7710 SR OS 5.0.R25
SOFTWARE RELEASE NOTES

These release notes are for Release 5.0.R25 of the 7710 SR OS software for the 7710 SR routers.

RELEASE NOTES ORGANIZATION

The following are the major topics covered in these Release Notes:

RELEASE 5.0.R25 DOCUMENTATION SET

The 7710 SR OS Release 5.0.R25 documentation set consists of Release Notes and the 7710 SR OS Release 5.0 manuals. The components of the Release 5.0.R25 documentation set are the following:

- 7710 SR OS 5.0.R25 Software Release Notes (Document Number: 93-0176-25 V5.0.R25)
- 7710 Service Router OS Basic System Configuration Guide 5.0 (93-0079-02)
- 7710 Service Router OS System Management Guide 5.0 (93-0080-02)
- 7710 Service Router OS Interface Configuration Guide 5.0 (93-0081-02)
- 7710 Service Router OS Router Configuration Guides 5.0 (93-0082-02)
- 7710 Service Router OS Routing Protocol Guide 5.0 (93-0083-02)
- 7710 Service Router OS MPLS Guide 5.0 (93-0084-02)
- 7710 Service Router OS Services Guide 5.0 (93-0085-02)

- 7710 Service Router OS Quality of Service Guide 5.0 (93-0086-02)
- 7710 Service Router OS Triple Play Guide 5.0 (93-0143-01)

RELEASE 5.0.R25 SUPPORTED HARDWARE

The following tables summarize the hardware supported in 7710 SR OS Release 5.0.R25. New hardware supported in 7710 SR OS Release 5.0 is printed in **bold**.

TABLE 1. Supported 7710 SR Chassis

Alcatel-Lucent Model #	Description
7710 SR-c12	7710 SR 12 CMA shelf with one Control and Forwarding Module, one Chassis Control Module, one fan filter, one fan tray and two Power Entry Modules (DC - 3HE01012AA; AC - 3HE01013AA)
7710 SR-c4	7710 SR 4 CMA shelf with one Control and Forwarding Module, one Chassis Control Module, one fan filter, one fan tray and two (of max three) Power Entry Modules (DC - 3HE02173AA; AC - 3HE02174AA)

summarizes the line cards supported in 5.0.R25.

TABLE 2. Supported 7710 SR Line Cards

Alcatel-Lucent Part #	Description
3HE01014AA	7710 SR-c12 12 Gbps Control and Forwarding Module (CFM)
3HE01019AA	7710 SR-c12 Chassis Control Module (CCM)
3HE02175AA	**7710 SR-c4 9 Gbps Control and Forwarding Module (CFM)**
3HE02181AA	**7710 SR-c4 Chassis Control Module (CCM)**

Table 3 summarizes the Media Dependent Adapters (MDAs) and Compact Media Adapters (CMAs) supported in 5.0.R25.

TABLE 3. Supported 7710 SR MDAs and CMAs

Alcatel-Lucent Part #	Description
3HE01020AA	8-port Channelized DS1/E1 CMA - RJ48c
3HE01021AA	**4-port DS3/E3 CMA – 1.0/2.3**
3HE01022AA	8-port 10/100TX Ethernet CMA - RJ45
3HE01023AA	1-port GigE CMA - SFP
3HE02185AA	**2-port OC-12c/STM-4c CMA - SFP**
3HE01024AA	MDA Carrier Module (MCM)
3HE00021AA	60-port 10/100TX MDA - mini-RJ21

TABLE 3. Supported 7710 SR MDAs and CMAs (Continued)

Alcatel-Lucent Part #	Description
3HE00023AA	20-port 100FX MDA - SFP
3HE00025AA	5-port GigE MDA - SFP
3HE01615AA	**5-port GigE MDA - SFP Rev B**
3HE00708AA	20-port GigE MDA - SFP
3HE00101AB	20-port 10/100/1000TX MDA - RJ45
3HE00032AA	8-port OC-3c/STM-1c MDA - SFP
3HE00043AA	**2-port OC-48c/STM-16c MDA - SFP**
3HE00071AA	**4-port ATM OC-12c/STM-4c MDA - SFP**

NEW FEATURES IN 5.0.R25

There are no new major features in 5.0.R25. See page 34 for a list of Enhancements in 5.0.R25 and page 81 for a list of Resolved Issues in 5.0.R25.

NEW FEATURES IN 5.0.R24

There are no new major features in 5.0.R24. See page 34 for a list of Enhancements in 5.0.R24 and page 82 for a list of Resolved Issues in 5.0.R24.

NEW FEATURES IN 5.0.R23

There are no new major features in 5.0.R23. See page 34 for a list of Enhancements in 5.0.R23 and page 83 for a list of Resolved Issues in 5.0.R23.

NEW FEATURES IN 5.0.R22

There are no new major features in 5.0.R22. See page 35 for a list of Enhancements in 5.0.R22 and page 87 for a list of Resolved Issues in 5.0.R22.

NEW FEATURES IN 5.0.R21

There are no new major features in 5.0.R21. See page 35 for a list of Enhancements in 5.0.R21 and page 87 for a list of Resolved Issues in 5.0.R21.

NEW FEATURES IN 5.0.R20

There are no new major features in 5.0.R20. See page 35 for a list of Enhancements in 5.0.R20 and page 88 for a list of Resolved Issues in 5.0.R20.

NEW FEATURES IN 5.0.R19

There are no new major features in 5.0.R19. See page 36 for a list of Enhancements in 5.0.R19 and page 89 for a list of Resolved Issues in 5.0.R19.

NEW FEATURES IN 5.0.R18

There are no new major features in 5.0.R18. See page 36 for a list of Enhancements in 5.0.R18 and page 89 for a list of Resolved Issues in 5.0.R18.

NEW FEATURES IN 5.0.R17

There are no new major features in 5.0.R17. See page 36 for a list of Enhancements in 5.0.R17 and page 90 for a list of Resolved Issues in 5.0.R17.

NEW FEATURES IN 5.0.R16

There are no new major features in 5.0.R16. See page 36 for a list of Enhancements in 5.0.R16 and page 90 for a list of Resolved Issues in 5.0.R16.

NEW FEATURES IN 5.0.R15

There are no new major features in 5.0.R15. See page 37 for a list of Enhancements in 5.0.R15 and page 91 for a list of Resolved Issues in 5.0.R15.

NEW FEATURES IN 5.0.R14

There are no new major features in 5.0.R14. See page 37 for a list of Enhancements in 5.0.R14 and page 92 for a list of Resolved Issues in 5.0.R14.

NEW FEATURES IN 5.0.R13

There are no new major features in 5.0.R13. See page 37 for a list of Enhancements in 5.0.R13 and page 93 for a list of Resolved Issues in 5.0.R13.

NEW FEATURES IN 5.0.R12

There are no new major features in 5.0.R12. See page 37 for a list of Enhancements in 5.0.R12 and page 94 for a list of Resolved Issues in 5.0.R12.

NEW FEATURES IN 5.0.R11

There are no new major features in 5.0.R11. See page 37 for a list of Enhancements in 5.0.R11 and page 95 for a list of Resolved Issues in 5.0.R11.

NEW FEATURES IN 5.0.R10

There are no new major features in 5.0.R10. See page 37 for a list of Enhancements in 5.0.R10 and page 97 for a list of Resolved Issues in 5.0.R10.

NEW FEATURES IN 5.0.R9

There are no new major features in 5.0.R9. See page 38 for a list of Enhancements in 5.0.R9 and page 98 for a list of Resolved Issues in 5.0.R9.

NEW FEATURES IN 5.0.R8

There are no new major features in 5.0.R8. See page 38 for a list of Enhancements in 5.0.R8 and page 100 for a list of Resolved Issues in 5.0.R8.

NEW FEATURES IN 5.0.R7

There are no new major features in 5.0.R7. See page 38 for a list of Enhancements in 5.0.R7 and page 101 for a list of Resolved Issues in 5.0.R7.

NEW FEATURES IN 5.0.R6

There are no new major features in 5.0.R6. See page 39 for a list of Enhancements in 5.0.R6 and page 102 for a list of Resolved Issues in 5.0.R6.

NEW FEATURES IN 5.0.R5

The following sections describes the major features in 5.0.R5. In addition, see page 39 for a list of Enhancements in 5.0.R5 and page 104 for a list of Resolved Issues in 5.0.R5.

IN-SERVICE SOFTWARE UPDATE (ISSU)

ISSU (in-service software update) allows in-service software updates for maintenance releases for systems with dual CFMs without requiring a reboot of the system which disrupts services. ISSU is comparable to performing a controlled High Availability switchover where the new image is loaded onto the standby CFM which becomes master, and then upgrading the image on the other CFM.

ISSU has the following limitations:

- ISSU is not supported across major versions, for example, from Release 5.0 to Release 6.0.
- The goal is to support ISSU across maintenance releases for up to six versions, for example, Release 5.0.R4 can be upgraded to Releases 5.0.R5 through 5.0.R10 using ISSU. If in the future, any changes in the code preventing ISSU from working across six versions will be release noted.
- If a firmware update is required, a chassis reboot will be required, and ISSU will not be supported for that version.
- Software downgrades are not possible with ISSU. A system reboot is required to downgrade to a prior version of 7710 SR OS.

MD5 AUTHENTICATION OF RSVP INTERFACE

When MD5 Authentication is enabled on an RSVP interface, authentication of RSVP messages operates in both directions of the interface.

A 7710 SR node maintains a security association using one authentication key for each interface to a neighbor. The following items are stored in the context of this security association:

- The HMAC-MD5 authentication algorithm
- Key used with the authentication algorithm
- Lifetime of the key: the user-entered key is valid until the user deletes it from the interface
- Source address of the sending system
- Latest sending sequence number used with this key identifier.

A 7710 SR RSVP sender transmits an authenticating digest of the RSVP message, computed using the shared authentication key and a keyed hash algorithm. The message digest is included in an Integrity object which also contains a Flags field, a Key Identifier field, and a Sequence Number field. The 7710 SR RSVP sender complies to the procedures for RSVP message generation in RFC2747.

A 7710 SR RSVP receiver uses the key together with the authentication algorithm to process received RSVP messages.

When a 7710 SR PLR node switches the path of the LSP to a bypass LSP, it does not send the Integrity object in the RSVP messages sent over the bypass tunnel. If the PLR receives an RSVP message with an Integrity object, it will perform the digest verification for the key of the interface over which the packet was received. If this fails, the packet is dropped. If the received RSVP message is a RESV message and does not have an Integrity object, then the 7710 SR PLR node will accept it only if it originated from the MP node.

A 7710 SR MP node will accept RSVP messages received over the bypass tunnel with and without the Integrity object. If an Integrity object is present, the proper digest verification for the key of the interface over which the packet was received is performed. If this fails, the packet is dropped.

The 7710 SR MD5 implementation does not support the authentication challenge procedures in RFC2747.

NEW FEATURES IN 5.0.R4

The following sections describe the major features in 5.0.R4. In addition, see page 41 for a list of Enhancements in 5.0.R4 and page 106 for a list of Resolved Issues in 5.0.R4.

SUBSCRIBER ROUTED REDUNDANCY PROTOCOL (SRRP)

SRRP (Subscriber Routed Redundancy Protocol) is actually a collection of functions and messaging protocols that allow a system to create a set of redundant gateway IP addresses shared by a local and remote node.

The 7710 SR platform supports special IES and VPRN subscriber IP interfaces that support system wide IP subnets. A subscriber IP interface operates like a system loop-back interface that supports adjacent IP hosts within the subnet on physical group IP interfaces. Group IP interfaces are created as unnumbered interfaces and do not have local subnets of their own. Instead, a subscriber IP interface contains unnumbered group IP interfaces providing a direct association between the subnet address space created on the subscriber interface and the group interfaces. Since many group IP interfaces may be created within a subscriber IP interface, the subscriber subnet address space may span across multiple physical IP interfaces.

A group IP interface supports the creation of many SAPs within the group interface. Each SAP must be on the same physical port or on the same Ethernet Link Aggregation Group (LAG). Each SAP is configured to support subscriber management which is the process used to manage the subscriber hosts. Multiple group IP interfaces may be created on the same physical port or LAG instance.

In a dual homing situation, normally, VRRP (Virtual Router Redundancy Protocol) would be used to provide gateway redundancy between edge routers. Due to the unnumbered nature of the group interfaces providing gateway routing to the subscriber hosts, VRRP cannot be used for this purpose. SRRP has been created to provide redundant IP gateway operation between group IP interfaces on different chassis.

Subscriber IP interfaces are also capable of providing wholesale routing for retail VPRN instances. In a wholesale/retail association, the SAPs on the group IP interfaces for the wholesale subscriber IP interface may support subscriber hosts that are associated with subscriber subnets in other routing contexts. There can be multiple subscriber subnets supported on a single group IP interface, and the group interface provides redundant gateway routing for each subnet.

DHCP Relay Subscriber Identification using Option 82 VSOs

In prior releases, the DHCP Relay Agent in the 7710 SR could insert subscriber information within the DCHP packets by appending sub-options 1 and 2 in the Option 82 fields. Release 5.0 continues to support this method of inserting subscriber information and provides an alternative method using DHCP Vendor Specific Options (VSOs) in sub-option 9 as defined in RFC4243. The alternative VSO method can be used in situations where the Option 82 sub-option 1 and 2 fields are already being used by other devices in the network.

VSOs exist for specifying the System ID, Client MAC address, Service ID, SAP ID and a text string which can be configured on a VPLS SAP, IES/VPRN IP interface, IES Subscriber Group interface or VPRN Subscriber/Subscriber Group interface. [55873]

Class-based Forwarding over RSVP LSPs

Class-based forwarding over RSVP LSPs allows a service packet to be forwarded over a specific RSVP LSP, part of an SDP, based on its ingress-determined forwarding class. The LSP selected depends on the operational status and load-balancing algorithms used for ECMP and LAG spraying.

This feature allows service providers to dedicate specific LSPs with a determined level of traffic engineering and protection to select service packets. For example, packets of a VoIP service are assigned the “ef” class to expedite their forwarding but are also sent over carefully traffic-engineered and FRR-protected LSP paths across the service provider network.

Class-based forwarding is not supported on a spoke SDP used for termination on an IES or VPRN service. All packets are forwarded over the default LSP.

Note that only user packets are forwarded based on their forwarding class. OAM packets are forwarded in the same way as if the SDP had Class-based Forwarding disabled. In other words, LSP Ping and LSP Trace messages are queued in the queue corresponding to the forwarding class specified by the user and are forwarded over the LSP being tested. Service and SDP OAM packets, for example, Service Ping, VCCV Ping and SDP Ping, are queued in the queue corresponding to the forwarding class specified by the user and forwarded over the first available LSP.

Manual Bypass Tunnels

The 7710 SR implements dynamic bypass tunnels as per RFC 4090, Fast Reroute Extensions to RSVP-TE for LSP Tunnels. When an LSP is signaled and the local protection flag is set in the session_attribute object and/or the FRR object in the path message indicates that facility backup is desired, the 7710 SR PLR will establish a bypass tunnel to provide node and link protection. If a bypass LSP which merges in a downstream node with the protected LSP exists, and if this LSP satisfies the constraints in the FRR object, then this bypass tunnel is selected.

With the manual bypass tunnels, an LSP can be pre-configured from a 7710 SR PLR which will be used exclusively for bypass protection. When a path message for a new LSP requests bypass protection, the 7710 SR will first check if a manual bypass tunnel satisfying the path constraints exists. If one is found, the 7710 SR will select it. If no manual bypass tunnel is found, the 7710 SR will dynamically signal a bypass LSP in the default behavior. Users can disable the dynamic bypass creation on a per node basis.

PORT-BASED EGRESS SCHEDULING POLICIES

Port-based Egress Scheduling Policies is a Quality of Service (QoS) feature where the policy defines a) the port's overall egress scheduling rate limit and b) up to eight (8) priority level schedulers on egress with options for a Committed Information Rate (CIR) and an overall rate limit for the priority level scheduler.

Port-based Egress Scheduling Policy can be either:

- the parent scheduler for service or subscriber schedulers providing scheduling at the port level, or
- the direct parent scheduler associated with service or subscriber polices providing the scheduling for these queues at the port level.

When a Port-based Egress Scheduling Policy is directly associated with service or subscriber queues, the port scheduler hierarchy allocates bandwidth on a per-class or per-priority basis to each service or subscriber queue. This allows a provider to manage the available egress port bandwidth on a "service tier" basis ensuring that during egress port congestion a deterministic behavior is possible on an aggregate perspective. While this provides an aggregate bandwidth allocation model, it does not inhibit per service or per subscriber queuing. The benefit of this single scheduler policy model is that the bandwidth is allocated per priority for all queues associated with the egress port. This allows a provider to preferentially allocate bandwidth to higher priority classes of service independent of service or subscriber instance. In many cases, a subscriber may purchase multiple services from a single site (VoIP, HSI, Video, etc.) and each service may have a higher premium value relative to other service types. If a subscriber has purchased a premium service class, that service class should get bandwidth before another subscriber's best-effort service class. An Aggregate Rate Limit can be defined, allowing per-service instance or per-subscriber instance aggregate SLA and a class-based port bandwidth allocation function.

When a Port-based Egress Scheduling Policy is used as the parent scheduler for multiple service or subscriber scheduler policies, multiple services or subscribers can have independent scheduler policy definitions while the independent schedulers receive bandwidth from the Port-based Egress Scheduler at the port level. By associating the two scheduler policies, available egress port bandwidth may be allocated fairly or unfairly depending on the desired behavior. When using the Port-based Egress Scheduling Policy in conjunction with service/subscriber schedulers, bandwidth is allocated on a per-service or per-subscriber basis as opposed to a per-class basis. A common use of this model is for a carrier-of-carriers mode of business. The goal of a carrier is to provide segments of bandwidth to providers that purchase that bandwidth as services. While the carrier does not concern itself with the interior services of the provider, it does however care a great deal in how congestion affects the bandwidth allocation to each provider's service. As an added benefit, this two policy approach provides the carrier with the ability to preferentially allocate bandwidth within a service or subscriber context through the service or subscriber level policy without affecting the overall bandwidth allocation to each service or subscriber.

Note Port-based Egress Scheduling Policies are supported on Ethernet and LAGs. Support for all other port types, including SONET/SDH, has been added in Release 5.0.R5.

NEW FEATURES IN 5.0.R3

The following sections describe new features added since 3.0(.R1) to the Release 5.0.R3 of 7710 SR OS. Note that for the 7710 SR OS, the first 5.0 version is 5.0.R3; there were no 5.0.R1 or 5.0.R2 releases.

FEATURES FROM 7750 SR R4.0

In Release 5.0, the 7710 SR features are being released in coordination with the Alcatel-Lucent's 7750 SR OS for the 7750 SR and 7450 ESS OS. The numbering of the 7710 SR releases is being coordinated with the 7750 SR and the 7450 ESS, so that a feature will be introduced on all three product lines at the same time and in the same release.

In addition, 7710 SR OS Release 5.0 also inherits all of the 7750 SR OS Release 4.0 features subject to limitations imposed by hardware such that the 7710 SR and 7750 SR have software feature parity.

Summary descriptions for the 7750 SR OS Release 4.0 features incorporated into the 7710 SR can be found in the 7750 SR OS Release 4.0.R4 Software Release Notes (Document Number: 93-0137-04) or later with more complete feature descriptions found in the 7710 SR OS software guides listed in Release 5.0.R25 Documentation Set on page 4 above.

Following is the list of features incorporated into 7710 SR OS from Release 4.0 of 7750 SR OS.

- Services
  - IP Interworking VLL
  - Default SAP on a Dot1Q Port
  - Layer 2 Protocol Tunneling (L2PT)
  - BPDU Translation
  - Multiple Spanning Tree Protocol (MSTP)
  - Interface-Triggered MAC Flush
  - Inter-AS IP-VPN (Model B)
  - OSPF CE-PE for IP-VPNs

  - IP-VPN Direct Route Comparison of BGP and MP-BGP Learned Routes
  - RFC2684 RBE Termination on ATM IES and IP-VPN SAPs
  - Draft Rosen Data MDT Support
  - Spoke Termination on IP-VPNs
- Triple Play Services Delivery Architecture (TPSDA)
  - Efficient Egress Multicast Replication
  - Enhanced Subscriber Management
  - Residential Split Horizon Groups
  - IP Multicast on Residential SAPs
  - Multicast VPLS Registration
  - Secure MAC Learning
  - MAC Learning Protection
  - ARP Reply Agent
  - Multipoint Shared Queuing
  - Policy-Based Forwarding for VPLS
  - QoS Policy Runtime Instantiation
  - Subscriber-Host Connectivity Verification (SHCV)
  - Routed CO and Layer 3 Group Interfaces
  - Layer 3 Subscriber Interfaces
  - Per-VC/Per-SDP Octet Counters
  - RADIUS Authentication of DHCP Sessions
  - Web Portal Redirection
- System
  - Non-Stop Routing for PIM and IGMP Graceful Restart
  - APS Multi-Chassis Support
  - LACP Enhancements
  - Network Time Protocol (NTP)
  - Reliance File System Integration
- Management
  - Generic Command Scheduling
  - SSHv2 Support
  - SNMP Support for ICMP Ping and Trace-Route
  - ATM ILMI 3.1 and 4.0 Support
- Routing
  - IPv6 Support (includes: OSPFv3, IS-IS for v6, IPv6CP, IPv6 BGP, 6 over 4 tunnel)
  - Multicast Extensions to BGP (MBGP)
  - Anycast RP for PIM-SM
- MPLS
  - Equal Cost Multi-path (ECMP) Support for LDP

- OAM
  - Bi-Directional Forwarding Detection (BFD)
  - SAA Enhancements
  - Virtual Circuit Connectivity Verification (VCCV) Ping for VLL
  - LDP Status Signalling
  - ATM-PING OAM Loopback Enhancement

HARDWARE

The following sections describe new hardware supported in 5.0.

- Common Equipment
  - 7710 SR-c4 Integrated Shelf on page 15
  - 7710 SR 9 Gbps Control and Forwarding Module (CFM) on page 15
  - 7710 SR-c4 Power Entry Modules (PEMs) on page 16
- MDAs
  - 5-port Gigabit Ethernet MDA Rev B on page 16
  - 2-port OC-48c/STM-16c MDA on page 16
  - 4-port ATM OC-12c/STM-4c MDA on page 16
- CMAs
  - 4-port DS3/E3 CMA on page 16
  - 2-port OC-12c/STM-4c CMA on page 17

7710 SR-c4 Integrated Shelf

- 7710 SR 4 CMA integrated shelf (3HE02177AA)

The 7710 SR-c4 integrated shelf supports 4 CMA interface positions and includes one air filter (3HE02183AA) and one fan tray (3HE02182AA). The shelf accepts one Control and Forwarding Module, one Chassis Control Module and up to three (3) AC or DC Power Entry Modules (PEM).

Both CMAs and MDAs (on MCMs) are supported in this shelf. The shelf supports up to four (4) CMAs, one (1) MDA and two (2) CMAs or two (2) MDAs.

7710 SR 9 Gbps Control and Forwarding Module (CFM)

- 7710 SR 9 Gbps Control and Forwarding Module (CFM) (3HE02175AA)

The 7710 9G CFM has one 9 Gbps flexible fast path and one control CPU. Due to the centralized architecture, there is no requirement for a switching fabric in the 7710 SR. The ingress forwarding path is directly connected to the egress forwarding path on the CFM.

7710 SR-c4 Chassis Control Module (CCM)

- 7710 SR-c4 Chassis Control Module (CCM) (3HE02181AA)

The 7710 SR-c4 Chassis Control Module (CCM) provides forward facing user interfaces for system functions such as system LEDs, console and management ports, compact flash, alarm indicators and alarm cut-off. The CCM interfaces are controlled by the CFM. The console port is DTE wired.

7710 SR-c4 Power Entry Modules (PEMs)

- 7710 SR-c4 DC Power Entry Module (PEM) (3HE02179AA)
- 7710 SR-c4 AC Power Entry Module (PEM) (3HE02180AA)

7710 SR-c4 AC or DC optionally redundant Power Entry Modules (PEMs) are inserted from the rear of the 7710 SR-c4 shelf. The system may remain fully functional with only one (1) PEM operating. Up to three (3) PEMs may be installed for power redundancy.

5-port Gigabit Ethernet MDA Rev B

The 5-port GigE MDA - SFP Rev B (3HE01615AA) is a 1000BASE Ethernet MDA which supports five (5) pluggable SFP optics. The Rev B MDA is a replacement for the existing 5-port GigE MDA - SFP (3HE00025AA). The Rev B MDA has on-board prioritization logic that can prioritize traffic based on IEEE 802.1p bits or DSCP bits.

The MDA supports all of the existing GigE SFPs. When an optical SFP is used, only the 1000BASE-T operation is supported. When a copper GigE SFP is used, the 10/100/1000BASE-T operation is supported. This multi-speed support with copper SFPs is a new feature of this Rev B MDA.

2-port OC-48c/STM-16c MDA

The 2-port OC-48c/STM-16c MDA - SFP (3HE00043AA) is a two (2) port OC-48c/STM-16c SONET/SDH MDA which supports two (2) pluggable SFP optics. The MDA supports the same OC-48/STM-16 SFPs common to the Alcatel-Lucent Service Router portfolio.

4-port ATM OC-12c/STM-4c MDA

The 4-port ATM OC-12c/STM-4c MDA - SFP (3HE00071AA) is a four (4) port ATM MDA which supports four (4) pluggable SFP optics. The MDA supports the same OC-3/STM-1 and OC-12/STM-4 SFP optics modules supported on the existing OC-12/STM-4 SONET/SDH cards. Note that 7710 SR OC-12 SFP modules are multirate and can support both OC-3/STM-1 and OC-12/STM-4 rates.

Both OC-12/STM-4 and OC-3/STM-1 rates are supported and can be selected on an individual port basis.

The ATM SAR function is performed within the Media Dependant Adapter (MDA). The MDA supports a 2.5 Gbps SAR (bi-directional). SAR hardware provides functions which include:

- ATM layer and AAL5 layer support
- Traffic Management support with per-VC queueing and ATM layer shaping
- F5 OAM support

4-port DS3/E3 CMA

The 4-port DS3/E3 CMA (3HE01021AA) is a four (4) port DS3/E3 CMA with eight (8) 1.0/2.3 BNC connectors (4 Tx, 4 Rx). DS3/E3 ports provide 75 Ohm impedance. Clear channel operation is supported (i.e., no channel group support). Frame Relay, PPP and Cisco-HDLC encapsulations can be configured on a port-by-port basis.

DS3/E3 CMA ports can be configured as either access or network ports. To be configured as network ports, PPP encapsulation must be selected. DS3/E3 ports must be access ports on the 7750 SR. Therefore, 7710 SR DS3/E3 ports configured for network mode must be connected to another 7710 SR or third party device supporting MPLS over DS3/E3. DS3/E3 ports configured for network mode cannot be connected to 7750 SR channelized MDAs as these cards currently only support access mode.

2-PORT OC-12c/STM-4c CMA

The 2-port OC-12c/STM-4c CMA - SFP (3HE02185AA) is a two (2) port OC-12c/STM-4c SONET/SDH CMA which supports two (2) pluggable SFP optics. The CMA supports the same OC-3/STM-1 and OC-12/STM-4 SFP optics modules supported on the existing OC-12/STM-4 SONET/SDH cards. Note that 7710 SR OC-12 SFP modules are multirate and can support both OC-3/STM-1 and OC-12/STM-4 rates.

Both OC-12c/STM-4c and OC-3c/STM-1c rates are supported and can be selected on an individual port basis.

SERVICES

The following sections describe new services features in Release 5.0:

RADIUS AUTODISCOVERY FOR VPLS

In prior releases, the VPLS configuration was manually provisioned via CLI or NMS. Whenever a new Provider Edge (PE) device was added to a VPLS, not only did the new PE have to be provisioned, but also all other PEs had to be re-provisioned to add the new PE. Even though a good NMS system solves this problem, many service providers still require an auto-discovery mechanism for VPLS.

IETF draft-ietf-l2vpn-radius-pe-discovery-02 proposes a mechanism using RADIUS to discover other PEs in the same VPLS, so that PE-to-PE pseudowires (PWs) can be established automatically with minimal or zero provisioning.

High-level features of the Release 5.0 implementation:

- PE discovery only
- Supports VPLS services.
- Auto-SDP is used for dynamic creation of GRE or LDP SDPs.

SDP/PSEUDOWIRE (PW) STANDBY BACKUP

In deployments where a backup pseudowire (PW) terminates on a different peer, MPLS protection mechanisms cannot address these scenarios.

7710 SR OS allows the creation of pre-provisioned backup SDPs/pseudowires where the switching to the backup is controlled by the following triggers:

- T-LDP session to the peer PE times out
- PW label withdrawal or PW status message received from peer PE

The status of the VLL stays UP during the switching operation, so no AIS or LMI status messages are generated by the PE for ATM or Frame Relay SAPs, respectively.

VLL Spoke Switching/Manual Pseudowire Stitching

The VLL Spoke Switching feature, also called Manual Pseudowire Stitching, provides the user with the ability to create a VLL service by cross-connecting two spoke SDPs. The objective of this feature is to allow the scaling of VLL and VPLS services in a large network in which the otherwise full mesh of PE devices would require thousands of Targeted LDP sessions per PE node.

Services with one SAP and one spoke SDP are created normally on the PE; however, the target destination of the SDP is the 7710 SR PW switching node instead of what is normally the remote PE. In addition, the user configures a VLL service on the PW switching node using the two SDPs.

The PW switching node acts in a passive role with respect to signalling of the PWs. It waits until one or both of the PEs sends the label mapping message before relaying it to the other PE. This is because it needs to pass the Interface Parameters of each PE to the other.

A PW switching point TLV is inserted by the switching PW to record its system address when relaying the label mapping message. This TLV is useful in a few situations:

- It allows for troubleshooting of the path of the PW especially if multiple PW switching points exist between the two PEs.
- It helps in loop detection of the T-LDP signalling messages where a switching point would receive back a label mapping message it had already relayed.
- The switching point TLV is inserted in PW status messages when they are sent end-to-end or from a PW switching node towards a destination PE.

PW OAM is supported for the manual switching PWs and allows the 7710 SR PW switching node to relay end-to-end PW status notification messages between the two PEs. The 7710 SR PW switching node can generate a PW status notification and to send it to one or both of the PEs by including its system address in the PW switching point TLV. This allows a 7710 SR PE to identify the origin of the PW status notification message.

IGMP on CE-PE links for IP-VPNs

Prior to Release 5.0, it was assumed that CEs in an IP-VPN environment were routers and that hosts were attached to those routers and not directly to the PE. Release 5.0 supports the IGMP protocol on the PE-CE interface of VPRNs.

Control Word Support for PWE3

The Control Word for PWE3 feature provides the option to add a Martini control word as part of the PW encapsulation for PW types for which the control word is optional. These are Ethernet PW, ATM N:1 cell mode PWs (VCC and VPC) and VT PW.

The control word is needed because when ECMP is enabled in the core network, packets of a given PW may be spread over multiple ECMP paths if the hashing router mistakes the PW packet payload for an IPv4 or IPv6 packet. This occurs when the first nibble following the service label corresponds to a value of 4 or 6.

The T-LDP control plane has been enhanced to allow the peer 7710 SR PEs to indicate whether the control word should be used or not on these types of PWs.

Prior to Release 5.0, a control word was only used on PW types for which it is mandatory, that is, FR PW, and ATM AAL5 SDU VCC PW.

IPv6 L3 Filter on L2 Service

For Layer 2 services, SAP/SDP ingress and egress filters have been extended to support IPv6 filters.

Restrictions:

- IPv6 ACLs are not supported on an L2 SAP which has a QoS mac-criteria associated
- MAC ACLs are not supported on an L2 SAP which has a QoS ipv6-criteria or an IPv6 ACL associated.

ATM Mirroring on ATM SAP

ATM Mirroring on an ATM SAP extends the service mirroring feature to include mirror sources with SAP type of ATM. Mirroring is supported on the following services:

- IES
- VPRN
- VPLS
- Epipe
- Ipipe
- Apipe VLL service with the AAL5 SDU mode, that is, atm-sdu spoke-sdp type.

This feature is supported on ATM MDAs.

Mirror destinations for ATM mirroring must be ATM SAPs and must not be part of an APS group.

ATM SAPs of an Apipe with N:1 cell mode, that is, atm-vcc, atm-vpc, and atm-cell spoke-sdp types, cannot be ATM mirror sources.

TPSDA

The following features are new to the Triple Play Service Delivery Architecture (TPSDA) in Release 5.0.

VPRN Routed CO

The TPSDA Routed CO model allows a provider to resell services while providing direct DSLAM connectivity. The operator can create a subscriber interface to identify the subscriber subnets and attach a collection of group interfaces that terminate one or more SAPs.

In 7710 SR OS Release 5.0, VPRN services are now supported using the Routed CO model. Because of the nature of VPRN, the service to each subscriber is limited to the VPN.

Wholesale/Retail VPRN Routed CO

With the VPRN Routed CO model, there two different modes of operations. In the first mode, the wholesaler creates a VPRN for the retailer and configures access from the subscribers as well as to the retailer network. Any further action will be as if the VPRN is a standalone router running Routed CO. This provides maximum flexibility to the retailer while minimizing the work and involvement of the wholesaler. Access cannot be shared among retailers unless a subscriber per SAP is used. This requires that the wholesaler maintains a different access node (DSLAM) for each retailer; this method does not scale well.

The second wholesale/retail model defines a wholesaler VPRN. The connections that are common to the access nodes are distributed to many retail instances (each in its own VPRN). Upstream subscriber traffic ingresses into the wholesaler instance, and after identification, it is then forwarded into the retail instance. The reverse will occur for traffic in the other direction. The wholesale/retail traffic flow is controlled with minimal communication with RADIUS. A RADIUS policy is defined in the wholesaler instance. The RADIUS response used during the subscriber instantiation provides the service context of the retailer VPRN. If the wholesaler has a retail business, the operator can configure a separate VPRN for their retail services.

DHCP Proxy Server

Within the Service Provider environment, the majority (if not all) of subscriber management systems rely heavily on RADIUS (RFC2865) as the means for identifying and authorizing individual subscribers into the network.

The Proxy DHCP Server capability will enable the deployment of DHCP into a providers network, by acting as a proxy between the downstream DHCP devices and the upstream RADIUS-based subscriber management system.

- Supported for L2 VPLS, L3 IES and L3 VPRN configurations
- Interacts with downstream DHCP Client devices (and DHCP Relay Agents in the path)
- Interfaces with RADIUS to authenticate DHCP requests
- Interfaces with RADIUS to receive all the necessary IP information to properly respond to a DHCP Client
- Ability to override the allocated IP address lease time, for improved IP address management

The Proxy DHCP Server is standards based and interfaces with the DHCP Clients and upstream DHCP Server using IETF defined standards RFC2131, RFC2132 and RFC3203.

The Proxy DHCP Server is tightly integrated with the 7710 SR feature set, supporting the following features:

- RADIUS-based authentication and re-authentication.
- the existing persistence/High Availability solution for per subscriber Lease state

- the existing Lease State dependent set of features
  - Dynamic ARP population
  - ARP Reply Agent
  - Anti-Spoofing Filters
  - MAC Pinning

The Proxy DHCP Server supports the following six different deployment scenarios:

1. DHCP Client to a RADIUS Server for authentication and all its IP information.
2. DHCP Client to a DHCP Server for all its IP information; no RADIUS Server involvement.
3. DHCP Client to both a RADIUS Server for authentication and a DHCP Server for its IP information.
4. Change of Authorization (CoA) initiated from a RADIUS Server, which can trigger a DHCP FORCERENEW message to the specified DHCP Client.
5. DHCP FORCERENEW initiated from an upstream DHCP Server.
6. CoA initiated from a RADIUS Server containing the full configuration for the given host.

DHCP Client/Server Spoofing

Enhanced Subscriber Management (ESM) uses information from DHCP to establish subscribers and hosts. The 7710 SR can now send a DHCP FORCERENEW message which will force a client to send a DHCP renew message and allow the operator to refresh the policy state (for example, the SLA-Profile) in the node.

In some situations, the 7710 SR will drop the final DHCP Ack coming from the DHCP server. In such situations, the provided IP address is normally not returned to the DHCP server's IP pool. Spoofing the DHCP release to the server (on behalf of the client) allows the server to return the IP to the pool.

DHCP Client Mobility

DHCP Client Mobility allows the node to use host monitoring to remove network and server state when a host is removed locally. This allows for MAC addresses learned and pinned to move based on policy parameters.

DHCP Client Mobility allows clients to move from one SAP to another SAP in the same service. This is only applicable in a VPLS service and group interfaces.

The first DHCP message on the new SAP with the same MAC address (and IP address for group interfaces) will trigger SHCV and will always be discarded.

SHCV will check that the host is no longer present on the SAP where the lease is currently populated to prevent spoofing. When SHCV detects that the host is not present on the original SAP, the lease state will be removed. The next DHCP message on the new SAP can instantiate the host.

RADIUS Framework Improvements

In the 7750 SR OS Release 4.0 DHCP authentication process, the 7750 SR elements accepted only an accept/deny indication from the RADIUS server. In many broadband aggregation networks, RADIUS is used for more than authentication; it is also used for subscriber-related configuration.

In addition to the above functionality, the 7710 SR OS Release 5.0 subscriber management for RADIUS has been enhanced to allow an operator to configure subscribers with the following:

- RADIUS-based assignment of subscriber-id, subscriber-profile and sla-id

  In 7750 SR OS Release 4.0, ESM extracted these objects using the subscriber-identification scripts as part of Enhanced Subscriber Management. These objects can now be assigned as a part of the RADIUS authentication process and provide an alternative mechanism for subscriber identification.

- RADIUS-based IP configuration (IP-address, default gateway, dns)

  These configuration parameters can be obtained from the RADIUS server, providing an alternative to DHCP server assignment of these parameters. This can be very useful in distributing DHCP server functionality in large networks. In this case, the 7710 SR element performs the DHCP server role towards the client. The selection of the IP address is however driven by RADIUS.

  This enhancement provides an alternative to a full DHCP server-based solution. Many providers making the switch between PPP and DHCP, do not have operational experience with provisioning and operating a DHCP server. This feature allows providers to use their current mode of IP assignment utilizing RADIUS. With this feature, the 7710 SR platform will behave as a DHCP server without local storage and management of the IP pools.

All of the above scenarios can be combined. The required information is returned by RADIUS response message in the form of VSAs providing a very flexible way of managing subscribers.

In addition to the above enhancements, RADIUS-based subscriber accounting is supported where an accounting record is created on the RADIUS accounting server for every newly created subscriber host. Updated accounting information can be regularly sent by the 7710 SR to the RADIUS server. At the termination of the subscriber session, an update with the latest accounting information is sent and the accounting record on the RADIUS accounting server is closed.

MULTICAST CAC POLICIES

In Triple Play aggregation networks, the available Broadband TV (BTV) channels frequently exceeds the capacity of BSA-DSLAM link (2nd mile), and it is not feasible to distribute all available channels to DSLAMs. Therefore, the IGMP protocol is used to register joins and send only requested channels.

7710 SR OS allows limiting the maximum number of channels which can be distributed on a given IP interface or VPLS SAP through the IGMP Snooping feature. However, this level of control, which is basically working on first-come-first-serviced basis, is not sufficient in an environment where not all channels are equal in their “priority” and bandwidth usage.

The BTV offerings typically consist of channel bundles with different levels of priority (for example, some channels should always be available in the DSLAM as they are part of mandatory offering). To reflect these aspects, sophisticated admission control is required.

7710 SR OS Multicast CAC Policies allow grouping of multicast group addresses into bundles and characterizing each channel in terms of the bandwidth it represents on the link. In addition, a Multicast CAC Policy defines constraints on the amount of channels and bandwidth per bundle which is allowed on a given IP interface or VPLS SAP. Based on these constraints, the 7710 SR IGMP function will accept/refuse individual IGMP Joins received.

The 7710 SR supports Multicast CAC Policies on VPLS SAPs, spoke SDPs, mesh SDPs and Routed CE interfaces. Within a VPLS, Multicast CAC policies are linked through IGMP Snooping.

GSMP Support to Adjust H-QoS Policies

General Switch Management Protocol version 3 (GSMPv3) is a generic protocol that allows a switch controller node to establish and maintain connections with one or more nodes to exchange operational information. Several extensions to GSMPv3 exist in the context of broadband aggregation. These extensions were proposed to allow GSMPv3 to be used in a broadband environment as additional information is needed to synchronize the control plane between Access Nodes (e.g. DSLAM) and Broadband Network Gateways (e.g. BRAS).

In the TPSDA framework, 7710 SR nodes fulfill some of the BRAS functionality, where per subscriber QoS enforcement is one of the most important aspects. To provide accurate per subscriber QoS enforcement, the network element not only knows about the subscriber profile and its service level agreement, but it is aware of the dynamic characteristics of the subscriber access circuit.

The most important parameter in this context is the subscriber-line capacity (dsl sync-rate). This information can be then used to adjust parameters of aggregate scheduling policy.

GSMPv3 is also used to convey other OAM information between switch controller and access switch. In this context, the 7710 SR can operate in two roles:

- Intermediate controller - 7710 SR extracts only relevant information and maintains the communication between DSLAM and BRAS intact (“man-in-the-middle” role).
- Terminating controller - the 7710 SR fulfils full role of the BRAS

The DSL forum working documents recommend that a dedicated L2 path (e.g. a VLAN in an Ethernet aggregation network) be used for this communication to provide a certain level of security. The actual connection between the DSLAM and BRAS is established at TCP level, and then individual L2CP messages are transported.

The 7710 SR element intercepts this communication and extracts all relevant information. For “man-in-the-middle” configurations, this means terminating the TCP connection with the DSLAM and setting up a new connection with the BRAS (similar to a proxy function). Alternatively, the DSLAM can maintain two connections with two controllers (BRAS and 7710 SR) and report respective information to both assuming this functionality is supported in the DSLAM.

In order to allow adjustment to scheduling policies, there is a correlation between line-ids in L2CP messages and scheduling policies configured in 7710 SR elements.

Multi-chassis LAG

Multi-chassis LAG (MC-LAG) is an extension to the LAG feature to provide not only link redundancy but also node-level redundancy. This feature is not defined in any IEEE standard, but Alcatel-Lucent has developed a proprietary solution.

A proprietary messaging between redundant-pair nodes supports coordinating the LAG switchover.

Multi-chassis LAG supports LAG switchover coordination: one node connected to two redundant-pair peer nodes with the LAG. During the LACP negotiation, the redundant-pair peer nodes act like a single node using active/stand-by signalling to ensure that only links of one peer nodes is used at a time.

Multi-chassis LAG is expected to be used primarily for supporting SAPs connecting to a VPLS service with pseudowires on the network core side, so that MAC Flush messages can be used to ensure a loop-free and consistent end-to-end topology.

Please note, that using MC-LAG for dual homing is not usable in ring topologies in the access network. The solution for such network configurations is not part of this feature.

Multi-Chassis Synchronization

Multi-Chassis Synchronization is an Alcatel-Lucent proprietary mechanism to synchronize dynamic state information (IGMP snooping and subscriber lease state information) between redundant pair nodes (inter-node state synchronization). This mechanism is not a standalone feature, but rather a mechanism required to provide a solution for redundant access (dual homing) in combination with subscriber management and Multi-chassis LAG. This mechanism requires that the involved nodes are running at the same system times (NTP/SNTP). It is also required that the local service configuration of the involved nodes are similar.

RADIUS Authentication of DHCP Sessions for VPRN SAPs

7750 SR OS Release 4.0 introduced RADIUS Authentication of DHCP Sessions that provided the functionality necessary to authenticate a DHCP-based subscriber requesting a connection to service provider's network via a RADIUS server on SAPs associated with IES and VPLS services.

In addition to the above, 7710 SR OS Release 5.0 adds support for this feature to SAPs associated with VPRN services. The same “subscriber policy” mechanism is used for authenticating DHCP sessions on VPRN SAPs as is already used for IES and VPLS SAPs.

DHCPv6 Relay

In the stateful address auto-configuration model, hosts obtain interface addresses and/or configuration information and parameters from a DHCPv6 server. Servers maintain a database that keeps track of which addresses have been assigned to which hosts. The stateful auto-configuration protocol allows hosts to obtain addresses, other configuration information or both from a server. This is especially important for service providers who want to use IPv6 as the triple play service delivery architecture.

The implementation complies with RFC 3315 “Dynamic Host Configuration Protocol (DHCP) for IPv6”.

IPv6 Local Proxy Neighbor Discovery

Local Proxy Neighbor Discovery (ND) is similar to local proxy ARP. This feature is useful in a residential bridging environment where end users are not allowed to communicate to each other directly.

Local Proxy ND can be enabled on an IPv6 router interface, and when enabled, the router will behave as follows:

- Respond to all neighbor solicitation messages received on the interface for IPv6 addresses in the subnet(s) unless disallowed by policy.
- Forward traffic between hosts in the subnet(s) of the interface.
- Drop traffic between hosts if the link-layer address information for the IPv6 destination has not been learned.

DHCPv6 Prefix Delegation

DHCPv6 Prefix Delegation is an 7710 SR OS implementation of the prefix delegation options described in RFC3633 that provides a mechanism for automated delegation of IPv6 prefixes using DHCP. The mechanism is intended for delegating a long-lived prefix from a delegating router to a requesting router, across an administrative boundary, where the delegating router does not require knowledge about the topology of the links in the network to which the prefixes will be assigned. For example, the delegating router can get a /48 prefix via DHCPv6 and assign /64 prefixes to multiple requesting routers.

In 7710 SR OS, multiple IPv6 prefix ranges can be configured for DHCPv6 Prefix Delegation where each prefix range is configured with the following:

- A matching client DUID.
- An optional IAID for the requesting router.
- The length of time a prefix should be preferred.
- The length of time a the prefix is valid.

STATIC IP-ONLY HOST

The Static IP-only Host feature allows for a static host definition in such a way that the MAC address becomes an optional parameter.

If the MAC address is not specified for a static host, the system learns the respective MAC address dynamically from ARP packets (ARP request or a gratuitous ARP) generated by the host with the specified IP address. On a VPLS service, this happens only if the ARP Reply Agent function is enabled on a given SAP. On L3 services (IES or VPRN), the ARP packets are always examined and hence, no additional configuration is needed.

The learned MAC address is handled as a MAC address of the static host with an explicitly defined MAC address. This means the following:

- The MAC address is not aged by a MAC aging or an ARP aging timer.
- The MAC address is not moved to another SAP as a consequence of a re-learning event, for example, if the same MAC address is seen on another SAP which would normally mean the MAC is re-learned and moved.
- The MAC address is not flushed from FDB due to SAP failure or STP flush messages.

Every time the given static host uses a different MAC address in its ARP request, the dynamic MAC learning process is performed; the old MAC address is overwritten by the new one.

The learned MAC address is not persistent (that is, the static host is not saved in the persistency file) which means that the “service discontinuity” of such host is proportional to its ARP cache timeout.

Static IP-only Hosts will interact with other features as follows:

- Anti-spoofing Filters (all services)

  When a Static IP-only Host is configured on a given SAP, both anti-spoof types IP and IP-MAC are supported. Static hosts for which the MAC address is not known will not have an anti-spoofing entry. The anti-spoofing is only added after the corresponding MAC has been learned. As a consequence, all traffic generated by the host before sending any ARP packets is most likely dropped.

- Enhanced Subscriber Management (all services)

  ESM is supported in a combination with a Static IP-only Host. It is assumed that IP-MAC anti-spoofing is enabled. The resources (queues, etc.) are allocated at the time such a host is configured, although they are used only after the anti-spoofing entry has been installed.

- MAC Pinning (VPLS service only)

  The dynamically learned MAC address of the Static IP-only host is considered as a static MAC and as such is not affected by a “no mac-pinning” command.

- ARP Reply Agent (VPLS service only)

  The ARP Reply Agent may be enabled on a SAP where the Static IP-only Host is configured. Besides the regular ARP Reply Agent functionality (reply to all ARP requests target-

ing the given host's IP address), learning of the host's MAC address (as described above) is performed. Obviously, as long as no MAC address has been learned, no ARP replies on behalf of such a host should be expected. Enabling of ARP Reply Agent is optional for a SAP with Static IP-only Hosts.

Following are some corner cases when using the Static IP-only Host feature:

- A host has statically configured ARP entry for its default gateway, so the host is most likely not going to send out any ARP packets, no learning of its MAC address is performed.
- The default gateway sends periodically gratuitous ARP packets which populates host's ARP cache, so the host will likely not originate any ARP packets, so no learning of its MAC address is performed.

Support for IP Multicast on Residential SAPs

7710 SR OS now supports IP multicast on residential SAPs and provides scaling on VPLS up to 2K SAPs per tree per Flexible Fast Path Complex. This feature is useful in TPSDA deployments where the DSLAM does not perform any multicast replication.

Quality of Service

The following are new Quality of Service features in Release 5.0.

- Time-of-Day Policies on page 26
- QinQ Dot1P QoS Selection on page 26

Time-of-Day Policies

In a Triple Play environment and in an effort to compete with BRAS functionality, the 7710 SR allow the providers to create time-based policies for ACL and QoS policies. The user will be able to create either a full policy change which can be handled by using multiple policies and controlling their assignment to the subscriber by means of CRON service. Time-of-Day Policies allow multiple ACL and QoS policies to be applied to a SAP, interface or Multi-Service Site. Statistics are collected and maintained for all policies. Alternatively, the user will be able to utilize a time stamp matching criteria for ACL and QoS policies.

Examples of use:

- The provider will be able to prevent P2P access during peak hours or alternatively rate limit it during peak hours.
- The provider will be able to allow new customers 30 days of Internet access as a trial without the need to re-provision their service to disallow it. If the customer signs up for the service a new ACL will be installed.

QinQ Dot1P QoS Selection

The 7710 SR can now be configured to optionally re-mark the top dot1p bits only at the Ethernet egress SAP with QinQ encapsulation. By default, both dot1p values (outer and inner) are re-marked with the same value.

SYSTEM

The following are new system features in Release 5.0.

7710 SR-c12 MCM and CMA Support

In Release 3.0, the 7710 SR-c12 supported two (2) MCMs in slots 1-2 and 3-4. In Release 5.0, MCMs can be used in slots 5-6, 7-8, 9-10, 11-12 to fill an entire 7710 SR-c12 chassis with MCMs and MDAs.

In Release 3.0, CMAs were supported in slots 5 through 12 with an overall system limit of support for eight (8) CMAs. In Release 5.0, there is still a system limit of eight (8) CMAs in the 7710 SR-c12, but CMAs can now be used in slots 1 through 4 in addition to slots 5 through 12.

Hardware/CPM Filters and Per-Peer Queuing

Hardware/CPM filters and per-peer queuing allows users to allocate dedicated CFM hardware queues for certain traffic destined to the CPUs on the 7710 SR and set a corresponding rate-limit for the queues. These queues then feed the appropriate CPM/DMA queues.

A CPM filter is a hardware filter operation performed by the Flexible Fast Path complex on the CFM that applies to all the traffic going to the CFM CPUs. These filters can be used to drop or accept packets.

CPM filters are applied before IP reassembly. All encapsulation types are supported, e.g., Ethernet, PPP, etc.

Note that CPM filters take precedence over per-peer-queuing. CPM filters are evaluated before the per-peer-queuing feature is executed and also prior to Management Access Filters.

CPM filters and per-peer queuing to both the 7710 SR-c4 and 7710 SR-c12 routers.

The following CFM traffic management features are supported:

- Traffic classification using cpm-filters

  The packets going to the CFM are classified by the ingress Flexible Fast Path complex into classes, which defines the CPM/DMA queues for a given packet. The CPM filter can be used to further classify the packets.
- Queue allocation

  Queues 1 to 8 are the default queues, and must not be configured/removed by users. There are no rate limits for the default queues.

TCP Authentication Extension

The TCP Authentication Extension is a new TCP Option specified in draft-bonica-tcp-auth-05.txt. The feature allows an operator to change session authentication keys without resetting the TCP session. The option also allows new, more secure, MAC algorithms to be used. The new TCP Option specifies a Key ID and an algorithm type, in addition to the hash data. In 7710 SR OS Release 5.0, the hash is over the TCP data and TCP options. There are two hash types specified in the draft, AES-CMAC-96 and HMAC-SHA1-96.

DHCPv6 Lease Persistency

DHCPv6 lease state information is reconciled with the standby CFM, so the lease state information is maintained across a High Availability switchover.

In addition, DHCPv6 lease information is stored on the Compact Flash, so DHCPv6 lease information can be restored from the Compact Flash after a system reboot. Note that the neighbor database entries resulting from handling by DHCPv6 Relay will also be restored if neighbor-resolution is enabled.

IPv6 Applications: FTP and TFTP

FTP and TFTP are now supported for IPv6.

MANAGEMENT

The following sections describe new management features in Release 5.0.

Cflowd Support for VPRNs

Cflowd flow statistics collection are now supported within the context of VPRNs (RFC2547 IP-VPNs). In prior releases, cflowd was only supported on network and IES interfaces.

For VPRNs, the IP flow statistics maintain the VPRN context along with the IP flow index to separate IP flows with overlapping IP address space. Flow statistics will be advertised in v5 format and the interface index can be used to identify the VPRN instance.

IPv6 Management Enhancements

The following management capabilities are supported under IPv6:

- IPv6 TACACS+ (IPv6 server addresses are supported)
- IPv6 RADIUS (IPv6 server addresses are supported)
- IPv6 Syslog (target syslog host address is supported)
- IPv6 SNMP
- IPv6 SNMP ping and traceroute

MIB Changes

Support for the following MIBs have been added to Release 5.0:

- DISMAN-SCHEDULE-MIB (RFC3231)
- DISMAN-SCRIPT-MIB (RFC3165)
- TIMETRA-NTP-MIB
- TIMETRA-OSPF-NG-MIB (replacing TIMETRA-OSPF-MIB)
- TIMETRA-SCHEDULER-MIB
- TIMETRA-MC-REDUNDANCY-MIB
- TIMETRA-MSDP-MIB
- TIMETRA-MCAST-CAC-MIB
- TIMETRA-GSMP-MIB
- TIMETRA-DOT3-OAM-MIB

- TRANSPORT-ADDRESS-MIB (RFC 3419)
- MSDP-MIB (extracted from draft-ietf-mboned-msdp-mib-01.txt)
- DOT3-OAM-MIB

ROUTING

The following sections describe new routing features in Release 5.0:

MSDP

Using PIM-SM, multicast sources and receivers, by means of the multicast router closest to them, register with their local Rendezvous Point (RP). The RP maintains the information with regard to the sources and receivers for any particular group. RPs in other domains do not have any knowledge about sources located in other domains.

MSDP is a mechanism that allows RPs to share information about active sources. When RPs in remote domains hear about the active sources, they can pass on that information to their local receivers and multicast data can then be forwarded between the domains. A useful feature of MSDP is that it allows each domain to maintain an independent RP that does not rely on other domains, but it does enable RPs to forward traffic between domains. PIM-SM is used to forward the traffic between the multicast domains.

MSDP is required to provide inter-domain multicast services using ASM.

ANYCAST-RP FOR MSDP

AnyCast-RP for MSDP allows a customer to have fast convergence when a PIM/MSDP Rendezvous Point (RP) router fails by allowing receivers and sources to Rendezvous at the closest RP. The fundamental problem that is solved by RFC 3446 is that it relaxes an important constraint in PIM-SM, namely, that there can be only one group to RP mapping active at any time. The single mapping property has several implications, including traffic concentration, lack of scalable register decapsulation (when using the shared tree), slow convergence when an active RP fails, possible sub-optimal forwarding of multicast packets, and distant RP dependencies.

As a result, ISP backbones require a mechanism that allows definition of multiple active RPs per group in a single PIM-SM domain. Further, any such mechanism should also address the issues listed above.

The mechanism is intended to address the need for better fail-over (convergence time) and sharing of the register decapsulation load (again, when using the shared-tree) among RPs in a domain. It is primarily intended for applications within those networks using MBGP, Multicast Source Discovery Protocol and PIM-SM protocols, for native multicast deployment, although it is not limited to those protocols. In particular, Anycast-RP is applicable in any PIM-SM network that also supports MSDP.

Note: A domain deploying Anycast-RP is not required to run MBGP. The anycast address is used as the RP address in the RP's SA messages.

IS-IS Multi-topology

Using IS-IS IPv6 TLVs for IPv6, routing has been supported since 7750 SR OS Release 4.0. This is so-called “native IPv6 routing with IS-IS” and has a limitation that the IPv4 and IPv6 topologies need to be congruent; otherwise traffic may be silently discarded. The service provider needs to ensure the IPv4 topology and IPv6 topology are the same which can be difficult to manage.

In addition to the above, 7710 SR OS Release 5.0 IS-IS Multi-topology allows the service provider to have different topologies for IPv4 and IPv6.

The 7710 SR OS implementation is compliant with draft-ietf-isis-wg-multi-topology-11.txt.

The following MT topologies are supported:

- MT ID #0 - Equivalent to the “standard”' topology
- MT ID #2 - Reserved for IPv6 routing topology

6PE

6PE allows IPv6 domains to communicate with each other over an MPLS IPv4 core network. This architecture requires no backbone infrastructure upgrades and no reconfiguration of core routers, because forwarding is based on MPLS labels rather than on the IP header itself. It provides a very cost-effective strategy for IPv6 deployment.

From a control plane perspective, the main principles of the solution are:

1. The 6PE MP-BGP routers are dual-stack: IPv6 towards the CE and IPv4 towards the MPLS core.
2. MP-BGP is used between 6PE routers to exchange IPv6 reachability information. A 7710 SR router uses “IPv6 Explicit Null” label (label 2) for all the IPv6 prefixes that it advertises, but it also accepts an arbitrary label from its peers. When a labeled packet is received and the “IPv6 Explicit Null” label indicates that an additional IPv6 route lookup is required.
3. At recursion time, the ingress 6PE extracts the IPv4 address contained in the IPv4 mapped IPv6 address. Then the ingress 6PE resolves this IPv4 address (using the IPv4 routing table) in order to get the label associated to the LSP for this destination. This “IPv4 label” has been populated in the IPv4 tables through regular IPv4 MPLS control plane procedures using IPv4 IGP and IPv4 label distribution protocols, such as LDP. This “IPv4 label” is then stored along with the BGP label (2 labels) for the destination IPv6 subnet in the IPv6 forwarding table of the ingress PE router.

6PE requires that MP-BGP be extended to be able to bind an MPLS label to the IPv6 route. The extension is described in draft-ooms-v6ops-bgp-tunnel-04.txt. The 7710 SR also supports inter-AS models. The inter AS model will work when LDP is enabled between the two ASBRs and the IPv4-mapped addresses are used for next hops, or the two ASBRs are dual-stack routers, in which case IPv6 packets are forwarded natively between them.

Multiple OSPF Instances within an RTM

The ability to create multiple instances of OSPFv2 within a given RTM/VRF is accomplished by assigning an instance number to OSPF when creating the protocol instance. The absence of an instance number implies the initial (0) instance allowing existing configurations to function without modification.

Route distribution when there are multiple OSPF instances is handled as follows:

- The preference for internal and external preference for each OSPF instance is configurable.
- As a general rule, distribution remains as is where the lowest preference value wins, and in the case of a tie, the lowest cost wins, and finally the lowest instance ID wins.
- Each OSPF instance can be configured for inclusion or exclusion in the unicast, multicast or both routing tables.

Multiple OSPF instances is supported in the main RTM context only (not VPRN VRFs) and is not supported for OSPFv3.

NSR for OSPFv3

Non-Stop Routing (NSR) for OSPFv3 makes OSPFv3 hot redundant on a 7710 SR with two CFMs. When the active CFM fails, the standby takes over and all of the OSPFv3 states (neighbors, LSA database, etc.) remain intact.

The implementation for OSPFv3 NSR is similar to OSPFv2 NSR.

The following functionality is supported:

- OSPFv3 adjacencies remain intact after a CFM switchover. For example, all the OSPFv3 neighbors in the “UP” state will continue to be “UP”. Neighbors will not detect that the switchover occurred.
- OSPFv3 databases (LSA database, TE database if applicable) are be preserved across the switchover.
- All the OSPFv3 related statistics are preserved after a CFM switchover.

IGMP Snooping High Availability (HA)

IGMP Snooping HA makes IGMP snooping hot redundant on a 7710 SR with two CFMs. When the active CFM fails, the standby takes over and all the IGMP snooping functions and maintains state.

The following functionality is supported:

- IGMP snooping state remains intact after a CFM switchover
- Multicast traffic forwarding is not affected by a CFM switchover
- IGMP v1, v2, and v3 are supported

Sub-second VRRP Timer Support

Sub-second VRRP timers allow the primary 7710 SR to support sub-second VRRP transitions to a secondary router in 600 msec or less.

If the VRRP process has not run for an extended period of time (longer than the message-interval), the router checks the receive queue for messages before declaring the primary router down to mitigate any false failure detections.

Sub-Second VRRP timers are supported on the 7710 SR-c12 and 7710 SR-c4.

MPLS

The following sections describe the new MPLS features in Release 5.0:

LDP-over-RSVP-TE

The primary reason for LDP over RSVP-TE is to provide end-to-end tunnels that have two important properties, fast reroute in under 50 msec and traffic engineering. Neither of these two properties are available for LDP (LDP FRR still has limited coverage, particularly in ring topologies, and traffic engineering using LDP, that is, CR-LDP, has been deprecated in the IETF).

The solution is focused at large networks, where there are over 100 nodes in the network. Simply using end-to-end RSVP-TE tunnels will not scale. While an LER may not have that many tunnels, any transit node will potentially have thousands of LSPs, and if each transit node also has to deal with detours or bypass tunnels, this number can make the LSR overly burdened.

In addition, without inter-area RSVP-TE, the architecture of the network also becomes an issue. For scalability, isolation of local flaps, etc., it is prudent to design these networks with a core network connecting multiple smaller regional or metro networks. RSVP-TE in multi-area environments is significantly more complicated.

NSR for RSVP-TE

NSR for RSVP-TE makes RSVP-TE hot redundant on a 7710 SR with two CFMs. When the active CFM fails, the standby takes over and all the RSVP-TE states (sessions, neighbors, etc.) remain intact.

The following functionality is supported:

- RSVP-TE interfaces and neighbors remain intact and neighbors do not detect a CFM switchover.
- RSVP-TE sessions remain intact after a CFM switchover. For example, all the RSVP-TE LSPs in “UP” state continue to be “UP” and functional.
- SDPs using RSVP-TE remain intact after a CFM switchover.
- Traffic forwarding for RSVP-TE LSPs and SDPs is not affected by a CFM switchover (except traffic loss due to switch fabric change).
- Statistics preserved after a CFM switchover include: LSP up time, path up time and MPLS interface transmit and receive counters.

Note that NSR using fast reroute (FRR) one-to-one method LSPs are not hot-redundant in this release.

LDP Graceful Restart Helper

7710 SR OS implements Graceful Restart (GR) Helper for LDP to allow the 7710 SR to act as a GR helper when a neighboring router performs GR.

The 7710 SR implementation supports both LDP and T-LDP and complies with RFC3478.

RFC4379 LSP Ping and LSP Trace

The 7710 SR OS LSP Ping and LSP Trace now comply with RFC4379.

LSP Ping will interoperate with LSP Ping in earlier versions of 7710 SR OS based on draft-ietf-mpls-lsp-ping-ping-03.txt, but LSP Trace will not interoperate due to major changes made to the Downstream Mapping TLV plus other minor changes such as return codes.

ECMP Tree Building and Multi-path Exercising

With ECMP Tree Building, the LER builds the ECMP tree from ingress to egress for each destination LDP FEC. The various paths from the ingress to the egress for a FEC are determined using LSP Ping and LSP Trace multi-path TLVs support as described in RFC4379. The LSR performing the ECMP hash replies with the range of 127/8 addresses that can be used to exercise each of the possible paths out of the LSR for this FEC. The tree building process runs every 60 minutes by default. The interval is configurable in minute units from 60 to 1440.

Once the tree is built, the 7710 SR PE uses a periodic LSP Ping to exercise the paths discovered in the tree building step. The LSP Ping frequency is on the order of a minute.

OAM

The following sections describe the new OAM features in Release 5.0:

- IEEE 802.3ah Ethernet OAM on page 33
- VCCV Ping Enhancements on page 33
- BFD Enhancements on page 34

IEEE 802.3ah Ethernet OAM

IEEE 802.3ah clause 57 defines the Ethernet First Mile (EFM) OAM sub-layer. IEEE 802.3ah is a link-level Ethernet OAM that allows network operators the ability to monitor link operation and quickly determine the location of failing links or fault conditions.

In Release 5.0, 7710 SR OS supports the following IEEE 802.3ah features:

- OAM Capability Discovery
- Active (default) and Passive modes
- OAM configurable transmit interval - the number of ports running OAM is limited to 30 per system with 100 ms timers and 120 ports with timers greater than or equal to 1 second
- OAM Events: Link Fault
- OAM Remote Loopback
- OAMPDU Tunneling on port-based services

VCCV Ping Enhancements

Virtual Circuit Connectivity Verification (VCCV) Ping has been enhanced in Release 5.0 as follows:

- The format of the echo request and echo reply messages have been modified to comply with RFC4379. The Release 5.0 message formats are backward compatible with the Release 4.0 implementation.
- Supports the optional draft-martini OAM control word to encapsulate VCCV within a VLL. The use of the OAM control word assumes that the draft-martini control word is also used on the user packets. In Release 3.0, 7710 SR OS makes use of the draft-martini control word only when it is mandatory, that is, in an ATM AAL5 SDU VLL and in a FR VLL. In Release 5.0, the 7710 SR OS adds support for the optional draft-martini control word.

- VCCV Ping can be used across manually switched pseudowires. The Release 5.0 implementation is based on draft-hart-pwe3-segmented-pw-vccv-00.txt and uses the OAM control word for single-segment pseudowires.

BFD ENHANCEMENTS

Bi-Directional Forwarding Detection (BFD) support has been enhanced to include the following:

- Support for BFD for IGPs and static routes on IES interfaces.
- Support for BFD for PIM on VPRN interfaces.
- BFD support on ATM and Frame Relay interfaces.

All encap types supporting IPv4 are supported as all BFD packets are carried in IPv4 packets, including Frame Relay and ATM.

ENHANCEMENTS

RELEASE 5.0.R25

There are no new enhancements in Release 5.0.R25.

RELEASE 5.0.R24

The following section describes the new enhancement added since 5.0.R23 to 5.0.R24 of 7710 SR OS.

LDP

- LDP static-FEC "swap" entries will no longer be activated unless the next-hop in the routing table matches the configured next-hop of the static-FEC entry. [73836]

RELEASE 5.0.R23

The following section describes the new enhancements added since 5.0.R22 to 5.0.R23 of 7710 SR OS.

HW/PLATFORM

- It is no longer possible to downgrade chassis firmware using the CLI command “admin reboot upgrade”. [78316]

SYSTEM

- The “file version” command can now determine the version of the boot.ldr file. [24775]
- There are firmware updates for the following Ethernet MDAs:
  - m20-1gb-tx
  - m20-1gb-sfp
  - m5-1gb-sfp-b

  The firmware update includes support for K28.5 comma detection. [59338]

LDP

- The LDP tunnel-down-damp-time feature has been enhanced to be triggered when an unresolved FEC satisfies the following two conditions:
  - The route must have at least one next-hop
  - The FEC must have a label binding from at least one peer

  The route table next-hop and the FEC peer do not need to match as in previous releases. This enhancement allows a faster reactivation of the FEC or LDP tunnel without the additional churn of withdrawing the FEC label and re-originating. In all other cases, the tunnel-down-damp-time feature will not activate. [73891]
- This release introduces the ability to delay sending a label withdrawal following the deactivation of an LDP FEC to dampen the effect of network flapping events.

  When the label withdrawal delay is configured using the config>router>ldp>label-withdrawal-delay command, LDP delays the withdrawal of the <FEC, label> binding it distributed to its neighbors after it deactivates the FEC. LDP may deactivate an FEC as the result of a number of events such as link down, a route change impacting the FEC prefix, a loss of LDP adjacency, etc. The default value of the timer is zero, which means that the timer is disabled and LDP sends a label withdrawal immediately after deactivation of the FEC.

  Packets received on an LSR node for dampened FECs will be discarded by the ingress IOM. When the timer expires, LDP will send a label withdrawal for the FEC to all of its neighbors. If a new next-hop becomes available such that the FEC gets resolved again before the timer expires, LDP activates the FEC and resets the timer.

  The same FECs on a node in the ingress LER role will be dampened via a different timer, config>router>ldp>tunnel-down-damp-time. [74424]

RELEASE 5.0.R22

The following section describes the new enhancements added since 5.0.R21 to 5.0.R22 of 7710 SR OS.

SYSTEM

- The system-wide limit on the number of SDP keepalives has been increased from 256 to 1000 sessions. [75481]

RADIUS

- Prior to Release 5.0.R22, the 7750 SR operator RADIUS server access method was changed from direct to round robin and the direct access method could no longer be configured. Starting in Release 5.0.R22, the RADIUS server access method can be configured to be either round robin or direct. [68478]

RELEASE 5.0.R21

There are no new enhancements in Release 5.0.R21.

RELEASE 5.0.R20

The following section describes the new enhancement added since 5.0.R19 to 5.0.R20 of 7710 SR OS.

LDP

- Release 5.0.R20 introduces an enhancement to the TTL processing of a Penultimate Hop Pop (PHP) LSR or an egress LER when an LDP or RSVP LSP is used as a shortcut to forward packets to a BGP next-hop of a route or to the indirect next-hop of a static route. Traceroute messages received at the ingress of the LSP will be properly propagated at the egress of the LSP.

  A PHP LSR that forwards an IP packet after popping the label stack sets the TTL in the IP header to MIN(MPLS_TTL-1, IP_TTL).

  An egress LER that forwards an IP packet after popping the label stack sets the TTL in the IP header to MIN(MPLS_TTL-1, IP_TTL-1) when the context for the packet requires uniform mode TTL processing. Uniform mode TTL processing applies to IP packets which are not received in a service context, for example, packets received on an LSP used as a shortcut.

  MPLS_TTL refers to the TTL in the outermost label in the popped stack. [71104]

RELEASE 5.0.R19

There are no new enhancements in Release 5.0.R19.

RELEASE 5.0.R18

The following section describes the new enhancement added since 5.0.R17 to 5.0.R18 of 7710 SR OS.

MANAGEMENT

- The system will continuously check for log files with expired retention periods once every hour and delete as many files as possible during a ten (10) second interval. [65208]

RELEASE 5.0.R17

There are no new enhancements in Release 5.0.R17.

RELEASE 5.0.R16

The following section describes the new enhancements added since 5.0.R15 to 5.0.R16 of 7710 SR OS.

SYSTEM

- An alarm notification has been added to notify the administrator when the ingress or egress P-chip error rate has exceeded its threshold. In addition to the optional log message and SNMP trap, the timestamp of the last occurrence of the event and number of times the threshold was crossed can now be seen in the “show card detail” command. The number of events and last timestamp are not cleared in the event of an IOM or MDA reset. The timestamp of the last event should be used to make sure the IOM or MDA state has not changed since the last occurrence. Further investigation is necessary to determine the root cause. [57589, 67271, 67274]
- An alarm notification has been added to notify the administrator when the Transmit XPL Data Error counter has exceeded its threshold. In addition to the optional log message and

SNMP trap, the timestamp of the last occurrence of the event and number of times the threshold was crossed can be seen in the “show mda detail” command. The number of events and last timestamp are not cleared in the event of an IOM or MDA reset. The timestamp of the last event should be used to make sure the IOM or MDA state has not changed since the last occurrence. Further investigation is necessary to determine the root cause. [67251]

RELEASE 5.0.R15

The following section describes the new enhancements added since 5.0.R14 to 5.0.R15 of 7710 SR OS.

IP MULTICAST

- New CLI commands have been introduced for Multicast CAC (MCAC). The commands “tools perform service id *x* mcac sap sap-id recalc”, “tools perform service id *x* mcac sdp sdp-id recalc” and “tools perform router *x* mcac recalc” allow MCAC re-evaluation on its users. This is necessary when the MCAC configuration has changed and the new values are not reflected in the user database. The command “show router mcac policy *policy-name* users” has been added to display all interfaces applied to a given MCAC policy. [59130, 62370]

RELEASE 5.0.R14

There are no new enhancements in Release 5.0.R14.

RELEASE 5.0.R13

The following section describes the new enhancements added since 5.0.R12 to 5.0.R13 of 7710 SR OS.

ROUTING

- Traceroute is now supported over LDP shortcuts when the source IP address of the traceroute is known on every router on the LDP-signaled MPLS path. [60426]

RELEASE 5.0.R12

There are no new enhancements in Release 5.0.R12.

RELEASE 5.0.R11

There are no new enhancements in Release 5.0.R11.

RELEASE 5.0.R10

The following section describes the new enhancements added since 5.0.R9 to 5.0.R10 of 7710 SR OS.

RADIUS

- In a RADIUS authentication policy, it is now possible to configure multiple RADIUS servers that have the same IP address but different UDP ports. [62340]

TPSDA

- Two extra user name formats have been added to the subscriber management authentication policies: “ascii-converted-circuit-id” and “ascii-converted-tuple”. These user name formats output the same information as circuit-id and tuple, respectively, but the information is now encoded as a hexadecimal string instead of a null-terminated ASCII string. [61379]

RELEASE 5.0.R9

The following section describes the new enhancements added since 5.0.R8 to 5.0.R9 of 7710 SR OS.

CLI

- When ping requests from the CLI time out, the message displayed now includes the ICMP sequence number of the individual ICMP echo request that timed out. [61729]
- To correct truncations of port strings and service IDs in the output of the “show service sap-using” command, some columns have been removed from the output of the basic form of this command and are now available in the output when one of two additional parameters are specified with this command. The new usage of this command is: “show service sap-using [sap *sap-id*] [vlan-translation | anti-spoof]”. [63275]

MPLS/RSVP

- In RSVP fast-reroute interoperability, the 7710 SR now accepts the system address of the tail-end router as a valid “node to avoid” address in the path message detour object when the head-end and tail-end router of an MPLS detour path are both Juniper routers and one or more transit routers are 7710 SR router(s). [62324]
- The RSVP path message TSPEC MTU of the primary path will no longer decrease in value after the bypass tunnel has been signaled on the head-end router in RSVP fast-reroute with facility backup. [62862]

RELEASE 5.0.R8

The following section describes the new enhancement added since 5.0.R7 to 5.0.R8 of 7710 SR OS.

VPRN/2547

- Subscriber host routes are no longer always advertised into VPRNs by default when using export routing policies. This allows configuration control as to when /32 host routes are to be advertised using routing policies. [57745]

RELEASE 5.0.R7

The following sections describe new enhancements added since 5.0.R6 to 5.0.R7 of 7710 SR OS.

CLI

- The “show router tunnel-table” command now supports “rsvp” as a protocol to filter the display for LDP-over-RSVP tunnels. [60338]

- The “oam ancp” command now supports a subscriber identification string in addition to the ancp-string already supported. [61696]

BFD

- BFD is now supported on TDM ports (T1, E1, DS3, E3). [59197]

RELEASE 5.0.R6

The following sections describe new enhancements added since 5.0.R5 to 5.0.R6 of 7710 SR OS.

CLI

- The “show card” output has been enhanced to display the bootrom version for CFMs. [59425]

SYSTEM

- Non-stop routing for RSVP-TE now supports detour-based Fast-ReRoute (FRR).

QOS

- Port-based Egress Scheduling Policies are now supported on all other port types, including SONET/SDH ports. This enhancement was added in Release 5.0.R5.

TPSDA

- Inter-node State Synchronization and Multi-Chassis LAGs can now support up to four peer groups. [58377]

RELEASE 5.0.R5

The following sections describe new enhancements added since 5.0.R4 to 5.0.R5 of 7710 SR OS.

HW/PLATFORM

- In 5.0.R3, support was added to allow field replacement of an MDA in a slot provisioned with older, non-Rev B GigE MDAs with a Rev B MDA. In this release, support has been added to allow field replacement of an MDA in a slot provisioned for a Rev B MDA with a non-Rev B MDA with the constraints described below.

  If no Rev B features are enabled on an MDA in a slot provisioned for a Rev B GigE MDA, a non-Rev B GigE MDA can be used as a direct Field Replaceable Unit (FRU) for Rev B MDA:

  - 5-port GigE MDA - SFP (3HE00025AA) can be used as a FRU for a 5-port GigE MDA - SFP Rev B (3HE01615AA) and

  If a non-Rev B GigE MDA is inserted into an MDA slot that is provisioned for the Rev B variant, the system will try to demote the MDA slot provisioning to the non-Rev B variant,

preserving all other provisioning, provided no Rev B-only features are configured on the MDA.

The following settings are required for the Rev B features for demotion to succeed:

- The Rev B pre-classifier must be disabled ("config>port *port-id*> ethernet>no ingress-rate").
- Ports with multirate 10/100/1000BASE-TX Copper SFPs must be configured for GigE operation ("config>port *port-id*>ethernet>speed 1000").
- The port duplex for all ports must be full duplex ("config>port *port-id*>ethernet>duplex full").

If any of the above conditions are not true, the MDA will be in a FAILED state in the system with the reasons for the failure noted in the system log.

Note that both the non-Rev B and Rev B MDAs support "auto-negotiation", "auto-negotiation limited" and "no auto-negotiation". The older non-Rev B MDAs only support 1000BASE operation, so "auto-negotiation limited" is the same as "auto-negotiation" on those MDAs. Since the Rev B MDAs support 10/100/1000BASE operation, it is possible that a Rev B port is working after auto-negotiating to speed of 10 or 100. If a Rev A is then used as a FRU, the port may not establish link because it does not run at speeds 10 or 100.

If a non-Rev B MDA is in the FAILED state because of unsupported settings, once the configuration errors noted in the system log have been corrected, the user can re-initiate the MDA initialization by re-inserting the MDA or performing a reset on the MDA ("clear mda *mda-id*"). [59443]

MANAGEMENT

- The DHCP Lease State Persistency files created in Release 3.0 of 7710 SR OS can now be read in and converted to the Release 5.0 format for these files. With this enhancement, it is no longer necessary to perform an intermediate upgrade to Release 4.0 to convert the persistency file format. [59340]

OSPF

- Controls have been added under OSPF to VPRN instances to control the inclusion/suppression of the DN bit for external LSAs sent to a CE device and to control the interpretation of the DN bit in LSAs received from a CE device.

  The "[no] suppress-dn-bit" command is used to control whether or not the DN bit is set for external LSAs sent by the 7710 SR. The default setting is to set the DN bit in external LSAs sent to the CE ("no suppress-dn-bit").

  The "[no] ignore-dn-bit" command is used to control whether the DN bit should be ignored in LSAs received from a CE device for the purposes of inclusion in the SPF calculation for the VPRN. The default setting is to not ignore the DN bit in LSAs received from the CE ("no ignore-dn-bit"), so these LSAs will not be used in the SPF calculation by the 7710 SR. [59603, 59623]

VRRP/SRRP

- SRRP will now report an issue when it detects a conflict between the IP address of a local subscriber interface and the IP address of the remote subscriber interface. When the conflict occurs, a new event tmnxSrrpDuplicateSubIfAddress is generated. [59300]

MPLS/RSVP

- When MPLS is operationally down, "show router mpls status" now indicates the reason why the protocol is operationally down. [60132]

SERVICES GENERAL

- SDPs will now monitor and limit the rate at which RSVP and LDP events are generated. If the rate for an LSP (or LDP tunnel) exceeds 10 events in a 10 second window, the LSP (or LDP tunnel) is held down for a period of 20 seconds to stabilize the flapping. A new SDP operating flag called TranspTunnUnstable has been added so this can be seen in event logs and via the show command output. [59864]

QoS

- The Port-based Egress Scheduling policies introduced in Release 5.0.R4 for Ethernet and LAG ports are now supported for all port types.
- An option has been added to control the CLP Tagging of non-expedited egress ATM traffic on an IES or VPRN interface. By default, the CLP bit is always set to zero on ATM traffic egressing an IES or VPRN SAP, but for traffic that traverses a third-party network, the option has been added to set the CLP bit to one (1) for non-expedited traffic. The command "config>qos>atm-td-profile *profile-id*>[no] clp-tagging" within the Traffic Descriptor Profile controls whether the CLP bit should be one (1) when applied to an IES or VPRN ATM SAP. [57768]

OAM

- In the following OAM commands, it is now possible to specify the forwarding class of the OAM requests:
  - mac-ping
  - mac-trace
  - cpe-ping
  - vprn-ping
  - vprn-trace. [42917]

RELEASE 5.0.R4

The following sections describe new enhancements added since 5.0.R3 to 5.0.R4 of 7710 SR OS.

SYSTEM

- SCP now uses the user's "home-directory" as the default destination relaxing the requirement to explicitly specify a destination path. [45928]
- Alarm and log events are now raised when an excessive clock skew condition exists between MCS peerings. These events are tmnxMcSyncClockSkewRaised and tmnxMcSyncClockSkewCleared. [58918]

LAG

- A hold up timer has been added to LAGs to allow control over the revertive behavior of a LAG subdivided into sub-groups and to delay switchover between the active/standby states. [59192]

MANAGEMENT

- An event log conversion mechanism has been added to support renumbering of old events that are replaced by new events and saving them using the newer event numbers. The two events handled in this release are:
  - security event 2013 (old) changed to security event 2026 (new)
  - security event 2017 (old) changed to security event 2025 (new). [59125]

ROUTING

- Unnumbered IP interfaces are now supported on the following:
  - Access mode SONET/SDH IES/VPRN interfaces
  - Access mode Frame Relay IES/VPRN interfaces
  - Access mode ATM IES/VPRN interfaces. [40362]
- The ECMP hashing algorithm has been changed to provide a better load distribution across the ECMP next-hop interfaces. [51376]

IS-IS

- When coming up for the first time, IS-IS will now send hello packets with the Suppress Adjacency (SA) bit set in the initial hello and then clear it in subsequent hello packets. This assists in cases where the router was previously on the network and remote routers start sending traffic immediately upon adjacency formation by using the previous copies of the LSP. This could result in lost traffic for tunneling topologies such as BGP or LDP-over-RSVP. Neighboring routers must have Graceful Restart Helper support enabled in order for this change to have any effect. If they do not, the neighboring routers will operate as before. [59024]

IP MULTICAST

- Efficient Egress Multicast Replication has been enhanced to support QinQ encapsulated SAPs. For a given egress multicast group, all SAPs must have the same inner and outer Ethertypes. [56614]
- The Multicast CAC Policy feature has been extended to be supported on Link Aggregation Groups (LAGs). Previously, a Multicast CAC Policy could only be applied to a single, physical port. [57816, 59236]

MSDP

- The “show router msdp peer” output has been enhanced to indicate the default peer. Previously, the default peer was only indicated in the detailed command output. [58177]

SUBSCRIBER MANAGEMENT

- Debugging has been enhanced to allow 1) debugging DHCP packets for a Layer 3 SAP (“debug service id <id> dhcp sap”) and 2) debugging based on a subscriber MAC (“debug router ip dhcp”). [56549]

VPLS

- The Multi-Chassis Synchronization (MCS) feature has been extended to maintain the state for Multicast VPLS Registration (MVR).

VPRN/2547

- VRRP policies are now supported on VRRP instances on VPRN IP interfaces. This enhancement was actually added in 5.0.R3. [53115]

RELEASE 5.0.R3

The following sections describe new enhancements added since 3.0.R1 to 5.0.R3 of 7710 SR OS.

HW/PLATFORM

- The image directory names on the Compact Flash devices now include the platform (7710).
- The MD5 checksums for all files in a software distribution set are now included on the Compact Flash in the file md5sum.txt.
- After the incomplete transfer of files with sizes that exceed the amount of available free space, the system will delete the partial file that is stored on the Compact Flash drive as a result of the incomplete transfer. [44052]

- A “clear mda mda_id” CLI command has been added to reset an individual MDA. [45981]
- A Rev B GigE MDA can be used as a direct field replacement unit (FRU) for the older (non-Rev B) GigE MDA:
  - 5-port GigE MDA - SFP Rev B (3HE01615AA) can be used as a FRU for 5-port GigE MDA - SFP (3HE00025AA) and

  If a Rev B GigE MDA is inserted into an MDA slot that is provisioned for the non-Rev B variant, the system will promote the MDA slot provisioning to the Rev B variant and all other provisioning will be preserved.

  In prior releases, Rev B and non-Rev B MDAs were considered distinct types, and the non-Rev B MDA would need to be unprovisioned and re-provisioned after replacement with a Rev B MDA.

  Note that:
  - Because a Rev B MDA supports more features than the non-Rev B equivalent, the reverse process is not supported, that is, a non-Rev B cannot be used as a FRU for a Rev B variant. [49468]
- The 7710 SR OS “version” command run against a 7710 SR OS (*.tim) file will check the image file for possible corruption. This is useful after transferring the file to the node and prior to using the file. [54268, 54269]
- A new parameter (“hi-bw-mcast-src”) has been added for MDAs that will optimize system resource allocation for MDAs with high-bandwidth ingress IP multicast. [52488]

BOF

- If the BOF is edited and a referenced configuration file does not exist, a console warning message is generated indicating the system may not reboot properly. In the case of a “redundant system” with “automatic sync boot-env” enabled and a referenced config file does not exist on the standby, the BOF is saved but a warning will indicate syncing of the configuration files has failed. For non-redundant systems or when “automatic sync boot-env” is disabled, the BOF will save with a warning message because the node may fail to reboot properly until there is an “admin save” on the node. [48040]

RADIUS

- The port numbers for RADIUS Authentication, Authorization and Accounting can all be individually configured. [31565]
- The “debug>radius” command has two new options “detail” and “hex” to facilitate debugging RADIUS connections. [38893]
- The 7710 SR now supports a single configurable “template user account” from which all RADIUS authenticated users will inherit user-level settings if the RADIUS server returns an Auth-Accept with an empty VSA. The default user account is enabled using the “use-default-template” command in the “config>system>security>radius” context. If the template is not enabled, Auth-Accept with empty VSAs is dropped. The default RADIUS template is configured in the “config>system>security>user-template radius_template” context. [56339]
- To ensure that accurate and up-to-date RADIUS dictionaries are readily available, the RADIUS dictionary is included in FreeRADIUS format on the distribution media of 7710 SR OS. [58162]

TACACS+

- In 3.0 releases, TACACS+ users could not be granted access to features available to locally defined users, such as:
  - Inbound FTP access
  - A login exec script (useful for defining aliases, etc.)
  - Home directories

  To allow for this capability, the 7710 SR now supports a single configurable “template user account” from which all TACACS+ authenticated users will inherit user-level settings.

  Example:

  config>system>security>user-template tacplus_default

  access console ftp

  home-directory cf3:/data

  console

  login-exec “cf3:/scripts/alias-login.cfg”

  exit

  Note that a member profile cannot be assigned to the default TACACS+ user template.

  The above example is consistent with the existing “default” authorization profile assigned that are not explicitly made members of another profile. [28268]
- Upon receiving password expiration warning messages from a CiscoSecure server, those expiration messages are now printed to the user upon login. Use the “password” command to change the password for telnet and SSH users. FTP users will still need their passwords changed on the CiscoSecure server. [30671]

CLI

- The CLI now supports “piping” (using the “|” character) and redirecting (using the “>” character) of output. Piping allows one command to be used as input into a grep-like “match” function which can filter the output to find the text or regular expressions. Using redirection, the command output can be sent to a File URL. [2741]

  Synax Examples:

  show router | match *regex*

  show router > *file-url*
- The “show router interface” command now displays the Port ID for the IP interface. [22075]
- A “*” character can now be displayed at the beginning of each command line prompt (for example, “*router_A>config>system#”) when either the BOF or the configuration has been changed and not yet been saved. The environment variable “environment>saved-ind-prompt” controls whether prompt should be prefixed with the “*” character; the prefix is displayed by default. The user may save changes or revert to the previously saved configuration when the “*” is present. [34219]
- CLI commands can be executed out of context by specifying the full path from the CLI root using a forward slash (“/”) or backward slash (“\”) at the beginning of the command. [38018]

  For example:

```
router_A>config>router# interface system address 1.2.3.4
router_A>config>router# /admin save
router_A>config>router# /clear router ospf database
```

- Commands entered in abbreviated form now display the entire command when the <Enter> key is pressed. For example, “sh ver<**Enter**>” now produces “show version”. [39168]
- The ability to translate DNS hostnames to IP addresses is now supported. Previously, IP addresses only reconciled with DNS names, not both directions. [39495]
- The “show service fdb-mac” command has a new “expiry” parameter that displays the how long the MAC entry will remain in the forwarding database if not refreshed. [43414]
- The “show port port_id” and “show port port_id detail” commands have been enhanced to display the distance information programmed in the SFP/XFP pluggable transceiver. The SFP/XFP type displayed in the “show port” command is the “type” programmed in the transceiver based on the industry Multi-Source Agreement (MSA) definitions. The type values defined in the MSA are generally limited to types defined by standards organizations, and in particular do not include ad-hoc industry standard types. For example, the MSA currently defines the longest distance 1000BASE SFP as an “LX” type which is the 10 km standard defined by the IEEE. Ad-hoc industry standards like “EX” (40 km) and “ZX” (70 km) are not defined in the MSA and are defined as “LX” types in the MSA fields but with longer distances. These distance values can now be viewed in the CLI. Previously, one would have to correlate the Alcatel part number to determine distance for the SFP/XFP. [44656]
- A text description can now be added to the following:
  - Physical ports (“config>port *port-id*>description”)
  - Router interfaces (“config>router>interface *if-int-name*>description”)
  - SNMP Trap Groups (“config>log>snmp-trap-group *log-id*>description”). [45615, 51746]
- The “show system information” command output now includes system/chassis type, software version and active CFM slot. [47353]
- The “show users” command now indicates whether users are connected using “SSHv1” or “SSHv2”. [49099]
- Prior to Release 5.0, debug output line termination was inconsistent for debug messages sent to the Console and to Compact Flash. Debug message lines are now always terminated with a carriage return (<CR>) and line feed (<LF>). [53106]
- The maximum length of port, LAG and IP interface descriptions has been increased from 80 characters to 160 characters. [53285]
- The SSH command syntax has been modified to be more consistent with other 7710 SR OS command. The new syntax to initiate an SSH session is “ssh *username@host*”; the old syntax was “ssh -l *username host*”. [54372]
- The “show config” command has been deprecated in this release since it could be run by users without administrative security permissions. The “admin>display-config” command replaces the “show config” command. [55057]

SYSTEM

- Syslog messages now indicate the Virtual Route Forwarding (VRF) instance/VPRN instance ID for logged events.
- The “summer” command to configure additional time zone shifts has been deprecated and replaced with new commands. The CLI now allows creation of new “dst-zone zone-name”

command contexts under config>system>time to configure summer/daylight savings time parameters for the time zone. Each “dst-zone” context has the following new commands:

- end - Configures the end date and time for the daylight savings time zone.
- offset - Configures the daylight savings time offset in minutes for the daylight savings time zone.
- start - Configures the start date and time for the daylight savings time zone.

Existing configuration files will be saved using the new commands. See the configuration guides for the complete command syntax and descriptions. [29332]

- The following time stamps are now preserved after a High Availability switchover:
  - System up time
  - OSPF status last enabled and interface last enabled
  - IS-IS status and adjacency up time
  - BGP neighbor up time and adjacency up times. [30362]
- An optional keyword “regexp” has been added to the “config>log>filter *filter-id*>entry *entry-id*>match>subject” and “show log log-id *log-id*” CLI commands to support regular expression matching criteria. When configuring match criteria, the “eq” and “neq” (equal or not equal) operators determine the type of string comparison to use to determine if the log event matches the regular expression. [39465]
- When saving a configuration with “admin>save” command, the dates and times of the backup configuration files will now be preserved rather than being updated to the current date and time. [44051]
- Free space is verified before a file is copied to a Compact Flash. Previously, if there was insufficient space, only a part of the file would be copied to Compact Flash. [47728]
- File destination references can now be removed or changed if the application referencing the file is administratively shutdown. For example, changing the file referenced in an Accounting Policy previously required removing the Accounting Policy before the destination file location could be changed; now, the destination can be changed or removed if the policy is shutdown. [49278]
- The following improvements have been made to the CRON feature:
  - A new line “next scheduled run” has been added to “show cron schedule” output to display the day, date and time of the next scheduled run.
  - A new parameter “count” has been added to “config>cron>schedule *schedule-name*” to configure the number of times to repeat a periodic schedule run.
  - A new line “repeat count” has been added to “show cron schedule *schedule-name*” to display the number of times a periodic schedule run has repeated.
  - An “end-time” parameter has been added to “configure>cron>schedule *schedule-name*” to stop a schedule at the specified time.
  - A new keyword “run-history” has been added to “show cron action” to show the run histories of all configured actions. [50839]
- New Zealand Standard Time (NZST) and New Zealand Daylight Time (NZDT) time zones have been added to the system. [52967]
- The minimum log file rollover period (config>log>file-id *log-file-id*>rollover) has been reduced from 15 minutes to 5 minutes. The default rollover period remains at 1440 minutes. [54356]

- The system now maintains administrative and operational CIR and PIR values for CFM queues. [55036]
- In accordance with recent US changes to the start and end dates of the four US daylight savings time zones beginning in 2007, those changes are now supported if the current configuration is saved and reloaded. Another way for the changes to take effect dynamically is to remove the dst-zone setting and add it back. This only applies to the four pre-defined US summer time zones (EDT, CDT, MDT and PDT). [55368]
- Atlantic Daylight Time (ADT) and Alaskan Daylight Time (AKDT) will now start and end in accordance with recent US and Canadian changes to the daylight savings time zones beginning in 2007. Also, the support for Newfoundland Standard Time (NST) (GMT - 3:30 hours) and Newfoundland Daylight Time (NDT) have been added to the system. [56682]

SONET/SDH

- SONET/SDH paths can now be configured to send a trace-signal of all zeroes, without a terminating <CR><LF> sequence for SONET or CRC-7 for SDH by with the “trace-signal zeros” command. [46928]

ATM

- ATM cell level policing is now supported on an ingress SAP which is in the UBR/UBR+MIR service category. When the policing option is enabled in the “atm-td-profile”, and this profile is applied to SAP ingress, a single GCRA instance with the SAP ATM traffic parameters PIR and CDVT is applied to police the traffic. Cells which do not comply to the policer are discarded. [39144]

LAG

- For consistency with other port types, LAGs must be shutdown before changing the encapsulation or mode. [41377]

MANAGEMENT

- The request count for the “ping” command has been increased to 100000 from 10000 to allow for higher count extended ping tests. [35590]
- The IP source address used for the management applications like ping, traceroute, dns, sntp, cflowd, radius, tacplus, telnet, SSH, SNMP traps, NTP and Syslog can now be user configured. The configured source address will be used for all unsolicited egress packets for the given application. [37433]
- The enumerations for the tmnxPppCpProtocol attribute have been changed to all lower-case notation. [45345]
- The 7710 SR OS event control and event logger have been enhanced with an event throttling mechanism that allows a maximum event count per unit time to be applied to events in addition to the existing “generate” and “suppress” options. A global event throttling rate based on <n> events per second(s) is defined and is optionally applied to events governed by event control. Any events in excess of the throttling rate are dropped and counted in the event drop counter. If any event is throttled, a TIMETRA-SYSTEM-MIB::tmnxTrapDropped notification is sent.
- The ifSpeed and ifHighSpeed MIB attributes are now populated with valid values for router interfaces. [45612]
- A timestamp of when the statistics were last cleared is now displayed in the “show service id *service-id* sap *sap-id* detail” and “show port *port-id* detail” output. [46034]

- Unsuccessful attempts to FTP a file to the router now generate a log event (SYSTEM: 2034 - tmnxFtpClientFailure) which can be forwarded to trap destination. [46736]
- A log event is generated when a SONET/SDH loopback test is provisioned to help differentiate a test from actual error conditions. [46734]
- The forwarding plane service statistics for a SAP (or L3 Spoke SDP) for IES and VPRN service IP interfaces are also mapped to the interface statistics in the standard MIB2 ifTable and ifXTable statistics. [51337]
- A new log event (VRTR: 2015 - tmnxVRtrBfdSessionUp) and associated SNMP trap have been added to indicate when a BFD session transitions to the up state. [56370]
- The format of SNMP messages in change logs has been revised to be more succinct but with more information in each log message. The more succinct format also eliminates erroneous buffer overflow messages that could occur with the old format. [58347]

ROUTING

- New optional parameters have been added to the "show router arp" command:
  - Parameter type local, dynamic or static to display ARP entries of a specific type
  - A mask-length option ("/nn") for the IP address parameter to only display ARP entries matching the IP address/mask combination. [9252]
- The "show router interface detail" command now displays Last Oper Chg indicating the timestamp of the last change in the operational state of the router interface. [9253]
- Entry IDs for route policy entries ("config>router>policy-options>policy-statement *name*>entry *entry-id*") can now be expressed using a 32-bit integer instead of an integer in the range of 1 to 65535. The change in the Entry ID numbering does not affect the overall scale of the number of policy entries supported; the number of policy entries is still 64K. [52742]

IPV6

- 7710 SR OS now supports IPv6 router advertisements of a prefix and other parameters on a LAN so hosts can auto-configure a global IP address.

PROXY ARP

- The Proxy ARP CLI has been changed as follows:
  - The command "config>[router]/[service>ies/vprn]>if>proxy-arp" has been deprecated.
  - The command "config>[router]/[service>ies/vprn]>if>remote-proxy-arp" replaces "config>[router]/[service>ies/vprn]>if>proxy-arp"
  - The command "config>[router]/[service>ies/vprn]>if>proxy-arp-policy" replaces "config>[router]/[service>ies/vprn]>if>proxy-arp>policy-statement"
- See the configuration guides for the complete command syntax and descriptions. [42824]

DHCP RELAY

- The source IP address for DHCP relay packets on IES and VPRN standard and group interfaces can now be configured. [44732]
- The new commands "detail-level" and "mode" have been added to the "debug>router>ip>dhcp" context to enhance debug tracing. [50460]
- DHCP Relay can insert an Option-82, sub-option "remote-id" into a DHCP packet expressed as an ASCII string. The ASCII string format can now be used in addition to previously supported format using the client hardware address format. [45221]

- A new format of circuit ID can be inserted into DHCP packets by DHCP Relay. This format includes the VLAN tag (P-bits as well as the VLAN ID) as received from the DHCP client. This format is primarily required for interworking with DSLAMs in a “QoS-aware-cross-connect” mode. The applicability and behavior for the various SAP types are as follows:
  - null-encapsulated SAPs - not applicable.
  - Dot1Q SAPs - dot1p bits from the service delimiting tag (along with the VLAN ID) will be placed into the Option 82 field in the upstream direction, and the service delimiting tag will be marked with dot1p bits from Option 82 field in downstream direction.
  - QinQ SAPs - dot1p bits of the inner tag (along with the inner VLAN ID) will be placed into the Option 82 field in the upstream direction, and both service delimiting tags will be marked with dot1p bits from Option 82 field in the downstream direction. [45373]

IS-IS

- A “run-manual-spf” command is available in the tools>perform>router>isis context to manually trigger a Shortest Path First (SPF) calculation. [20289]
- In 3.0 releases, point-to-point IS-IS interfaces were controlled by the Level 1 parameters according to standards. Therefore, interface parameters like hello-interval needed to be configured under Level 1 in order to be persistent. This was true for point-to-point interfaces configured for Level 2 only. The behavior now is that if a point-to-point interface is configured for Level 2 only, the Level 2 parameters will be used. [42721]
- The hello-authentication-type and hello-authentication-key for point-to-point interfaces configured for Level 2 capability only are now controlled at the interface Level 2 context as well. [52158]

OSPF

- OSPF timers granularity has been reduced to approximately 10 ms. The previous “spf-timers” command has been deprecated and replaced by commands under a new “timers” context under config>router>ospf. Configuration files read in with the “spf-timers” command will be written back with the new commands. If values of zero are configured under the “spf-timers” command, the new minimum values will be written back out. The new OSPF timer commands are:
  - lsa-arrival - Configures the minimum delay that must pass between receipt of the same LSAs arriving from neighbors.
  - lsa-generate - Configures the throttling of OSPF LSA-generation.
  - spf-wait - Configures the maximum interval between two consecutive SPF calculations.

  Note that the new “lsa-arrival” and “lsa-generate” commands are used to override the RFC-defined constants MinLSArrival and MinLSInterval.

  See the configuration guides for the complete command syntax and descriptions. [40406]
- The default OSPF metric for the system interface is now 0 instead of 1 to be more in line with industry standard defaults. [48901]
- An optional “remain-down-on-failure” parameter can now be specified on OSPF interfaces that are BFD-enabled to keep OSPF from reaching the state FULL if the BFD session to

that neighbor cannot be established. This option is disabled by default and should only be used if OSPF full adjacency is not desired when no BFD session present. [50984]

BGP

- BGP peer tracking, if enabled, will cause a BGP peer to be torn down as soon as all routes, resolving a BGP peer's address, are removed from the system. This allows BGP to respond immediately to BGP peer unreachability. If this feature is not enabled, then the BGP peer will remain up until the BGP hold timer expires. The default state for this feature is disabled. [48502]

MPLS/RSVP

- The LSP ID increments every time the LSP needs to be re-signaled as a result of a number of possible events: for example, an RSVP interface shutdown/no shutdown, a failure of the primary LSP when the head-end is also the PLR, or a change to any of the LSP configuration parameters. In a previous releases, the LSP path was re-signaled with the same LSP ID which could cause state issues in the network, as the PSB in the rest of the network may not have timed out. [19566, 36994]
- The auto-generated path names for node/link bypass tunnels have been shortened to “bypass_link_ip_addr” and “bypass_node_ip_addr” to improve the output of the “show router rsvp session [bypass-tunnel]” command. [34676]
- Per-LSP debug logging for make-before-break (MBB) LSPs is now supported. When debugging the LSP ID of a MBB LSP, LSP setup events for both the original and the MBB LSP path are logged. If the user specifies the LSP ID of the MBB LSO path, then only LSP setup events for the MBB LSP path will be logged. [36992]
- The FRR implementation has been enhanced to provide better strict capacity planning capability in the following scenario: the 7710 SR is the ingress LER with a primary path and secondary path(s) for an LSP and FRR (either detour or bypass tunnel) is enabled for the LSP. If the primary path fails and the PLR is exactly the ingress LER, the new behavior is to switch to an operational hot standby secondary path immediately, unlike the old behavior where the switchover to a secondary path only occurred when the detour had failed. If no hot standby secondary path is available but a cold secondary is configured, then that path is first signaled. After it becomes operational, the switchover from the detour path is made to the operational secondary path. The failure may happen at any point with the path that is protected by a fast reroute path. [37353]
- The output of the command “show router mpls lsp path” has been enhanced as follows:
  - The next-hop address and interface when an FRR backup LSP is activated at a PLR node.
  - For facility bypass LSPs, the command displays the bypass for next-hop and interface.
  - For one-to-one LSPs, the command displays the detour LSP next-hop address and interface. [42297]
- 7710 SR OS now supports the global revertive mode as defined in RFC4090 when the 7710 SR is the head-end LER. (Note that support for local revertive mode has been deprecated in Release 5.0.)

  In global revertive mode, the head-end LER will attempt a re-optimization of the LSP shortly after it knows the PLR has activated local protection. If CSPF finds a path, the head-end LER will switch the LSP to the new optimized path. If CSPF fails to find a path, the LSP stays on the detour/bypass or standby path until the next re-optimization attempt

upon the expiry of the LSP retry-timer, that is, in 30 seconds by default. However, the 7710 SR OS head-end LER will continue global revertive re-optimization attempts even after a successful local reverting; it does so until a new path is found or the LSP is torn down. [46128]

- In prior releases, both Path and Resv messages were sent out precisely at the expiry of their timers. Both messages are now jittered +/- a few seconds from the expiry timer setting, so if there are many messages, the messages are spread out over a few seconds rather than all being sent in a burst. [52010]
- In prior releases, when a 7710 SR PLR node associate a bypass tunnel with the protected LSP primary path which requested ‘node-protect’ protection type, it was possible that a ‘link-protect’ bypass tunnel was associated with the LSP path. This can happen either because a suitable 'node-protect' bypass was not found by CSPF at LSP establishment time or that the associated ‘node-protect’ bypass went down and no other one was found by CSPF. If a suitable ‘node-protect’ bypass tunnel became available later, the 7710 SR PLR would not attempt to associate the protected LSP path with this bypass tunnel. The workaround to this issue was to manually re-signal the protected LSP primary path from the head-end node.

  In Release 5.0, the PLR keeps track of LSP paths for which “node-protect” bypass is requested and which are associated with a “link-protect” bypass. As soon as a manual or dynamic “node-protect” bypass which meets the constraints of protected LSP becomes available, the PLR will revert its association back to this ‘node-protect’ bypass tunnel. [54488]
- The “show router mpls static-lsp” output has been enhanced and now includes the following information:
  - Static LSP ID
  - Static LSP up/down time
  - Label Map up/down time (for transit and terminate LSPs)
  - In/Out Port. [54723]
- A hold timer has been added to control the reprogramming of an active path. During a MBB process, the hold timer defines a discreet period of time 7710 SR OS will wait before switching to the new MBB path to allow a downstream node to program its forwarding table with the labels for the new LSP. The hold timer can also be used to induce a delay when switching from a detour path to a signaled secondary path.
- The hop-limit parameter is used to determine the maximum hops that a detour/bypass path can traverse. In Release 5.0 7710 SR OS has changed the default value of the hop-limit from 6 to 16, as many ring topologies do not have protection paths available within 6 hops.
- As global revertive is the default behavior for Facility Bypass LSPs, support for FF-style signalling for LSPs requesting the FRR facility has been disabled. MBB for FF-style signaled LSPs would lead to double booking of resources in the network.
- LSP IDs are now randomized to prevent some LSP ID collisions that had occurred in previous releases where the same LSP ID was incorrectly reused for a different path upon a High Availability switchover. This behavior is different from all earlier releases. [59308]

LDP

- The 7710 SR will now accept LDP hello packets from the remote router’s “system address” to be compatible with Huawei’s LDP implementation. Previously, the7710 SR strictly adhered to RFC1112 enforcing that the IP source address of outgoing IP multicast

datagrams must be one of the individual addresses corresponding to the outgoing interface. [54408]

MSDP

- The “show router msdp statistics” output has been enhanced to display the number of source addresses rejected as a result of import and policies at the group, peer or source levels. In addition, the peer statistics have been enhanced to display the number of received unknown and reserved messages. [56492, 57967]

IP MULTICAST

- Prior to Release 5.0, some inconsistencies existed with regard to configuring “*,g” “starg”, or “star-g” in the CLI. “starg” is now used consistently in all contexts. [56495]
- A new parameter “mfib-allowed-mda-destinations” has been added to “config>service>vpls *svc-id*>spoke-sdp *sdp-id:vc-id*>egress” and “config>service>vpls *svc-id*>spoke-sdp *sdp-id:vc-id*>mesh-sdp *sdp-id*>egress” to specify which MDAs may participate in the multicast tree. [57254]

IGMP

- IGMP query response behavior has been modified to allow for a faster startup. Four queries are sent on startup of an IGMP interface in the first 10 seconds, each with a query response interval of 1 second. After this initial “burst”, the configured values are then used in subsequent queries.
- A “max-groups” command for IGMP has been added to limit the number of groups allowed to join on an interface. IGMP keeps track of only the IGMP downstream groups that are created as a result of IGMP Joins using the IGMP configured value. If the configured value is dynamically changed to a smaller value than the number of active groups present, existing groups are not pruned. [40047]

PIM

- A “max-groups” command for PIM has been added to limit the number of groups allowed to join on an interface. PIM keeps track of only the PIM downstream groups that are created as a result of PIM Joins using the PIM configured value. If the configured value is dynamically changed to a smaller value than the number of active groups present, existing groups are not pruned. [40047]
- A PIM log event is generated when a data Multicast Distribution Tree is reused (event # 2012 - vRtrPimDataMtReused). [50304]
- In 3.0.R3, a new feature was added to advertise a neighbor list TLV within PIM Hello packets. The TLV number assigned by the IANA for this list was 24. However, document draft-ietf-pim-sm-v2-new-xx had already specified TLV 24 for usage as a secondary address list without assignment by the IANA. IANA has now assigned TLV 25 for this feature.

  In order to be compatible with the 7750 SR OS 3.0.R10, 3.0.R11, 4.0.R2 and 4.0.R3 implementations and the 7710 SR OS 3.0.R3 implementation, a new “compatibility-mode” option has been added to the CLI to toggle between the two TLV values.

  In 7750 SR OS 3.0.R10, 3.0.R11, 4.0.R2, 4.0.R3 and 7710 SR OS 3.0.R3, the three-way-hello feature is enabled by default on VPRN multicast tunnels (MT) and spoke-SDPs and disabled on other interface types. In 7750 SR OS 3.0.R12 and later, 4.0.R4 and later and 7710 SR OS 3.0.R5 and later, the feature is disabled by default on all interfaces. [53238]

- When the 7710 SR acting as a Rendezvous Point (RP) is recovering from a down state, it will delay advertising itself as an RP to allow time for the multicast network tree to re-converge. [57281]

QoS

- Network Queue Committed Burst Size (CBS) and Maximum Burst Size (MBS) parameters can now be configured in hundreds of a percent. The configuration range is 0.00% to 100.00%. [52290]

FILTERS/PBR/TCS

- IP Filters now accept a wildcard character ("*") for a protocol keyword for an entry match criteria to indicate that the protocol can be either TCP or TCP (config>filter>ip-filter *filter-id*>entry *entry-id*>match protocol *). The wildcard character can also be used for the next-header of an IPv6 Filter (config>filter>ipv6-filter *filter-id*>entry *entry-id*>match next-header *) to the same effect. [20399]
- ACL filter logger now maintains state as to which filter entries most recently generated a logging event and will now throttle log entries by counting the occurrences of the hit rather then issue a message per hit. [44053]
- IP Filters now accept inverse masks for source and destination IP address match criteria (config>filter>ip-filter *filter-id*>entry *entry-id*>match>dst-ip | src-ip). [45720]

SERVICES GENERAL

- The qinq encapsulation is now supported for LAG service access points (SAPs).
- A "summary" keyword has been added to the "show service id *service-id* interface" CLI command which is identical in functionality to the "summary" keyword on the "show router interface" CLI command. [50289]
- A "summary" keyword has been added to the "show service id *service-id* arp" command to display the number of ARP entries for the given service ID. Previously, "show router arp summary" was used view this information. Also, a "summary" keyword has been added to the "show service id *service-id* host" command to display the total number of hosts associated with the service. [50314]
- Prior to 3.0.R6 when routing into an IES spoke, the MTU of the egress port the spoke traverses is used to determine fragmentation/discard decisions. Problems can arise when the far-end termination of the spoke has a restricted Maximum Receive Unit (MRU) (for example, VPLS).

  An administrative IP-MTU parameter has been added to IES and VPRN IP interfaces. The IP-MTU on the L3 access interface is used to override the egress port's MTU which was the operational MTU used in previous versions. The IP-MTU is useful when a provider is using the metro network for multiple service types that do not all support the same MTU value. [51143]
- An event is now generated (SVCMGR: 2208 - sapReceivedProtSrcMac) when the system shuts down a SAP due to a source MAC being received that is a protected MAC on another SAP. [52595]
- Since the entire service ID was not displayed in the output of the "show service sap-using" and "show service id <id> sap" commands, the accounting policy column has been removed. The accounting policy field is still displayed under the output of the "show service id <id> sap <sap-id>" command. [54665]

- SDPs can now be configured with a static “metric” which serves as a deterministic tie-breaking mechanism when the system needs to choose between multiple SDPs to the same destinations. [55053]

SUBSCRIBER MANAGEMENT

- A log event is generated when an attempt to write subscriber billing statistics fails, for example, if the accounting policy writes to a non-existent Compact Flash. [48119]
- The “brief” keyword for the “show service active-subscribers” CLI command has been changed to “summary” to be more consistent with other CLI commands. [50289]

TPSDA

- GSMP is now supported for Enhanced Subscribers (in addition to being supported on SAPs and Multi-Service Sites). [58377, 59370]

EPIPE/VLL

- An EPIPE service can transparently transport LACP packets. Tagged LACP packets are always transparently tunneled. Untagged packets can optionally be transparently tunneled by configuring the “lacp-tunnel” parameter in the config>port port-id>ethernet context. [37111]
- The ARP cache can be cleared on an IP Interworking VLL (Ipipe) service to relearn the MAC address associated with the IP address of the attached CE device. The CLI command to clear an Ipipe ARP cache is the “clear service id *service-id* arp” command. The 7710 SR immediately sends an unsolicited ARP request after the cache is cleared. [47212]

VPLS

- The number of non-residential SAPs that can be configured within a single VPLS service has been increased from 32 to 127.
- When a SAP or SDP is operationally down due to “relearnLimitExceeded”, an additional line “Time to come up/Retries left” has been added to the “show service id *service-id* sap” and “show service id *service-id* sdp” output to display a countdown until the SAP/SDP brought up again and how many retries remain. [49275]
- A Management VPLS (mVPLS) is now supported for default-SAP and NULL VLAN tag SAPs. In 3.0 releases, only explicit VLAN tag ranges were supported within mVPLS. [51195]
- A new keyword “alarm-only” has been added to the “restrict-protected-src” commands under the VPLS command tree. Adding this parameter modifies the behavior of the MAC protection feature in a multi-subscriber environment so the SAP is not shutdown but an alarm event is generated. [52103]
- New events are generated when a SAP (SVCMGR: 2209 - sapTlsMacMoveExceeded) or an SDP (SVCMGR: 2314 - sdpBindTlsMacMoveExceeded) is brought operationally down after the fourth MAC move event, indicating to the operator that a loop may exist. These events are generated each time a SAP or SDP is blocked when the relearn limit has been exceeded indicating the number of retires left and a countdown until the SAP/SDP again becomes operational. [54299]
- In compliance with the latest revisions of draft-ietf-l2vpn-vpls-ldp-xx, receipt of an LDP address withdraw message for MAC withdraws (also known as MAC flush) containing an empty (not included) address list TLV is now supported. In particular, note that a 7710 SR VPLS will now correctly interoperate with a 7250 SAS VPLS for MAC withdrawals. [54794]

IGMP SNOOPING

- An explicit IGMP version to use for sending IGMP queries and for receipt of reports can now be configured at the VPLS SAP (config>service>vpls *service-id*>sap *sap-id*>igmp-snooping>version) and SDP level (config>service>vpls *service-id*><mesh | spoke>-sdp *sdp-id*>igmp-snooping>version). [51917]

VPRN/2547

- The 7710 SR VPRN Next Hop Label limit has been increased from 16K to 64K.
- A log event is now generated when more Data MDTs groups are used than the configured MDT data range (config>service>vprn *service-id*>pim>mdt>data). [50304]
- When a Service Provider (SP) provides VPRN/IP-VPN service to a customer and the customer uses OSPFv2 to advertise its routes to the SP, the Customer Edge (CE) Router and a Provider Edge (PE) Router are OSPF peers and customer routes are sent via OSPFv2 from the CE to the PE. The customer routes are converted into BGP routes, and BGP carries them across the backbone to other PE routers. The routes are then converted back to OSPF routes sent via OSPF to other CE routers. As a result of converting the routes from OSPF to BGP and back to OSPF, some of the information needed to prevent loops may be lost.

  The high-order bit of the LSA Options field (a previously unused bit) is used to solve the loop prevention problem described above. This bit is referred to as the DN bit. When a type 3, 5, or 7 LSA is sent from a PE to a CE, the DN bit MUST be set. The DN bit MUST be clear in all other LSA types. When the PE receives, from a CE router, a type 3, 5, or 7 LSA with the DN bit set, the information from that LSA MUST NOT be used during the OSPF route calculation. As a result, the LSA is not translated into a BGP route. The DN bit is ignored in all other LSA types. [51272]
- The number of VRFs into which a route can be redistributed (leaked) into has been increased from 64 to 2K.
- Hub and Spoke IP-VPNs are VPNs configured such that spoke sites must route all traffic through a hub site. This logical topology has been positioned by several service providers as having security advantages over standard mesh IP-VPNs. In previous releases, the 7710 SR could support this type of VPRN topology through the use of multiple VRFs (one per spoke site) and the configuration of route policies. The configuration of Hub and Spoke VPRNs has been simplified by allowing network administrators to configure a single VPRN instance and identify the necessary VPRN site as a spoke site and internalize the configuration complexity.

VRRP

- This enhancement adds a new timer to VRRP to allow a delay to the initialization of VRRP. This ensures other protocols have converged before the local node tries to assume the primary VRRP role. [54121]

BFD

- BFD support has been added to POS interfaces. [47863]
- The CLI has been enhanced to allow the administrator to display all BFD-enabled interfaces and the configured timer values on the system. [48151]

OAM

- Service Assurance Agent support for vprn-ping and vprn-trace is supported.
- Traceroute now supports displays DNS names in the display output. The ‘no-dns’ command line option disables the DNS lookup of IP addresses. [30662]

- Transparent forwarding of IEEE 802.3ah frames for VLLs is now supported. Transparent forwarding is configured using the “tunneling” command in the “configure>port *port-id*>ethernet>efm-oam” context. [44059]
- The “debug>oam>lsp-ping-trace” command has been added to allow decoding and displaying of the LSP Ping and LSP Trace packet information on the console. The command enables displaying the packet decode information from the transmit and receive sides or both directions. [45967]
- VCCV-Ping now supports the use of the OAM control word encapsulation over a VLL service. The OAM control word option adds a 4-byte control word to the packet with the first nibble set to a value of 0x1 as per draft-ietf-pwe3-vccv-12.txt. This option is useful in networks where LSR nodes perform ECMP and may be forwarding the VCCV-Ping packet on a different path than the user packets. [47211]
- DNS name resolution is now performed for a target address in “ping” and “traceroute” commands at the CLI. [49778]
- When a rapid ping is running, an exclamation mark (“!”) indicates a successful ping reply and a dot (“.”) indicates a the ping request failed. Previously, there was no indication of failed pings. [51793]

USAGE NOTES

The following information supplements or clarifies information in the manuals for Release 5.0.R25 of 7710 SR OS.

COMPACT FLASH DEVICES

- Only Alcatel-Lucent-sourced Compact Flash devices for the 7710 SR are supported.

MANAGEMENT

- It is highly recommended that the management ports be on protected and controlled network segments not directly accessible from the Internet to prevent unwanted Denial-of-Service attacks. [52314]

TCP AUTHENTICATION EXTENSION

- Keychains with no active entries will keep LDP and BGP peerings down. [57917]

DISALLOWED IP PREFIXES

- The following IP address prefixes are not allowed by the unicast routing protocols and the Route Table Manager and will not be populated within the forwarding table:
  - 0.0.0.0/8 or longer
  - 127.0.0.0/8 or longer
  - 224.0.0.0/4 or longer (used for multicast only)
  - 240.0.0.0/4 or longer

  Any other prefixes that need to be filtered can be filtered explicitly using route policies.

IP FILTERS AND NON-IP PACKETS

- IP filters with a default-action of discard will not discard non-IP packets (such as ARP and IS-IS). [40976]

CLI

- The special characters | and > can no longer be used inside environment alias strings. Additionally, the special characters / and \ cannot be used as the first character inside an alias string.

APS

- It is recommended the lb2er-sd and lb2er-sf alarms be enabled for SONET/SDH ports belonging to APS groups to better understand some APS group switchovers between the working and protect circuits.
- For SONET/SDH ports belonging to APS groups that have a very large difference in the transmission delay between the working and protect circuits, it is recommended that the hold down timers be increased from their default values.

ATM

- 7710 SR allows configuration of user traffic on reserved ATM Forum UNI specification VCI values (VCIs from 0 to 31 inclusive). It is recommended not to configure any user traffic on those VCIs on any VP as other equipment may treat that traffic per the defined usage reserved to a given VCI value. [53205]

MLPPP

- When a MLPPP bundle is out of service (oos), the Oper MTU and Oper MRRU are derived from the configured MRRU.
- Currently, LCP echo ids from 0 - 255 are separated into two ranges:
  - 0-127 is used for keepalive function
  - 128-255 is used for differential delay detection.

  Keepalive statistics only count echo packets with IDs from 0-127.

ROUTING

- Reducing the interval/timeout timers much below default values is not recommended for OSPF, IS-IS, PIM, BGP, LDP and RSVP to ensure stability under transitional events like a CFM switchover. [56792, 58891]

IS-IS

- IS-IS authentication is not activated at any given level or interface unless both the authentication key and type are added at that level. For instance, if hello-authentication-type is set to password for an interface, it is not activated until a key is added at the interface level. [34256]

IS-IS TE

- The protocol sends advertisements with the IS-IS Traffic Engineering (TE) Router ID TLV when traffic engineering is disabled. [17683]

LDP

- On LDP interfaces and targeted session keep-alive commands, it is recommended that the “factor” setting be set to a value greater than one or it may lead to unexpected drops in LDP peerings. [67153]

VPRN/2547

- Modifying local VPRN policies might have unexpected side effects if there exist one or more VPRN import polices that refer to a community list name that does not exist. One side effect is that a route refresh message is sent to all BGP-VPN peers when this is not required. To prevent this, route policies must never refer to non-existing objects (prefix lists, community lists, etc.). [60879]

TIME-OF-DAY SUITES

- In a TOD suite, items can be defined that cannot be applied to all SAP types: for instance, an IP filter in the TOD suite that is then assigned as the TOD suite to a VPLS SAP. When the IP filter becomes active, the system will detect that it is not possible to assign the suite to the SAP and generate a log event.
- When a TOD suite is applied to a SAP, there may be conflicts that make it impossible to install all of the current TOD suite defined values. The conflicts can be between the TOD suite defined values or between SAP configured values and TOD suite defined values. A log event is always generated when a conflict occurs. The possible conflicts are:
  - An ingress MAC filter cannot be installed with an ingress IP filter, ingress Ipv6 filter or ingress QoS policy which has Ipv6 criteria. The MAC filter will not be installed.
  - An egress MAC filters cannot be installed with an egress IP filter or egress Ipv6 filter. The MAC filter will not be installed.
  - An ingress Ipv6 filter cannot be installed with ingress an QoS policy which has MAC criteria. The filter will not be applied.
- At system boot, it is possible that the “intended value” (be it from the TOD suite or a configured value) of a policy assignment cannot be applied due to resource unavailability; at that time, there is no previous state to which to revert, and the SAP (or multi-service site (MSS)) ends up with a default policy assignment. In this situation, the SAP (or all of the MSS's SAPs) is (are) placed in an operationally down state with the appropriate flag set.
  - “SapTodResourceUnavail” indicates that the SAP has a TOD suite and could neither apply it nor revert to the previous state. The SAP will have a default policy configured.
  - “SapTodMssResourceUnavail” indicates that the SAP has a Multi-Service Site that uses a TOD suite, and the MSS could neither apply the TOD suite nor revert to its previous state. The SAP will have a default scheduler policies configured, i.e. none.

  These flags get cleared whenever a subsequent application of the TOD suite is successful and the intended policies can be configured.
- When the QoS and scheduler policy assignment of a SAP or MSS is changed by action of its TOD suite, packet loss may occur, just like when the configuration is modified directly by CLI or SNMP.
- The number of assignments in a given TOD suite is implicitly limited to 100 (10 types of parameters each with 10 possible priority values).

SOFTWARE UPGRADE PROCEDURES

The following sections contain information for upgrading to the 5.0.R25 software version. In particular, there are sections that describe the following:

- Software Upgrade Notes on page 60

  Information on upgrading the router from previous versions of 7710 SR OS software including rules for upgrading firmware and any special notes for upgrading from specific earlier versions.
- ISSU Upgrade Procedure on page 62

  Procedure for performing an ISSU to 5.0.R25 including information on applicability of ISSU for earlier versions.
- Standard Software Upgrade Procedure on page 66

  Procedure for performing a standard, service-affecting upgrade including updating of firmware images.

Note:

In-Service Software Upgrade (ISSU) is not supported between a release earlier than 5.0.R23 and a release later than 5.0.R23. ISSU is supported between 5.0.R23 and later releases.

SOFTWARE UPGRADE NOTES

The following sections describe notes for upgrading from prior versions of 7710 SR OS to 5.0.R25.

In the sections below, the following terminology is used:

- Deprecated commands are not flagged as errors upon reading a configuration file with deprecated commands, but these commands will not be written to a saved configuration file.
- Modified command are read using the old format, but they are written out with the new format in a configuration file; so a configuration file saved with modified commands is not compatible with earlier releases.
- Modified parameters are supported when they are read, but the modified parameters will be converted to new minimums or maximums when saved in a configuration file.

5.0.R25 FIRMWARE UPDATE RULES

7710 SR OS 5.0.R25 requires a mandatory firmware update from version 0x1A introduced in 3.0.R1 to the new version 0x1B.

Following the steps in the Standard Software Upgrade Procedure on page 66 will update the firmware images.

5.0.R9 OR EARLIER TO 5.0.R25

The following note applies to upgrading from releases prior to 7710 SR OS 5.0.R9 to 7710 SR OS 5.0.R25.

VCCV-PING OPTIONS

The loading or execution of the configuration file will fail and abort if an SAA vccv-ping is configured with an invalid value (for example, zero) for any of the following parameters: src-ip-address, dst-ip-address or pw-id. This will also occur if one or more of the parameters are

missing in a configuration file saved on a system running software version 5.0.R9 or earlier. This issue is resolved by updating the parameters to valid, non-zero values and ensuring that all three parameters are present before loading or executing the configuration file.

3.0.R5 or Earlier to 5.0.R25

The following notes apply to upgrading from releases prior to 7710 SR OS 3.0.R5 to 7710 SR OS 5.0.R25.

Autonegotiation and Link Aggregation Groups

The Ethernet ports in a Link Aggregation Group must have autonegotiation disabled or set to limited autonegotiation or the configuration will fail to load. Older configuration files must be manually edited to correct the port autonegotiation settings before upgrading to the 5.0.R25 release. [46342, 47100]

3.0 to 5.0.R25

The following notes apply to upgrading from 7710 SR OS 3.0 to 7710 SR OS 5.0.R25.

Management

- The “trap-destination” command has been changed to the “trap-target” command to support name targets as used in SNMPv2 MIB structures. The default name for a converted trap-destination is the IP address, a colon, and the port number (for example, 192.1.1.1:162).
- Trap targets with a notify-community string of exactly 32 characters are no longer supported. The new maximum length for the notify-community parameter on the “trap-target” command (which replaced “trap-destination”) is 31 characters. [57822]

Event-Control Changes

- NTP events 2006 and 2007 (tmnxNtpUtcOffsetExThreshold and tmnxNtpUtcOffsetInThreshold) were never implemented and are removed in 5.0.R3. Saved configurations files containing these events will deprecate them on load and not re-save them. [47996]

Network Policy Limit on Router Interfaces

When upgrading to 5.0.R25, a hard limit of 255 network policies for router interfaces is now enforced. Previously, this limit could be exceeded upon loading the configuration but excess QoS network policies were not operational. When upgrading to 5.0.R25, any additional network policies will now be rejected and must be removed prior to upgrading. [51437]

RSVP Message Pacing

The minimum values for the RSVP “msg-pacing” commands “max-burst” and “period” parameters have been changed from 10 milliseconds to 100 milliseconds. Configurations containing values less than 100 milliseconds will be converted to the new minimum when loaded. [56350]

VRRP Authentication

Because only password authentication is supported at this time, the VRRP “authentication-type” command has been deprecated. [42531]

BFD Minimum Multiplier

The new minimum multiplier value for BFD-enabled interfaces is now 3. Configuration files saved with the old minimum value of 2 will be continue to be loaded but will be converted to an operational value of 3. [49034]

PIM Three-way Hello Defaults

In 7750 SR OS 3.0.R10, 3.0.R11, 4.0.R2, 4.0.R3 and 7710 SR OS 3.0.R3, the PIM three-way-hello feature (config>router>pim>interface *ip-int-name*>three-way-hello) is enabled by default on VPRN multicast tunnels (MT) and spoke-SDPs and disabled on other interface types. In 7750 SR OS 3.0.R12 and later, 4.0.R4 and later and 7710 SR OS 3.0.R5 and later, the feature is disabled by default on all interfaces.

After an upgrade, the operational value will be the default for the updated version, meaning 1) the PIM three-way-hello feature will be disabled, 2) "compatibility-mode" required for interoperability with the above versions will be disabled and 3) the PIM adjacencies may be dropped. Manual setting of the three-way-hello parameter may be required to restore proper operation.

TACACS+ and RADIUS Source Address

The "source-address" entries for RADIUS and TACACS+ will need to be re-entered under the config>system>security>source-address context using the "application" command for these two protocols.

ISSU Upgrade Procedure

This section describes the ISSU Upgrade Procedure which can be used:

- When no firmware update is required
- On routers with redundant CFMs
- On routers running 5.0.R23 to 5.0.R24.

If any of the above criteria do not apply, the Standard Software Upgrade Procedure on page 66 must be performed.

Note:

Although the software upgrade can be performed using a remote terminal session, Alcatel-Lucent recommends that the software upgrade procedure be performed at the system CONSOLE device where there is physical access to the 7710 SR as remote connectivity may not be possible in the event there is a problem with the software upgrade. Performing the upgrade at the CONSOLE with physical access is the best situation for troubleshooting any upgrade problems with the help of the Alcatel-Lucent Technical Assistance Center.

Step 1 Backup Existing Images and Configuration Files

New software loads may make modifications to the configuration file which are not compatible with older versions of the software.

Note:

Configuration files may become incompatible with prior releases even if no new features are configured. The way in which a particular feature is represented in the configuration file may be updated by the latest version of the operating software. The updated configuration file would then be an unknown format to earlier software versions.

Alcatel-Lucent recommends making backup copies of the BOOT Loader (`boot.ldr`), software image and configuration files, should reverting to the old version of the software be required.

STEP 2 Copy 7710 SR OS Images to cf3:

The 7710 SR OS image files must to be copied to the cf3: device on the 7710 SR. It is good practice to place all of the image files for a given release in an appropriately named subdirectory off the root, for example, "cf3:\5.0.R25". Copying the `boot.ldr` and other files in a given release to a separate subdirectory ensures that all files for the release are available should downgrading the software version be necessary.

STEP 3 Copy boot.ldr to the Root Directory on cf3:

The BOOT Loader file is named `boot.ldr`. This file must be copied to the root directory of the cf3: device.

STEP 4 Modify the Boot Options File to Point to the New Image

The Boot Options File (`bof.cfg`) is read by the BOOT Loader and indicates primary, secondary and tertiary locations for the image file.

- The `bof.cfg` should be modified as appropriate to point to the image file for the release to be loaded.
- Use the "`bof save`" command to save the Boot Options File modifications.

STEP 5 Synchronize Boot Environment

Once the Boot Options File has been modified, the active and standby CFM boot environments must be synchronized.

- Use "`admin redundancy synchronize boot-env`" to synchronize the boot environments between the active and standby CFMs.

STEP 6 Reboot the Standby CFM

In the sample output below, the active CFM is in Slot A and the standby CFM is in Slot B. Before the start of ISSU, the card will look like the following:

```
A:router1# show card

===============================================================================
Card Summary
===============================================================================
Slot      Provisioned         Equipped            Admin      Operational
          Card-type           Card-type           State      State
-------------------------------------------------------------------------------
1         iom-12g             iom-12g             up         up
A         cfm-12g             cfm-12g             up         up/active
B         cfm-12g             cfm-12g             up         up/standby
===============================================================================
```

- Use "`admin reboot standby now`" to reboot the standby CFM and start the ISSU process.

```
A:router1# admin reboot standby now
A:router1# show card

===============================================================================
Card Summary
===============================================================================
Slot      Provisioned         Equipped            Admin      Operational
          Card-type           Card-type           State      State
-------------------------------------------------------------------------------
1         iom-12g             iom-12g             up         up
A         cfm-12g             cfm-12g             up         up/active
B         cfm-12g                                 up         down/standby
===============================================================================
```

STEP 7 Wait for Standby CFM to Synchronize

After the ISSU has been initiated, the card status of the standby CFM (in Slot B in this example) will show as “synching”.

```
A:router1# show card

===============================================================================
Card Summary
===============================================================================
Slot      Provisioned         Equipped            Admin      Operational
          Card-type           Card-type           State      State
-------------------------------------------------------------------------------
1         iom-12g             iom-12g             up         up
A         cfm-12g             cfm-12g             up         up/active
B         cfm-12g             cfm-12g             up         synching/standby
===============================================================================
```

When the standby CFM has completely synchronized, the standby CFM will indicate a state of “ISSU”.

```
A:router1# show card

===============================================================================
Card Summary
===============================================================================
Slot      Provisioned         Equipped            Admin      Operational
          Card-type           Card-type           State      State
-------------------------------------------------------------------------------
1         iom-12g             iom-12g             up         up
A         cfm-12g             cfm-12g             up         up/active
B         cfm-12g             cfm-12g             up         ISSU/standby
===============================================================================
```

Step 8 Reboot the Active CFM

After the standby CFM has synchronized and indicates a card status of “ISSU”, the active CFM (in Slot A in this example) must now be rebooted.

- Use “`admin redundancy force-switchover`” to reboot the active CFM.

In the sample output below, the active CFM in Slot A is rebooted from the CONSOLE on Slot A and the boot up messages from 7710 SR OS are displayed:

```
A:router1# admin redundancy force-switchover

TiMOS-C-5.0.Rx both/hops ALCATEL SR 7710 Copyright (c) 2000-2008 Alcatel-
```

```
Lucent.
All rights reserved. All use subject to applicable license agreements.
Built on ddd mmm d hh:mm:ss PST 2007 by builder in /rel5.0/panos/main

<...>
```

STEP 9 If Necessary, Re-establish a Console Session

If the ISSU is performed from the serial port CONSOLE on the CFM, the console session must be re-established on the newly active CFM.

STEP 10 Wait for Standby CFM to Synchronize

Before continuing with the ISSU procedure, the standby CFM must re-synchronize by transitioning from the “down”, “synchronizing” and finally to the “up” states. Use the command “show mda” to monitor the status of the MDAs and CMAs. Note that the MDAs and CMAs now have an “ISSU” status indicating that the active CFM is running the new image.

```
B:router1# show mda

===============================================================================
MDA Summary
===============================================================================
Slot  Mda    Provisioned           Equipped              Admin     Operational
             Mda-type              Mda-type              State     State
-------------------------------------------------------------------------------
1     2      c1-1gb-sfp            c1-1gb-sfp            up        ISSU
      3      m5-1gb-sfp            m5-1gb-sfp            up        ISSU
      5      c2-oc12/3-sfp         c2-oc12/3-sfp         up        ISSU
      6      c8-chds1              c8-chds1              up        ISSU
      9      c1-1gb-sfp            c1-1gb-sfp            up        ISSU
      11     c8-10/00eth-tx        c8-10/100eth-tx       up        ISSU
===============================================================================
```

Step 11 Reset the MDAs and CMAs to Load the New Image

The MDAs and CMAs must now be reset to load the new image.

Note:

The system does not allow cards to run in an ISSU state indefinitely; the system automatically resets the MDAs and CMAs after 2 hours. The “Comments” field in the “show card state” output displays the time until the system resets the MDA/CMA in the ISSU state.

The timing and order of the MDA and CMA resets should be sequenced to maximize the effectiveness of any redundant interfaces (LAGs, VRRP, etc.) spanning MDA/CMA slots.

- Use “`clear mda 1/<n>`” to reset an MDA/CMA

The sample output below shows the operational state transitions for a single MDA.

```
B:router1# clear mda 1/3
B:router1# show mda

===============================================================================
MDA Summary
===============================================================================
Slot  Mda   Provisioned           Equipped              Admin     Operational
            Mda-type              Mda-type              State     State
-------------------------------------------------------------------------------
1     2     c1-1gb-sfp            c1-1gb-sfp            up        ISSU
      3     m5-1gb-sfp                                  up        provisioned
      5     c2-oc12/3-sfp         c2-oc12/3-sfp         up        ISSU
      6     c8-chds1              c8-chds1              up        ISSU
      9     c1-1gb-sfp            c1-1gb-sfp            up        ISSU
      11    c8-10/100eth-tx       c8-10/100eth-tx       up        ISSU
===============================================================================
```

When the MDA/CMA is in the “up” state, it will have the new image so it will no longer have an “ISSU” operating state.

```
B:router1# show mda

===============================================================================
MDA Summary
===============================================================================
Slot  Mda   Provisioned           Equipped              Admin     Operational
            Mda-type              Mda-type              State     State
-------------------------------------------------------------------------------
1     2     c1-1gb-sfp            c1-1gb-sfp            up        ISSU
      3     m5-1gb-sfp            m5-1gb-sfp            up        up
      5     c2-oc12/3-sfp         c2-oc12/3-sfp         up        ISSU
      6     c8-chds1              c8-chds1              up        ISSU
      9     c1-1gb-sfp            c1-1gb-sfp            up        ISSU
      11    c8-10/100eth-tx       c8-10/100eth-tx       up        ISSU
===============================================================================
```

When all of the MDAs and CMAs have been rebooted, the ISSU is complete.

STANDARD SOFTWARE UPGRADE PROCEDURE

This section describes the Standard Software Upgrade Procedure which is service-affecting and must be used:

- When a firmware update is required
- On routers with non-redundant CFMs
- When ISSU is not supported between the old and new versions of software.

Each software release includes a BOOT Loader (`boot.ldr`). The BOOT Loader performs two functions:

1. Initiates the loading of the 7710 SR OS image based on the Boot Options File (`bof.cfg`) settings
2. Reprograms the boot ROM and firmware code on the CFM cards to the version appropriate for the 7710 SR OS image.

This section describes the process for upgrading the software and, if necessary, the boot ROM and firmware images with the BOOT Loader.

The software checks the firmware images on the CFM and reports any mismatch. If the loaded version is earlier than the expected version, the firmware may need to be upgraded; a console

or log message will indicate if a firmware upgrade is required. If the firmware version loaded is later than the expected version, no firmware programming is required.

The following steps describe the software upgrade process using the automatic firmware upgrade procedure introduced in 3.0.R1 of the 7710 SR OS image and BOOT Loader.

Note:

Although the software upgrade can be performed using a remote terminal session, Alcatel-Lucent recommends that the software upgrade procedure be performed at the system CONSOLE device where there is physical access to the 7710 SR as remote connectivity may not be possible in the event there is a problem with the software upgrade. Performing the upgrade at the CONSOLE with physical access is the best situation for troubleshooting any upgrade problems with the help of the Alcatel-Lucent Technical Assistance Center.

STEP 1 Backup Existing Images and Configuration Files

New software loads may make modifications to the configuration file which are not compatible with older versions of the software.

Note:

Configuration files may become incompatible with prior releases even if no new features are configured. The way in which a particular feature is represented in the configuration file may be updated by the latest version of the operating software. The updated configuration file would then be an unknown format to earlier software versions.

Alcatel-Lucent recommends making backup copies of the BOOT Loader (`boot.ldr`), software image and configuration files, should reverting to the old version of the software be required.

If the firmware version loaded is later than the expected version reported by the BOOT Loader, no firmware programming is required.

STEP 2 Copy 7710 SR OS Images to cf3:

The 7710 SR OS image files must to be copied to the cf3: device on the 7710 SR. It is good practice to place all the image files for a given release in an appropriately named subdirectory off the root, for example, “cf3:\5.0.R25”. Copying the `boot.ldr` and other files in a given release to a separate subdirectory ensures that all files for the release are available should downgrading the software version be necessary.

STEP 3 Copy boot.ldr to the Root Directory on cf3:

The BOOT Loader file is named `boot.ldr`. This file must be copied to the root directory of the cf3: device.

STEP 4 Modify the Boot Options File to Boot the New Image

The Boot Options File (`bof.cfg`) is read by the BOOT Loader and indicates primary, secondary and tertiary locations for the image file. The `bof.cfg` should be modified as appropriate to point to the image file for the release to be loaded. Use the "`bof save`" command to save the Boot Options File modifications.

STEP 5 [Redundant CFMs] Synchronize Boot Environment

On systems with Redundant CFMs, copy the image files and Boot Options File to the redundant CFM with

- "`admin redundancy synchronize boot-env`".

STEP 6 Reboot the Chassis

If no firmware update is required based on the rules in 5.0.R25 Firmware Update Rules on page 60, the chassis should be rebooted with the "admin reboot" command. If a 7710 SR is rebooted with the "admin reboot" command (without the "upgrade" keyword), no automatic firmware upgrades are performed by the BOOT Loader.

If a firmware update is required, use the "admin reboot upgrade" command. The "admin reboot upgrade" command in 7710 SR OS 3.0.R1 sets a chassis flag for the BOOT Loader and on the subsequent boot of 7710 SR OS on the chassis, any firmware images on CFMs requiring upgrading will be upgraded automatically. In addition, any CFMs which are inserted in the chassis until the next reboot will be upgraded automatically if possible. On rare occasions, the firmware of an inserted card cannot be upgraded automatically. Such cards will have an operational state of failed and a log message will be generated indicating an FPGA version mismatch.

Warning:

Do not power off, reset the system, insert cards or remove cards when firmware programming is being performed or the cards may be rendered inoperable requiring return for resolution.

Issuing and confirming the "admin reboot upgrade" operation will cause the 7710 SR to reboot, upgrade firmware as necessary and then reboot loading the configuration file.

The sample output below shows a 7710 SR-c12 running 3.0R12 performing an automatic firmware upgrade with the 5.0.R3 BOOT Loader and software image.

```
B:routerB# admin reboot upgrade
  ***********************************
  ** --->    W A R N I N G    <--- **
  ** ALL REQUIRED CHASSIS FIRMWARE **
  ** UPGRADES WILL BE DONE ON THIS **
  ** REBOOT AND MAY TAKE SEVERAL   **
  ** MINUTES. THE CHASSIS MUST NOT **
  ** BE RESET OR POWERED DOWN, NOR **
  ** CARDS INSERTED OR REMOVED,    **
  ** DURING THIS PROCESS. ANY OF   **
  ** THESE PROHIBITED ACTIONS MAY  **
  ** CAUSE CARDS TO BE RENDERED    **
  ** INOPERABLE!                   **
  ***********************************
Are you sure you want to reboot (y/n)? y

sysToMonitor: Resetting...


Alcatel 7xx0 Boot ROM. Copyright 2000-2008 Alcatel-Lucent.
All rights reserved. All use is subject to applicable license agreements.
Build: X-3.0.R12 on Tue Apr 24 10:55:17 EDT 2007 by builder
Version: 0x1A
Starting 7710 CFM B
COLD boot on processor #1
CPU Control FPGA version is 0x24
?Starting bootrom RAM code...
Boot rom version is v26
Booted from Control PROM 1
>>>Testing mainboard FPGA chain...
>>>Validating SDRAM from 0x7ff00000 to 0x80000000
>>>Clearing SDRAM from 0x02200000 to 0x7ff00000
>>>Testing Compact Flash 1... Slot Empty
>>>Testing Compact Flash 2... Slot Empty
>>>Testing Compact Flash 3... OK (SILICONSYSTEMS INC 256MB)
CFMCTL FPGA version is 0x24
CCM FPGA version is 0x99
Board Serial Number is 'NS070380006'
Chassis Serial Number is 'NS070380007'
Searching for boot.ldr on local drives:
Searching cf3 for boot.ldr...
************************************************
CFM Control FPGA mismatch. Expecting version 0x29, but have 0x24.

Total Memory: 2016MB  Chassis Type: sr-c12  Card Type: canada_r1
TiMOS-L-5.0.R3 boot/hops ALCATEL SR 7710 Copyright (c) 2000-2008 Alcatel-
Lucent.
All rights reserved. All use subject to applicable license agreements.
Built on Mon Apr 16 18:16:55 PDT 2007 by builder in /rel5.0/b1/B1-34/panos/main
```

```
TiMOS BOOT LOADER
Time from clock is TUE APR 24 20:40:13 2007 UTC

Unexpected but acceptable bootrom version; found v26, expected v27

** CPU FPGA is too old; is v36, should be v40 **
Chassis firmware will be upgraded automatically

Switching serial output to sync mode...   done

Checking for firmware upgrades...

Embedded BOOT ROM: 27/X-5.0.R3 on Mon Apr 16 17:11:07 PDT 2007 by builder

Sl Type      Serial #        FPGA up? BOOT ROM                              up?
 A iom-12g   *none*           24=>29  26/X-3.0.R12 on Tue Apr 24 10:55:17 E YES
 B iom-12g   NS070380006      24=>29  26/X-3.0.R12 on Tue Apr 24 10:55:17 E YES

********************************
** ROM PROGRAMMING IN PROGRESS **
**   DO NOT RESET CHASSIS,     **
**   INSERT OR REMOVE CARDS    **
**   WHILE IN PROGRESS, OR     **
**   CARDS MAY BE RENDERED     **
**   INOPERABLE.               **
********************************

STOPPING CPM IN SLOT A DURING PROGRAMMING...
Slot A not ready to download. Clearing...OK
OK

***** Programming Slot A *****

PROGRAMMING FPGA ROM UPGRADE SPACE
   Erasing 70000 bytes at address 1f200000: 70000 OK
Programing 67173 bytes at address 1f200000: 67173 OK

PROGRAMMING BOOT ROM UPGRADE SPACE
   Erasing 40000 bytes at address 1f3c0000: 40000 OK
Programing 31338 bytes at address 1f3c0000: 31338 OK

***** Programming Slot B (this slot) *****

PROGRAMMING FPGA ROM UPGRADE SPACE
   Erasing 70000 bytes at address 1f200000: 70000 OK
Programing 67173 bytes at address 1f200000: 67173 OK

PROGRAMMING BOOT ROM UPGRADE SPACE
   Erasing 40000 bytes at address 1f3c0000: 40000 OK
Programing 31338 bytes at address 1f3c0000: 31338 OK

CPM in slot A will now clear


******************************************
** CARD PROGRAMMING COMPLETED SUCCESSFULLY **
******************************************

Resetting...OK
```

```
Alcatel 7xx0 Boot ROM. Copyright 2000-2008 Alcatel-Lucent.
All rights reserved. All use is subject to applicable license agreements.
Build: X-3.0.R12 on Tue Apr 24 10:55:17 EDT 2007 by builder
Version: 0x1A
Starting 7710 CFM B
COLD boot on processor #1
CPU Control FPGA version is 0x24
?Starting bootrom RAM code...
Boot rom version is v26
Booted from Control PROM 1
>>>Testing mainboard FPGA chain...
>>>Validating SDRAM from 0x7ff00000 to 0x80000000
>>>Clearing SDRAM from 0x02200000 to 0x7ff00000
>>>Testing Compact Flash 1... Slot Empty
>>>Testing Compact Flash 2... Slot Empty
>>>Testing Compact Flash 3... OK (SILICONSYSTEMS INC 256MB)
CFMCTL FPGA version is 0x24
CCM FPGA version is 0x99
Board Serial Number is 'NS070451368'
Chassis Serial Number is 'NS070380007'
Upgrading BOOT rom to v27 / X-5.0.R3 on Mon Apr 16 17:11:07 PDT 2007 by buil
Programming 369304 bytes at address 0xbfc00000
!........!........!........!........!........!........

rebooting...

Alcatel 7xxx Boot ROM. Copyright 2000-2008 Alcatel-Lucent.
All rights reserved. All use is subject to applicable license agreements.
Build: X-5.0.R3 on Mon Apr 16 17:11:08 PDT 2007 by builder
Version: 0x1B
Starting 7710 CFM B
COLD boot on processor #1
CPU Control FPGA version is 0x24
?Preparing for jump to RAM...
Starting bootrom RAM code...
Boot rom version is v27
Booted from Control PROM 1
>>>Testing mainboard FPGA chain...
>>>Validating SDRAM from 0x7ff00000 to 0x80000000
>>>Clearing SDRAM from 0x02200000 to 0x7ff00000
>>>Testing Compact Flash 1... Slot Empty
>>>Testing Compact Flash 2... Slot Empty
>>>Testing Compact Flash 3... OK (SILICONSYSTEMS INC 256MB)
CFMCTL FPGA version is 0x24
CCM FPGA version is 0x99
Board Serial Number is 'NS070380006'
Chassis Serial Number is 'NS070380007'
Upgrading FPGA rom to v29
File name:               cfmctl.ncd
FPGA name:               2vp40ff1152
Compile date:            2006/11/22
Compile time:            14:43:47
Image size (bytes):      0x1E4424
Programming...0x1E4440
Verify0x1E44400354A0
fpgaProgramPROM() status: 0

rebooting...
```

```
Alcatel 7xxx Boot ROM. Copyright 2000-2008 Alcatel-Lucent.
All rights reserved. All use is subject to applicable license agreements.
Build: X-5.0.R3 on Mon Apr 16 17:11:08 PDT 2007 by builder
Version: 0x1B
Starting 7710 CFM B
COLD boot on processor #1
CPU Control FPGA version is 0x29
?Preparing for jump to RAM...
Starting bootrom RAM code...
Boot rom version is v27
Booted from Control PROM 1
>>>Testing mainboard FPGA chain...
>>>Validating SDRAM from 0x7ff00000 to 0x80000000
>>>Clearing SDRAM from 0x02200000 to 0x7ff00000
>>>Testing Compact Flash 1... Slot Empty
>>>Testing Compact Flash 2... Slot Empty
>>>Testing Compact Flash 3... OK (SILICONSYSTEMS INC 256MB)
CFMCTL FPGA version is 0x29
CCM FPGA version is 0x99
Board Serial Number is 'NS070380006'
Chassis Serial Number is 'NS070380007'
Searching for boot.ldr on local drives:
Searching cf3 for boot.ldr...
*************************************************

Total Memory: 2016MB  Chassis Type: sr-c12  Card Type: canada_r1
TiMOS-L-5.0.R3 boot/hops ALCATEL SR 7710 Copyright (c) 2000-2008 Alcatel-
Lucent.
All rights reserved. All use subject to applicable license agreements.
Built on Mon Apr 16 18:16:55 PDT 2007 by builder in /rel5.0/b1/B1-34/panos/main

TiMOS BOOT LOADER
Time from clock is TUE APR 24 20:46:51 2007 UTC
Switching serial output to sync mode...   done

Looking for cf3:/bof.cfg ... OK, reading

Contents of Boot Options File on cf3:
    primary-image    cf3:/TiMOS-5.0.R3
    primary-config   cf3:/7710SR.cfg
    address          192.168.10.3/24 active
    address          192.168.10.4/24 standby
    primary-dns      192.168.10.1
    dns-domain       domain.com
    autonegotiate
    duplex           full
    speed            100
    wait             4
    persist          off
    console-speed    115200

Hit a key within 1 second to change boot parms...

Primary image location: cf3:/TiMOS-5.0.R3
Initializing management port tme0 using IP addr 138.120.193.135.
Loading image cf3:/TiMOS-5.0.R3/both.tim
Version B-5.0.R3, Mon Apr 16 18:14:17 PDT 2007 by builder in /rel5.0/b1/B1-
34/panos/main
text:(26582684-->61715760) + data:(1609174-->14006224)
Executing TiMOS image at 0x2800000
```

```
Total Memory: 2016MB  Chassis Type: sr-c12  Card Type: canada_r1
TiMOS-B-5.0.R3 both/hops ALCATEL SR 7710 Copyright (c) 2000-2008 Alcatel-
Lucent.
All rights reserved. All use subject to applicable license agreements.
Built on Mon Apr 16 18:14:17 PDT 2007 by builder in /rel5.0/b1/B1-34/panos/main
```

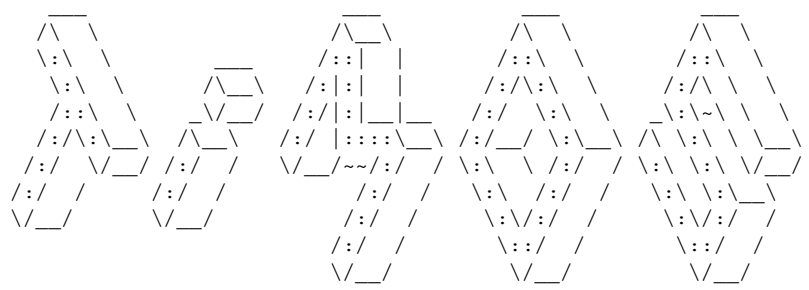

```
Time from clock is TUE APR 24 20:47:11 2007 UTC

Attempting to exec primary configuration file:
   '7710SR.cfg' ...
System Configuration
Log Configuration
System Security Configuration
Card Configuration

TiMOS-B-5.0.R3 both/hops ALCATEL SR 7710 Copyright (c) 2000-2008 Alcatel-
Lucent.
All rights reserved. All use subject to applicable license agreements.
Built on Mon Apr 16 18:14:17 PDT 2007 by builder in /rel5.0/b1/R3/panos/main


Login:
```

STEP 7 Verify the Software Upgrade

Allow the boot sequence to complete and verify that all cards come online.

Note:

If any card fails to come online after the upgrade, contact the Alcatel-Lucent Technical Assistance Center for information on corrective actions.

KNOWN LIMITATIONS

Following are the known limitations for Release 5.0.R25 of 7710 SR OS.

AUX PORT

- The AUX serial port on the CFM is currently not supported in software. The 7710 SR OS software does not provide a means of configuring the device.

GIGABIT ETHERNET

- In certain interoperability situations, it may be necessary to disable autonegotiation on Gigabit Ethernet ports.

DS1/E1

- Via SNMP, a value of 0 will be returned for tmnxDS1BERTTotalBits as this function is not supported on the DS1/E1 CMA. This value is properly shown as “N/A” in the CLI. [bz1400]

SONET/SDH

- OC-48c/STM-16c SONET/SDH interfaces only run in CRC32 mode. CRC16 mode cannot be configured for these interfaces.
- The ports on OC-3c/STM-1c SONET/SDH MDA are serviced in groups of four (1-4, 5-8) by a single framer chip, and as such, all must have the same framing across a given group. If framing on one port is changed, all four ports in a group must be shutdown and the framing will be changed on all four ports.
- The framer on 2-port OC-48/STM-16c MDA supports a single software reset for all transmit subsystems, so changes to the transmit clock source on a single port will result in a short traffic interruption on all ports on the MDA. As a result, a short interruption will be experienced on all ports on the MDA when the transmit clock source for any one port is changed, for example from line to node timed. Also, traffic will be interrupted on all ports on the MDA when the port loopback mode on a port also configured with loop timing are transitioned in any of the following ways:
  - from ‘no loopback’ to Internal
  - from Internal to ‘no loopback’
  - from Internal to Line
  - from Line to Internal.
- If the H1 and H2 bytes are set to 0xFF but the H3 byte is not set to 0xFF, an AIS-P condition is not reported but an LOP-P condition is reported. [30498]
- When the clock-source on an OC-48 port is changed, the transmit clock on all OC-48 ports on the same MDA are reset, resulting in a brief loss of traffic on the other ports. [31401]

APS

- APS implementation supports 1+1 signaling and 1:1 data path (data is TX on active circuit only). The far-end node must be compliant with bi-directional or uni-directional APS signaling specifications to interoperate with the 7710 SR APS implementation.
- ITU-T G.783 Annex B is not supported.
- MLPPP cannot be configured on APS-protected ports
- Uni-directional APS is not supported
- In some cases of RDI-L, the transmitted K1/K2 bytes on the wire may differ from those maintained by the CFM's APS controller (as displayed in the CLI). [36537]

ATM MDAS ACCESS MODE ONLY

- The ATM interfaces only support the customer-facing access mode.

ATM AND IS-IS

- IS-IS is not supported on IES and VPRN interfaces with ATM PVC SAPs in this software release.

ATM Traffic Management Limitations

The following applies to OC-12c/STM-4c ATM MDA.

In the context of multiple services using an ATM MDA, the following two criteria must be met in order to satisfy the QoS guarantees:

- VC fairness
- COS fairness

VC fairness implies that each VC gets its due share of bandwidth relative to the other VCs and COS fairness implies that within each VC, each COS gets its due share of bandwidth. What is considered the “due share” is very specific to the configuration. (For example, for two VCs of the same ATM service category, the due share will be proportionate to the configured rates of the VCs; for 2 VCs with different ATM service categories, the due share will depend on the priority of the service category and the configured rate, etc.)

A minor loss of throughput (< 2% of line rate) may occur if an OC-12 port is configured with small number of shaped PVCs, the difference in the configured ATM rates of the PVCs is large, and the sum of the shaped rates is equal to port rate. The loss of packet throughput occurs in the highest traffic parameter VC and only. [28869]

The ATM layer shaping in the MDA schedules cells of the high-priority Forwarding Class queues with strict priority over cells of low-priority Forwarding Class queues within a SAP. This is performed such that packet delay and jitter are minimized on the high-priority forwarding class queues. As a result in some traffic loading scenarios, the lower priority forwarding class queues may not achieve their fair share of bandwidth. This is the case when the high-priority Forwarding Class queues have an offered traffic to the ATM MDA per-VC queue equal or higher than the PIR of the ATM VC. The user can alter this behavior and trade delay performance for forwarding class fairness in this specific scenario configuring H-QoS schedulers to limit the total offered load out of the forwarding class queues to the ATM MDA per-VC queue to the PIR of the ATM VC. [30819]

ATM Traffic/Statistics Limitations

The following limitations apply to the OC-12c/STM-4c ATM:

- OC-12/STM-4 latency increases when applying a new ingress SAP policy that adds more queues. The latency increases from around 22.2 μs to 24.8 μs over a 1 min period. Traffic loss does not occur during this period.
- Port input statistics do not increase when terminating e-t-e AIS cells are received.
- PVC admin state is not applicable - There is no command that can administratively disable a PVC. In order to disable a PVC, the user must disable the applicable service or service interface.

Class of Service Fairness Affected on Shaped VCs

The following applies to OC-12c/STM-4c ATM MDA.

In the case of ATM VCs configured with more than two classes of service where one queue, queue A, is allowed no burst beyond CIR and another queue of the same priority, queue B, is allowed to burst up to line-rate; the traffic offered from queue B might prevent queue A from achieving its CIR. The problem has a lesser degree of impact if there is an increased number of ATM VCs on the port and can also be addressed by lowering the configured PIR of queue B. [35224]

Shared Queueing QoS

- In a SAP Ingress QoS Policy with shared queueing, high-priority packets dropped will be counted in the low-priority drops in the SAP ingress service queue statistics. [32335]

HW/PLATFORM

- On the 8-port SONET/SDH OC-3c/STM-1c MDA and 2-port OC-12c/STM-4c CMA, only the first 16 bytes of the 62 byte trace string can be unique for each group of four (4) ports (for example, for ports 13 through 16) for ports operating in SONET mode at OC-3. The last 48 bytes of the trace string will be the same for all ports and will be the last value set. Basically, a unique trace string per port is not possible if the unique part of the string is longer than 14 characters.
- On the 8-port SONET/SDH OC-3c/STM-1c MDA, 4-port ATM OC-12c/STM-4c MDA and 2-port OC-12c/STM-4c CMA, the normal range for the SONET/SDH line signal failure Bit Error Rate (BER) threshold configured using the configure>port *port-id*>sonet-sdh>threshold command is 3 to 6. For these MDAs, the allowed threshold values are 3 to 5. The SNMP variable for this exponential threshold is tmnxSonetBerSfThreshold.

RADIUS

- If the system IP address is not configured, RADIUS user authentication will not be attempted for in-band RADIUS servers unless a source-address entry for RADIUS exists.
- The NAS IP address selected is that of the management interface for out-of-band RADIUS servers. For in-band RADIUS servers if a source-address entry is configured, the source-address IP address is used as the NAS IP address, otherwise the IP address of the system interface is used.
- SNMP access cannot be authorized for users by the RADIUS server. RADIUS can be used to authorize access to a user by FTP, console or both.
- If the first server in the list cannot find a user, the server will reject the authentication attempt. In this case, the 7710 SR router does not query the next server in the RADIUS server list and denies access. If multiple RADIUS servers are used, the software assumes they all have the same user database.

TACACS+

- If the TACACS+ start-stop option is enabled for accounting, every command will result in two commands in the accounting log.
- If TACACS+ is first in the authentication order and a TACACS+ server is reachable, the user will be authenticated for access. If the user is authenticated, the user can access the console and any rights assigned to the default TACACS+ authenticated user template (config>system>security>user-template tacplus_default) (see TACACS+ enhancement on page 45). Unlike RADIUS, TACACS+ does not have fine granularity for authorization to define if the user has just console or FTP access, but 7710 SR OS supports a default template for all TACACS+ authenticated users.

  If TACACS+ is first in the authentication order and the TACACS+ server is NOT reachable, authorization for console access for the user is checked against the user's local or RADIUS profile if configured. If the user is not authorized in the local/RADIUS profile, the user is not allowed to access the box.

  Note that inconsistencies can arise depending upon combinations of the local, RADIUS and TACACS+ configuration. For example, if the local profile restricts the user to only FTP access, the authentication order is TACACS+ before local, the TACACS+ server is UP and the TACACS+ default user template allows console access, an authenticated TACACS+ user will be able to log into the console using the default user template because TACACS+ does NOT provide granularity in terms of granting FTP or console access. If the TACACS+ server is DOWN, the user will be denied access to the console as the local profile only authorizes FTP access. [39392]

CLI

- The CLI allows the user to specify a tftp location for the destination for the “admin save” and “admin debug-save” commands which will overwrite any existing file of the specified name. [18554]
- The firmware limits ICMP packet to be generated at the rate of 100 packets/sec. However when configuring an interface in the CLI, the user is allowed to configure ICMP packets to be generated at rates up to up to 1000 packets/sec. [46767]
- The system does not prevent the user from using the same IP address of its BGP peer on one of its router interfaces. [57198]

SYSTEM

- The 7710 SR-c12 chassis cannot differentiate between a missing and non-functioning fan tray. [17756]
- Dropped incoming packets due to a packet processing error are not being counted in the ifInErrors SNMP counter. Examples of packets such as this include any packet with a malformed IP header. [27699]
- When a fan is removed from a 7710 SR-c12, an erroneous “fan high temperature alarm” is generated that is cleared when the fan is replaced. [36112]
- Remapping of control plane traffic from a default CFM queue to a different queue is not supported on the 7710 SR. [59438]

ATM

- On the OC-12c/STM-4c ATM MDA, ATM Apipes configured with a vc-type of atm-vpc drop all ATM OAM F4 segment cells and pass through the ATM OAM F4 end-to-end cells. The PTI field of the forwarded ATM OAM F4 end-to-end cells is set to five and might cause interoperability issues if the third-party equipment expects the PTI field to be zero. [40451]
- The 7710 SR does not support bi-directional FR PVC management procedures over an ATM VC part of an FRF.5 VLL. When doing FRF.5 interworking between the 7710 SR and other products, the bi-directional network PVC management over the ATM VC must be disabled on the other products. [49696]

MLPPP

- In order to interoperate with other vendors’ MLPPP implementations, the MLPPP sub-layer will accept packets with or without leading zeros in the protocol field even though the 7710 SR does not advertise the protocol field compression (PFC) option during LCP negotiation. [25996, 29923]

MANAGEMENT

- Collision events detected on a CFM management Ethernet port are reported as CRC/Alignment errors. [30205]
- The SNMP MIB attribute snmpEngineTime is restarted upon a High-Availability switchover or system reboot. Additionally, the snmpEngineBoots attribute is restarted on a system reboot. [46179]
- Source address configuration applies only to the Base routing instance, and where applicable, to VPRN services. As such, source address configuration does not apply to unsolicited packets sent out the management interface.
- The SSHv2 implementation does not support the RC5 cryptographic algorithm. [47122]
- After 497 days, system up-time will wrap around due to the standard RFC 1213 MIB-II 32-bit limit. [51129]

TCP AUTHENTICATION EXTENSION

- It is not possible to delete an authentication keychain if that keychain was recently removed from a BGP neighbor while BGP was operationally down. BGP has to become operationally active before the keychain can be deleted. [57277]

ROUTING

- Routes exported from one protocol to another are redistributed with only the first ECMP next-hop. Therefore, if BGP routes having multiple next-hops are exported to a VPRN client, only one next-hop for the route will be exported. The one chosen is the lowest IP address of the next-hop address list. [40147]

IPv6

- MAC filtering does not match on IPv6-enabled IES interfaces. [44897]
- When “debug router ip packet” is enabled, packets received on a 6over4 tunnel do not display the IPv4 header information and packets sent on the tunnel do not display the IPv6 header information as the encapsulation and decapsulation is performed on the line card. [45606]
- A change in any IS-IS Multi-topology and/or level will cause the SPF to be run in all levels and/or topologies. [56527]

OSPF PE-CE

- The following OSPF features are not supported in OSPF PE-CE instances in this release:
  - stub areas
  - graceful-restart
  - traffic engineering

BGP

- If BGP transitions to the operationally disabled state, the “clear router bgp protocol” command will not clear this state. The BGP protocol administrative state must be shutdown/no shutdown to clear this condition. [12074]
- Under rare circumstances, multiple BGP peer sessions between the same two systems may not be successfully established without a “local-address” explicitly configured for each peer. [77884]

MPLS/RSVP

- The “no rsvp” command in the config>router context has no effect as the state of RSVP is tied to the MPLS instance. The “no mpls” command deletes both the MPLS and RSVP protocol instances. [8611]
- Shutting down a port on an OC-3c/STM-1c MDA may not provide sub-50 ms failover for an RSVP path signaled over that port. This issue does not occur if the fiber is disconnected or if the path is shut down. [39973]
- Fast failover times of less than 100 ms cannot be achieved for fast reroute protected LSPs if the failed link is detected by copper Ethernet SFPs. Sub-second failover times are achieved, but the failover times with copper Ethernet SFPs are inherently longer based on how the system communicates with the SFP. [49003].
- There are scenarios where the bypass optimization does not ensure that a node-protect manual bypass will be selected over a node-protect dynamic bypass tunnel. This is because

the manual bypass may be unavailable when the association of a bypass LSP is made with the primary LSP.

The bypass optimization feature only changes the association for an LSP which requested node protection but is currently associated with a link-protect bypass.

To ensure this selection when using manual bypass, dynamic bypass must explicitly be disabled. [60261]

IP MULTICAST

- The Router Alert IP option is not included in mtrace queries that are unicast to the last-hop router in the trace as defined by the IETF draft. Note that this causes no known interoperability issues since this packet is still destined for an IP address on this last-hop router. [37923]
- Cisco routers that incorrectly send mtrace queries to the group multicast address rather than the ALL-ROUTERS.MCAST.NET address (as defined by the IETF draft) will be discarded. Additionally, some Cisco routers do not fill in the “oif” field in the response block, and some do not accept an mtrace query that comes in on the “oif” interface. A workaround in this last case is to use the RPF as the destination address for the query. [39070]

QOS

- QoS and IP filter matches on IP frames are limited to Ethernet Type 2 IP frames. In particular, Ethernet SNAP IP frames will not be matched with IP match criteria. [15692]
- Small amounts of packet loss may occur on queues with an MBS of 4 KB or lower when the traffic rate through that queue is large or if that queue is subject to a lot of back pressure from the egress port. This packet loss is possible without congestion on queues where the traffic rate is lower than the PIR. To avoid this type of packet loss, the MBS of a queue should not be configured to a value lower than 5 KB. [66687]

FILTERS/PBR/TCS

- IP filters will not match against IS-IS control packets since these packets do not contain IP headers. [40976]
- An IP address must be assigned to the system interface and the interface operationally up in order for Web Portal Redirection to operate. [46305]

SUBSCRIBER MANAGEMENT

- Subscriber Management is only supported on Ethernet Layer 2 VPLS SAPs and for Layer 3 only on Routed CO IES and VPRN interfaces.
- Dynamic subscribers learned (via DHCP) while sub-sla-mgmt is shutdown will continue to use the SAP-level ingress and egress filter rules. Once the subscriber is relearned (renewed), the subscriber profile filters will then be used. This does not apply to static subscribers. [47167]

TPSDA

- ATM RBE SAPs are not supported for the Routed CO model.

IES

- In the saved configuration for IES services, the IES instance and interfaces will appear twice: once for creation purposes and once with all of the configuration details. This allows configuration items such as DHCP server configuration to reference another IES interface without errors. [56086]

- If two IES interfaces are connected back-to-back through a 2-way spoke-SDP connection with SDPs that have keep-alive enabled and IGP is enabled on the IES interface with a lower metric as the network interfaces, the related SDPs will bounce due to SDP keep-alive failure. The GRE encapsulated SDP ping reply will be ignored when it is received on an IES interface. [68963]

VPRN/2547

- The use of auto-bind and spoke-sdp within a VPRN is mutually exclusive. [21529]
- Dynamic Multipath changes might not work in the case of VPN-IPv4 routes and might require a restart of the service. [31280]
- Each MP-BGP route has only one copy in the MP-BGP RIB, even if that route is used by multiple VRFs. Each MP-BGP route has system-wide BGP attributes and these attributes (preference) cannot be set to different values in different VRFs by means of vrf-import policies. [34205]

VRRP/SRRP

- The MAC address displayed for an SRRP gateway IP in the “show router arp” output on a subscriber interface does not show the MAC address of the Virtual Router but is that of the interface. Use the “show srrp” command to see the VR MAC address actually in use. [57838]

CFLOWD

- Cflowd is not supported on subscriber SLAs.

MIRROR SERVICE

- Simultaneous Filter Logging and Service Mirroring on egress is not supported. When simultaneously performing filter logging and service mirroring at egress, the service mirroring operation takes precedence over the filter logging operation. This behavior was introduced in Release 2.0. In Release 1.3 and earlier releases, the filter logging takes precedence and the service mirroring of the packet is not performed.

SPANNING TREE

- The RSTP and MSTP Spanning Tree Protocols operate within the context of a VPLS service instance. The software allows for the configuration of an STP instance per VPLS service instance (1K for a 7710 SR); however, operation within the STP specification limits is not always guaranteed as it is highly dependent on 1) the number of SAPs/SDPs per VPLS and 2) the number of MAC addresses active within a VPLS.

OAM

- OAM-VPRN ping and traceroute for VPRN in a hub and spoke topology using hairpin routing does not work. If a hub and spoke topology is used, the spoke site must be associated with the hub VRF or the default route created must point to the hub site not a blackhole. If not, some sites will not be reachable from the spoke site.
- OAM-VPRN ping and traceroute does not work in a hub and spoke network topology with the 7710 SR as the Customer Edge (CE) hub. As a workaround, the 7710 SR will send a control plane response from the hub to the requester Provider Edge (PE) to confirm connectivity to the hub PE.
- LDP Treetrace is not supported for LDP-over-RSVP tunnels as these tunnels do not support ECMP at this time. [58297]

- An OAM Service Ping request for a VPRN service is always sent over the data plane (over the spoke SDP) and not through the control plane. A VPRN Ping should be used to send a ping request using the control plane for a VPRN instance. [58479]

RESOLVED ISSUES

RESOLVED IN 5.0.R25

Following are specific technical issues that have been resolved in Release 5.0.R25 of 7710 SR OS since Release 5.0.R24.

HW/PLATFORM

- Ethernet ports now correctly renegotiate Ethernet parameters (speed, duplex, pause support). [82939]

RADIUS

- A system configured to use "exit-on-reject" that also had RADIUS authentication as the first of two or more authentication methods might not have attempted subsequent methods even if it did not receive an explicit REJECT from the RADIUS server (for example, first response time-out). This issue has been resolved. [85048]

TACACS+

- When configuring multiple authentication methods with the “exit-on-reject” option and if “tacplus” is not the last authentication method, the system will only exit and not try the next authentication method if it explicitly receives a reject from the TACACS+ server. [63012]

MANAGEMENT

- Creating tech-support files on a system that has IMA bundle protection groups (bgrp-ima) configured no longer causes system instability. [86859]

ROUTING

- When using triggered-policy, changing the import and export policy of a particular neighbor at the same policy-configuration commit was only notifying BGP of changes on the import policy for the neighbor and changes on the export policy were ignored. This issue has been resolved. [85984]
- Using regular expressions with the format “[A B...] {X}”, where X is greater or equal to two (2), in route policy as-path lists or in the “show router bgp routes aspath-regex” command no longer causes system instability. [86420]

BGP

- When a BGP next-hop address was resolved through a default route, it would incorrectly show as unresolved in the "show route bgp next-hop" CLI command output. This issue has been resolved. [80043]
- When loop-detection is set to "discard-route" on a BGP neighbor, the loop detect parameter in the "show router bgp neighbor detail" command output is no longer incorrectly set to “BADVAL”. [83298]

MPLS/RSVP

- An MPLS-enabled interface that flaps for an extended period of time (for example, days) no longer results in system instability. [84596]
- An MPLS FRR bypass tunnel that is signaled over more than 16 hops will no longer result in a standby CFM reset. [85001]

SERVICES GENERAL

- The active/standby pseudowire SDP's operational state for Epipes might have been displayed incorrectly as "Pruned" in the "show service sdp-using" output. This issue has been resolved. [85200]

VPLS

- Having an address configured on a VPLS management interface that is used by another host on the same subnet (duplicate address) no longer causes system instability. [84785]

VPRN/2547

- Changing the Route Distinguisher in a VPRN residing on a node with a large number of VPRNs and VPRN routes might have, in rare cases, resulted in a CFM High Availability switchover. This issue was resolved in 5.0.R9. [59814]

RESOLVED IN 5.0.R24

Following are specific technical issues that have been resolved in Release 5.0.R24 of 7710 SR OS since Release 5.0.R23.

CLI

- The output of the "show router bgp routes" command no longer truncates AS-path values for BGP routes when the total number of characters including spaces is greater than 83 characters and the last AS entry starts before the 83rd character and ends after. [80755]

SYSTEM

- Performing an "ssh server-shutdown" while the system is in the process of closing out recently terminated SSH connections no longer causes the session to lock-up. [77069]
- The system will now accurately generate an event at the end of the sampling period during which a throttled event has exceeded its pre-configured throttle rate and the system had discarded the additional entries. This applies to both SNMP and log-only events. [80974]
- Configuring an interface with the "allow-directed-broadcast" option no longer results in system instability when processing directed broadcast packets. [81145]

NTP

- Systems configured to use NTP or SNTP to synchronize their clocks will now regularly update the standby CFM's real-time clock. [82379]

ROUTING

- Broadcast ICMP ping packets received on a group interface were not replied correctly and caused the system to raise an error message in the log: "A:ATIC_L3:UNUSUAL_ERROR icmp_EchoReqRx: couldn't find the prime address for if". This issue has been resolved. [64560]
- The use of the regular expression {m,} to match community values in route policies might have caused the system to take more time than usual to process the condition and potentially slowed down protocol convergence. This issue has been resolved. [83847]

IS-IS

- The system can now establish more than 42 IS-IS adjacencies over a single broadcast IS-IS interface. [80820]

OSPF

- If the 7710 SR receives its own router LSA with MaxSequenceNumber, it max ages as expected and now correctly sets the sequence number of that max aged LSA to 0x7fffffff. [73931]
- An ABR configured as “nssa originate default-route” and “no summaries” no longer generates a default-route when all of its interfaces to Area 0 are “Down” as long as the system interface is administratively up and is configured in Area 0. [80044]

BGP

- Sharing policy communities that use regular-expressions between BGP policies and OSPF policies might have caused system instability. This issue has been resolved. [76697]

QoS

- When using H-QoS in conjunction with low traffic rates, if the last operational rate of a given queue was 1Kbps and subsequently that queue goes into an inactive state (for example, not sending more packets), the scheduling algorithm now correctly offers no bandwidth to that queue. [81002]
- QoS network-queue policy values "mbs 100" and "high-prio-only default" for queues one (1) and nine (9) were not saved correctly with "admin save". After a High Availability switchover or upgrade, those values were incorrectly changed to "mbs 50" and "high-prio-only 10". This issue has been resolved. [82922]

CFLOWD

- When monitoring fragmented TCP or UDP traffic, a cflowd flow entry is no longer created for every fragment. [81356]

VPLS

- ARP-reply packets are no longer incorrectly remarked with a dot1p value of zero (0) when egressing a residential split-horizon group SAPs that has ARP reply agent enabled. [81040]

VPRN/2547

- Configuring an unnumbered interface as RIP neighbor no longer results in a CFM High Availability switchover. [82752]

RESOLVED IN 5.0.R23

Following are specific technical issues that have been resolved in Release 5.0.R23 of 7710 SR OS since Release 5.0.R22.

HW/PLATFORM

- The MDI/MDX state of an Ethernet port is now displayed correctly when using a copper SFP. [70720]
- The system may have failed to boot from a 2GB compact flash disk that was formatted using CLI because of its inability to locate the BOOT.LDR file on the compact flash disk. This did not occur on compact flash disks that had been formatted using a Windows PC. This issue has been resolved. [74239]

CLI

- RMON threshold entries now correctly load after upgrading to a newer revision of SR OS where those entries refer to obsolete MIB objects. [73209]
- A 7710 SR no longer incorrectly increments the “Int MAC Tx Errs” count for certain traffic patterns on m20-1gb-tx MDA ports operating at 1Gbps. [75076]
- The CLI command “show router bgp route ipv6 xx::/32” now correctly displays all matching IPv6 prefixes with a mask length of 32. [76192]
- The report-alarm command is no longer allowed on non-10GE ports. [76752]

SYSTEM

- Running several SNMP walks of the log event table on a system where both the security and main source streams go to the same log-id might have caused the system to run at a higher than normal CPU utilization. This issue has been resolved. [66291, 71508]
- A continuously flapping port that belongs to a LAG might have caused high CPU utilization. This mainly affected a scaled configuration with a large number of SAPs configured on the LAG. The workaround was to configure a timer (“hold-time up”) on ports belonging to a LAG to prevent the flapping from affecting the CPU. This issue has been resolved. [71377]
- A faulty compact flash disk that is used as a destination for a system’s DHCP persistency files no longer results in a High-Availability switchover or in a standby CFM reset depending on where the faulty compact flash disk is located. [73766]
- Sending traffic streams at, or close to, line rate with a mix of jumbo Ethernet frames and small Ethernet frames to 1GE interfaces no longer result in a small packet loss due to MAC transmit errors (counted as “Int MAC Tx errors” in the port statistics) or, on rare occasions, in a lock-up of the port on 7710 chassis. [75076, 75240, 77601, 78873]
- The system will now correctly poll an SNTP server that was not available during the first 10 minutes after initial boot-up within the standard poll period. [77407]
- The “first record” and “last record” values displayed in the header of any log file were one number less than the actual first and last records in the event log. This issue has been resolved. [80041]
- The system is designed to keep a log file open past its rollover period if there are no new events to log past the rollover period. If an event that needed to be logged occurred after the rollover period had passed, then that event was incorrectly logged as the last record in the open log file, after which it was closed; a new file was created to collect all subsequent logs. This issue has been resolved. [80041]

TDM

- Some constant BERT patterns received on a DS3 or DS1 port no longer results in a sync when the port is disconnected. [74318]

MANAGEMENT

- An SNMP GET is now able to obtain the description of a LAG from the standard interface MIB after a CFM switchover. [70750]
- The SYSTEM #2009 tmnxStateChange event is now the default event for LDP interface status changes. The LDP #2003 vRtrLdpIfStateChange event that was previously used to log these changes still exists, but is suppressed by default. [76302]
- Loading the TIMETRA-CAPABILITY-7710-V5v0 MIB file into HP OpenView no longer fails due to missing references. [78271]

- SNMP sets of the snmpCommunityStatus MIB variables no longer result in system instability when the createAndWait value is used. [78528]

ROUTING

- When ping packets were originated from a node and sent out over a tunnel (for example, LDP shortcuts or VPRNs), it was possible that those ping packets were dropped after the interface that was used as the source IP address for the ping packets had bounced. This issue has been resolved. [76260]

IS-IS

- Receiving a full-packet purge LSP instead of a header-purge LSP for an LSP that is not available anymore in the system's IS-IS database will no longer cause system instability. [77727]

OSPF

- Values of zero (0) could incorrectly be selected for transit-delay and retransmit-interval via the CLI and/or SNMP. The correct minimum value is one (1) for each of these, and the setting is now correctly enforced. [76075]
- In a setup with redundant Area Border Routers (ABRs) belonging to the same stub area with conflicting summary settings configured, ABRs without any summaries configured remained in the loading state. This issue has been resolved. [76130]
- Virtual links to the ABR via Dot1Q-tagged router interfaces might have, at times, become operationally down when ECMP was enabled in the configuration. This issue has been resolved. [78288]
- The OSPF preference can no longer be set to an invalid value of zero (0) with SNMP. [78416]
- BFD now functions correctly on all OSPF instances. Previously, BFD was only supported on interfaces in the main OSPF instance. [79012]
- In some cases, when using lsa-generate to customize the throttling of OSPF LSA-generation, the system might have incorrectly reset the lsa-wait timer back to the lsa-initial-wait time. This issue has been resolved. [80107]

BGP

- When exporting local 6over4 static routes, the next-hop contained in the BGP update message is now encoded correctly; remote BGP peers can now resolve the next-hop and install such routes. [78592]

MPLS/RSVP

- An ERO object is no longer added to a path refresh message if the original path message did not include one. [76503]
- IS-IS TLV 134 (Router ID) is now used in the CSPF calculation of an avoid-node bypass tunnel. [76689]

LDP

- Ambiguity in RFC 5036 regarding the maximum LDP PDU length might have created some interoperability problems with other vendors' LDP implementation that concatenate multiple LDP messages in a single PDU with size over 4096 bytes. The maximum PDU negotiated by the 7710 SR is 4096 bytes and was encoded in the LDP header as 4092. Receiving any LDP packet with a PDU length greater than 4092 in the LDP header caused

the 7710 SR to send a notification and tear down the LDP session to the peer sending the oversized PDU. This issue has been resolved. [79610]

SERVICES GENERAL

- The Ethernet SAP of an Ipipe did not respond to a gratuitous ARP request from the remote CE. This was required if the remote CE was backed up by a standby CE and a protocol such as VRRP or HSRP was running between the active and standby CE. This issue has been resolved. [77164]
- Router originated ARP-reply packets from IES or VPRN services sent over spoke-SDPs were not padded to 64 bytes at the service PDU level. This issue has been resolved. [77813]
- Broadcast DHCP OFFER and DHCP ACK messages from a server to a client behind the same VPLS SAP with only snooping enabled would broadcast those DHCP messages in the VPLS but also out of the SAP on which the messages were initially received. The workaround was to enable lease-populate to ensure that the DHCP reply messages were dropped immediately. This issue has been resolved. [79396]
- Modifying the SDP keepalive hold-down-time after it had been configured to a value of zero (0) and then administratively disabling SDP keepalives no longer causes a High Availability switchover. [80213]

DHCP

- In DHCP snooping on a VPLS, the DHCP lease was terminated if a DHCP INFORM from the client was acknowledged by the DHCP ACK from the server with the “Your IP Address” field in the DHCP ACK set to a non-zero value. This was an illegal DHCP ACK according to RFC 2131 paragraph 4.3.5 but is now allowed and the DHCP lease will no longer be terminated. [79128]

CFLOWD

- Cflowd packets were always sent with sample mode and sample rate set to zero (0) regardless of the configuration. A workaround was to manually configure the sample rate in the collector. This issue has been resolved. [77814]
- Cflowd packets are now sent to the collector every minute even if there are less than 30 records. [78270]
- The correct srcAS and dstAS are now populated in the cflowd PDU. [78926]

CRON

- Using the “tools perform cron action stop” command to abort cron actions that were in the executing state and that launch CLI commands that take a long time to execute, such as displaying a large route table or a ping with a large number of packets, might not have terminated those actions until they are completed. This issue has been resolved. [79746]

OAM

- An IPv6 traceroute from a Cisco router over a 7750 SR router no longer fails. [75818]
- When the OAM cpe-ping command was used or SHCV was enabled on a VPLS, the chassis MAC address was populated in the VPLS FDB as type OAM. In rare circumstances, non-OAM packets that ingress the VPLS SAP with the chassis MAC address as the destination might have caused system instability. To avoid attracting non-OAM packets to the CFM, the source IP address in the cpe-ping command or in the SHCV configuration should be completely unrelated to the IP addresses used in the network. This issue has been resolved, but traffic loss may still occur if cpe-ping is executed with a source IP address that is known to the network. [76573]

RESOLVED IN 5.0.R22

Following are specific technical issues that have been resolved in Release 5.0.R22 of 7710 SR OS since Release 5.0.R21.

SYSTEM

- Attempts to simultaneously copy multiple files with SCP (SSH copy) no longer result in a CFM High-Availability switchover when the user had no access privilege to the files because they were located outside of the user's home directory. [73717]
- If an SSH client that uses the "Diffie-Hellman group 1" or "Diffie-Hellman group 14" key exchange algorithms had set up a connection, no subsequent SSH connections were possible even after the termination of the first connection. This issue has been resolved. [74040]
- Systems with a large number of log files stored in the Compact Flash devices might have taken longer than usual to execute the configuration file. This was caused by the system deleting old log files with expired retention periods. This issue has been resolved. [76067]

MANAGEMENT

- The system no longer becomes unstable when creating a tech-support snapshot with the "admin tech-support *url*" command where *url* is an FTP URL on an FTP server that was very slow to respond. [71372]
- The system now correctly returns a value when an SNMP GET operation uses the specific OID for entries in the "sapPortIdIngQosSchedStats" and "sapPortIdEgrQosSchedStats" tables (found in TIMETRA-SERV-MIB). [75189]

OSPF

- A setup with redundant Area Border Routers (ABRs) in the same stub area and with one ABR losing its connection to area zero (0) no longer causes the default LSA to max-age in the stub area. This issue was actually resolved in 5.0.R17. [69189]
- When multiple networks are summarized using the "area-range" command, the system no longer sets an incorrect cost for the summary LSA. [74363]

RESOLVED IN 5.0.R21

Following are specific technical issues that have been resolved in Release 5.0.R21 of 7710 SR OS since Release 5.0.R20.

HW/PLATFORM

- Ingress traffic on a GigE port will no longer cause the MDA to fail Self-Test. This issue was actually resolved in 5.0.R19. [69792]

CLI

- Router "vpls-management" is now stored correctly in the configuration under the SAA context. [74529]

SYSTEM

- When the configured time zone is different from UTC, daylight savings time (DST) was not taken into account for timestamps of saved BOF and configuration files. This issue has been resolved. [72319]
- Using a Java SSH client to connect to a system no longer results in a High Availability CPM switchover if the client is designed to execute CLI commands shortly after logging into the system in a non-interactive mode. [73713]

MANAGEMENT

- If an entry was added to a CPM filter with a lower number than existing entries, all entries that were numbered higher than the newly added entry would no longer be evaluated. The workaround was to re-execute the complete CPM filter configuration. This issue has been resolved. [68808]

OSPF

- Modifying the “area-range” command no longer incorrectly allows LSAs to be originated if multiple area-range statements exist for the same IP address with different prefix lengths. [74306]

MPLS/RSVP

- If the primary path of an MPLS LSP that was protected by FRR facility backup was switched to a bypass tunnel and the primary path could not be global-revertive-resignaled, the LSP might have gone down after a few minutes. This only occurred if the shortest IP path from the merge point to the point of local repair (PLR) arrived at a non-MPLS-enabled interface on the PLR. This issue has been resolved. [73353]
- A 7710 SR that is a transit LSR for a number of MPLS LSPs will now only forward an RESV-ERROR message for the LSPs that match the filter-spec of that RESV-ERROR message. [73817]

PIM

- When running PIM in a VPRN that has both Bootstrap Router (BSR) and Candidate Rendezvous Point (CRP) configured for a local interface address in the VPRN, RP addresses are now correctly learned. [74418]

RESOLVED IN 5.0.R20

Following are specific technical issues that have been resolved in Release 5.0.R20 of 7710 SR OS since Release 5.0.R19.

SYSTEM

- SSH login to a 7710 SR node may have failed with an error message "RSA_public_decrypt failed: error:0407006A:rsa routines:RSA_padding_check_PKCS1_type_1:block type is not 01" after several days when there was a lot of SSH activity on the node. This issue has been resolved. [59424]

MANAGEMENT

- The snmpInASNParseErrs counter now correctly increments for an invalid length field in SNMP requests. [70842]

OSPF

- External OSPF routes were sometimes incorrectly refreshed in certain networks after an OSPF summary LSA was received for an unrelated route. This issue was actually resolved in 5.0.R17. [68087, 72430]
- When a router was configured for OSPF graceful restart helper mode, an OSPF adjacency formed over a BFD-enabled OSPF interface may have unexpectedly bounced while an adjacent router was performing a graceful restart. This issue has been resolved. [71712]
- A system that has both a unicast routing table and a multicast RPF routing table no longer leaks routes from both tables into its IGP protocols. The system will only leak routes from the unicast routing table unless specified otherwise with either the unicast-import-disable or multicast-import command. [72945]

BGP

- Enabling BGP peer-tracking on dual CFM systems no longer results in the standby CFM to reset several days after enabling BGP peer-tracking. [70363]

LDP

- When the operational state of an LDP interface changed, the system may have only generated a “vRtrLdplfRowStatus” and not a “vRtrLdplfStateChange” trap. This issue has been resolved. [57624]
- Some LDP adjacencies may have gone down after a certain type of corrupted UDP packet was received on an LDP socket. This issue has been resolved. [72599]

SERVICES GENERAL

- Modifying the configuration of an SDP with a script in a scaled setup no longer results in some of those SDP bindings being disabled. [71890]

VPLS

- Enabling BPDU-translation STP on a VPLS SAP no longer causes a forwarded PVST BPDU to strip off the PVID TLV. [71028]

RESOLVED IN 5.0.R19

There are no new resolved issues in Release 5.0.R19 of 7710 SR OS since Release 5.0.R18.

RESOLVED IN 5.0.R18

Following are specific technical issues that have been resolved in Release 5.0.R18 of 7710 SR OS since Release 5.0.R17.

CLI

- Performing an SNMP GET on an interface description now returns the correct description of that interface if the same CLI command was used to create the interface and its description (for example, “configure router interface <ifName> description <ifDesc>”). This was not an issue if the interface was created first, followed by the description. [65622]

SYSTEM

- Clearing, reseating or replacing an MDA will now clear any XPL error event counts that were recorded prior to the MDA being cleared, reseated or replaced. [69906]
- When using the “exit-on-reject” option for authentication, the authentication method following RADIUS in the authentication-order will now be consulted if the RADIUS server is not reachable. [70081]

MPLS/RSVP

- If an interface address was modified prior to the interface becoming operationally “up” or being added to MPLS, it was possible that RSVP would use an incorrect address when the interface was added to MPLS. This issue has been resolved. [66530]

OSPF

- Self-originated OSPF LSAs will no longer age-out after multiple High-Availability CFM switchovers. [69010]

- Next-hop addresses are no longer incorrectly resolved for Type 7 NSSA LSAs when there are equal cost parallel links between the stub router and the ABR with ECMP enabled. [70298]
- Removing and adding an OSPF area-range summarization may have resulted in the summary prefix not being installed in the routing table and longer prefixes still to be advertised. This issue has been resolved. [70745]
- An OSPF LSA that contains a network with a multicast address of 240.0.0.0/8 or higher is now processed properly; the LSA is accepted into the OSPF database and the multicast route is rejected for the route table. [71002]

BGP

- BGP import policies that modified a BGP attribute of a prefix received from multiple BGP peers to different values may have resulted in a CFM High-Availability switchover. This issue has been resolved. [69875]
- For local imported routes into a BGP confederation AS, the original AS number may have been added multiple times when advertised to an eBGP peer. This only occurred when the BGP confederation AS was the same as that configured under group or neighbor. This issue has been resolved. [70315]

LDP

- An LDP withdrawal message that contains a MAC flush TLV with certain parameters received from a Cisco router no longer results in a standby CFM reset. [70186]

PIM

- If the reverse-path for an (S,G) and (*,G) diverge before merging back onto a router, that router no longer incorrectly encodes the Prune(S,G,rpt) along with the Join(*,G) and Join(S,G) and no longer sends the composed message to the RP. [69785]

RESOLVED IN 5.0.R17

There are no new resolved issues in Release 5.0.R17 of 7710 SR OS since Release 5.0.R16.

RESOLVED IN 5.0.R16

Following are specific technical issues that have been resolved in Release 5.0.R16 of 7710 SR OS since Release 5.0.R15.

HW\PLATFORM

- When 100FX SFPs (3HE00266AA) with version ICS02 were hot-inserted into an MDA, these SFPs would not appear in the CLI but the data path would be active and the SFPs would function correctly. The workaround was to reset the MDA with the SFPs inserted. This issue has been resolved. [69211]
- The system no longer generates a critical “PHY failed to init” error on a port with a defective SFP when that port has been administratively shutdown. [69308]

LAG

- When the first LAG port is removed from a LAG, the other LAG port entries are no longer incorrectly removed from the tLagMemberTable as viewed using SNMP. [68171]

IS-IS

- The output of the CLI command “show router isis database *lsp-id* detail” now correctly displays the admin tag information after receiving an LSP with subTLV1 (part of TLV 135) from a Cisco router. [65778]

VPLS

- The “show service id *svc-id* fdb detail expiry” command output no longer displays the value “429496*” in the expiry field when it has reached the value zero “0”. [69309]

OAM

- The Ethernet OAM (802.3ah) Loopback PDU was being transmitted incorrectly as a 68-byte PDU. This may have potentially caused interoperability issues and has now been corrected to transmit a 64-byte packet. [69753]

RESOLVED IN 5.0.R15

Following are specific technical issues that have been resolved in Release 5.0.R15 of 7710 SR OS since Release 5.0.R14.

HW/PLATFORM

- System stability is no longer affected by a tech-support file request on a system with a defective SFP present in the m20-1gb-sfp MDA (3HE00708AA). [67781]

CLI

- The “show router interface” command output now correctly truncates the “Port/SapId” field for any interface when the string length of this field is more than 14 characters. [66751]
- When configuring a trap target with a name that has a period (“.”) in it, an extra line of configuration no longer appears when configuring another trap target under a lower snmp-trap-group id. [68501]
- The CLI no longer incorrectly allows an entry in the CPM filter with “match src-ip 0.0.0.0/0” which does not match anything and makes no sense. An error message is now displayed for this issue. [68809]

LAG

- An SNMP trap notification will now be generated when the link state of a standby LAG member port is restored. [64444]

ROUTING

- Configuring an IP address on an interface where that IP address is already configured on an existing interface now results in an error message and no longer results in the original IP address on that interface to become unavailable. [59610]
- IP traffic with options received on POS interfaces is now correctly forwarded. [65457]
- The aggregate route is now removed from the routing table when all of the more specific BGP component routes become invalid. [68060]

IS-IS

- Under certain circumstances, the IS-IS link-up convergence time was higher than expected in a network due to a delay in sending out an ARP request over the link after IS-IS adjacency had been restored on that link. This issue has been resolved. [68137]

OSPF

- In an OSPFv3-enabled network with multiple OSPF areas that are interconnected with multiple ABRs, and with ASBRs that inject external IPv6 routes into OSPFv3, a link failure that triggers a next-hop change of these external routes could in some cases result in a High Availability CFM switchover. This issue has been resolved. [67921]

MPLS/RSVP

- An MPLS LSP that requires CSPF no longer fails to establish if a static route exists that matches the far end address of the LSP. The failure code was “noCSPFRouteOwner”. [65302]
- Removing an interface that has active RSVP neighbors no longer results in a High Availability CFM switchover. [65903]
- IS-IS TLV 134 (Router ID) is now used to compare the LSP end address in a CSPF calculation. [67834]

VPLS

- Valid BOOTP packets are no longer incorrectly logged as “suspicious DHCP packets” when they are received on a VPLS SAP with DHCP snooping enabled. [68555]

RESOLVED IN 5.0.R14

Following are specific technical issues that have been resolved in Release 5.0.R14 of 7710 SR OS since Release 5.0.R13.

CLI

- The match function now works properly after it is used against the “file type” command when the text is wrapped on some lines within the file. The workaround was to log out and log back in to restore the match function. [67656]
- The “Software boot (rom) version” field under the card detail and in the output of the “boot-messages” no longer incorrectly prefixes the letter ‘L’ to the boot ROM version instead of the letter ‘X’. [67749]

SYSTEM

- The system time MIB object stiDateAndTime is the UTC time and no longer includes the time zone offset in SNMP get and set requests. This issue was erroneously reported. [66553]
- Synchronization of configuration files and the boot environment between CFMs may have failed on a redundant node if the compact flash boot order was changed from the factory setting. This was indicated by the following message in the event log: “Boot order on CFM set to unsupported value [213] at factory”. An invalid compact flash boot order may have been caused by a terminal server which sends back the characters it receives from the console port. The workaround was to always disable echo on the terminal server. This issue has been resolved. [67130]
- An HTTP client connected to a web-redirect enabled interface that rejected the web-redirect “MOVED” response and rapidly retried the HTTP request no longer results in the system to become busy. This did not cause any service impact but some measures have been implemented to prevent the system to become busy in any case. [68119]

OSPF

- The command “tools dump router ospf route-table inter-area” no longer results in a High Availability CFM switchover. [68088]

MPLS/RSVP

- If a node on a broadcast network does not include the remote IP address sub-TLV in its OSPF-TE LSA and an exclude constraint exists for that advertising router, the MPLS CSPF algorithm no longer includes the excluded node in the computed path. [67262]
- Receiving PATH and RESV messages with an ERO and RRO object that contain a total count of more than 196 sub-objects no longer results in a High Availability switchover. [68009]

VPRN/2547

- Convergence times between dual-homed CE routes are no longer slowed when a large amount of sequential VPRN routes are advertised with a label per prefix of length 26 or longer. [67655]

RESOLVED IN 5.0.R13

Following are specific technical issues that have been resolved in Release 5.0.R13 of 7710 SR OS since Release 5.0.R12.

SYSTEM

- Accounting with record type “combined-network-ing-egr-octets” will now log the correct value for the egress in-profile dropped packets counter. [64713]

IS-IS

- The system now correctly discards an LSP with zero checksum and a non-zero lifetime. [67005]

OSPF

- Virtual links to the ABR will now become operationally up when multiple ECMP paths exit through the transit area. [67025]
- The 7710 SR will now interoperate properly with Juniper routers for OSPFv3 Graceful Restart. [67418]

PIM

- When PIM is enabled on an IES or VPRN interface, one or more ingress multicast queues may be created on the SAP of this interface depending on the SAP ingress policy. After a High Availability switchover, these multicast queues will incorrectly no longer be displayed in the “show service <id> sap” output. This issue has been resolved. [62715]
- Multicast traffic will no longer be impacted by a CFM High Availability switchover on nodes that have the multicast source directly connected if the source address of the multicast stream is in a different subnet as the incoming interface address and if the incoming interface has multicast-sender-always enabled. [67411]

QoS

- No error message is given at the CLI when an exclusive SAP ingress policy is applied to multiple SAPs. This issue has been resolved. [64304]

IES

- An IES instance no longer stops sending ARP requests out of a specific interface if a script is used to swap the primary and secondary IP addresses of that interface without allowing sufficient time for the deletion operation to be processed first. [65675]

VPRN/2547

- When subscriber management host routes are leaked between two local VRFs on a PE router, the leaked routes will no longer be installed in the other VRF routing table with an invalid next hop interface. [65472]

VRRP/SRRP

- In a VRRP setup where the router is in “non-owner” mode and where “standbyforwarding” is enabled, a ping to the VRRP backup IP now works from the backup router(s) and from any remote interfaces in the same VPRN after a High Availability switchover is executed on the backup router(s). [67422]

RESOLVED IN 5.0.R12

Following are specific technical issues that have been resolved in Release 5.0.R12 of 7710 SR OS since Release 5.0.R11.

SYSTEM

- Network instability would sometimes cause route cache corruption after the network was stabilized. This issue was actually resolved in 5.0.R5. [58296]
- CLI through SSH sessions on systems with large numbers of TCP connections (BGP peer sessions and/or LDP peer sessions) will no longer experience a delay in the system’s response time to commands. [63217]
- The active CFM will no longer reset when MIB table vRtrlgmpGrpSrcTable is read via SNMP with an index (virtual router id) of 4096, which is an invalid index. [65819]
- An SNMP GET on the RMON MIB object “etherStatsTable.etherStatsDataSource” now correctly returns the ifIndex instead of the ifEntry of the source that the etherStatsTable entry is configured to analyze. [66449]
- In rare cases, a system that is very busy before and during a High Availability switchover may not properly synchronize the new standby CFM with the active CFM. This issue has been resolved. [66834]

ATM

- ATM OAM loopback cells that do not contain “101” in the PTI field are now processed properly and are no longer dropped. This issue was actually resolved in 5.0.R10. [63039]

MLPPP

- Removal of one or more links from a multi-link bundle while passing traffic no longer causes DRAM errors to be reported in the event log and does not cause the system to stop forwarding traffic. [65837]

ROUTING

- HTTP packets received on interfaces with web-redirect enabled are now processed properly if the HTTP does not have a “\r\n” pattern at the end of the HTTP field. [65783]

DHCP RELAY

- A DHCP relayed packet to a DHCP server in a VPRN will now use the ingress interface as the source IP address instead of the interface with the lowest IP address. This issue was actually resolved in 5.0.R11. [64561]

LDP

- If one of the downstream ABR nodes is turned off in an LDP-over-RSVP network in which IS-IS is used as IGP, the upstream ABR nodes did not always switch their LDP-over-RSVP path to the redundant downstream ABR node. This only occurred with P-nodes in between the upstream and downstream ABR nodes so that the next-hop of the IS-IS route on the upstream ABRs towards the downstream ABRs remained unchanged. This issue has been resolved. [66542]

QOS

- Adding a QoS policy to two or more services with “scope exclusive” will now work properly. [66713]

SERVICES GENERAL

- GRE-encapsulated SDPs no longer stop forwarding traffic if the route for the far-end address of the GRE SDP is indirect and the route flaps. [65732]

IGMP SNOOPING

- Static IGMP groups configured under Multicast CAC policy bundles are not being taken into account at bootup when the Multicast CAC policy groups are the same as those configured under IGMP snooping. The bandwidth for these groups is not populated and remains at zero. This issue has been resolved. [63711]

RESOLVED IN 5.0.R11

Following are specific technical issues that have been resolved in Release 5.0.R11 of 7710 SR OS since Release 5.0.R10.

HW/PLATFORM

- Multicast traffic was impacted for 200 to 400 milliseconds when a soft reset “admin reboot standby” of the standby CFM was executed. Multicast traffic impact for soft reset of CFMs has now been reduced to a few milliseconds. [63255]

RADIUS

- The system now includes the correct NAS-IP-Address in RADIUS authentication requests. [65334]

TACACS+

- The system incorrectly tested connectivity to the TACACS+ server every 15 seconds even if health-check is disabled. This issue has now been resolved. [64394]

CLI

- The “show cflowd interface <ifName>”, “show service id <vprnId> interface detail” and “show router <vprnId> interface detail” commands now show the cflowd information for the specified VPRN interfaces configured for cflowd collection and analysis. [61317]
- Improper termination of CLI sessions in the process of displaying a large system fib table are now handled properly. [65246]

SYSTEM

- A management access filter entry with “src-port cpm” no longer incorrectly matches packets received from inband ports. [64559]
- When the system was unable to read the MDA type from the EEPROM of a newly inserted MDA, it showed “Unsupported” instead of “Unknown” in the MDA Equipment type field. It also incorrectly produced a log event indicating that the MDA is “Unsupported”. This issue has now been resolved. [64966]
- A loop in the “vRtrLdpAddrFecMapTable” object in TIMETRA-LDP_MIB has now been resolved. [65409]
- The TOS field of ourgoing SSH packets would sometimes change from a value of 0x88 (DSCP AF41) to 0x10 (DSCP CP4) during an originating SSH session. The TOS field now correctly stays at value 0x88. [65880]
- The SDP metric value used by the system will now match the configured SDP metric value after a CFM switchover. [65964]

LAG

- Pseudowire status bits are now set correctly when a LAG is removed from an MC-LAG configuration. [65381]

ROUTING

- When using LAG links with configured port speeds that are different from the default port speeds, the bandwidth of the LAG may change after a CFM switchover. As a result, OSPF and IS-IS metric values derived based on the reference-bandwidth may change. This issue has now been resolved. [64861]

IS-IS

- When a point-to-point IS-IS interface is configured with default level-capability (for example, level-1/2) and different costs for Level 1 and Level 2, the lowest cost is advertised regardless of the level even if global IS-IS level-capability is set to “level-1 only”. This issue has now been resolved. [64704]

OSPF

- OSPF router IDs configured within the OSPF level of a VPRN are not activated upon configuration as OSPF must be shutdown then restarted. This will also occur upon a restart of the system as the VPRN level router ID will be used again instead of the one configured at the OSPF level. This issue has now been resolved. [61100]

BGP

- If there are multiple BGP route reflectors for BGP routes, the system will install routes in the FIB with multiple next-hops: one for each route reflector. This issue has now been resolved. [63969]

MPLS/RSVP

- In a network with duplicated IP addresses (for example, different interfaces with the same IP address), signalling of MPLS FRR bypass tunnels now work correctly. [64752]
- For some protected LSPs in link-bypass scenario cases, the output of the “show router mpls bypass-tunnel protected-lsp detail” command displays the system address of the avoid node in the Avoid Node/Hop field instead of the IP address of the interface. This issue has now been resolved. [65056]

SERVICES GENERAL

- Adding LAG ports on a LAG that belongs to a routed CO group-interface no longer affects outgoing traffic on that LAG. The workaround was to always shut down the LAG prior to modifying ports within the LAG. [64618]

VPLS

- Policy-based forwarding of multicast traffic in a VPLS now works properly when combined with IGMP snooping. [64597]

VPRN/2547

- In a system with two CFMs, the standby CFM may reset when multiple VRFs on the same PE router import each other's routes in a way that results in control plane routing loops and rapid route table updates on all VRFs involved. This issue has now been resolved. [60814]
- In a VPRN instance with multiple spoke-SDP interfaces, a ping (without specifying a source IP address) to an IP address that is not directly connected to the VPRN no longer causes ICMP packets to be sent out with the source IP address of the last interface that became operationally up. [61304]
- Modifying a vrf-import policy in policy-options now works properly and no longer temporarily removes prefixes in other VPRNs. [65879]
- A CFM High Availability switchover could occur on VPRN PE routers that have ECMP enabled on one or more VPRNs when a large amount of BGP-VPN route updates are received at the same time. The workaround was to disable ECMP (ECMP value 1) on all VPRNs. This issue has now been resolved. [66384]

RESOLVED IN 5.0.R10

Following are specific technical issues that have been resolved in Release 5.0.R10 of 7710 SR OS since Release 5.0.R9.

CLI

- Closing an SSH client window in the authentication stage before receiving a response back from the system no longer causes the system to deny future SSH-based login attempts. Systems that had the SSH server disabled were not affected. [63787]
- For security reasons, the "info" command is no longer allowed at the root CLI level. The "admin display-config" command can be used instead. [63840]
- The system no longer displays an asterisk '*' indicating an unsaved configuration change after reboot of a node configured with BGP-VPN groups. [64149]
- A CLI session can now be terminated by closing the client window or by using the "admin disconnect <user>" command in the middle of a "show" or "info" command output with the "match" option specified without resulting in system instability. [64392]

SYSTEM

- The system will now raise a "power lost" alarm instead of a "power removed" alarm against the affected input feed when detecting an input power source failure. [61162]
- If the default action of a CPM filter is "reject", this action did not work properly after a system reboot. This issue has been resolved. [62416]

MANAGEMENT

- Configuring a Multi-Service Site via SAM could result in a reset of the standby CFM. This issue has now been resolved. [64142]

OSPF

- Configuring an authentication type with no authentication key on an OSPF interface no longer results in system instability. [64131]

BGP

- If ORF is enabled on a BGP-VPN route reflector and a BGP import policy is configured on this route reflector, routes will now always be advertised correctly to the ORF enabled route reflector clients. [63932]

MPLS/RSVP

- The 7710 SR will now properly process a received RRO object containing a node-id address in the RRO IPv4 sub-object as defined in RFC 4561. For instance, a 7710 SR PLR will be able to associate a primary LSP path with an FRR bypass or to signal a detour backup path by inspecting the RRO object containing either the node-id address sub-object, the interface-id address sub-object, or both. When refreshing a Resv message upstream, the router will always append an interface-id address sub-object to the RRO object received from the downstream node. [63945]
- RSVP packets are now transmitted in the correct interval when RSVP message pacing is enabled. [64260]
- A non-redundant CFM reset could result in a protected LSP using the detour path when the primary path is up. This issue has now been resolved. The workaround was to bounce the egress port on the primary path. [64431]

LDP

- Performing a “clear router ldp instance” on a heavily scaled, busy system no longer causes a delay in the LDP protocol returning to a stable state. This issue was resolved in 5.0.R7. [57973]
- LDP-based SDPs no longer take one to three seconds to go operationally down after the related interfaces go operationally down when the link LDP transport address was configured as an interface address instead of a system address. [58443]

PIM

- If (P-S, P-G) was the multicast distribution tree (MDT) of a VPRN tunnel, then at the root of the tunnel (the ingress PE), “show router pim group <P-G> source <P-S> detail” did not show the statistics of traffic carried over the tunnel. This issue has been resolved. [63176]
- Network instability in a network with hundreds of PIM multicast groups no longer results in a race condition between PIM joins and IGP route updates which would leave some of the PIM groups unresolved. [64399]

OAM

- If a bypass LSP were signaled to a 7710 SR with an implicit or explicit null label (where the sending router supported PHP), then an LSP ping issued over that bypass tunnel could fail to reach its destination. This issue has been resolved. [61532]
- In a multi-vendor network, the 7710 SR now replies correctly to ATM OAM periodic loopback cells. [62930]

RESOLVED IN 5.0.R9

Following are specific technical issues that have been resolved in Release 5.0.R9 of 7710 SR OS since Release 5.0.R8.

RADIUS

- In a system configured for RADIUS authentication, a user logging in and logging out of multiple sessions now works properly. [61619]

SYSTEM

- The system will no longer incorrectly generate and clear false fan failure alarms when the power source is momentarily disrupted. This issue was actually resolved in 5.0.R1. [55129]
- Inaccuracies with the “complete-service-ingress-egress” service accounting policy have now been corrected. [62236]
- A High Availability switchover will no longer result in traffic loss if the following configuration is present in the BOF: an active IP address, no standby IP address, and any number of DNS addresses. [62751]

CLI

- The “admin disconnect” command now works properly for FTP users when an IPv6 address is given. [60349]
- The “show card detail” command now correctly shows the complete serial number of the compact flash device. [62095]
- When using the “show router route-table” command with a route prefix and the “longer” keyword for multicast IPv4 routes, matching entries in the IPv4 unicast routing table are no longer incorrectly displayed. [62411]
- Subscriber management authentication and accounting policies that contain spaces in the name are now saved correctly. [63214]
- Unnumbered interfaces that contain spaces in the name are now saved correctly. [63659]

RIP

- RIP debugging now works properly when the RIP protocol is disabled and then re-enabled. [60211]

IS-IS

- The “show router isis lsp detail” command will now properly display the full content of an LSP that includes Sub-TLVs 250 and 251 which are part of Cisco-specific extensions in support of MPLS-TE. Huawei routers also support these extensions. [62464]
- A system configured for LSP authentication now preserves the authentication TLV received in an LSP when an LSP is aged out, and the authentication TLV will be included in the flooded LSP. [63572]

OSPF

- The “show router ospf *ospf-instance* opaque-database” CLI command will now correctly display entries from the backbone (area ID 0) and other areas for the specified OSPF instance. Previously, the user would have to specify the area ID in the command in order to see entries for areas other than the backbone area. [60468]
- OSPF neighbors now remain stable after a High Availability switchover with scaled numbers of services and when OSPF timers are set to minimum values (hello-interval 1, dead-interval 4). [61101]

MPLS/RSVP

- The endpoint address of the bypass tunnel is now encoded in the ERO when the primary path is sent over the bypass tunnel as required by RFC4090. [60918]

- Path messages with no “Controlled-Load Service” object present in the Adspec object are now decoded properly. [63030]
- The 7710 SR will now skip RRO sub-objects received in RSVP Reservation messages from a Redback router configured to use RFC4561. [63429]

PIM

- If all of the last-hop routers in the C-instance are configured not to switch from the shared tree to the shortest path tree, the ingress PE will not stop and restart forwarding traffic on the "improved-assert" enabled MDT. [62564]
- Changing the hello interval (“hello-interval”) value from its default value of 30 seconds for a PIM interface will be committed and applied to the first hello sent when the interface comes up. [62717]

SERVICES GENERAL

- Overriding the advertised VC-type MTU on an SDP (“adv-mtu-override”) is now committed as soon as it is configured. [62484]

VPLS

- Authentication-policy strings are now always quoted in the configuration file. Policy names containing a space or other special characters can now be saved and restored. [61158]
- In a VPLS with mesh SDPs, traffic forwarding is no longer affected when the STP state is quickly toggled from enabled to disabled. [63352]

VPRN/2547

- For a specific VPRN that is configured with the maximum number of routes and a default threshold value of 0, the system no longer incorrectly raises a threshold crossing alarm if the number of routes exceeds 50% of the maximum number of routes configured. A threshold value of 0 disables threshold crossing alarms. [60912]
- VPRN-related IP packets originating from the CFM sent over an implicit null MPLS tunnel are now correctly encoded. Previously, incorrect encoding could occur in the case where 7710 SR routers were interoperating with third-party equipment because 7710 SR routers never require implicit null labels. For example, an ICMP ping within a VPRN could fail if implicit null MPLS tunnels were used. [61785]
- A High Availability switchover no longer results in routes being temporarily withdrawn from the VPRN when the VPRN is configured with SDPs that have signaling enabled. [62792]

OAM

- OAM sdp-ping requests sent over an LDP-over-RSVP tunnel that has moved to a facility-bypass LSP will now correctly transit across the network. [61385]
- VCCV ping requests sent over an Epipe service with an SDP binding of vc-type “vlan” now work properly. [62760]

RESOLVED IN 5.0.R8

Following are specific technical issues that have been resolved in Release 5.0.R8 of 7710 SR OS since Release 5.0.R7.

SYSTEM

- Under conditions where the active CFM's memory buffers are heavily utilized, CFM processing of protocol packets is no longer impacted. [61977]

DHCP RELAY

- When the system receives DHCP discover messages with a null source MAC address and DHCP snooping is enabled, the system will no longer incorrectly learn and install the null MAC address into the forwarding database. [61538]

IS-IS

- In IS-IS graceful restart interoperating with Huawei routers, the system did not explicitly clear the SA bit in the second IS-IS hello packet that was sent out after a reboot of the node. This issue has been resolved. [62167]

PIM

- PIM adjacencies will now be formed on interfaces that have multicast-senders configured to "always". [61866]
- A few milliseconds of multicast packet drop could occur every 2 minutes in certain types of VPRN topologies. This issue has been resolved. [62006]

QOS

- Operational WRED slope values are recalculated when the slope policy is applied to a port and a new network queue policy changes the shared pool memory size. [60919]

RESOLVED IN 5.0.R7

Following are specific technical issues that have been resolved in Release 5.0.R7 of 7710 SR OS since Release 5.0.R6.

HW/PLATFORM

- If an upgrade was performed from 3.0.R13 to R5.0 prior to 5.0.R7, communication to the standby CFM would be lost. This issue has been resolved. [61496]

RADIUS

- When a RADIUS server is used for authentication, the number of "rejected logins" under authentication statistics is no longer incorrectly incremented after health-check attempts. [61270]

ROUTING

- ICMP checksum failures are now properly detected. [61279]

OSPF-TE

- A node that runs multiple OSPF instances can now terminate TE LSPs in all OSPF instances. [61361]

BGP

- IPv6 BGP peering is now established when interoperating with routers from other vendors when MD5 authentication is configured. [61056]

MPLS/RSVP

- Path messages received with an ERO with non-local first address are now rejected. [54879]
- Sending of RSVP hello packets is no longer delayed on MD5-enabled LSP tunnels where there is a high frequency of flapping. [61079]

LDP

- System instability could occur if an LDP neighbor was performing a restart (for example, the 7710 SR was performing the helper function for Graceful Restart) on a targeted session while an SDP was being configured to use that session or a reconciliation was being completed after a High Availability switchover. This issue has been resolved. [60890]
- LDP graceful restart helper can now be disabled. [61776]

PIM

- When using Anycast RP and the multicast sources are directly connected and dual-homed to two or more RP nodes, the PIM group state is now correctly removed after the multicast sources stop transmitting data. [61014]

VPLS

- When switching from vc-type ether to vc-type vpls, the VLAN VC tag is no longer incorrectly overwritten with a value of zero. [60107]

VPRN/2547

- The system's CPU utilization is no longer abnormally high when there are two VPRNs with BGP configured that share the same interface with multiple non-BGP routes that are shared between the two services. [59618]
- Leaked VPN-IPv4 routes with different route distinguishers (RDs) and/or route attributes are now properly compared when ECMP is enabled in the VPRN instance. [60480]
- Modifying a vrf-import policy in the policy-options now works properly and no longer temporarily removes prefixes in other VPRNs. [61111, 60996]
- Modifying a vrf-import policy on a redundant route reflector now works properly and no longer removes prefixes from other VPRNs. [61179]
- Quickly toggling the administrative state of a BGP VPRN peering with a large number of routes now works properly. [61273]

OAM

- MAC trace requests sent on the control plane are now being processed correctly and will no longer indicate erroneous failures. [58815]

RESOLVED IN 5.0.R6

Following are specific technical issues that have been resolved in Release 5.0.R6 of 7710 SR OS since Release 5.0.R5.

HW/PLATFORM

- When an unsupported MDA is installed, the equipped field now reads "Unsupported" instead of being blank as if no device were installed. [55535]

CLI

- The "admin debug-save" command now saves debug commands for services. [40694]
- The "show card detail" command now displays the string "Software boot (rom) version" instead of "Software boot version" before the bootROM version. [59425]
- Performing an SSH server shutdown on a node when there are a significant number of SSH login attempts now works properly. [56227]
- The commands "admin disconnect ssh" and "admin disconnect user" no longer incorrectly disconnect SAM-initiated SSH sessions. [60471]

- The CLI output after executing a ping within a VPRN instance that uses hub and spoke VPN interfaces no longer incorrectly reports negative round-trip delay times. [60550]
- Redirection of CLI output reports to a file now works properly for telnet and SSH console sessions in addition to the console serial port. [60556]
- Attempts to log into the node using SSH with a null login name and null password followed by issuing “show users” command are now properly processed by the router. [60661]

MLPPP

- MLPPP packets dropped due to fragmentation and reassembly errors are now correctly counted as Errored Packets. [60580]

MANAGEMENT

- When combined-network-ing-egr-octets was configured for accounting, the iod (in-profile octets dropped) counter was incorrectly being populated with the octets forwarded data. This issue has been resolved. [60634]
- SNMP conformance issues with the RMON alarmTable, eventTable and logTable have been resolved in this release. [60639, 60640]

BGP

- BGP sessions can now be established if the neighbor advertises BGP capability code 0x42 (BGP MDT SAFI) in the BGP open message. [60646]

MPLS/RSVP

- If the primary path and the ingress detour of a one-to-one backup both fail at the same time on a CSPF-signaled LSP, the LSP will now correctly be resignaled with a secondary path. [50583]
- A label re-use error condition is now correctly processed by the standby CFM. [58709]
- The system will now correctly ignore path messages for which the sender's IP address starts with “127.” [60968]

LDP

- MD5 key changes no longer require the LDP peering session to be cleared. The resolution to this does mean that MD5 key changes on active LDP peering must now be coordinated on each end of the peering session to occur before the expiration of the session keep-alive timer. [58403]
- After changing the metric for some LSPs, LDP bindings are now correctly updated to select the best available LSP based on the metric changes. [59949]

IP MULTICAST

- If a Multicast CAC policy were deleted while active on an IGMP interface, the interface still had references to the deleted policy. If a new policy were added to that interface, the data for bandwidth was incorrectly inherited from the previously deleted policy. This issue has been resolved. [60154]
- Multicast CAC policies are now supported on a LAG group. [60798]

SERVICES GENERAL

- L2TP, STP, IGMP snooping and OAM packets are now processed properly when received over a double encoded MPLS packet, such as, LDP-over-RSVP tunnels and one-to-one backup tunnels.

  Another manifestation of this now resolved issue is that the Spanning Tree Protocol did not function correctly on an SDP if that SDP was based on an RSVP LSP and the last part of that LSP was switched to its FRR bypass tunnel. This only occurred for FRR facility protection and not for FRR one-to-one protection. [60422]

VPLS

- The Spanning Tree Protocol now works properly on a VPLS with mesh SDPs with VC type “vlan”. [60505]

OAM

- The router no longer displays “!” for echo replies received as a result of unsuccessful rapid pings to an unreachable host. This issue was actually resolved in 5.0.R4. [59455]

RESOLVED IN 5.0.R5

Following are specific technical issues that have been resolved in Release 5.0.R5 of 7710 SR OS since Release 5.0.R4.

RADIUS

- RADIUS server authentication statistics now increment properly. [58899]

SYSTEM

- Logging out of the console of a chassis system running with a single CFM no longer generates the following benign trap: “The active CFM card {A | B} is operating in singleton mode. There is no standby CFM card.” [59271]
- When the time zone was set to something other than UTC, timestamps of saved BOF and configuration files were not correctly updated. This issue has been resolved. [60182]

MANAGEMENT

- The terminal length now defaults to 24 lines for an SSHv2 session. Previously, the terminal length default for SSHv2 was 100 lines. [48101]

OSPF

- With a scaled number of services and OSPF timers set to smallest possible values (hello-interval 1, dead-interval 4), OSPF neighbors are now stable after a High Availability switchover. [60179]
- In an OSPF multi-area configuration, CSPF will now choose the best route from all available OSPF areas. [60199]
- The summary LSA is now refreshed on an OSPF ABR that is in more than two areas where summarization is used for a prefix that is in multiple areas. [60360]

MLPS/RSVP

- The RSVP TTL on a one-to-one detour at the ingress node would incorrectly be sent with a value of zero if the frr-hop-limit were configured with a value of 255. This issue has been resolved. [60137]

LDP

- For VPLS and Epipe interoperability with the Juniper ERX, the 7710 SR no longer sends a fatal notification message when an LDP label release message is received with an invalid interface MTU sub-TLV included. [60489]

PIM

- During a BSR election, incoming bootstrap messages with a hash-mask length of 0 are now processed properly. [60325]

QoS

- Aggregate Rate Limits for Port-based Egress Scheduling Policies are now supported on ANCP nodes. [60160]

SERVICES GENERAL

- An SDP may no longer be incorrectly associated with an LSP that has no far-end IP address configured. [60312]

SUBSCRIBER MANAGEMENT

- If a DHCP discover message is received on a service with DHCP snooping enabled and the DHCP discover message includes option 50, the 7710 SR DHCP proxy server will no longer check if the option 50 IP address is already in use, before forwarding the DHCP discover message to the remote DHCP server. This only applies if the 7710 SR DHCP proxy server authenticates the subscriber via a RADIUS server. This change was made to be compatible with certain DHCP clients that include option 50 in the DHCP discover message. [60439]

VPLS

- In a user VPLS with spoke SDPs that are managed by an STP-enabled VPLS where the user spoke SDP VC IDs are lower than the spoke SDPs in the management VPLS, the user spoke SDPs will correctly continue to forward traffic after a High Availability CFM switchover. [60368]

VPRN/2547

- In rare instances, a High Availability switchover on a 7710 SR provisioned with a large number of VPRN instances and BGP peers could cause the PE-CE BGP peering to flap. This issue has been resolved. [59242]
- Specifying a valid source IP address when using the traceroute command within a specific VPRN instance now works properly and no longer fails with the following error message displayed “MINOR: OAM #0000 - The address is invalid or does not match the address type”. [59909]

VRRP/SRRP

- SRRP no longer transitions state after a “clear card” command is executed. [59805]
- In a VRRP setup where the router is in “non-owner” mode and where “standby-forwarding” is enabled, a ping to the VRRP backup IP now works from the backup router(s) and from any remote interfaces in the same VPRN. [60369]

OAM

- Starting LSP pings in one telnet session and then stopping the telnet session from another console session with the “admin disconnect” command could result in a High Availability switchover. This issue has been resolved. [59275]
- lsp-ping, lsp-trace and vccv-ping now support the “Copy PAD TLV to reply” option defined in Section 3.4 of RFC4379. [60389]

RESOLVED IN 5.0.R4

Following are specific technical issues that have been resolved in Release 5.0.R4 of 7710 SR OS since Release 5.0.R3.

HW/PLATFORM

- Depending on where the BOOT.LDR file physically resided on the Compact Flash, the boot ROM on the CFM would fail to locate the BOOT.LDR file on CF3 and the router would fail to boot. This issue has been resolved. [59910]

CLI

- The output of the “show router tunnel-table summary” command now displays the summarized number of RSVP tunnels. [58833]
- The “show port a/1” and “show port b/1” commands that are used to display the status of the Ethernet management ports will now display the correct duplex mode for the specified management port. [59664]
- An asterisk (“*”) that is displayed after modifying the static route entry in the BOF, is now cleared after saving the BOF. [59983]
- An asterisk (“*”), indicating an unsaved configuration change, is no longer incorrectly displayed at the prompt after the execution of a rapid ping command. [60139]

SYSTEM

- DNS resolution now works properly when no system IP address is configured in the BOF. [59777]

LAG

- Removing an active LAG member port from a LAG and re-assigning it to the same LAG now works properly. [59740]

MANAGEMENT

- Various SNMP conformance issues with get-next processing of large index values on IP address indices have been resolved. [45227]

OSPF

- OSPF hello packets were occasionally not being sent at their intended intervals, especially under scaled conditions or when using small hello and dead intervals. This issue has been resolved. [59174]
- When an OSPF area is reconfigured to be a stub area, the Area Border Router (ABR) now correctly withdraws the associated ASBR Summary LSAs (Type 4) for the ASBRs that were in the area from the OSPF database. [59391]
- Changing the system time to a value in the past no longer affects OSPF adjacencies formed over interfaces that are MD5-authenticated. [59431]
- If the OSPF reference bandwidth was set to a value lower than 100,000, and a given metric was configured on an OSPF interface and later removed, then the OSPF metric of that interface would incorrectly become the value 0. This issue has been resolved. [59725]
- Translation of NSSA external LSA’s to external LSA’s on an OSPF area border router would in rare instances not work properly when there were area-ranges configured. This issue has been resolved. [60194]

BGP

- A BGP route reflector would incorrectly advertise BGP-VPN routes to its route reflector clients that it received with the same cluster ID as its own cluster ID. This issue was actually resolved in 5.0.R3. [59514]
- BGP next-hops are no longer improperly resolved through aggregate routes. [59920]

MPLS/RSVP

- If a Bypass to the Merge Point is already setup, the hop limit is now enforced for subsequent LSP associations. [57913]
- Setting the value of the vRtrRsvpIfHelloInterval MIB attribute to values that are not multiples of 1000 is now properly restricted. [59133]
- Under scaled LSP conditions, CPU utilization has been improved in this release. [59655]

LDP

- High Availability switchovers now correctly maintain the correct LSP choice for the LDP-over-RSVP tunnelling. In some cases, the tunnelling would fall back to link-level LDP, and in other cases, an operationally down LSP was incorrectly chosen. These issues have been resolved. [60313]

PIM

- In a network with multicast traffic originating from two different PIM sources whereby the network address of one source is a subnet of the network address of the other source, PIM would incorrectly not send a prune message to the source address with the higher netmask when the source with the lower netmask became known in IGP. This would result in duplicate multicast traffic for a short period of time. This issue has been resolved. [59837]

QoS

- The H-QoS scheduler rate limit accuracy for low rates around 1 Mbps has been improved to better than 5%. [60027]

FILTERS/PBR/TCS

- MPLS-encapsulated IP packets are no longer incorrectly matched against the default action of an egress IP filter. [59980]

SUBSCRIBER MANAGEMENT

- Change of Authorization statistics are now supported. [57147]

VPLS

- There is no longer any DCHP packet loss during a High Availability CFM switchover on a 7710 SR that has a VPLS instance configured with DHCP snooping enabled on thousands of SAPs. [58001]

VRRP

- Sub-second VRRP timers are now supported on SDP interfaces. [58839]

RESOLVED IN 5.0.R3

Following are specific technical issues that have been resolved in Release 5.0.R3 of 7710 SR OS since Release 3.0.R1.

HW/PLATFORM

- An m5-1gb-sfp, m60-10/100eth-tx or m20-100eth-sfp MDA in an MCM could exhibit some lost or truncated packets with particular traffic patterns in the packet payload at high traffic rates. This issue has been resolved. [bz2227]
- If a DS1 channel with a PPP encapsulation is set to accept a remote loopback request and the remote side sends a loopback request, the PPP session is now properly renegotiated when the remote loopback mode is exited. [41225]
- Issuing the “no loopback” command twice in succession on a DS-1 port will no longer place the port into a “diag” state. [42138]
- The CLI command “file dir <string>*” for which “<string>*” matches a range of file names was incorrectly reporting “File Not Found” if only directories matched. This issue has been resolved. [42291]
- CPM queues defined with non-default CBS or MBS values were not reliably created properly and hence would let flow traffic sent to them without shaping/policing. The rate steps for these queues start at 1000 Kbps, so PIR or CIR set to less than this value will be defaulted to 0. This issue has been resolved. [43696]
- Changing from chassis mode A to B could cause some MDAs to lose their port configuration. This issue has been resolved. [43813]
- Issuing repeated “admin save” commands will no longer result in system instability. [44004]
- A low-level firmware issue could cause packet loss on the CFM Flexible Fast Path in very rare instances. This issue is now fixed, and code has been added to automatically detect and correct similar issues. [44528, 45214, 46009]
- The syntax for accessing or managing the directory structures on the Compact Flash card is not case sensitive. In addition when specifying the path to the config file or the image files in the BOF, it is no longer necessary to properly specify the case used for the target directory on the Compact Flash for the config/boot-env synchronization process between the two CFMs to succeed. [44807]
- When connected to a 7710 SR node using FTP, an “ls” command will return the files in the current directory and will no longer incorrectly add a “File Not Found” message at the end of the file list. [45652]
- Performing a transfer of large files to/from the system could result in difficulty managing the device or in system instability. This issue has been resolved. [46358]
- Transferring files using tftp to a destination reachable in-band now works properly. [46586]
- BERT configuration data is no longer reset to defaults when the port state goes operationally down. [46643]
- The system no longer reports error messages during the insertion of some Alcatel Gigabit Ethernet SFPs. This issue only affected one Alcatel SFP supplier’s parts. [47291]
- During the five (5) seconds after a Compact Flash has been improperly removed (without first administratively shutting the Compact Flash device down), the MIB table will be empty for that device and then will correctly report a non-present Compact Flash. This issue has been resolved. [47870, 48666]

- An initialization timing issue has been resolved where some SFPs sourced by Sumitomo/Excelight with vendor part number 6Fxx would not be recognized upon insertion. [50498]
- An initialization issue has been resolved where 3HE00867AA GigE EX SFPs sourced by OCP would not enable the transmit laser when installed in some GigE MDAs: 5-port (3HE01615AA) and 20-port (3HE00708AA). [50498, 51448]
- Traffic corruption could be experienced when the MDA configuration of slots 1/1 or 1/3 were changed. This issue has been resolved. [51592]
- Rapidly configuring and de-configuring MDAs could occasionally cause the active CFM to reset. This issue has been resolved. [56772]
- If the DHCP persistent index files could not be saved to the Compact Flash due to insufficient space on the Compact Flash, the Compact Flash no longer becomes unreadable. [52644]
- Log file rollovers on CF2: were sometimes incorrectly removing the last files added rather than the oldest files. This issue has been resolved. [54529]
- Executing sync-if-timing force-reference to the second qualified reference now works properly. [55739]
- After a “clear mda” on the 20-port GigE SFP MDA (), some SFPs with part number 3HE00027AA manufactured by OCP would not come up once the MDA was out of reset. This issue has been resolved. [55956]
- A High Availability switchover while an MDA was booting could result in the MDA not coming online. This issue has been resolved. [58439]

BOF

- If the BOF.CFG file is not present or corrupted on the Compact Flash, the console port will be the last configured speed value when the system comes up (after manual user intervention to enter minimum boot parameters). Previously, the speed would default to 115200 when the BOF.CFG were not present or corrupted. [51264]

RADIUS

- If a RADIUS server is configured and used for authentication and authorization, under specific conditions (two active CFM switchovers with users logging in and an 'admin save') the config file could have become un-executable due to some unexpected data written to it. This issue has been resolved. [41166]
- RADIUS authentication requests are now generated if the system interface is operationally down. [49713]

TACACS+

- If a TACACS+ server is used for authentication and authorization, unexpected stack trace messages were displayed on user sessions after logging in or executing CLI commands. This would occur when the TACACS+ server is not working properly and closing sessions too early. This issue has been resolved. [42216]

CLI

- If an APS group is being monitored via the CLI while an APS switchover occurs, the monitor rate will no longer incorrectly show a value greater than 100% for one monitoring interval. [41089]

- DHCP information is now correctly not displayed in the output of the “show router interface detail” command for network interfaces. Support for DHCP on network interfaces has been removed in this release. [42214]
- The CLI no longer allows one to associate the default filter log (101) with a syslog destination. [42269]
- The range of ingress service labels for a mirror destination's remote-source far-end address is now correctly reported in the help text as 2048..18431. [42272]
- The online help for the OAM command “cpe-ping” used to incorrectly display a “send-count” option where the correct parameter should have been “count”. The online help for that OAM command has been modified accordingly. [42337]
- The copy command at the config>system>security level to copy profiles no longer incorrectly allows destination profile names with spaces. [42684]
- The graceful-restart command is no longer available as a command under the BGP context of a VPRN. Graceful-restart helper commands appeared in previous versions of the software even though it was not a supported feature. The command has been removed until that feature is supported. [43013]
- The “admin tech-support” command now correctly appears in change logs. [43628]
- Issuance of the “tools perform router mpls cspf” command was not always causing the command to execute. This issue has been resolved. [44098]
- The CLI command “show router bgp damping” now properly displays damped routes. [44417]
- Using a wild card in the source file of the “copy” command when trying to copy a file from one Compact Flash device to another Compact Flash device on the standby CFM now works properly. [44677]
- An IES interface with a VRRP association now requires the VRRP association be removed before the IP address can be deleted. [44962]
- The help text for the commands in the config>router>ospf>timers context now correctly indicates the values entered are in milliseconds. [45470]
- The SSH version is now displayed in the “show users” output. [49099]
- Minor CLI error messages are no longer generated if the “show service id <id> igmp-snooping all” command is performed on a service configured with residential split horizon groups. [49194]
- If the command “no login-banner” was entered under configure>system>login-control and later the configuration was saved, the node would stop executing the config file at that command when the system was rebooted. This issue has been resolved. [54048]
- Traceroute messages that are received after 100 ms are no longer incorrectly displayed as an asterisk in the traceroute output. [52011]
- Executing the command “show router bgp” while a CFM synchronization is in progress now works correctly. [52981]
- An erroneous, benign error message about duplicate socket closes no longer appears on the console during an FTP copy operation. [54344]
- A configuration file will now load properly if a “%” is used in the text string of a description field. [54355]
- Using the “ssh” or “scp” command on a 7710 SR system to access a FreeBSD server now works properly. [54977]

- The “admin disconnect telnet” command no longer incorrectly disconnects SSH users. [55467]
- The SSH server can now be shutdown at system initialization in the configuration file. [56235]
- If the system was busy and the user issued a rapid ping command, the system could generate the following benign error message “UNUSUAL_ERROR cliIcmpPing: sia_tmnxOamPingHistoryEntryGet failed for HistIdx" in the log. This issue has been resolved. [57630]
- The “show router route-table n.n.n.n/n longer” command now works properly and no longer includes routes matching the base subnet that are shorter than the mask supplied. [58010]

SYSTEM

- The receipt of malformed SSH packets will no longer cause the SSH subsystem to generate critical unusual error log event entries. The receipt of these packets are now properly classified with a lower trace severity. [42667]
- All image and config file locations configured via BOF statements must use absolute paths. Relative paths used to be accepted which could result in a failure of boot synchronization between an active and standby CFM. [44619]
- Accounting policies saved to files on Compact Flash were not properly logged following a High Availability switchover. The workaround was to perform a “shutdown” and “no shutdown” on each accounting policy. This issue has been resolved. [44822]
- Under rare circumstances, SNTP request packets are misdirected and lost generating an SNTP event in the log-id 99. This issue has been resolved. [45917]
- Instances where neither, both or an incorrect CFM is providing a qualified synchronous timing clock will now generate a log error message. [46523]
- The list of configurable standard time zones no longer includes zone names for daylight and summer time zones. [45079]
- LACP and BPDU control packets no longer share access resources to the CFM mitigating protocol interaction. [49449]
- Creating or deleting a user while any user is attempting to login to the node via an SSH session could cause system instability. This issue has been resolved. [52440]
- On a redundant system when the sync mode was set to “Configuration”, the system could report a configuration sync failure in the sync status field. This issue has been resolved. [52686]
- Upon a High Availability switchover, card, MDA, and port insertion messages are no longer erroneously displayed in the event logs. [53172]
- The 1-port GigE CMA and 20-port GigE MDA interfaces with SFPs installed no longer experience an up to a one second outage on CFM activity switches. [54182]
- CMA initialization failures are now consistently logged by the system. [54491]
- Performing an SSH server shutdown on an SSH version 1 capable node when there is a significant number of SSH login attempts now works properly. [56227]
- When a custom time zone offset included a minutes offset in addition to an integral number of hours, the minutes portion of the offset was being added/subtracted in the wrong direction. This issue has been resolved. [56717]

- Summer time periods that cross year boundaries, as used in the Southern Hemisphere, now work properly. [57891]
- The system now gracefully handles situations where event logging and statistics collection are configured to use local files on a Compact Flash device that has reached full capacity. [58260]

ETHERNET

- Ethernet ports configured for half duplex operation no longer report a higher than expected number of excess collisions when operating near link capacity. [bz2079, bz2217]
- The MCM preclassifier for the 5-port GigE MDA, 20-port 100FX MDA and 60-port 10/100 MDA now correctly classifies system interface and port interface traffic. [bz2219]

DS1/E1

- When a remote line loopback is engaged, the LED now flashes yellow to indicate that the port is in a maintenance state. [bz1262]
- When receiving AIS, the 7710 SR will no longer use the port as a valid synchronization source. [bz1613, bz1620]
- On DS1/E1 CMAs, a frequency shift of -10 to -15 ppm is no longer observed when switching from holdover back to reacquiring a valid signal. [bz1725]
- The DS1/E1 statistics are displayed correctly after a shutdown at the channel group / DS1 level. [bz1933]
- On DS1 and E1 ports in LOS state, RAI is now transmitted in this state. [bz2139]
- DS1 ports now pass jitter tolerance test (per ITU-T G.824 / G.813) when configured as source of sync for the system. [bz2187]
- Continuous large packets on a low bandwidth channel group no longer experience potential packet loss. [bz2223]
- Frame relay packets larger than 4736 bytes are no longer discarded. [bz2261]
- The MTU for DS1/E1 and DS3/E3 ports now have a configurable range of 512 to 9208 octets. [bz2262]
- Both the far-end and near-end would loopback when fdl-ansi loopback is initiated. This is issue now resolved. [bz2271]
- CLI is now functioning properly for Cisco HDLC keepalive and up-count commands. [bz2277]
- Full line-rate on individual channelized interfaces is now supported on all ports at the same time. [49298]
- AIS is now transmitted at DS1/E1 port shutdown. [50993, bz388]
- Large packets on one link no longer affects small packet streams on another link. [51216]
- When an internal loopback is present, the link LED will now flash amber. [51380]
- The DS1/E1 CMA now remains in-service following CFM activity switches. [51458]
- A CFM activity switch no longer results in reference clock source qualification issue. [51461]
- For FDL initiated loopbacks, the status LED now flash amber to indicate maintenance activity and the loopback persists across CFM activity switchovers. [52062, 52092]
- Remote loopback is now deactivated when switching directly to line loopback. [52099]

- If a channel group was removed from an MLPPP bundle and then a CFM switchover occurred, the PPP session for the channel group could go down. This issue has been resolved. [52193]
- Shutdown of a channel group now properly clears octet statistics. [52194]
- The payload-ansi loopback now functions properly to pass BERT tests. [52307]
- When a channel group is configured for network mode, the ability to configure accounting policies and statistics collection is now available. [52515]
- Link-level protocols, such as PPP and LMI running on a channel group no longer need be re-established immediately after a CFM switchover. [54406]
- When a channel group is deleted on a configured, but not present DS1/E1 CMA, subsequent channel group configuration on another CMA may cause the active CFM to restart. This issue has been resolved [54869]
- On DS1 ports, LOS is now properly cleared after 10 seconds. [55773]
- When the DS1/E1 port is the system timing source, some wander tests would fail. This issue has been resolved. [56975]
- Following a CFM activity switch, SAP egress statistics would temporarily stop incrementing. Traffic was not affected but statistics would not be properly counted temporarily following the activity switch. This issue has been resolved. [57455]
- If a port is shutdown and no shutdown while an internal port loopback is active, the port transmit now works properly without having to reset the CMA. [58776]
- For MLPPP ports configured for network mode, accounting policies and statistics collection now work properly. [59331]
- Requesting remote ANSI loopbacks via FDL now works properly [59493]

DS3/E3

- MDL strings are now properly detected upon a CMA reboot or clear operation. [54032]
- BERT tests using ones/alternating pattern now work on the first test. [54053]
- AIS or RDI alarms are now correctly cleared following a CFM activity switch when the condition was present prior to the activity switch. [54050]
- DS3 MDL string information is now correct during loopback conditions. [54175]
- Line Errored Seconds and Line Code Violations will no longer increment while the line has an internal loopback. [54225]
- When an internal loopback is set in a DS3 CMA port, the AIS signal will now be sent with the correct P-bit values. [54229]
- DS3 M23 framing now works properly. [54240, 54326]
- The DS3/E3 port octet statistics no longer temporarily go negative under certain scenarios. [54258]
- After nodal reboot, the MDL string is now properly transmitted and received. [54421]
- DS3/E3 CMAs would not boot reliably in slots 1-4 of a 7710 SR-c4 chassis. This issue has been resolved. [54471]
- BERT tests on DS3 ports now work properly when the framing is set to M23. [55955]

SONET/SDH

- The 8-port OC-3c/STM-1c SONET/SDH MDA no longer experience packet corruption with specific sequences of large and small packets. [bz2265, bz2274]

- During AIS-L conditions, the M1 error counters are no longer incorrectly incremented (which resulted in the reporting of REI-L alarm). [30474]
- When a channelized SONET/SDH port is in a Path-AIS or Path-RDI state, the received path trace (J1) string displayed is no longer incorrectly showing the last good value. [51025]
- The 8-port OC-3c/STM-1c SONET/SDH MDA no longer experiences packet loss with a sequence of packets of very small packet size at a high traffic rate. [51863]
- The 8-port OC-3c/STM-1c SONET/SDH MDA no longer experiences packet corruption with specific sequences of large and small packets. [51968, 51972]
- The 7710 SR OS CLI help will now displays the help line “string ‘zeros’ will send all zeros in the J1 bytes”. [52504]

FRAME RELAY

- The service-mtu on an Fpipe service no longer needs to be set much higher than the SAP's MTU to become operationally up. The service MTU only needs to be 2 bytes larger than the MTU of the port containing the SAP. [55820]
- Small size Frame Relay frames could be dropped when traversing an Fpipe if ingressing on sub-DS3 SAPs. Frames down to 1 byte will now be passed. [56529]

LAG

- The qinq encapsulation for LAGs configured as SAPs is now supported in this release. [42189]
- A Link Aggregation Group (LAG) with Link Aggregation Control Protocol (LACP) enabled can now correctly be looped back to a different LACP-enabled LAG group on the same system. Previously, the LAG would not come up when configured this way. [47246]
- Administratively disabling a LAG will now bring down all of the port members. [52513]
- The application list for the “show debug” command now includes “lag”. [55450]

MLPPP

- An MLPPP LFI bundle no longer must be shutdown before removing/adding a member. [bz2186]
- LFI traffic is now scheduled with the correct priority. [bz2186]
- An MLPPP bundle now forwards packets with a size of up to 9208 bytes when the MTU is configured for 9208 bytes. [bz2256]
- When an MLPPP bundle is in a high congestion state, the excess packets will be effectively dropped by the appropriate Network Egress Queue. [53103]
- Setting the Fragment Threshold for an MLPPP multilink bundle using SNMP will return an error if the value is outside the supported values of 128 to 512. [54423]
- The encap-type of an MLPPP bundle can no longer be incorrectly set to Cisco HDLC which is not supported on MLPPP bundles. [55713]
- The encap-type on the primary member of an MLPPP bundle can no longer be changed when the bundle is assigned to a service. [55800]
- If a far-end MLPPP group were changed from long-sequence IDs to short-sequence IDs, a console trace error message was generated, and traffic would not successfully pass. This issue has been resolved. [57253]

APS

- The CLI now allows the proper range of up to 16 for APS groups. [bz2268, bz2269]
- If an APS group is associated with a port-down VRRP policy statement, the VRRP policy statement now must be deleted prior to deleting the APS group. [41772]
- An APS switchover concurrent with a High Availability switchover in a scaled APS environment will no longer cause an increase in the switchover time (to maximum of 1.5 seconds). This issue has been resolved. [42375]
- In a fully loaded system, clearing, rebooting or quickly performing a shutdown/no shutdown on multiple MDAs simultaneously where APS groups were configured would occasionally cause one APS path to be stuck in the “link up” state. This issue has been resolved. [48048]

MANAGEMENT

- SNMP trap-destinations used to require that either the system IP address or a management IP address be configured or SNMP traps would not be sent. This is no longer an issue as source addresses for SNMP applications can now be configured explicitly. [29802]
- The MIB object “ifOperStatus” for DS3 and DS1 and the CLI status are now in agreement. [46453]
- Passwords of exactly 8 and 16 characters saved in configuration files were not always being compared correctly when read back in following a reset of the system. This issue has been resolved. [40350]
- Values set for ifAlias MIB objects were not persistent or saved to the configuration file. Therefore, the ifAlias MIB has been changed to a read-only MIB. [40375]
- The variable bindings of APS-related trap events are now correct. [42091]
- If the 5620 SAM system was used to set the dot1p configuration on a MAC match criteria for an access ingress policy, the set of the mask parameter to “Default” was correctly setting the mask value to -1, but this value was incorrectly being saved via the “admin save” command resulting in an unloadable configuration file. This issue has been resolved. [44334]
- SNMPv1 trap destinations reachable via in-band routes no longer result in trace error messages displayed on the console shortly after rebooting the system. [44555]
- When a summer time zone starts, the “show time” command now correctly displays the summer time zone (instead of the non-summer time zone, for example, PST instead of PDT). [45070]
- The Western European Summer Time zone (WEST) is now correctly displayed in the list of pre-configured summer time zones. The system previously had WET as the acronym for that time zone. [45072]
- The three European Summer Time zones are now all correctly being triggered at 1 AM UTC rather than at midnight UTC when entering and exiting the summer time periods. [45074]
- A configuration file saved with SNMPv3 users (privacy keys enabled) now correctly restores the user authentication information and no longer displays error messages upon reloading of the configuration file. [46112]
- The SNMP MIB attribute snmpEngineBoots and snmpEngineTime are now correctly synchronized/reconciled between active and standby. After a High Availability failover, snmpEngineBoots will be incremented and snmpEngineTime reset to 0. Power-cycle or system reboot resets snmpEngineBoots=1. [46179]

- If a system threshold setting had been configured for either the cflash-cap-alarm or the memory-use-alarm and only the rising-threshold had been set, and afterwards a CFM switchover has occurred and the configuration has been re-saved, the resulting configuration file would be unloadable due to the addition of “falling-threshold 0” to the threshold configuration. This issue has been resolved. [46361]
- If the system interface is shutdown or not configured, SNMP traps and syslog events are now correctly being sent. [49251]
- In-band syslog messages no longer include the IP address of the management interface (if configured) in the text of the syslog message. [49709]
- If a source-address configuration was applied to a management protocol and the interface/address configured was operationally down, the following benign messages could appear on the console or in the user session: “Configured source IP address is not used.” This issue has been resolved. [50493]
- SSH login failures now correctly generate log events. [50788]
- If one configured the exponential-backoff command under login-control and no other commands were configured under that CLI context, the resulting configuration file would fail to load since the login-control statement was missing. This issue has been resolved. [50945]
- Configuring “src-port cpm” on a management-access-filter entry now correctly matches against out-of-band management traffic and does not match against in-band management traffic. [51436]
- If all RADIUS servers are down, incoming telnet sessions are no longer delayed when connecting. [52278]
- Exceeding the maximum number of allowed SSH sessions is now handled gracefully. [52444]
- It is no longer possible to configure some managed objects via SNMP with SNMP index strings that contain unprintable ASCII characters. [52540]
- SNMPv3 users can no longer be created incorrectly as “read-only”. [52736]
- The “router” match criteria under management-access-filters now works properly for VPRN instances. [53129]
- The enterprise MIB files now correctly define all intervening OID levels prefixing SNMP notifications. For each notification set, there was a level defined with a zero that had no explicit define of its own. [53202]
- The snmpEngineBoots attribute is now incremented properly when no engine-ID has been configured. [54388]
- SNMP get requests with an incomplete object identifier for tNetworkQueueTable and tSharedQueueTable (TIMETRA-QOS-MIB) now correctly return a “No Such Object” error message. [54449]
- Using the port ID and encapsulation values from a “vRtrIfEntry” (ref: TIMETRA-VRTR-MIB) of a spoke-terminated Layer 3 interface to index any parameter in the “sapBaseInfoTable” (ref: TIMETRA-SERV-MIB) no longer causes system instability. [56077]
- Changing the address or port of a syslog entry that is assigned to a filter log no longer affects the logging of messages. [56271]

ROUTING

- Local-proxy-arp is no longer dependent on proxy-arp being configured. [37480]
- If a static-arp entry is configured under one router interface, the static-arp entry is already configured under a different router interface and the static-arp MAC address differs, then the MAC address would incorrectly be changed under the router interface where the static-arp was originally configured. This has been resolved. [42251]
- A static-route where an indirect next-hop is resolved by an LDP tunnel is now properly programmed in the forwarding and routing table for traffic received on network interfaces destined to that prefix. [44027]
- TTL expired ICMP messages were incorrectly sent out using the egress interface IP address as the source address instead of using the ingress interface IP address on which the original packet needing an ICMP reply was received. The correct IP address is now used. [45174]
- When configuring static host, the sub-option “subscriber-ID” now works properly. By configuring the subscriber-ID, traffic between hosts connected to a bridged residential gateway will stay local and no longer have to be looped by the access node. [45390]
- IP multicast packets within a VPLS with checksums equal to 0xFFFF are no longer discarded at ingress interfaces. [47056]
- A system interface that is configured as a remote proxy ARP server was not replying to incoming ARP requests for IP addresses of local router interfaces. The interface now replies to those ARP requests. [47711]
- IP packets destined to an address with a /31 network mask are no longer incorrectly forwarded by the CFM. [48023]
- Modifying a routing policy (or adding an export statement for such a policy) on a router with a large routing table and a large number of policy entries could have resulted in delays in processing the transmission and reception of the hello packets of protocols such as OSPF and IS-IS. In cases of extremely low hello and dead interval times, this could have resulted in adjacency flapping. This issue has been resolved in this release. [50609]

DHCP RELAY

- DHCP offer packets are being dropped incorrectly on IES interfaces configured with /31 subnet masks. This issue was reported in earlier releases, but it was found not to be an erroneously reported issue. [43764]
- DHCP packets are now correctly being filtered by cpm-filters. [43846]
- Clearing DHCP lease states via SNMP is now supported in this release. [43964]
- The DHCP ACK packet sent by the server answering a DHCP Inform packet from a client was improperly dropped at IES interfaces or VPLS SAPs where the “lease-state populate” feature was enabled. These ACK packets are now correctly forwarded by the DHCP relay on these interfaces. [45616, 45617]

RIP

- A RIP route import policy based on matching specific route tags now properly accept or reject routes with those tag values. [56221]

IS-IS

- The IS-IS graceful-restart helping statistics and status are now properly restored after a High Availability switchover. [42257]
- Scaling point-to-point adjacencies in a narrow metric environment no longer cause LSP fragments to contain partial data at the end (incomplete TLVs) which, depending upon the

IS-IS implementation, could cause the LSP update to be dropped. This issue has been resolved. [42630]

- An LSP will now come up even if its terminating node is in the IS-IS overload state. [45428]
- The CLI command “show router isis database detail” could in previous versions lead to system instability if a Juniper router in the IS-IS routing domain sent LSPs containing unicast IPv6 multi-topologies TLVs. This issue has been resolved. [47672]
- Multi-topology (MT) TLVs are now decoded in the output of the “debug router isis packet detail” command. [48108]
- The metric for the system interface configured under IS-IS is now correctly advertised in the IP reachability TLVs. [48795]
- A log event is now generated when an IS-IS adjacency toggles from init to down on a broadcast interface and from up to down on a point-to-point interface. [49704]
- If an IS-IS packet is received on a VPRN interface, it was incorrectly being processed by the IS-IS protocol in the base routing instance. This issue has been resolved. [50591]
- IS-IS hello packets are no longer incorrectly egressing on a level 2 IS-IS point-to-point interface that is administratively shutdown. [54924]
- An IS-IS complete sequence number PDU (CSNP) packet greater than 1492 bytes is now handled properly. [56139]
- A High Availability switchover no longer causes a change in the checksum value of an IS-IS self-generated LSP. [58249, 58388]

OSPF

- If a virtual-link was computed via a transit-area with the exact number of paths that equals the ECMP value, only one ECMP path was installed and used. This issue has been resolved. [41504]
- Overlapping OSPF area-ranges no longer cause the refreshing of a subsumed range matching exactly the aggregated route to be suppressed. This is in addition to 39815 that had addressed the same problem for ranges including (but not matching) aggregated routes [42763]
- If a summary or external LSA was received with a non-zero route portion but a zeroed netmask, the router was being installed this as a default route. This is valid according to the OSPF RFCs, but this value used to interfere with the reception of a true default route of 0.0.0.0/0. This issue has been resolved. [44553]
- NSSA address ranges containing solely NSSA type-1 LSAs now correctly take the highest cost of the NSSA Type-1 LSAs. [46317]
- If 1) the OSPF spf-wait timers for initial-spf-wait and/or second-spf-wait were set to values lower than 1000 ms (one second), 2) the standby CFM reset and 3) a High Availability switchover was later initiated, the newly active CFM incorrectly reported zero for those two timers resulting in a subsequently unloadable configuration file if saved. This issue has been resolved. [47569]
- Deleting an OSPF area using a mixture of CLI and SNMP commands now works properly. [47661]
- If an NSSA router is originating external routes into the NSSA area, the forwarding address is currently set to the router ID, which may not be a reachable IP address within the NSSA area. This issue has been resolved. [47923]

- If an IP address is modified on an OSPF-enabled interface, after a High Availability switchover, the original IP address may incorrectly still be active at the OSPF level. This issue has been resolved. [48992]
- If OSPF PE-CE is used in a dual-homed configuration, it was possible for an operator to create a routing loop in the network by leaking routes from AND to BGP within the network. This issue has been addressed with the enhancement [51272] on page 56. [49283]
- When an Area Border Router (ABR) connected to more than one NSSA area becomes operationally up, the default route is now correctly advertised into the NSSA areas. [50207]
- If an OSPF VPRN instance were created and the VPRN was deleted without first deleting the OSPF configuration, the previous areas could not be configured again under a new OSPF instance. This issue has been resolved. [50446]
- If the IP subnet mask of an OSPF-enabled IP interface was modified to a different value (but the IP address remained the same), the OSPF hello packets incorrectly contained the original subnet mask after a High Availability switchover. This issue has been resolved. [50513, bz2211]
- The removal of component routes for an NSSA route range while the NSSA ABR is running an SPF is now a benign operation. [50608]
- Redistribution of direct routes into OSPF that match existing component routes of an aggregated area range now works properly. [54475]
- OSPF adjacencies will now remain stable when short hello-interval and dead-timers are configured on OSPF interfaces and the limit of routes set in the “external-db-overflow” command is high. [54833]
- Running OSPF over IES interfaces using spoke-sdp tunnels now works properly even if a tunnel’s operational state is flapped at a high frequency. [54835]
- Debugging of the OSPF RTM updates to a disk file or syslog server when there are thousands of OSPF route updates now works without resetting OSPF neighbors. [56039]
- OSPF will now correctly install local IP addresses in the RTM if they come from a peer with a /32 subnet mask. [56371]
- When no metric was explicitly assigned to an IES (or VPRN) spoke interface, the OSPF metric was sometimes set to a value of 1 and sometimes set to a value of 65535. Now, an OSPF-enabled spoke interface receives a default metric as if it were a 10 Mbps interface (10000 if the reference-bandwidth has not been modified). [56831]
- Flapping an OSFP interface at specific intervals could prevent some routes from being added to the route table. This issue has been resolved. [56936]
- Interface down events were handled with some delay in OSPF when the interface was a LAG with more than one active member, and either the LAG was shutdown or its last link member went operationally down. This issue has been resolved. [57736].
- When an OSPF interface detects a mismatch between its operational MTU and the OSPF MTU of the far-end interface, the system now generates an event message to warn the user of this misconfiguration. [57800]
- When an ABR connected to more than one NSSA area becomes operationally up, the default route is now correctly advertised into the NSSAs. [57950]
- An ABR configured as “nssa originate default-route” and “no summaries” may continue to generate a default-route even if all of its interfaces to Area 0 are “Down” as long as the system interface is administratively up and is configured in Area 0. [80044]

OSPF-TE

- If OSPF-TE is enabled after LSPs have already been established, the value of maximum bandwidth in the TE LSAs is now set correctly. [57887]

BGP

- The object bgpPeerState in the SNMP trap bgpBackwardTransition now correctly reflects the current state of the BGP peer. [44100]
- Several successive CFM switchovers between active and standby could occasionally cause a PE-CE BGP session to flap. This issue has been resolved. [44561]
- When the advertise-inactive feature is enabled, a non-BGP route (static, local or IGP) that is redistributed into BGP and not advertised to peers because it is also received from a BGP peer will not be advertised if the BGP route is withdrawn. This issue has been resolved. [47457]
- The output of the “show router bgp routes” command no longer incorrectly displays routes from the local RIB. [48731]
- Regular expression-based policies matching on communities now work properly for received BGP routes that contain more than 64 communities. [50219]
- BGP will now correctly install /32 routes for the router’s own interface addresses if learned from a BGP peer. [50955]
- Upon removal of a router reflector cluster-id configuration, the previous cluster-id is no longer incorrectly used in iBGP route selection. [53263]
- BGP routes received with no MED value are no longer selected incorrectly over equivalent routes containing a MED value. [54031]
- Route refresh messages for unrecognized address families were incorrectly being processed. This issue has been resolved. [54365]
- Using eBGP as the PE-CE protocol for a scaled number of VPRN instances, an eBGP session could bounce on rare occasions when the standby CFM came online. This issue was actually resolved in the 5.0.R1 release. [58896]
- If BGP is configured with "loop-detect discard-route" and a prefix is learned and installed in the routing table, a subsequent BGP update for the same prefix with an AS-loop will be discarded, but the original route will incorrectly remain. [78520]

MPLS/RSVP

- An unknown Class Number will now generate the correct error code. [12746]
- When two LSPs with different MTUs are signaled in the same tunnel and the LSP with the active MTU is subsequently torn down, the merge point LSR now correctly sends the updated MTU value to the upstream node. [16158]
- If an LSP is configured with a primary path and several secondary paths, none of which would setup end-to-end, all of the secondary paths are now signaled. [17356]
- RSVP Resv errors returned by the 7710 SR no longer contain erroneous data in the error specification for the node, flags, error code or error value. [30562]
- For LSPs configured with fast reroute (FRR), the “show router mpls lsp path detail” CLI command now contains complete information in the “Actual Hops” part after the LSP moves to a detour path. [30982]
- If the metric cost of the active path of an LSP was increased after the LSP was created in the operationally up state and the new metric cost value was higher than other paths in the network, resignalling would not bring the LSP onto other better cost paths in the network. This issue has been resolved. [31109]

- When OSPF was used as IGP, LSP re-signalling would not be triggered after OSPF-TE (traffic engineering) has been turned off/on. This issue has been resolved. [31198]
- When an LSP is on the ingress detour/bypass tunnel and the LSP is resignaled, the ‘tools dump router rsvp psb’ command would not display the old LSP path before the resignalling was triggered although the resignalling worked properly. This cosmetic issue has been resolved. [31940]
- For LSPs protected by facility bypass fast reroute (FRR) that have switched to the ingress node tunnel detour, a make-before-break (MBB) initiated by a change of reservation bandwidth now works properly. [32233]
- Very rarely, traffic carried by an LSP using one-to-one FRR protection and established on a strict primary path takes two hits (instead of one) when switching to the detour path. This issue has been resolved. [33894]
- If a 7710 SR was the Merge Point (MP) and a Juniper router was the Point of Local Repair (PLR) for an LSP with Fast-Reroute facility bypass method enabled and using the local revertive mode, traffic was lost when the link was restored and the Juniper PLR switches the path of the LSP to the primary. Since the 7710 SR MP no longer supports local revertive mode, this is no longer an issue. [40295]
- Two active CFM switchovers will no longer incorrectly reset the RSVP refresh-timer for LSPs back to the default value (30 sec.). [42320]
- LSPs set up with explicit paths to a router one hop away no longer incorrectly bounce on the addition of a black-hole static-route that encompasses the address range used in the path. [42348]
- When a fast re-route (FRR) LSP is active on a detour at the ingress node, manual resignalling now works properly. [43081]
- For the facility fast-reroute method, the PathTear message for a protected LSP that is currently on the bypass tunnel on the head end node now correctly has the egress interface IP address, which is the same as used in the Path Message sent over the tunnel, as the Sender IP address in the sender Template. [43187]
- With the resignalling timer turned on for a fast re-route LSP where the new primary path and the old detour path share the same link at the Merge Point (MP), the RESV messages of those two paths are now merged properly. [43974]
- A ResvTear which does not contain the FlowSpec object is now correctly accepted as valid. Previously, those messages were incorrectly rejected as invalid, and the LSP would not be torn down or rerouted until the RSVP session expired (up to 3 refresh timers). [45009]
- If the Merge Point of a Bypass protected LSP did not change the Sender IP address in the Sender Template to the Source Address of the LSP, the egress node could stop sending Resv back for this LSP. This issue has been resolved. [45482]
- When a ResvTear is received on the bypass return path, the message is now properly propagated upstream. [45594]
- A label change in the Resv message for a bypass protected LSP now correctly causes the forwarding table to be updated with the new label expected at the merge point during FRR protection. [45639]
- A ResvTear received for a bypass protected LSP at the PLR will no longer incorrectly trigger a switch to the bypass tunnel. [45853]
- For facility FRR LSPs, a bypass path is now searched beyond the previous 7 hops limit. [45913]

- A PathTear received at a Merge Point for the LSP being protected by a bypass tunnel will no longer incorrectly clean up the primary LSP. [46183]
- LSPs were not being activated properly in a non-backbone OSPF area if the backbone area (area 0) had been configured without any interfaces. This issue has been resolved. [45574, 46665]
- If a PathError were received for an unmapped LSP, sometimes existing CSPF data was used instead of the LSP recomputing the path from scratch. This issue has been resolved. [46711]
- An LER that has an LSP with “adspec” and “fast-reroute” configured now correctly uses the MTU size of the negotiated MTU when trying to bring up the detour path. [47196]
- Static LSPs setup on several network interfaces using VLAN (802.1Q) tags on the same port now come up properly when these interfaces are shutdown/no shutdown. [47404]
- Occasionally in specific configurations, an LSP reverting from a bypass tunnel/detour to its primary path could experience a short traffic disruption. This issue has been resolved. [47491]
- IP packets originating from the CFM and encapsulated with a single MPLS label were incorrectly sent out with an MPLS Ethertype on Ethernet interfaces instead of an IP Ethertype when the next hop router requires an implicit null label. Traversing traffic will carry the correct Ethertype. This would occur only with equipment from other vendors as 7710 SR routers never require implicit null labels. This issue has been resolved. [48135]
- A protected LSP (FRR facility or one-to-one method) with one or more standby secondary LSPs that share a common path with the primary LSP will no longer go down momentarily if one of its standby LSPs goes down. [48892]
- When setting the FRR property on an LSP to disabled, the LSP no longer undergoes a make-before-break (MBB) as was the case in earlier releases. [49046]
- A one-to-one FRR LSP tunnel configured to use CSPF is now correctly signaled even if traffic-engineering is disabled. [49152]
- LSPs configured to a far-end loopback address can now be correctly established. [49434]
- PathTear messages for facility-bypass LSPs are now properly sent on the correct egress interface upon a path shutdown. [49523]
- Receiving labels for a FEC from two different LDP neighbors at the same time could result in rare cases in an incorrect next-hop for that FEC. This issue has been resolved. [49622]
- If a primary path configured with bypass tunnels were not initially signaled when the LSP becomes active (it is signaled using a standby path), the bypass tunnels were not being signaled when the primary path subsequently became active. This issue has been resolved with the addition of the global revertive mode functionality as per RFC 4090. [50237]
- If a 7710 SR was the ingress LER for an LSP with Fast-Reroute one-to-one method enabled, the node had re-signalling enabled and the re-signal timer went off while traffic was flowing over the ingress node detour, an incorrect PUSH entry operation for the new computed primary path would cause traffic on that LSP to stop. This issue has been resolved. [50395].
- On a RESV timeout, a facility bypass LSP will now correctly bring down the LSP and resignal rather than switch to the bypass tunnel. [50992]
- RSVP-TE LSPs are no longer incorrectly bounced after 479 days. This issue has been resolved. [51129]

- If an explicit null label is enabled on the egress node of a Fast-Reroute Facility LSP, traffic is no longer affected if switched to the bypass tunnel terminating on that egress node. [51692]
- If the MTU on a network port used for an LSP were reconfigured, the affected LSPs dynamically changed to use the lowest of the MTU values configured, old and new. This issue has been resolved. [51706]
- A Juniper router acting as point of local repair (PLR) node for an LSP that is concurrently doing a local reversion on the Juniper router and a global reversion on the 7710 SR head-end could cause system instability of the 7710 SR. This issue has been resolved. [53261]
- When the link along an LSP path recovers after being operationally down for an extended period of time, the path would not recover unless the path was removed and re-added to the LSP. This issue has been resolved. [54230]
- Traffic is no longer affected for a global revertive make-before-break scenario with multiple failures. [54418]
- Manually resignalling an LSP that is currently attempting timer-based resignalling will no longer cause the LSP to go down. Previously, the LSP could go down in rare instances. [55966]
- If a tail-end node receives a detour or standby PATH message for an existing LSP containing an Adspec object that has a lower MTU than the existing LSP, the node now replies with the correct MTU value (equal to the negotiated value of the Adspec object) in the Flowspec object of the RESV message. [56230]
- Admin group attributes of MPLS interfaces were not properly synchronized to the standby CFM causing the IGP to not advertise these attributes after a CFM switchover. These attributes are now properly synchronized. [57528]
- An LSP with a path containing strict hops where the first hop is learned with multiple next-hops is now setup properly. [57581]
- If the port MTU is increased, the MTU of existing LSPs that use that port are now increased. [58302]
- A flapping link could cause merging one-to-one detours to go operationally down if the IGP SPF timers were large enough (>1 sec.). This issue has been resolved. [58402]

LDP

- If an interface address is deleted on a router interface, adjacent LDP peers that are still connected to that interface properly re-advertise that FEC back to the router that originally advertised it. [38211]
- Label withdraw messages sent with a valid label and with a wildcard address prefix are now accepted and processed. [38279]
- If there is a routing loop in the network and there are static FECs configured on the interfaces involved in the routing loop, the LDP FECs are no longer continually advertised, then withdrawn. [40073]
- If traffic is forwarded over static FECs and a High Availability switchover occurs, traffic is forwarded without interruption on the LDP sessions using those static FECs. [41249]
- The “configuration sequence number” in LDP Hellos is now correctly incremented when a session is restarted only if a local configuration change was made. [47731]
- The “show router ldp binding” command now displays the qtag values properly on tagged interfaces. [48370]

- Upon shutting down an SDP signalled using targeted LDP and if “debug router ldp peer <n.n.n.n> label detail” is enabled, incorrect data will no longer appear in the output of the Label Release message. [49330]
- Under high loads of label withdrawal and re-advertisement, some FECs would not be re-advertised properly. This issue has been resolved. [51092, 51355]
- Upon a series of label advertisements and withdrawals where labels for the router's own interfaces are advertised back to the router, labels will no longer incorrectly be reused. [53050]
- LDP adjacencies will now come up properly on network interfaces where the IP address has been set and modified (configured as IP x then as IP y) before a port was assigned and LDP enabled. [55122]
- Receipt of a label for FEC 0.0.0.0/0 followed by a High Availability switchover where a route for that FEC is subsequently installed in the RTM is now handled properly. [56288]
- Entering the “info” command in CLI and paging through the LDP peers configuration section could cause some instability in that protocol if authentication keys were configured. This issue has been resolved. [56876]
- A FEC will now be correctly programmed if the label withdraw for the FEC is sent out and the label release message coming back from an LDP peer is received after the FEC is re-advertised. [57262]
- LDP static FECs (fec-originate) can now be configured for routes 0.0.0.0 with mask lengths /0 through /7. [57839]
- LDP label removal in an ECMP configuration now works properly. [58160]
- The value of LDP graceful restart state is always “capable”, even when the remote side did not signal that it is capable of performing graceful restart. [79430]

IGP SHORTCUTS

- Static routes configured with an LDP tunnel to be the indirect next-hop (for example, as an IGP shortcut) would incorrectly add two labels onto tunneled packets rather than just a single label. This issue has been resolved. [47961]
- Subscriber interfaces are now properly resolved over IGP shortcuts and appear in the route-table at the far-end of the IGP shortcut. [50762]

IP MULTICAST

- An mtrace request is no longer incorrectly sent to all PE routers in a VPRN. Instead, it is correctly only being sent to the VPRN reachable via the RPF address. [42172]
- IP multicast to IES interfaces bound to spoke-sdp's is now supported. [43207]
- If an IGMPv3 membership report message is received that contains multiple multicast groups including invalid groups in the range of 224.0.0.0/24, the invalid groups are not dropped and the valid groups are not processed. This is applicable for Layer 3 services with IGMP termination, as well as Layer 2 services with IGMP snooping. [78632]

IGMP

- If a second source is configured for an SSM-translated group (*,G) on the same node, IGMP now correctly refreshes all sources instead of just the first source. [48974]
- An SNMP GET received for an (S,G) entry on a non-existent interface referenced by a large index value could result in system instability. This issue has been resolved. [49837]

PIM

- PIM Join and Hello packets that are sent across a GRE-based MDT are now correctly being sent as network-control (CS6) packets. [42136]
- Previously, mismatched MTUs on each end of a VPRN multicast tunnel could have resulted in dropped PIM control signals being sent across the tunnel. To correct this, the maximum MTU for such PIM control signals is limited to 1400 bytes in this release when being sent across such a tunnel. [42563]
- On spoke IES interfaces where PIM was enabled, there would occasionally be some delay in bringing up a PIM adjacency with the far-end node after the tunnel transporting the IES spoke-SDP was brought down and up. This issue has been resolved. [50286]
- When the sticky-dr priority on an interface is modified, the designated router (DR) election is now run immediately rather than waiting for the remote hello. [54223]
- A loop no longer exists in the vRtrPimSptSwitchoverThdTable. [55798]
- Incoming PIM hello packets with an incorrect source IP address 0.0.0.0 on a PIM enabled VPRN will no longer result in system instability. [56083]
- Shutting down IS-IS now works properly for systems running PIM where at least one multicast source address is resolved with two route table entries learned from IS-IS where one uses a generic prefix (i.e. 0.0.0.0/0) and another uses a more specific prefix. [56229]

QoS

- Once the total number of ingress queues has been exhausted on an MDA, it was no longer possible to remove a QoS policy from a SAP on that MDA. This issue has been resolved. [39293]
- Traffic flowing across a service () and rate limited at the ingress SAP by a QoS policy (PIR not equal to max) would be affected (not hitless) during an LSP revertion from standby/detour to primary path (assuming that the service is transported by an MPLS SDP using the aforementioned RSVP LSP). This issue has been resolved. [43596]
- If an ingress or egress sub-port rate had been set on a port that was greater than the operational speed of the port, if the port speed subsequently changed, the correct sub-port rate was not taking effect. Also, the sub-port-rate setting was not being taken into account when calculating pool sizes and queue rates. This issue has been resolved. [56649]

FILTERS/PBR/TCS

- Clearing only the egress counters on IP and MAC filters is now supported. [45265]
- IP filter entries with no action specified are incorrectly being assigned an “action drop” when a “renum” command is issued. Also, the following commands for IP filter entries are not being configured properly when the filter entry is renumbered:
  - filter-sample
  - interface-disable-sample
  - action forward next-hop interface
  - action forward redirect-policy
  - action drop

  This issue has been resolved. [45355, 45529]
- IP filter entries with an action configured for “next-hop interface” cannot incorrectly be assigned to VPRN interface SAPs. [45517]

- An IP filter configured on the egress side of a DHCP relay enabled IES interface was not able to filter DHCP replies (Acks/Offers). These packets will now be processed by the filters. [47488]
- Web redirection is now supported when a Web Portal Redirection filter entry is applied to a VPLS mesh SDP. [48717]
- The action “forward next-hop interface” is no longer allowed for numbered IPv4 interfaces. This filter action only is supported on unnumbered interfaces. [50118]
- An IP frame containing a single IP option no longer is incorrectly matched against an egress IP filter configured with “multi-option”. [51449]

SERVICES GENERAL

- SDPs for all services will remain operationally up even when a local SAP is operationally down. This maintains the VC label mappings to the far end.
- For ATM and Frame Relay pipes, VC label mappings are now sent or withdrawn based on the SAP state ensuring proper OAM notifications are sent to the far end CE devices. [40666]
- IP addresses of IES/VPRN SAPs that are administratively and operationally up were not displayed in the routing table when one of the interfaces of the IES/VPRN was down. The system now correctly displays those addresses in the routing table. [46618]
- Sdp-keepalive used to send packets set to the maximum size of the tunnel (equal to the MTU negotiated by the transport layer). If the transport mechanism configured is using LSPs protected by the fast-reroute facility method, sdp-keepalive would fail when an LSP switched to a bypass tunnel because of MPLS label stacking. Sdp-keepalive packet have now a smaller size (40 bytes) in order to avoid the issue. A workaround in prior releases is to manually decrease the packet size of sdp-keepalive messages. [46850]
- IP packets routed through IES/VPRN spoke-terminated interfaces where the size is larger than the spoke SDP tunnel path MTU are now fragmented correctly. [50294]

SUBSCRIBER MANAGEMENT

- Deleting a configured subscriber host entry can now be performed using only the host IP address. [48590]
- Anti-spoofing will no longer drop received multicast packets, so features such as VRRP and multicast streaming now work properly when received on a subscriber interface with anti-spoofing enabled. [48780]

VPLS

- Ingress SAP VPLS billing statistics are now collected for all traffic types: unicast, broadcast, multicast and destination unknown. [16206]
- When creating more VPLS entries with ‘stp no shutdown’ commands than the maximum number of STP instances allowed, warning messages are generated indicating this limit has been exceeded and the illegal configuration will not be saved to the configuration file. [27756]
- If the VPLS FDB table size has reached the value assigned by the fdb-table-size, MAC addresses were not being relearned until five MAC addresses or more have been removed from the system. This issue has been resolved. [40014]
- Combining Multicast VPLS Registration (MVR) SAPs with MROUTER SAPs now functions properly. [42970]

- When saving the configuration of a SAP which has MVR in the IGMP snooping node, the SAP was saved in an invalid format and would fail to load. This issue has been resolved. [44734]
- Creation of more than one m-vpls instances are now correctly assigned unique Spanning Tree bridge IDs. [45266]
- Creating a new SAP and shutting it down simultaneously using a single SNMP SET command now works properly and does not affect the standby CFM. [50489]
- The restrict-protected-src alarm-only option is now working properly in this release. [52103]
- Large root STP path costs are no longer incorrectly displayed as negative numbers. [51234]
- Cisco PVST PDUs ingressing on a LAG are now processed instead of being discarded. [51440]
- An event is now generated if a SAP associated with an mVPLS is in the downstream loop operational state. [51874]
- FDB aging of MAC addresses are more precisely controlled and will no longer drift as much as 6 seconds for every minute of aging. [52145]
- In a VPLS spanning more than one node, a misconfiguration of RSTP (enabled on one node and not on others) no longer causes system instability on the node running RSTP if there is a loop in the non-RSTP-enabled part of the L2 domain. [52838]

IGMP SNOOPING

- Changing the LAG configuration when IGMP snooping is configured on the LAG now works properly. [48181]

IES

- Packets dropped by an egress filter on a spoke-SDP IES interface are now correctly being counted in the SDP binding statistics. [40528]
- Modifying the path-mtu of an SDP linked as a spoke-sdp in an IES interface will now correctly modify the OSPF operational MTU. [42164]
- Traffic being routed across an IES spoke interface (traffic not destined for the directly attached subnet) is now being properly forwarded. [43242]
- An egress IP filter applied to a spoke SDP bound to an IES interface no longer requires the filter default action be toggled to make the filter work properly. [43284]
- Packets ingressing a spoke-sdp interface are now “trusted” or “untrusted” based on the “tos-marking-state” setting and with TOS preserved or re-marked accordingly on network egress. [51473]
- If the “ip-mtu” is manually configured on an IES interface bound to a spoke-sdp with an MTU value that is larger than the SDP’s path-mtu, the SDP binding is now correctly being brought up with the smaller SDP path-mtu value. [52319]
- When using RSVP as a transport mechanism for spoke SDP IES interfaces, the SDP path-mtu no longer needs to be provisioned 4 bytes lower than the expected value (to account for an extra label being added if the underlying LSP moved to a bypass tunnel). [56590]

VPRN/2547

- For a VPRN, the confederation AS (configured in the provider core network) is no longer incorrectly passed to an external BGP peer (CE) as part of the AS-Path attribute of the route(s) being exchanged. [43819]

- The “longer” option for the “show router bgp routes” command now works for BGP-VPN routes. [47913]
- BGP update for a prefix advertised through the MP-BGP core no longer overwrites a previously installed route for the same prefix learned with a differing route-distinguisher. [50936]
- Large route distinguishers are no longer incorrectly displayed as negative numbers in the output of “show router bgp routes”. [51119]
- In the event logs, the ID shown for VPRN events now includes the service ID of the VPRN. [51869]
- If a /32 prefix for a local routing interface was received from a CE router, it would be installed in the local routing table, causing BGP peerings to toggle. This issue has been resolved. [52056]
- If the “ip-mtu” is manually configured on a VPRN interface bound to a spoke-sdp with an MTU value that is larger than the SDP’s path-mtu, the SDP binding is now correctly being brought up with the smaller SDP path-mtu value. [52319]
- In the “show router bgp next-hop detail” output, the egress label for remote PEs is now correctly displayed for all next-hops. [54026]
- Under some circumstances, the active CFM could reboot when more than 250 VPRNs were provisioned. This issue is now resolved. [55961]
- Changing the next-hop of bgp-vpn routes in a VPRN import policy now works properly. [57517]
- The origin attribute of a BGP route is now considered in the BGP route selection process of a VPRN between routes coming from the MP-BGP side and the PE-CE side. [57590]
- Configuring an IP address on a VPRN interface that is identical to an SDP far-end IP address no longer generates an error. [57864]
- BGP VPN routes that were previously redistributed into OSPF in VPRN instances were incorrectly advertised in OSPF LSAs if the same prefix were advertised in OSPF. This issue has been resolved. [59010]

VRRP

- When becoming master, the new VRRP master now correctly sends a gratuitous ARP. [42210]
- The VRRP master now correctly replies to local proxy ARP requests with the MAC address of the VRRP instance. [42316]
- Provisioning arp-populate on IES interfaces protected by VRRP no longer causes active DHCP originated ARP entries to disappear after a high availability CFM switchover. Also, unprovisioning that feature no longer causes the standby CFM to become unstable. [42571, 43215]
- Under some conditions, it would take one extra interval for VRRP to converge following a link failure as the first packet was not being sent on the wire. This issue has been resolved. [58248]

MIRROR SERVICE

- When mirroring traffic from a remote-source that is running a version of the 7750 SR OS, 7450 ESS OS or 7710 SR OS earlier than 3.0.R2, it was recommended that static labelling be used rather than using the dynamic labelling method. This prevented unsolicited label

releases from being incorrectly processed on the newer box which would result in LDP sessions flapping. This issue has been resolved. [42230]

- Enabling egress mirroring on VPLS SAPs that are forwarding replicated traffic (multicast, broadcast and unicast unknown) now works properly and will not affect system stability. [55505]

OAM

- OAM mac-ping and mac-trace commands sent to a remote APS group SAP now correctly respond to those OAM requests. [41223]
- SAA mac-trace commands are no longer erroneously reported as failed if the final MAC address is reachable and is less than 4 hops away. [41700]
- Configuring a DHCP host MAC on a VPLS that has dynamically learned the same MAC address will no longer result in a loss of synchronization and reset of the standby CFM. [41834]
- Trace error messages are no longer seen on the primary CFM of a High Availability system when OAM messages are received that cause a MAC address to be populated on a local VPLS. [41857]
- An OAM service ping issued from the VPLS side of a spoke SDP towards the IES spoke interface now works properly. [43099]
- Sending a “cpe-ping” through a SAP on which the arp-reply-agent now works properly. [50122]
- The trap notification for a CPE Ping no longer has the RttSumOfSquares value incorrectly set to zero. [52982]

KNOWN ISSUES

Following are specific technical issues that exist in Release 5.0.R25 of 7710 SR OS. Please also consult Known Limitations on page 73 as some known issues may have been moved to that section.

HW/PLATFORM

- When multiple power supply or fan failures are cleared, the clear event does not indicate which specific event is being cleared. [21221]
- A SONET/SDH port that is shutdown or in internal loopback is incorrectly being allowed as a valid synchronous timing reference. [36448]
- The sync-if-timing external reference on the standby CFM is not being monitored for validity as a qualified clock source. Thus, in cases of a CFM switchover, the sync-if-timing state may transition. [38071]

RADIUS

- In defining RADIUS Vendor Specific Attributes (VSAs), the TiMetra-Default-Action parameter is required even if the TiMetra-Cmd VSA is not used. [13449]

CLI

- There is currently no ‘show’ command to show the current values of the password hash settings. [32747]
- When logged in with the telnet client supplied with Windows and issuing an SSH to a remote server, an unexpected carriage return may be sent causing the SSH connection to fail. [50657]
- An asterisk (“*”), indicating an unsaved configuration change, may not be displayed after changing some of the parameters under the “configure>system” and “configure>log” contexts, among others. Additionally, the “Change Since Last Save” field in the “show system information” output may not be updated. [61271]
- When an SSH client is used to log in, long CLI commands may not display properly when the up-arrow key is used and may also render the backspace key unusable. [68252]
- The "show service *id* host summary" command for retailer services incorrectly reports zero hosts. The non-summary version of the same command will print the correct number of hosts. [70958]

SYSTEM

- When password aging option is enabled, the reference time is the time of the last boot and not the current time. Password expiry will also be reset on every reboot. [64581]
- A faulty compact flash disk that is used as a destination for a system’s DHCP persistency files may cause the system to generate a “sbmConfirmCB: Wrong magic num” error message in log 99. This is an internal message and does not affect functionality; it can be safely ignored. [73764]
- Line triggered FCS errors on POS ports may incorrectly result in “Ingress Pchip error” alarms. [76053]
- The system may not prevent the user from deleting a cron action that has running scripts even if it was not in shutdown state. This may result in the user’s inability to clear the running script’s entry from the cron action history list. [79794]

SONET/SDH

- On the OC-3c/STM-1c MDA, LOP-P defects received by the MDA are incorrectly reported as AIS-P events. [8658]
- The ‘show port’ command on a SONET/SDH interface will only display the bottom 4 bits of the S1 byte but will incorrectly display the bits as an entire byte. [17364]
- CV errors are incorrectly being incremented during a Severely Errored Seconds (SES) state. [29052]
- An OC-48/STM-16 port that is administratively shutdown is reporting LOS alarms to the event logger. [31786]
- If AIS-L (MS-AIS) or LOF alarms are detected on a SONET/SDH port used as a line timing reference, the reference incorrectly remains a valid reference when it should be disqualified as a timing source. [34387]
- If the statistics on a SONET/SDH port are cleared while bit error counters are incrementing on a section, line or path, the CV error counters may increment after the statistics are cleared. This leads to a condition where the CV counters have a value but no other error counters are non-zero. [41515]
- SONET/SDH line and path alarms are incorrectly being filtered in the console port output when the “no report-alarm” command is configured. The configuration parameter should only affect the logging of those alarms and not the CLI display. [61628]

TDM

- When a TDM channel is administratively disabled, the alarm statuses from “show port” are correct; however, the alarm log “Alarm RAI Set” is only reported when the condition is cleared. [58505]

PPP

- PPP is not preventing IPCP negotiation with a non-matching IP subnet address. [24475]
- The PPP dropcount parameter should not be reduced below 3. If set lower than 3, the PPP session may occasionally be dropped and re-established during CFM activity switches. [58684]

ATM

- ATM ports whose operational state toggle at a high rate (faster than both the up and down hold timers) may remain in a “Link Up” but not be in the operationally Up state. The workaround is to wait for the hold timer to expire before issuing the “no shutdown” command. [35066]
- ATM port statistics for AAL5 packets include all AAL type frames as well as ATM cells received on L2 ATM pseudowires (Apipes) on the OC-12c/STM-4c ATM MDA. [39089]
- If the receive side fiber of an ATM Apipe SAP loses link and that Apipe is also bound to an SDP, then remote OAM cells received on that SDP will be dropped since the Apipe service is locally in a down state. Additionally, ETE-RDI cells will be transmitted out the ATM SAP to the CE. [39571]
- If traffic is passing on an ATM OC-12 port and the port speed is changed to OC-3, “Unknown Protocol Discards” may be seen at the console although no such frames are actually being received. The OC-3 port's operational state is not affected, although some noise may be interpreted as end-to-end VC-RDI/AIS cells by newly configured ATM PVCs, which would cause those PVCs to go operationally down. The condition will clear as soon as ATM traffic passes once again through the port. [58197]
- ATM cells in a VPC connection with the GFC field not equal to zero will be discarded. [75387]

LAG

- A failure of the link holding the primary port of the LAG can sometimes very briefly impact (<10e-4 seconds) flows on other links of the same LAG. This is not the case for failures on other links (non-primary) of a LAG. [49698]

APS

- In exceptional cases, especially in a fully loaded node, where the occurrence of a high availability CFM switchover is exactly concurrent with an APS switch from Working to Protect (both unidirectional or bi-directional failures), PSBF may potentially be posted by the Far-End node during the APS K1/K2 byte exchange due to the increase latency response of the Near-End where the CFM switchover is occurring. [41192]

MANAGEMENT

- If a log file in the file system is manually deleted, the log may become operationally down as no new log file is created. The ‘clear log’ command will recreate the log file. [16317]
- Port-level and SAP-level statistics do not reflect packets processed by the CFM, for example, packets destined to a router IP address or a packet with the router alert options set. Another case is where DHCP relay packets ingress on a spoke-SDP bound to an IES

interface as these packets are first sent to the CPU, so the SDP does not reflect that these are ingressing packets. [16330]

- SNMPv3 user authentication and privacy keys in the ‘config>system>security>user username>snmp>authentication’ command must be entered as maximum length strings. [18314]
- Manual editing of SNMP persistent index files can cause errors in loading the configuration file. Persistent index files should only be created by the system. [24327]
- A log that is shutdown but is bound to a file-id will still create a file at the location for that file-id. [39044]
- A BOF static-route will remain active in the out-of-band management VRF even if the management Ethernet port is down. [40583]
- TIMETRA-PORT-MIB.mib does not include an entry for “Link Length support” as an attribute of a Gigabit Ethernet port. This prevents Alcatel-Lucent 5620 SAM from reporting the value even though this attribute is reported in the CLI. [46225]
- In a dual-CFM system where the management port is configured with a single IP address and auto-negotiation is disabled, the management port on the standby CFM may not operate in the correct duplex mode. In addition, the system may not bring the physical port down as a result of a duplex mismatch. [62142, 62143]
- The system may not correctly count the number of failed SNMPv3 authentication attempts in the event-control log. [64537]
- The system may not return a lexicographically higher OID than the requested OID in an SNMP GET-NEXT operation when incorrect values are used. This behavior is seen in the following tables:
  - inetCidrRouteTable [80593]
  - vRtrInetCidrRouteTable [80593]
  - tcpConnectionTable [80594]
  - vRtrInetStaticRouteTable [80598]
  - tmnxIPMafMatchTable [80600]
  - vRtrPimStaticRPTable [80605]
- After 497 days, the "Last Oper Chg" value in the output of the "show router interface" command and the "Last Mgmt Change" value in the output of the "show service service-using" command will wrap around. [83801]

ROUTING

- Setting a metric of zero in OSPF or IS-IS is not supported and causes the interface to fall back to the “reference-bandwidth” computed value instead of setting the value to zero. [17488]
- It is recommended that the preference value for BGP routes be set to a higher value than that of the internal (IGP) routes used to resolve the next-hop addresses of iBGP routes or routing instability can occur while the BGP routes are constantly re-learned. [31146]
- Reducing the interval/timeout timers much below default values is not recommended for OSPF, IS-IS, PIM, BGP, LDP and RSVP to ensure stability under transitional events like a CFM switchover. [56792, 58891]
- The use of long regular expressions (multiple terms separated by the “|” operator) in policy statements may cause the system to take more time than usual to process such regular expressions and potentially slow down protocol convergence. [85825]

- Policy statements referring to nonexistent community lists, prefix lists, as-path lists or damping profiles will be re-evaluated by protocols every time the route-policy configuration is committed, which may result in high CPU usage for a prolonged period of time. [86129]

DHCP RELAY

- If the addition of the Option 82 information to a DHCP packet would cause the maximum size of 1500 bytes to be exceeded, the DHCP packet is incorrectly not forwarding the original DHCP packet (without the additional Option 82 information). [37061]

RIP

- The RIP global statistics for all RIP instances is incorrectly being displayed for each VPRN instance. This has the effect of causing one to think that the VPRN instance has learned routes when in fact it has not. [26472]

IS-IS

- The granularity of the IS-IS hold timer is accurate only to within +/- 0.5s, so having a computed holdtime value of less than 2s may result in adjacencies being randomly dropped. It is recommended that hello-intervals and hello-multiplier values be adjusted accordingly, paying specific attention to the smaller hold-times computed on DIS systems. [29490]
- A POS-based interface will incorrectly not come up if it is configured as an IS-IS passive interface and the remote side does not respond to OSICP packets. [60100]
- Some types of corrupted CSNP PDUs are incorrectly acknowledged with a PSNP BPDU when they should be silently discarded. [67761]

OSPF

- The system may refresh self-originated LSA shortly after completing a CFM switchover which may mean the entry is refreshed before the expiration of the age-out period. [65195]
- Multiple links between two OSPF routers that regularly flap at the same time may cause the following invalid trap: “Conflicting configuration netMaskMismatch on interface”. [76131]
- The CSPF metric incorrectly toggles in situations where a point-to-point OSPF interface has asymmetric costs on both sides and an LSP is resignaled. [84404]

BGP

- The ‘remove-private’ command currently causes affected BGP peerings to be reset and re-established. [16727]
- If a 6PE prefix is received with two or more labels but for the same next-hop, the reference count in the “show router bgp next-hop” output will always show a value of 1. [56638]

MPLS/RSVP

- An invalid Class Number or C-Type in the Session Object does not cause a Path Error message to be generated. [12748]
- To disable OSPF-TE on a link, both ends of the link should be MPLS/RSVP disabled for CSPF to work correctly and be removed from the TE database. [15127]
- Loose hop LSPs are not re-signaled should a new equal-cost path become available. [17494]
- The bandwidth parameter is not supported on Path and ResV messages of one-to-one detour and facility bypass paths. [27394, 57847]

- For LSPs configured with fast reroute (FRR), on the transit nodes originating a detour path (PLR node), the detour's RESV message has the Local Protection Available (0x1) flag set even though the detour's PATH message does not have the Local Protection Desire flag set. [31486]
- For (rare) topologies in which the protected LSP and the detours are set up along parallel links across several hops (link protection only), fast reroute may take longer to restore traffic if the primary path is broken. [39808]
- Stress testing a large number of one-to-one protected LSPs may cause the CLI command "show router mpls lsp *lspname* path detail" to display incorrect information (egress label, outgoing interface, etc.) about detours. However, the forwarding table is properly programmed for these detours. [43782]
- For an SE-style one-to-one LSP with both primary and standby paths created, breaking the primary and standby path (by shutting down the port) and bringing up the standby path (by no shutting the port) may result in the generation of a harmless critical alarm at one of the LSRs. [46972]
- One-to-one detours for a protected LSP can be merged onto a locally originated detour in some network topologies, even if the latter traverses nodes that the one-to-one detours attempt to avoid. [48644]
- If a node is the point of local repair (PLR) for a large number of LSPs attempting to do global revertive after moving to a detour, it might take a few RSVP refresh timer periods before all LSPs complete the revertion process. [55185]
- If the MPLS config exec takes longer than two refresh cycles, the exec might fail due to contention between a new LSP and the bypass tunnel. [57809]
- A manual-bypass tunnel that terminates on the incoming interface IP address at the merge point will become operational but will not be properly associated with the primary LSP. The recommendation is to always use the IP address of the system interface to ensure reachability to the node. [59184]
- After changing the MTU size on the egress interface of an LSR on an LSP setup with strict hops associated with IP interfaces and protected with node-level facility bypass, the LSP may be incorrectly setup with link-protect bypass instead of node-protect bypass. [59738]
- RSVP LSPs cannot be signaled over a channelized DS1 or E1 interface if the channel group bandwidth is less than 1 Mbps. [59776]
- If a primary LSP path protected by FRR facility backup switches to a bypass tunnel and the total number of hops traversed by the primary path exceeds the configured primary path hop limit, then the MPLS LSP will go down after a few minutes. [72424]
- An MPLS LSR will incorrectly send RSVP RESV refresh messages to the upstream LSR after the link to that LSR had failed and recovered. This keeps the state for the LSP alive and the head-end LSR will be unable to set up a new LSP if it does not increase the LSP ID. [82284]

LDP

- When LDP is shut down on an intermediate ABR, an LDP-over-RSVP tunnel is not being resignaled to an alternate ABR. [58442]

IP MULTICAST

- If an 'rp static-address' is configured, the current PIM implementation will install an implicit deny-all for 224.0.0.0/4. To re-permit this address range, another static entry for this range must be installed. [38630]

- When a multicast CAC (MCAC) policy is applied under IGMP-snooping of a SAP with static-groups that are configured in the bundle of the same MCAC policy, the bandwidth used by the static groups on the SAP is not recalculated after the bundle is disabled, then re-enabled (commands “shutdown” and “no shutdown”). The used bandwidth remains at zero for the static groups. In addition, the MCAC recalculation command “tools perform service id *id* mcac sap *sap* recalc policy *policy*” fails to recalculate the used bandwidth and the use of the option “bundle” in the command returns an error. [71023]

PIM

- There is no CLI show command to see the SSM groups configured on PIM. The only way to see those SSM group is to use “info” in the config menu. [33746]

QoS

- All network egress traffic on Dot1Q tagged router interfaces has the dot1p values remarked even if the dot1p remarking option on that QoS network policy is disabled. [25850]
- When ler-use-dscp is enabled on network ingress and multicast vprn traffic is tunnelled through an SDP, ingress classification on network ingress will happen based on the TOS bits in the transport (outer) IP header as opposed to the customer IP packet. This behavior is seen strictly for multicast VPRN packets. [40348]
- Network control traffic (or other high-priority, expedited traffic) should not be configured to share a queue on a port scheduler policy with non-expedited or lower priority traffic or the queue could get into a state where the higher priority traffic will not be forwarded out the egress port. This can also occur if the traffic is on two separate queues that are mapped to the same level. [59298, 59435]

FILTERS/PBR/TCS

- When an IP filter is applied to a Layer-2 VPN service SAP with the action of remarking the DSCP or IEEE 802.1p bits, the router incorrectly does not issue a warning that these actions are not supported. [16850]
- If a Transparent Cache Switching (TCS) redirect-policy destination does not have a test clause defined, the operational state is reported as “Up”. [21227]

SERVICES GENERAL

- The CLI does not display an error when the user attempts to apply a filter log and a mirror-source to a given SAP at the same time. A filter log and mirror-source cannot be applied simultaneously to the same SAP. [22330]
- A static ip-mac host configured under the SAP of a group-interface where the port is shutdown will incorrectly switch from the “Not Forwarding” state to the “Forwarding” state upon a CFM switchover. [70345]
- The inactive CFM may go through a reset when ARP-populate is being enabled on an interface while the system is near the ARP table limit. [71052]
- The operational state for an SDP bound to a mirror service may show as “up” when the mirror service is administratively disabled. [81807]

SUBSCRIBER MANAGEMENT

- DHCP persistency should not be configured to use Compact Flash drives formatted with the newer Reliance file system in this release. [50940]
- Removing the first IP address from a Routed CO subscriber management interface can incorrectly result in the removal of this interface from OSPF, even if this interface has more than one IP address configured. [58225]

- Creating a static host with the same IP address as that of an existing host in another RCO IES service may cause downstream traffic to that IP address of that static host (and existing host) to be dropped. Removing the offending static host will restore all forwarding traffic to the existing host. [68394]
- If the subscriber management SAPs are on a LAG and a port scheduler is configured on the ports of that LAG, the port scheduler may not function correctly after the subscribers are created. In that case, the port scheduler will start to function correctly after one of the LAG ports is bounced. [81272]

VLL SPOKE SWITCHING

- If the control word is modified on a TPE device in a pseudowire switched environment with either a Cisco or an Alcatel-Lucent router running a previous software revision as the SPE device, it may be necessary to toggle the spoke binding status on the SPE device (l2vfi connection in the case of a Cisco). [57494]

CFLOWD

- Self-generated CFLOWD packets do not use the source IP address configured in the configure>system>security>source-address>application>cflowd context. [83165]

VPLS

- If a large number of MAC addresses exists in the VPLS FDB and the entire FDB is flushed and relearned, there may be a period of up to 5 seconds where RSTP BPDUs are not sent out. A partial workaround is to configure fdb-table-size limits. [40532]
- In a distributed VPLS configuration, it may take up to (2*(Max Age)-1) seconds to age a remote MAC address, and in cases of CFM switch over, it may take up to (3*(Max Age)-1) seconds. [48290]
- A user VPLS SAP might stop forwarding traffic after the SAP port bounces if that SAP is managed by a management VPLS (mVPLS) with Spanning Tree Protocol disabled. The workaround is to remove the mVPLS if the Spanning Tree Protocol is not required. If Spanning Tree Protocol is required, it should be enabled on the mVPLS. [60262]

VPRN/2547

- VPRN service traffic with the DF (Do Not Fragment) flag set and requiring fragmentation to be transported through an SDP tunnel is correctly discarded, but an ICMP Type 3 Code 4 (fragmentation needed and DF set) message is not issued. [18869]
- The total number of routes displayed in show router <vprn-id> summary is the total number of routes for all VRFs, not just the “best” routes installed. [21642]
- The service operational state of a VPRN might be displayed incorrectly as Up during its configuration while some mandatory parameters to bring it up have yet to be set. [31055]
- The “triggered-policy” feature does not apply to VRF import and export policies in a VPRN. One needs to reset the target VRF instance in order to re-evaluate these policies or to disable the “triggered-policy” feature. [43006]
- Executing a ping from a VPRN without a configured loopback address may fail with a “no route to destination” error message despite there being a valid route in the routing table. The error message is misleading and should state that the reason for the failure is not having a source address configured. [55343]
- When there are overlapping BGP and BGP-VPN routes in a VRF, after the BGP-VPN routes has flapped a number of times, it is possible that the USED flag is no longer set for the BGP routes. [82112]

VRRP/SRRP

- SRRP log events are labelled as VRRP events in the event log output. [57709]
- If VRRP or SRRP interfaces are configured for sub-second intervals lower than 300 ms, transmission of VRRP/SRRP packets could be delayed on a High Availability switchover, leading to toggling of VRRP/SRRP master states. [59695, 60009]
- If the in-use priority on each side an SRRP connection go to zero, both routers will incorrectly elect themselves as master. [60032]

MIRROR SERVICE

- If a Dot1Q SAP is being mirrored on an IES interface, DHCP responses from the server to the DHCP clients are not mirrored. A workaround is to mirror the port instead of the SAP. [40339]

BFD

- BFD is not supported on aggregated interfaces (LAG and APS), but configuring BFD is incorrectly not blocked in CLI. [57774]

OAM

- If a mac-ping or mac-trace request is sent with an unknown source MAC address and there are multiple SAPs, the user will see duplicated results because the request is flooded to each SAP and each SAP sends a reply to the request message. This is the expected behavior. [16298]
- OAM vprn-trace may sometimes display a “no route to destination” error when the route is in fact reachable via an in-band ping via the VPRN tunnel. [24272]
- When an “oam saa lsp-ping” test is configured with rising and falling thresholds, the system may erroneously generate threshold crossing alarms following a “clear saa” command or system. reboot. [54444]
- OAM DNS lookups are not working correctly if the full DNS name is not provided. [54239, 54689]
- OAM vprn-trace packets are incorrectly timing out when sent to ASBRs in an inter-AS configuration. [59395]

Document Part Number: 93-0176-25 V5.0.R25

93-0176-25 V5.0.R25