, The Embedded Control Solutions Company, EtherSynch, FlashTec, Hyper Speed Control, HyperLight Load, IntelliMOS, Libero, motorBench, mTouch, Powermite 3, Precision Edge, ProASIC, ProASIC Plus, ProASIC Plus logo, Quiet-Wire, SmartFusion, SyncWorld, Temux, TimeCesium, TimeHub, TimePictra, TimeProvider, Vite, WinPath, and ZL are registered trademarks of Microchip Technology Incorporated in the U.S.A.

Adjacent Key Suppression, AKS, Analog-for-the-Digital Age, Any Capacitor, AnyIn, AnyOut, BlueSky, BodyCom, CodeGuard, CryptoAuthentication, CryptoAutomotive, CryptoCompanion, CryptoController, dsPICDEM, dsPICDEM.net, Dynamic Average Matching, DAM, ECAN, EtherGREEN, In-Circuit Serial Programming, ICSP, INICnet, Inter-Chip Connectivity, JitterBlocker, KlearNet, KlearNet logo, memBrain, Mindi, MiWi, MPASM, MPF, MPLAB Certified logo, MPLIB, MPLINK, MultiTRAK, NetDetach, Omniscient Code Generation, PICDEM, PICDEM.net, PICKit, PICtail, PowerSmart, PureSilicon, QMatrix, REAL ICE, Ripple Blocker, SAM-ICE, Serial Quad I/O, SMART-I.S., SQI, SuperSwitcher, SuperSwitcher II, Total Endurance, TSHARC, USBCheck, VariSense, ViewSpan, WiperLock, Wireless DNA, and ZENA are trademarks of Microchip Technology Incorporated in the U.S.A. and other countries.

SQTP is a service mark of Microchip Technology Incorporated in the U.S.A.

The Adaptec logo, Frequency on Demand, Silicon Storage Technology, and Symmcom are registered trademarks of Microchip Technology Inc. in other countries.

GestIC is a registered trademark of Microchip Technology Germany II GmbH & Co. KG, a subsidiary of Microchip Technology Inc., in other countries.

All other trademarks mentioned herein are property of their respective companies.

© 2021, Microchip Technology Incorporated, All Rights Reserved.

ISBN:978-1-5224-8178-2

For information regarding Microchip's Quality Management Systems, please visit www.microchip.com/quality.

Worldwide Sales and Service

AMERICAS

Corporate Office
2355 West Chandler Blvd.
Chandler, AZ 85224-6199
Tel: 480-792-7200
Fax: 480-792-7277
Technical Support:
<http://www.microchip.com/support>
Web Address:
www.microchip.com

Atlanta
Duluth, GA
Tel: 678-957-9614
Fax: 678-957-1455

Austin, TX
Tel: 512-257-3370

Boston
Westborough, MA
Tel: 774-760-0087
Fax: 774-760-0088

Chicago
Itasca, IL
Tel: 630-285-0071
Fax: 630-285-0075

Dallas
Addison, TX
Tel: 972-818-7423
Fax: 972-818-2924

Detroit
Novi, MI
Tel: 248-848-4000

Houston, TX
Tel: 281-894-5983

Indianapolis
Noblesville, IN
Tel: 317-773-8323
Fax: 317-773-5453
Tel: 317-536-2380

Los Angeles
Mission Viejo, CA
Tel: 949-462-9523
Fax: 949-462-9608
Tel: 951-273-7800

Raleigh, NC
Tel: 919-844-7510

New York, NY
Tel: 631-435-6000

San Jose, CA
Tel: 408-735-9110
Tel: 408-436-4270

Canada - Toronto
Tel: 905-695-1980
Fax: 905-695-2078

ASIA/PACIFIC

Australia - Sydney
Tel: 61-2-9868-6733

China - Beijing
Tel: 86-10-8569-7000

China - Chengdu
Tel: 86-28-8665-5511

China - Chongqing
Tel: 86-23-8980-9588

China - Dongguan
Tel: 86-769-8702-9880

China - Guangzhou
Tel: 86-20-8755-8029

China - Hangzhou
Tel: 86-571-8792-8115

China - Hong Kong SAR
Tel: 852-2943-5100

China - Nanjing
Tel: 86-25-8473-2460

China - Qingdao
Tel: 86-532-8502-7355

China - Shanghai
Tel: 86-21-3326-8000

China - Shenyang
Tel: 86-24-2334-2829

China - Shenzhen
Tel: 86-755-8864-2200

China - Suzhou
Tel: 86-186-6233-1526

China - Wuhan
Tel: 86-27-5980-5300

China - Xian
Tel: 86-29-8833-7252

China - Xiamen
Tel: 86-592-2388138

China - Zhuhai
Tel: 86-756-3210040

ASIA/PACIFIC

India - Bangalore
Tel: 91-80-3090-4444

India - New Delhi
Tel: 91-11-4160-8631

India - Pune
Tel: 91-20-4121-0141

Japan - Osaka
Tel: 81-6-6152-7160

Japan - Tokyo
Tel: 81-3-6880-3770

Korea - Daegu
Tel: 82-53-744-4301

Korea - Seoul
Tel: 82-2-554-7200

Malaysia - Kuala Lumpur
Tel: 60-3-7651-7906

Malaysia - Penang
Tel: 60-4-227-8870

Philippines - Manila
Tel: 63-2-634-9065

Singapore
Tel: 65-6334-8870

Taiwan - Hsin Chu
Tel: 886-3-577-8366

Taiwan - Kaohsiung
Tel: 886-7-213-7830

Taiwan - Taipei
Tel: 886-2-2508-8600

Thailand - Bangkok
Tel: 66-2-694-1351

Vietnam - Ho Chi Minh
Tel: 84-28-5448-2100

EUROPE

Austria - Wels
Tel: 43-7242-2244-39
Fax: 43-7242-2244-393

Denmark - Copenhagen
Tel: 45-4485-5910
Fax: 45-4485-2829

Finland - Espoo
Tel: 358-9-4520-820

France - Paris
Tel: 33-1-69-53-63-20
Fax: 33-1-69-30-90-79

Germany - Garching
Tel: 49-8931-9700

Germany - Haan
Tel: 49-2129-3766400

Germany - Heilbronn
Tel: 49-7131-72400

Germany - Karlsruhe
Tel: 49-721-625370

Germany - Munich
Tel: 49-89-627-144-0
Fax: 49-89-627-144-44

Germany - Rosenheim
Tel: 49-8031-354-560

Israel - Ra'anana
Tel: 972-9-744-7705

Italy - Milan
Tel: 39-0331-742611
Fax: 39-0331-466781

Italy - Padova
Tel: 39-049-7625286

Netherlands - Drunen
Tel: 31-416-690399
Fax: 31-416-690340

Norway - Trondheim
Tel: 47-7288-4388

Poland - Warsaw
Tel: 48-22-3325737

Romania - Bucharest
Tel: 40-21-407-87-50

Spain - Madrid
Tel: 34-91-708-08-90
Fax: 34-91-708-08-91

Sweden - Gothenberg
Tel: 46-31-704-60-40

Sweden - Stockholm
Tel: 46-8-5090-4654

UK - Wokingham
Tel: 44-118-921-5800
Fax: 44-118-921-5820