

Alcatel-Lucent 5620

SERVICE AWARE MANAGER | RELEASE 8.0 R5

RELEASE DESCRIPTION

3HE 06054 AAAE TQZZA Edition 01

Alcatel-Lucent Proprietary
This document contains proprietary information of Alcatel-Lucent and is not to be disclosed
or used except in accordance with applicable agreements.

Copyright 2010 © Alcatel-Lucent. All rights reserved.

Alcatel-Lucent assumes no responsibility for the accuracy of the information presented, which is subject to change without notice.

Alcatel, Lucent, Alcatel-Lucent, the Alcatel-Lucent logo, and TiMetra are registered trademarks of Alcatel-Lucent. All other trademarks are the property of their respective owners.

Copyright 2010 Alcatel-Lucent.
All rights reserved.

Disclaimers

Alcatel-Lucent products are intended for commercial uses. Without the appropriate network design engineering, they must not be sold, licensed or otherwise distributed for use in any hazardous environments requiring fail-safe performance, such as in the operation of nuclear facilities, aircraft navigation or communication systems, air traffic control, direct life-support machines, or weapons systems, in which the failure of products could lead directly to death, personal injury, or severe physical or environmental damage. The customer hereby agrees that the use, sale, license or other distribution of the products for any such application without the prior written consent of Alcatel-Lucent, shall be at the customer's sole risk. The customer hereby agrees to defend and hold Alcatel-Lucent harmless from any claims for loss, cost, damage, expense or liability that may arise out of or in connection with the use, sale, license or other distribution of the products in such applications.

This document may contain information regarding the use and installation of non-Alcatel-Lucent products. Please note that this information is provided as a courtesy to assist you. While Alcatel-Lucent tries to ensure that this information accurately reflects information provided by the supplier, please refer to the materials provided with any non-Alcatel-Lucent product and contact the supplier for confirmation. Alcatel-Lucent assumes no responsibility or liability for incorrect or incomplete information provided about non-Alcatel-Lucent products.

However, this does not constitute a representation or warranty. The warranties provided for Alcatel-Lucent products, if any, are set forth in contractual documentation entered into by Alcatel-Lucent and its customers.

This document was originally written in English. If there is any conflict or inconsistency between the English version and any other version of a document, the English version shall prevail.

CONTENTS

Release 8.0 – At A Glance	10
<i>Target Schedule</i>	<i>10</i>
<i>Network Element Support</i>	<i>10</i>
<i>Feature List</i>	<i>10</i>
1. Scalability	25
<i>Scalability Targets</i>	<i>25</i>
<i>Performance Targets</i>	<i>26</i>
2. Platform.....	28
<i>Support x86 Intel 64-bit (NEBS compliant)</i>	<i>29</i>
<i>IPMP Support</i>	<i>29</i>
<i>Inband & OutOfBand Enhancements</i>	<i>30</i>
<i>Support for preferred management address SNMP table for SR</i>	<i>32</i>
<i>Preferred management address after un-manage</i>	<i>33</i>
<i>Rename Primary/Secondary address fields to OOB/IB</i>	<i>33</i>
<i>Management Ping versus IB and OOB</i>	<i>33</i>
<i>IPv6</i>	<i>34</i>
Configuring the address	34
Discovery	36
Determining the address used	37
Configuring IPv6 Trap Destinations	37
Trap Target	38
Limitations	38
Future Evolution (R5 Candidate)	38
<i>SAM Database and Server Security Enhancements</i>	<i>39</i>
<i>Network Element Backups & MME Perf Mgmt Files</i>	<i>39</i>
<i>One Build Support - Network Element Upgrades</i>	<i>40</i>
<i>GUI Framework Enhancements</i>	<i>40</i>
Reduce JMS events processed by SAM GUI client	40
“Reset” button on Configuration Forms	41
Last Search Time in Table Views	42
Changes in Layout of Configuration Forms	42
<i>Map Enhancements</i>	<i>43</i>
Map Filter Functionality	43
Make Map from Group	44
Show Only Selected Option	44
Info Table Usability	45
Show on Mouse Over option	45
Show Header option	45
500 Elements per Group	46
Icon Legend	46
Find Vertex/Edge Enhancements	47
Find by Label functionality	47
Show On Map	48
Logical Groups	48
Keyboard controls	48
Configurable labels for LLDP	49
<i>Span of Control/Scope of Command</i>	<i>49</i>
Span on Services & Customers	49
Services	50
Span Types	50

Span Rule.....	51
Creating Services	52
Customers	52
Restricting Read Only Access	53
Menu for Scope of Command & Span of Control	53
<i>System Administration</i>	53
Task Manager Enhancements	53
<i>Policy Infrastructure</i>	54
Explicit Policy distribution scalability.....	54
Global policy created by the first local policy discovery.....	55
Policy performance improvements	56
Switch Mode of policy from list	56
<i>VI for CLI</i>	56
Multiple NE Types for a Script.....	56
<i>Tools</i>	57
SAP Copy/Move Support for ATM.....	57
Tools CLI commands available for SAM.....	57
<i>Scripts</i>	58
XML U/S - Modify Templates incl def queries / curr	58
3. Equipment Management.....	58
<i>7450/7750 IOM3/IMM Mixed-Mode Chassis Support</i>	58
<i>New 7xx0 MDA support</i>	60
<i>Hybrid Port Support</i>	61
Business and Mobile Backhauling.....	61
Service Destination Based Shaping	62
Hybrid Port Mode Configuration.....	62
SAP Capability	63
Network IP Interface Capability	63
Hardware Support.....	63
QoS Requirements	63
<i>M1-10GB DWDM TUN with DWDM Wavelength Selection</i>	64
<i>Support for OUT changes</i>	64
<i>Synchronization Enhancements on the SR & ESS</i>	66
<i>ESMC (Ethernet Synchronization Message Channel) Configuration on an Ethernet Port</i>	66
<i>System Clock Configuration</i>	66
<i>Enhanced BITS in port redundancy</i>	67
<i>MultiService ISA support in the IOM3 for Video Services</i>	67
<i>Prioritization Mechanism for RET vs. FCC</i>	67
<i>ECMP fate sharing</i>	67
<i>E-LMI Protocol</i>	68
<i>Increase the LAG limit to 16</i>	68
<i>Enhancements to LAG Hashing for consistent per service forwarding</i>	68
<i>SAP and Service Configuration Enhancement for Frame Relay</i>	68
<i>Multi-Chassis PW Endpoint Support for VPLS</i>	69
MC Endpoint support in Node Redundancy management.....	71
MC Endpoint support on VPLS Site.....	72
Binding MC Endpoint with Spoke SDP	72
Alarm Support	73
Statistics Support	73
<i>7750 SR c12 Enhancements</i>	73
CPM Filter and CPMQ	73
Timestamp up to 128 bytes into the packet	73
Scalability Requirements	73
<i>OMNI Support</i>	74

OSPF.....	77
RIP.....	79
MPLS/LDP.....	80
VRF Instances.....	82
OMNI Scalability.....	82
7210 SAS Support.....	83
SFP Diagnostics and Monitoring.....	84
9500 Support.....	85
Network Architecture.....	86
UMTS Backhaul Architecture.....	86
CDMA Backhaul Architecture.....	87
Transport Tunnel.....	87
7705 Support.....	87
7705 SAR R3.0 R1 functionality supported in 5620 SAM 8.0 R1.....	87
4 Port DS3/E3 ASAP Daughter Card.....	87
8 Port Ethernet Daughter Card Enhancements.....	88
16 Port DS1/E1 Daughter Card Members Per IMA/MLPPP Bundle.....	88
12 Port Serial Data Interface Daughter Card Port Display Enhancement.....	88
802.1ag/Y.1731 Ethernet OAM Enhancements.....	88
VLL Redundancy – Active/Standby.....	89
VLL Switching Site.....	89
VPRN.....	89
Management IES.....	89
Routing Policy.....	90
IP Filter Policies.....	90
Performance Statistics.....	90
IP Service Tunnels.....	90
DHCP Relay On Network Interfaces.....	91
Routing Instance AS Number.....	91
Route Aggregation.....	91
IP Addresses.....	91
BGP.....	91
7705 SAR 3.0 R1 Functionality added in 5620 SAM 8.0 R3.....	92
7705 SAR 3.0 R1 Functionality NOT currently supported.....	92
GNE Support.....	92
Hardening of GNE profile.....	92
LLDP support (GNE).....	92
4. Assurance.....	93
Alarm Timestamps.....	93
Test Suite support for CFM tests.....	93
OAM Results on Map (including CFM).....	96
Ethernet SAP OAM (Mapping).....	96
Modify AIS selection.....	96
Add MEP Fault Propagation.....	96
Down MEP on Ipipe Ethernet SAP.....	96
Epipe Service & 3.5 MIP for Epipe.....	96
Enable MEP/MIP on Epipe Service.....	97
Auto MEP/MIP Creation.....	97
Service & Composite-Service OAM.....	98
MAC address provisioning for MIP.....	98
Add MAC property to MIP.....	98
Fast CCM timer for all MEPs.....	99
CCM Interval Enumeration:.....	99
Continuous SAA CFM OAM Tests.....	100

5. Service & Routing Enhancements.....	101
<i>VLL Redundancy Enhancements</i>	<i>102</i>
Network Redundancy	102
Access Redundancy	103
SAM Alarms.....	103
<i>Segmented Service</i>	<i>103</i>
<i>Service Provisioning Enhancements.....</i>	<i>103</i>
<i>Tunnel Selection Enhancements</i>	<i>103</i>
<i>Service Component Tree Navigation</i>	<i>104</i>
<i>Service Site Creation from Service Map.....</i>	<i>104</i>
<i>Service CAC.....</i>	<i>104</i>
Request for Admission Control	106
Topology change and BW usage re-calculation	106
Audit	107
Ethernet CAC	107
<i>IPVPN - VRF id based on strings rather than numbers</i>	<i>108</i>
<i>Diff-Serv Class type change during failures.....</i>	<i>109</i>
<i>BFD support of OSPF CE-PE adjacencies</i>	<i>111</i>
<i>Spoke termination for IPv6 IES & 6VPE.....</i>	<i>111</i>
<i>Support for MPLS hash label</i>	<i>111</i>
<i>RSVP LSP Primary and LDP LSP backup within a SDP</i>	<i>114</i>
<i>Re-signalling of primary RSVP-TE paths.....</i>	<i>121</i>
<i>Full IGP Shortcuts and Forwarding Adjacencies</i>	<i>122</i>
<i>Enable/disable the no-propagate-ttl capability</i>	<i>124</i>
<i>IS-IS and OSPF TE bandwidth updates triggered by threshold crossing events.....</i>	<i>125</i>
<i>IMPLICIT NULL label option support on egress LER.....</i>	<i>127</i>
<i>PBB (MMRP) Scalability for inter-domain services</i>	<i>128</i>
<i>Downstream on demand label for LDP (Tunnel only).....</i>	<i>129</i>
<i>PCP (dot1p) and DE bits transparency for PBB.....</i>	<i>130</i>
<i>LDP Shortcut for IGP Route Resolution.....</i>	<i>131</i>
<i>GR helper for PE-CE protocols.....</i>	<i>132</i>
<i>Precedence support for LSP secondary paths</i>	<i>134</i>
<i>Multiple LDP LSR-IDs and LDP Instances.....</i>	<i>135</i>
<i>IP Pseudowire L3 Termination.....</i>	<i>139</i>
<i>T-LDP status TLV (Active/Standby and Oper) support on VPRN</i>	<i>139</i>
<i>Option to Place IGP into Overload if Switch Fabric Fails</i>	<i>141</i>
<i>Load Sharing Multiple BGP Paths Even If AS-Path is Different.....</i>	<i>142</i>
<i>Support for RFC3107 BGP Label for L2 Services.....</i>	<i>143</i>
<i>RSVP Shortcut for BGP Next-Hop Resolution.....</i>	<i>144</i>
<i>PBB Ethernet Tunnel Enhancements.....</i>	<i>144</i>
‘Create CFM Continuity Check’	146
<i>Point to Multi Point LSP Enhancements</i>	<i>146</i>
<i>Tunnel Interface Changes in 8.0.....</i>	<i>147</i>
SAM 8.0 Changes to IGMP Tunnel Interface	147
<i>Multicast Tree BFD.....</i>	<i>147</i>
<i>Source Redundancy</i>	<i>147</i>
<i>LDP for P2MP.....</i>	<i>148</i>
<i>P2MP-Lsp-Ping for RSVP P2MP LSP</i>	<i>149</i>
<i>P2MP-Lsp-Trace for RSVP P2MP LSP.....</i>	<i>149</i>
<i>Hierarchical Policy Support.....</i>	<i>149</i>
Configuration.....	150
<i>Policer Control Hierarchies.....</i>	<i>151</i>
<i>Policer Control Policy.....</i>	<i>151</i>
Arbiter Entry	152
Priority Level Entry	152

Shared Policer Output Queue	152
<i>Queue Group Template and Port Queue Group Enhancements</i>	152
Queue Mbs unit changes.....	152
Forwarding Class for Egress Queue Group Template	153
Default Egress Queue Group Template	153
Egress Queue.....	153
Port Access Egress Queue Group Configuration.....	154
<i>Access Ingress/Egress Policy Enhancements</i>	154
Policer.....	154
Forwarding Class	155
<i>Subscriber Profile Enhancement</i>	155
Applying Policer Control Policy.....	155
Configure Policer Control Override	155
<i>SLA Profile Enhancement</i>	156
Policer Override.....	156
Service Access Point.....	156
Assign Policer Control Policy.....	157
Policer Control Override (8.0R4 candidate)	157
Policer Override (8.0 R4 candidate)	157
<i>Down MEP support for IES & VPRN Subscriber Group Interface Ethernet SAP</i>	158
<i>VPRN and IES support in Service creation for Global Maintenance Associations</i>	158
<i>Down MEP creation/discovery for VPRN/IES Subscriber Group Interface</i>	158
<i>FC Mapping based on Exp Bits at VLL/VPLS SAP</i>	158
<i>PPPoE for Residential & Business Wholesale</i>	160
Hardware Requirements	161
<i>Typical L2TP Configuration</i>	161
<i>L2TP Protocol and L2TP Protocol Site</i>	162
L2TP Group Configuration.....	162
L2TP Group Operations	162
L2TP Tunnel Configuration Tab	163
Tunnels Status Tab	163
<i>L2TP Tunnel Configuration</i>	163
<i>L2TP Tunnel Status</i>	163
<i>L2TP Tunnel Operations</i>	163
<i>L2TP Tunnels</i>	164
<i>Purging L2TP Tunnels</i>	164
<i>L2TP Peer</i>	164
<i>L2TP Peer Tunnels</i>	165
<i>L2TP Peer Operations</i>	165
<i>L2TP Sessions</i>	165
<i>L2TP Session Managed Routes</i>	165
<i>L2TP ISA LNS Group</i>	165
<i>ISA LNS Group Member Operations</i>	166
<i>L2TP Manager</i>	166
<i>Network Address Translation (NAT)</i>	166
<i>NAT Policies</i>	167
<i>Nat Configuration for IES and VPRN services</i>	169
NAT Pools	169
Nat Destinations	170
L2 Aware Ip Addresses	170
<i>ISA Groups and MDA's</i>	170
Assigning NAT Policies to Subscriber Policies.....	171
NAT and IP Filters	171
Statistics.....	171
Traps/Alarms	172
<i>tmnxNatIsaMdaSessionUsageHigh</i>	172

<i>tmnxNatPlBlockUsageHigh</i>	172
<i>tmnxNatPlBlockAllocationLsn</i>	172
<i>DOS Protection</i>	173
6. Triple Play Market	173
<i>ESM Enhancements for L2TP</i>	173
<i>IES/VP RN Group Interface Configuration Enhancements</i>	174
<i>Local User Database Configuration Enhancements</i>	174
<i>Subscriber Hosts and PPPoE Sessions Enhancements</i>	174
<i>L2TP Performance Statistics</i>	174
<i>ARP Host</i>	174
<i>ARP host VPLS configuration</i>	174
<i>ARP host VPLS stats</i>	175
<i>ARP host IES/VP RN Group interface configuration</i>	175
<i>ARP hosts</i>	175
<i>ARP host managed routes</i>	175
<i>IPSec VPNs</i>	175
<i>Static IPSec Tunnel</i>	176
<i>IPSec Tunnels on a Single Site</i>	176
<i>IPsec Tunnels on Multiple Sites</i>	177
<i>IPSec Application Function as part of the Corporate Network</i>	178
<i>IPSec Enhancements in SAM 8.0</i>	178
<i>IPSec Application Function</i>	178
<i>IPSec Application Function Discovery/Resync</i>	179
<i>IPSec Application Function Topology</i>	180
<i>Alarm Correlation and OAM</i>	183
<i>Service Test Management (STM)</i>	184
<i>Test Suites</i>	184
<i>Test Scheduling:</i>	184
<i>Test Policy</i>	184
<i>IPSec Session Management</i>	184
<i>Bidirectional Forward Detection (BFD)</i>	185
<i>Include DHCP-Options into RADIUS Authentication-Request</i>	186
7. Application Assurance	186
<i>ISA-AA Group CFLOWD implementation</i>	186
<i>Application Assurance Policy Enhancements</i>	188
<i>New Manage -> Application Assurance Window</i>	188
<i>Multiple ISA-AA Groups and Partitions Support</i>	188
<i>ISA-AA Group Configuration Enhancements</i>	189
<i>ISA-AA Group Partition Configuration</i>	189
<i>AA Group Policy Management Enhancements</i>	189
<i>AA Account Policy Enhancements in Release 8.0</i>	189
<i>AA Statistics/Debug Enhancements</i>	190
<i>AA Policer Enhancements</i>	190
<i>Object References and Filters to AA Group/Partition Policy Enhancements</i>	190
<i>Protocol Shutdown for New Signature Upgrade Support</i>	190
<i>Custom Pattern Based Protocol Support</i>	191
<i>Application Filter Expression Match Extensions</i>	191
<i>Capacity Based Load Balancing</i>	191
<i>Spoke-SDP aa-subs</i>	191
<i>Spoke SDP AA Performance and Real-time Stats</i>	191
<i>ISA-AA only Upgrades</i>	191
8. LTE	192

<i>Key Elements</i>	<i>192</i>
<i>5620 SAM Functions</i>	<i>194</i>
9. Deprecations	195
<i>Deprecations This Release</i>	<i>195</i>
5620 SAM LSP MAP	195
SAM-O Deprecations	196
<i>Deprecations in Future Releases.....</i>	<i>196</i>
Dynamic Shelf Drawings.....	196
SAM-O Deprecations	196
Platform Deprecations	196
10. Ordering information.....	197
<i>Sales Engineering Information.....</i>	<i>197</i>
<i>Licensing Information.....</i>	<i>197</i>
<i>Oracle Ordering Information</i>	<i>197</i>
11. References.....	200

RELEASE 8.0 - AT A GLANCE

Target Schedule

Four releases with new content are planned.

8.0 R1 - April 21, 2010

- 7x50 8.0 R1 features
- NSM content
- Architecture work to drive SAM scaling and performance upward
- 7.0 Scaling limits supported in 8.0
- Support for other NEs whose DR4 dates align
- CPAM 4.0

8.0 R3 - June 30, 2010

- 7x50 8.0 R3 features and R1 stretch features
- Support for other NEs whose DR4 dates align

8.0 R5 - Oct 27, 2010

- 7x50 8.0 R4/R5 features
- NSM content
- Support committed 8.0 Scaling Targets
- Support for other NEs whose DR4 dates align
- CPAM 5.0

8.0 R7 - January 2011

- Support for other NEs

There will also be maintenance releases scheduled in between content releases throughout the year.

Network Element Support

Please see the *5620 SAM Network Element Compatibility Guide* for information on release compatibility.

Feature List

The following table lists all candidate and committed nodal and NMS functionality to be supported in 5620 SAM along with expected target dates.

Note: This view is subject to change. Content targeted for releases later than R5 is not yet committed. View as of September 14, 2010. Ask your regional representative for an updated view if required.

Rel	Load	Feature Description	Status	Node Number	Node Rel	Node Load
Nodal Features						
8.0	R1	Link Loss Forwarding on SAS-E	Committed	7210	1.1	R6
8.0	R1	PIM object model change	Committed	7x50	8.0	R1
8.0	R1	Error Recovery Mechanism	Committed	9500 MPR	1.3 E	

8.0	R1	Trusted Manager Support	Committed	9500 MPR	1.3 E	
8.0	R1	IOM3 support with 200G SF/CPM2	Committed	7x50	8.0	R1
8.0	R1	Chassis mode B for L2TP LAC/PPP termination (IOM1)	Committed	7x50	8.0	R1
8.0	R1	L2TP AVP 22 (calling-number) must be configurable	Committed	7x50	8.0	R1
8.0	R1	LAC support for "address float" for L2TP tunnels	Committed	7x50	8.0	R1
8.0	R1	L2TP tag value zero (0)	Committed	7x50	8.0	R1
8.0	R1	Remove subscriber ID restrictions	Committed	7x50	8.0	R1
8.0	R1	16 Link bundles (ML-PPP)	Committed	7705	3.0	R1
8.0	R1	Two-tier scheduling on Ethernet MDA	Committed	7705	3.0	R1
8.0	R1	Mobile Service	Committed			
8.0	R1	ISA-AA Upgrade	Committed	7x50		
8.0	R1	OSPF shadow tables	Committed	7x50		
8.0	R1	Load-sharing support for video	Committed	7x50		
8.0	R1	IGMP Snooping	Committed	7210	1.1	R6
8.0	R1	Spoke-SDP aa-subs	Committed	7x50	8.0	R1
8.0	R1	5780 DSC R1.0 Beta Support	Committed	5780 DSC	1.0	
8.0	R1	LTE Performance Management	Committed			
8.0	R1	LTE Licensing	Committed			
8.0	R1	LTE Control Bearer Management	Committed			
8.0	R1	9471 MME Support	Committed	9471 MME	LM2.0.1	
8.0	R1	OmniSwitch 9000E (Fuji2) / part of AOS 6.4.2	Committed	OmniSwitch		
8.0	R1	OmniSwitch 6250 / AOS 6.6.1 stream support	Committed	OmniSwitch		
8.0	R1	16 link bundles (IMA)	Committed	7705	3.0	R1
8.0	R1	Y.1731 FM/PM	Committed	7705	3.0	R1
8.0	R1	APD (Automatic Discovery Protocol) on Ethernet Por	Committed	7705	3.0	R1
8.0	R1	Active/Standby mode for PW-redundancy	Committed	7705	3.0	R1
8.0	R1	SSM for Synch-E with Q/L	Committed	7705	3.0	R1
8.0	R1	VPRN	Committed	7705	3.0	R1
8.0	R1	Pseudo-wire switching	Committed	7705	3.0	R1
8.0	R1	IP tunnels	Committed	7705	3.0	R1
8.0	R1	4 Port T3/E3 Card (-48V & +24V)	Committed	7705	3.0	R1
8.0	R1	Support for 2XGE ports (SAS-MX)	Committed	7210	1.1	R6
8.0	R1	ATM counters (per VPI/VCI)	Committed	9500 MPR	1.3 E	
8.0	R1	IMA counters per group	Committed	9500 MPR	1.3 E	
8.0	R1	16-Port E1/T1 ASAP Access card	Committed	9500 MPR	1.3 E	
8.0	R1	UBR,UBR+ flows support	Committed	9500 MPR	1.3 E	

8.0	R1	CBR flows support	Committed	9500 MPR	1.3 E	
8.0	R1	Max number of VPI/VCI	Committed	9500 MPR	1.3 E	
8.0	R1	Max number of E1s in a IMA group	Committed	9500 MPR	1.3 E	
8.0	R1	VCC mode support	Committed	9500 MPR	1.3 E	
8.0	R1	VPC mode support	Committed	9500 MPR	1.3 E	
8.0	R1	PWE3 ATM N:1 cell mode (with N=1)	Committed	9500 MPR	1.3 E	
8.0	R1	ATM cell concatenation into a PWE3 word	Committed	9500 MPR	1.3 E	
8.0	R1	PWE3 encapsulation according to RFC 4717	Committed	9500 MPR	1.3 E	
8.0	R1	ATM2ETH	Committed	9500 MPR	1.3 E	
8.0	R1	ATM2ATM	Committed	9500 MPR	1.3 E	
8.0	R1	E1-IMA AF-PHY-0086.001	Committed	9500 MPR	1.3 E	
8.0	R1	ATM-IMA over E1/T1 - 75 Ohms, 1.0/2.3 panel	Committed	9500 MPR	1.3 E	
8.0	R1	ATM-IMA over E1/T1 - 75 Ohms, BNC panel	Committed	9500 MPR	1.3 E	
8.0	R1	ATM-IMA over E1/T1 - 75 Ohms, 1.6/5.6 panel	Committed	9500 MPR	1.3 E	
8.0	R1	ATM-IMA over E1/T1 - 100/120 Ohms, RJ45 panel	Committed	9500 MPR	1.3 E	
8.0	R1	ATM-IMA over E1/T1 - 100/120 Ohms, FW	Committed	9500 MPR	1.3 E	
8.0	R1	1+1 U-APS for CES	Committed	7x50	8.0	R1
8.0	R1	Increase LDP adjacencies	Committed	7x50	8.0	R1
8.0	R1	LSR FRR performance improvement	Committed	7x50	8.0	R1
8.0	R1	Router independent signature upgrade within a Majo	Committed	7x50	8.0	R1
8.0	R1	Capacity-Cost based load balancing	Committed	7x50	8.0	R1
8.0	R1	Fat SAP Support	Committed	7x50	8.0	R1
8.0	R1	Software use of 8G on AA ISA	Committed	7x50	8.0	R1
8.0	R1	cflowd volume and TCP performance reporting	Committed	7x50	8.0	R1
8.0	R1	Support multiple (7) AA ISA Groups	Committed	7x50	8.0	R1
8.0	R1	Policy: partition and scale	Committed	7x50	8.0	R1
8.0	R1	Policy: app-filter expression match extensions	Committed	7x50	8.0	R1
8.0	R1	Custom pattern-based protocols	Committed	7x50	8.0	R1
8.0	R1	Protocol shutdown for new signature upgrade	Committed	7x50	8.0	R1
8.0	R1	Protocol Signatures	Committed	7x50	8.0	R1
8.0	R1	NTP within a VPRN	Committed	7x50	8.0	R1
8.0	R1	Higher scaling numbers for OSPF	Committed	7x50	8.0	R1

8.0	R1	LDP Shortcut feature for IP forwarding over MPLS (Committed	7x50	8.0	R1
8.0	R1	SAP Packet Interleaving for FRF.12 e2e fragmentati	Committed	7x50	8.0	R1
8.0	R1	WRED per Queue/Forwarding class in Sparrow	Committed	7x50	8.0	R1
8.0	R1	BGP rapid update for all address families includin	Committed	7x50	8.0	R1
8.0	R1	Multiple instances of IS-IS	Committed	7x50	8.0	R1
8.0	R1	Local switching for Epipe in PBB access dual-homin	Committed	7x50	8.0	R1
8.0	R1	RSVP Shortcut for BGP Next-Hop Resolution (RFE 767	Committed	7x50	8.0	R1
8.0	R1	RFE 72052 - PCP (dot1p) and DE bits transparency f	Committed	7x50	8.0	R1
8.0	R1	Option to place IGP into overload if switch fabric	Committed	7x50	8.0	R1
8.0	R1	Precedence support for LSP secondary paths	Committed	7x50	8.0	R1
8.0	R1	RFC5095 (deprecations of type 0 routing headers)	Committed	7x50	8.0	R1
8.0	R1	RFC4552 - Authentication for OSPFv3	Committed	7x50	8.0	R1
8.0	R1	Downstream on demand label for LDP (Tunnel only)	Committed	7x50	8.0	R1
8.0	R1	Increase maximum LAG members limit to up to 16	Committed	7x50	8.0	R1
8.0	R1	support for RFC3107 BGP label for L2 services	Committed	7x50	8.0	R1
8.0	R1	IPv6 route table scaling (150K) - Discuss options	Committed	7x50	8.0	R1
8.0	R1	BGP Peering Scaling - 5K (w/ SFM3)	Committed	7x50	8.0	R1
8.0	R1	Traffic leaking to the GRT from a VPRN	Committed	7x50	8.0	R1
8.0	R1	VPRN indirection for NH in the core and access	Committed	7x50	8.0	R1
8.0	R1	BGP AD for Sparrow and Aragorn - feature parity, o	Committed	7x50	8.0	R1
8.0	R1	PBB (MMRP) Scalability for inter-domain services (Committed	7x50	8.0	R1
8.0	R1	Enhancements to LAG hashing for consistent per-ser	Committed	7x50	8.0	R1
8.0	R1	"as-path multipath-relax" to load-share multiple B	Committed	7x50	8.0	R1
8.0	R1	FC mapping based on EXP bits (on VLL/VPLS SAP)	Committed	7x50	8.0	R1
8.0	R1	IMPLICIT NULL label option support on egress LER	Committed	7x50	8.0	R1

8.0	R1	2K network ports	Committed	7x50	8.0	R1
8.0	R1	Hybrid Port Mode	Committed	7x50	8.0	R1
8.0	R1	Ingress LER FRR performance improvement	Committed	7x50	8.0	R1
8.0	R1	IS-IS and OSPF TE bandwidth updates triggered by t	Committed	7x50	8.0	R1
8.0	R1	Provide CLI command to enable/disable the no-propa	Committed	7x50	8.0	R1
8.0	R1	T-LDP status TLV (Active/standby and oper) support	Committed	7x50	8.0	R1
8.0	R1	IP pseudowire L3 termination	Committed	7x50	8.0	R1
8.0	R1	Increase VRRP scale to 2k/system; Increased VRRP s	Committed	7x50	8.0	R1
8.0	R1	Full IGP Shortcuts (including Cisco autoroute) and	Committed	7x50	8.0	R1
8.0	R1	make before break re-signalling of primary RSVP-TE	Committed	7x50	8.0	R1
8.0	R1	RSVP LSP Primary and LDP LSP backup within a SDP	Committed	7x50	8.0	R1
8.0	R1	Increase number of MC-LAG peers to 20	Committed	7x50	8.0	R1
8.0	R1	Multiple loopback for LDP, T-LDP (at least 16)	Committed	7x50	8.0	R1
8.0	R1	Support for MPLS hash label	Committed	7x50	8.0	R1
8.0	R1	Increase cflowd performance/scale	Committed	7x50	8.0	R1
8.0	R1	Netflow v9 - support flow stats on MPLS terminated	Committed	7x50	8.0	R1
8.0	R1	Spoke termination for IPv6 IES & 6VPE	Committed	7x50	8.0	R1
8.0	R1	BFD support of OSPF CE-PE adjacencies	Committed	7x50	8.0	R1
8.0	R1	Diff-Serv Class type change during failures	Committed	7x50	8.0	R1
8.0	R1	IPVPN - VRF id based on strings rather than number	Committed	7x50	8.0	R1
8.0	R1	BFD support within IPsec tunnels	Committed	7x50	8.0	R1
8.0	R1	NAT Support (applicable in MSE as wel)	Committed	7x50	8.0	R1
8.0	R1	DHCP/Radius enhancements - Translation between DHC	Committed	7x50	8.0	R1
8.0	R1	H-Pol support on IOM-3 (ingress/egress H-Pol on ES	Committed	7x50	8.0	R1
8.0	R1	PPPoE for Residential & Business Wholesale (L2TP -	Committed	7x50	8.0	R1
8.0	R1	CFM enhancements for e-pipes (MIP, UP MEP for e-pi	Committed	7x50	8.0	R1
8.0	R1	Enable vi editor in the CLI	Committed	7x50	8.0	R1

8.0	R1	CLI command parameter auto-completion	Committed	7x50	8.0	R1
8.0	R1	Enhance SAA to include Y.1731 and IEEE 802.3ah OAM	Committed	7x50	8.0	R1
8.0	R1	BFD for T-LDP	Committed	7x50	8.0	R1
8.0	R1	BFD Scalability improvements to 300 (BFD functiona	Committed	7x50	8.0	R1
8.0	R1	Fast CFM MEPs	Committed	7x50	8.0	R1
8.0	R1	ETH CFM Scalability improvements	Committed	7x50	8.0	R1
8.0	R1	802.1ag triggered interface down + tie in into 802	Committed	7x50	8.0	R1
8.0	R1	E-LMI w link loss forwarding at vlan level	Committed	7x50	8.0	R1
8.0	R1	CFM Down MEP on Ipipe Ethernet SAP	Committed	7x50	8.0	R1
8.0	R1	Ethernet SAP OAM (AIS, CC, ELMI) mapping: PW (Epip	Committed	7x50	8.0	R1
8.0	R1	802.1ag under group interfaces and R-SAPs	Committed	7x50	8.0	R1
8.0	R1	G.8031 Ethernet tunnels for L2/L3 SC-access (L2 P1	Committed	7x50	8.0	R1
8.0	R1	GR helper for PE-CE protocols (dropped from 7.0)	Committed	7x50	8.0	R1
8.0	R1	Multi-Services ISA for NAT and L2TP LNS (8GB SDRAM	Committed	7x50	8.0	R1
8.0	R1	ISA-Video support for IOM3-XP	Committed	7x50	8.0	R1
8.0	R1	WaveTracker and Full OTU support for Austria2 (10G	Committed	7x50	8.0	R1
8.0	R1	IP/VPN scaling improvements in Sparrow	Committed	7x50	8.0	R1
8.0	R1	VPLS scaling improvements in Sparrow	Committed	7x50	8.0	R1
8.0	R1	VLL/SAP scaling improvements in Sparrow	Committed	7x50	8.0	R1
8.0	R1	Sparrow time stamp up to 128 bytes in packet	Committed	7x50	8.0	R1
8.0	R1	Sparrow support for CPM filters and CPM queuing	Committed	7x50	8.0	R1
8.0	R1	Transient Immunity & BITS redundancy	Committed	7x50	8.0	R1
8.0	R1	ECMP Fate Sharing	Committed	7x50	8.0	R1
8.0	R1	7450/7750 IOM3/IMM mixed-mode chassis support	Committed	7x50	8.0	R1
8.0	R1	48x10/100/1000 copper MDA-XP (Saigon)	Committed	7x50	8.0	R1
8.0	R1	9500 2.0.1 A - - DS3, C-pipe, MPT	Committed	9500 MPR	2.0.1	
8.0	R1	Equipment Management Enhancements For S gateway	Committed	7750 MG	1.0	B2

8.0	R1	LSP-ping & LSP-trace for P2MP RSVP LSP	Committed	7x50	7.0	
8.0	R1	ARP Host	Committed	7x50	7.0	
8.0	R1	Synchronous Ethernet including SSM for 1-port 10GE	Committed	7x50		
8.0	R1	SSM support for Rev-C MDAs and Krakatoas	Committed	7x50		
8.0	R1	Multi-chassis Endpoint for VPLS	Committed	7x50	7.0	
8.0	R1	PPPoE for Residential/Business Wholesale (L2TP/LAC)	Committed	7x50		
8.0	R3	Alarm Propagation	Committed	7750 MG	2.0	
8.0	R3	Service Within a Service (DCR 582118)	Committed			
8.0	R3	Extensions to Drill Down	Committed	7750 MG		
8.0	R3	SAS M ETR Chassis support	Committed	7210	1.1	R8
8.0	R3	Link Loss Forwarding for SAS M	Committed	7210	1.1	R7
8.0	R3	Service MTU check is performed including the SAP encaps	Committed	7210	1.1	R7
8.0	R3	APN/IP pool configuration	Committed	7750 MG	2.0	
8.0	R3	Stats plotting - pre 2.0 and 2.0	Committed	7750 MG	2.0	
8.0	R3	Collapse ShelfPDU/SlotPDU/Ralarm.. to single form	Committed	9471 MME	LM2.0.1	
8.0	R3	Security Permissions (MME/DSC changes)	Committed	9471 MME	LM2.0.1	
8.0	R3	Software Download of MME binary	Committed	9471 MME	LM2.0.1	
8.0	R3	Access Point Network (APN) support on P-GW	Committed	7750 MG	1.0	B2
8.0	R3	PIM support in VRF mode support /OS9000E /Fuji 2	Committed	OmniSwitch		
8.0	R3	VLAN Uplink support for nodes with validated encaps	Committed	7210		
8.0	R3	Work for BT re: router compatibility	Committed	7x50	5.0	R1
8.0	R3	9500 2.2 A support	Committed	9500 MPR	2.2 A	
8.0	R3	9500 2.1 A support	Committed	9500 MPR	2.1 A	
8.0	R3	Cross Launch AWY J-USM Manager	Committed	9500 MPR	AWY 2.1.4	
8.0	R3	dot1q VLAN service for ETSI	Committed	9500 MPR	1.4 E	
8.0	R3	OmniSwitch 6250 AOS 6.6.2 Release Recognition Only	Committed	OmniSwitch	6.6.2	
8.0	R3	UDP Relay - DHCP Option Port 82 and Snooping	Committed	OmniSwitch	6.4.2	
8.0	R3	Cisco PVST	Committed	7210	2.0	R1
8.0	R3	LACP tunneling in VLL service (SAS M/X)	Committed	7210	2.0	R1
8.0	R3	Egress Queue Statistics	Committed	7210	2.0	R1
8.0	R3	Graceful Restart of OSPF / ISIS	Committed	7210	2.0	R1

8.0	R3	9500 1.4E	Committed	9500 MPR	1.4 E	
8.0	R3	DSC 5780 2.0 R1	Committed	5780 DSC	2.0	R1
8.0	R3	LI for LAC sessions	Committed	7x50	8.0	R4
8.0	R3	LI mirroring via RADIUS based on calling station ID or accounting session ID	Committed	7x50	8.0	R4
8.0	R3	Calling-station-id attribute must be configurable to being the remote-id	Committed	7x50	8.0	R1
8.0	R3	SLA override without sla-profile change (policer)	Committed	7x50	8.0	R4
8.0	R3	LTS in LNS	Committed	7x50	8.0	R4
8.0	R3	Default DNS in VRF for framed IP support (PPPeE)	Committed	7x50	8.0	R1
8.0	R3	Missing attributes in acct msges (acct-dely-time, acct-auth, nas-port)	Committed	7x50	8.0	R1
8.0	R3	Client MAC address in the calling station ID	Committed	7x50	8.0	R1
8.0	R3	Configurable NAS-PORT-TYPE	Committed	7x50	8.0	R1
8.0	R3	LAC in VRF (support of L2TP tunnels within VRF)	Committed	7x50	8.0	R4
8.0	R3	CLI for broadcast for IP interface (NE stretch goa	Committed	7210	2.0	R3
8.0	R3	QinQ Ethertype configuration per port - max of 4 v	Committed	7210	2.0	R1
8.0	R3	IGMPv3 support	Committed	7210	2.0	R1
8.0	R3	IES	Committed	7705	3.0	R2
8.0	R3	1:1 with FRR detour LSPs	Committed	7705	3.0	R1
8.0	R3	Ingress SAP Statistics	Committed	7210	1.1	R7
8.0	R3	Egress SAP Statistics	Committed	7210	2.0	R1
8.0	R3	Full AOS 6.4.2 support for OS 6850, 6855, 6400, 90	Committed	OmniSwitch	6.4.2	
8.0	R3	OmniSwitch 6855-U24X / part of AOS 6.4.2	Committed	OmniSwitch	6.4.2	
8.0	R3	6 Port E&M Card	Committed	7705	3.0	R2
8.0	R3	TDM support on 2p OC3 Card	Committed	7705	3.0	R1
8.0	R3	1+1 APS on network ports - OC3 PoS Card	Committed	7705	3.0	R2
8.0	R3	7210 2.0 R1 RFEs	Committed	7210	2.0	R1
8.0	R3	IP interface in a VPLS service (ping only)	Committed	7210	2.0	R1
8.0	R3	Dry contact relay inputs	Committed	7210	2.0	R1
8.0	R3	Y.1731 (SW implementation of time stamping)	Committed	7210	2.0	R1
8.0	R3	SFP diagnostics and monitoring	Committed	7210	2.0	R1
8.0	R3	Mesh SDPs in VPLS	Committed	7210	2.0	R1

8.0	R3	TE Graceful Restart of RSVP on MPLS uplinks / MPLS	Committed	7210	2.0	R1
8.0	R3	BGP VPLS	Committed	7x50	8.0	R1
8.0	R3	Fallback in RADIUS authentication (including msaps, RFE 65775)	Committed	7x50	8.0	R1
8.0	R3	Redundant (Active/Stand-by) PW for mirror/LI services	Committed	7x50	8.0	R1
8.0	R5	IP packet reassembly MDA and S1u/S5 ref pt support	Committed	7750 MG	2.0	R5
8.0	R5	GGSN - Mobility map adjustments	Committed	7750 MG	2.0	R1
8.0	R5	APN - adjustments for last MG mib import	Committed	7750 MG	2.0	R1
8.0	R5	Line Timing of Ethernet Ports Using SynchE	Committed	7210	2.0	R3
8.0	R5	Out of Band management (SAS M)	Committed	7210	2.0	R2
8.0	R5	PM Analysis of multiple pts on different NEs	Committed	9500 MPR		
8.0	R5	Radio Controls (Lpbk Cfg on ports)	Committed	9500 MPR		
8.0	R5	Y.1731 (SW implementation of time stamping) - SAS-M	Committed	7210	2.0	R3
8.0	R5	Latitude Longitude Based Node Location	Committed	9500 MPR		
8.0	R5	Analog Performance Management	Committed	9500 MPR		
8.0	R5	Timing Quality	Committed	7705	3.0	R1
8.0	R5	Y.1731 Dual-Ended Loss Measurement	Committed	7705	3.0	R1
8.0	R5	9500 MPT-sa Stand Alone Support (as GNE with x-launch)	Committed	9500 MPR	1.0	
8.0	R5	NAT Related Dynamic Objects DCR 577946	Committed	7x50		
8.0	R5	MVRP	Committed	OmniSwitch	6.4.3 / 6.6.2	
8.0	R5	L2 Protocol Mac Tunneling	Committed	OmniSwitch	6.6.2	
8.0	R5	802.ag & Y1731 Hybrid Configuration	Committed	OmniSwitch	6.4.3 / 6.6.2	
8.0	R5	Appdex Threshold Fields	Committed	7x50		
8.0	R5	TCP/UDP Application Flag	Committed	7x50		
8.0	R5	5620 SAM supported devices and RAN Releases (eNB models)	Committed	9412 eNodeB	LA3.0	
8.0	R5	5620 SAM Supervision (Demo only)	Committed	9412 eNodeB	LA3.0	
8.0	R5	5620 SAM eNodeB Call Trace Support	Committed	9412 eNodeB	LA3.0	
8.0	R5	5620 SAM eNodeB Self Configuration	Committed	9412 eNodeB	LA3.0	
8.0	R5	5620 SAM eNodeB Configuration Management	Committed	9412 eNodeB	LA3.0	

8.0	R5	5620 SAM eNodeB Performance Management	Committed	9412 eNodeB	LA3.0	
8.0	R5	5620 SAM eNodeB State Management	Committed	9412 eNodeB	LA3.0	
8.0	R5	5620 SAM eNodeB Element Management	Committed	9412 eNodeB	LA3.0	
8.0	R5	5620 SAM eNodeB Security Extension	Committed	9412 eNodeB	LA3.0	
8.0	R5	Support for the X.21 panel with the Serial Data Interface Card	Committed	7705	3.0	R3
8.0	R5	ISA-AA Fairness between SAPs in congestion state	Committed	7x50	8.0	R4
8.0	R5	Support for SDP binds in per-service-hashing-enabled services	Committed	7x50	8.0	R4
8.0	R5	QoS burst management	Committed	7x50	8.0	R4
8.0	R5	PBB support for RFC3107 BGP label for L2 services	Committed	7x50	8.0	R4
8.0	R5	ISSU Support	Committed	7x50	8.0	R4
8.0	R5	IPv6 local user database	Committed	7x50	8.0	R4
8.0	R5	Video Quality Monitoring (VQM)	Committed	7x50	8.0	R4
8.0	R5	Dual video stream	Committed	7x50	8.0	R4
8.0	R5	Dynamic computation of FCC duration and bandwidth	Committed	7x50	8.0	R4
8.0	R5	Audio reordering for FCC	Committed	7x50	8.0	R4
8.0	R5	GGSN (integrate into new RP and billing config into Bearer Management)	Committed	7750 MG	2.0	R1
8.0	R5	1+1 MG-ISM redundancy (S-GW only)	Committed	7750 MG	2.0	R1
8.0	R5	File based KPI/KCI (3GPP format) for SGW	Committed	7750 MG	2.0	R1
8.0	R5	L2TP LNS feaures deferred from 8.0 R1 (LNS IPv6, LNS over MPLS)	Committed	7x50	8.0	R6
8.0	R5	OC-768 DWDM IMM (Japan-based)	Committed	7x50	8.0	R5
8.0	R5	802.1ag - CFM clause 8.1	Committed	OmniSwitch	6.4.3 / 6.6.2	
8.0	R5	802.3ah - UNI Loopback	Committed	OmniSwitch	6.4.3 / 6.6.2	
8.0	R5	SLA / SAA stats / IP & ethernet / L2	Committed	OmniSwitch	6.4.3 / 6.6.2	
8.0	R5	Ports / Queue Stats for Tx / Disc pkts per queue	Committed	OmniSwitch	6.4.3 / 6.6.2	
8.0	R5	Egress Filtering/Adv Policy Cond & Actions / Rule Stats	Committed	OmniSwitch	6.4.3 / 6.6.2	
8.0	R5	Ingress Filtering/Adv Policy Cond & Actions / Rule Stats	Committed	OmniSwitch	6.4.3 / 6.6.2	
8.0	R5	IPM VLAN support	Committed	OmniSwitch	6.6.2	
8.0	R5	Y1731 - Fault Mgmt & Perf Reporting	Committed	OmniSwitch	6.4.3 / 6.6.2	
8.0	R5	9500 AWY 2.1.5 GNE support	Committed	9500 MPR	AWY 2.1.5	

8.0	R5	9500 MSS-1c Compact Chassis support (as GNE with NETO x-launch)	Committed	9500 MPR	2.2 E	
8.0	R5	9500 1.2 A support (complete support)	Committed	9500 MPR	1.2 A	
8.0	R5	9500 1.2 A support (eqmt mgmt only)	Committed	9500 MPR	1.2 A	
8.0	R5	350 channels per ISA / 700 channels per system spr	Committed	7x50	8.0	R4
8.0	R5	Port Segregation Enhancements	Committed	9500 MPR		
8.0	R5	LDP-Configurable transport address	Committed	7x50	8.0	R4
8.0	R5	Queue Parent Weight Override	Committed	7x50	8.0	R4
8.0	R5	ATM Aware QoS for Broadband Network Gateway	Committed	7x50	8.0	R4
8.0	R5	Support disconnected subscribers	Committed	7x50		
8.0	R5	1+1 U-APS for Sicily	Committed	7x50	8.0	R4
8.0	R5	Aux alarm card	Committed	7705	3.0	R3
8.0	R5	ESM: Support mid-session changes for DHCP/PPOE	Committed	7x50		
8.0	R5	H-Pol support on Sparrow	Committed	7x50	8.0	R4
8.0	R5	AIS Generation	Committed	7210	2.0	R3
8.0	R5	Out of Band management (SAS E)	Committed	7210	2.0	R3
8.0	R5	Ipipe divert to AA	Committed	7x50	8.0	R4
8.0	R5	Subscriber attribute override of app-profile value	Committed	7x50	8.0	R4
8.0	R5	NAT static port forwarding	Committed	7x50	8.0	R4
8.0	R5	IPoE user name option 61 and option 60	Committed	7x50	8.0	R4
8.0	R5	Multiple RADIUS authentication based on domain nam	Committed	7x50	8.0	R4
8.0	R5	Remove/add/modify domain names in RADIUS authentic	Committed	7x50	8.0	R4
8.0	R5	Subscriber multicast	Committed	7x50	8.0	R4
8.0	R5	CFM frame-rate limiting support	Committed	7x50	8.0	R5
8.0	R5	128K labeled BGP routes	Committed	7x50	8.0	R5
8.0	R5	RFC3107 BGP-to-LDP label stitching	Committed	7x50	8.0	R4
8.0	R5	mVPN scaling increase (to 500)	Committed	7x50	8.0	R4
8.0	R5	G8032 ring protocol	Committed	7x50	8.0	R4
8.0	R5	H-Pol support on IOM-3 for all non-subscriber over	Committed	7x50	8.0	R4
8.0	R5	SDIC multi-drop data bridge	Committed	7705	3.0	R3
8.0	R5	1+1 U-APS for MLPPP network port and ASAP	Committed	7x50	8.0	R5
8.0	R5	Magma IMM (1-port 100GE)	Committed	7x50	8.0	R5
8.0	R5	FRR LSP Scaling	Committed	7x50	8.0	R4

8.0	R5	LSP Automatic Bandwidth Adjustment	Committed	7x50	8.0	R4
8.0	R5	Enable ESMC on WAN ports	Committed	7x50	8.0	R4
8.0	R5	9500 2.1 E	Committed	9500 MPR	2.1 E	
8.0	R5	LDP support (NE stretch for R1)	Committed	7210	2.0	R3
8.0	R5	T1 / E1 MDA for CES	Committed	7210	2.0	R3
8.0	R5	Dynamic port buffer allocation (Named pools) on La	Committed	7x50	8.0	R1
8.0	R5	Soft reset support for IOM3/IMM and Ethernet MDAs	Committed	7x50	8.0	R4
8.0	R5	Magma IMM (12-port 10GE)	Committed	7x50	8.0	R5
8.0	R5	IPv6 Support on Ipipe VLL Service	Committed	7x50	8.0	R5
8.0	R5	Uni-directional APS service parity for Sonet/SDH channelized and non-channelized	Committed	7x50	8.0	R4
8.0	R5	SSM/ESMC - comprehensive	Committed	7x50	8.0	R4
8.0	R5	7750 SR1/c4 combo product (Sicily)	Committed	7x50	8.0	R4
8.0	R5	Single chassis APS support for MLPPP network ports	Committed	7x50	8.0	R5
8.0	R5	HSDPA offload fallback solution: pw-redundancy wit	Committed	7x50	8.0	R4
8.0	R5	multi-homing for L3 services	Committed	7x50	8.0	R5
8.0	R5	VPLS Multi-homing using BGP AD	Committed	7x50	8.0	R1
8.0	R5	Network Domain Queue Optimisation	Committed	7x50	8.0	R1
8.0	R5	Routed VPLS on IOM3 with unicast support	Committed	7x50	8.0	R4
8.0	R5	RSVP P2MP I-PMSI using dynamic template	Committed	7x50	8.0	R5
8.0	R5	Multicast P2MP LDP (mLDP) for GRT	Committed	7x50	8.0	R4
8.0	R5	FCC/RET Statistics - ALU SQM MIB Additions	Committed	7x50	8.0	R1
8.0	R5	RET Payload Format	Committed	7x50	8.0	R1
8.0	R5	FCC Hybrid mode (Burst + Dent)	Committed	7x50	8.0	R1
8.0	R5	Prioritization Mechanism for FCC vs. RET	Committed	7x50	8.0	R1
8.0	R5	Time & Volume based accounting with local policy o	Committed	7x50	8.0	R4
8.0	R5	Diameter	Committed	7x50	8.0	R4
8.0	R5	IPv6 subscriber management for routed CO (DHCP, PP	Committed	7x50	8.0	R1
8.0	R5	BNG RADIUS - Over-ride SLA, filters, etc	Committed	7x50	8.0	R1
8.0	R5	TPSDA scaling improvements in Sparrow (subscribers	Committed	7x50	8.0	R4
8.0	R5	BITS-out support:	Committed	7x50	8.0	R4

8.0	R5	Routed subscriber with PPP and DHCP support for BG	Committed	7x50		
8.0	R5	Failover support for DHCP server	Committed	7x50	7.0	
8.0	R5	New SAS-X Platform (Includes XGE, HQOS)	Stretch	7210	2.0	R4
8.0	R6	Short keep-alive time for limited number of PPPoE sessions	Committed	7x50	8.0	R6
8.0	R7	support for sap-ingress policy with num-qos-resources set to arbitrary values	Candidate	7210	3.0	R1
8.0	R7	Supervision step 2	Candidate		LA3.0	
8.0	R7	Support of eNB PM catch-up	Candidate	9412 eNodeB	LA3.0	
8.0	R7	Support of eNB TLA3.0 and of eNB LA4.0.0	Candidate	9412 eNodeB	LA4.0.0	
8.0	R7	MME / SAM trap destination config	Candidate	9471 MME	LM2.1	
8.0	R7	MME MDS Service Split	Candidate	9471 MME	LM2.1	
8.0	R7	MME Gr and luPS Interface	Candidate	9471 MME	LM3.0	
8.0	R7	MME MPH Service	Candidate	9471 MME	LM3.0	
8.0	R7	MME LM3.0 General Support	Candidate	9471 MME	LM3.0	
8.0	R7	5620 SAM Dimensioning and KPI for LE3.0	Candidate	9412 eNodeB	LA3.0	
8.0	R7	5620 SAM migration from XMS LA2.0 (for mgt of eNB LA2.0)	Candidate	9412 eNodeB	LA3.0	
8.0	R7	Timezone Management	Candidate		LA3.0	
8.0	R7	eNodeB Licensing	Candidate	9412 eNodeB	LA3.0	
8.0	R7	MME Multishelf Support	Candidate	9471 MME	LM3.0	
8.0	R7	QinQ on Access	Candidate	7210	3.0	R1
8.0	R7	2x10G MDA support	Candidate	7210	3.0	R1
8.0	R7	Link Layer Discovery (LLDP or 802.1ab)	Candidate	7210	3.0	R1
8.0	R7	BFD for L3 protocols	Candidate	7210	3.0	R1
8.0	R7	QinQ enet config per port - max 4 values	Candidate	7210	3.0	R1
8.0	R7	Support CESoP for CES MDA	Candidate	7210	3.0	R1
8.0	R7	Ethernet APS (G.8031, G.8032) and Ring Protection	Candidate	7210	3.0	R1
8.0	R7	FRR Facility Backup with PHP	Candidate	7210	3.0	R1
8.0	R7	Dry Contacts	Committed	1830		
8.0	R7	Backup and Restore	Committed	1830		
8.0	R7	Topology and Logical View - generated by WebUI	Committed	1830		
8.0	R7	SAM-O	Committed	1830		
8.0	R7	Regen Service	Committed	1830		

8.0	R7	Tag Alarms as service or non-service affecting	Committed	1830		
8.0	R7	Routing Config	Committed	1830		
8.0	R7	Port - Timeslot Assignment	Committed	1830		
8.0	R7	GUI Services	Committed	1830		
8.0	R7	Protected and Unprotected Services	Committed	1830		
8.0	R7	Wavelength Services	Committed	1830		
8.0	R7	APS Group	Committed	1830		
8.0	R7	Span of Control	Committed	1830		
8.0	R7	Topology Map	Committed	1830		
8.0	R7	9500 2.2.1 A Support at 2.2 A level	Committed	9500 MPR	2.2.1 A	
8.0	R7	5620 SAM eNodeB Auto Neighbor Relation	Committed	9412 eNodeB	LA3.0	
8.0	R7	5620 SAM eNodeB Physical Cell ID	Committed	9412 eNodeB	LA3.0	
8.0	R7	Photonics management extension	Committed	1830		
8.0	R7	Integration of Wavelength Tracker functions	Committed	1830		
8.0	R7	Basic FCAPS mgmt of PSS-32 R2.5	Committed	1830	PSS-32	2.5
8.0	R7	Basic FCAPS mgmt of 1830 PSS-16 R2.5 End office shelf	Committed	1830	PSS-16	2.5
8.0	R7	Basic FCAPS mgmt of 1830 PSS-1 AHP 1.0	Committed	1830	PSS-1 AHP	1.0
8.0	R7	Basic FCAPS mgmt of 1830 PSS-1 MD4H Release 1.5	Committed	1830	PSS-1 MD4H	1.5
8.0	R7	Basic FCAPS mgmt of 1830 PSS 1 GBEH Release 2.5	Committed	1830	PSS-1 GBEH	2.5
NMS Features						
8.0	R1	Change timing of timestamps on alarms	Committed			
8.0	R1	Map Enhancements	Committed			
8.0	R1	LTE NBI Support	Committed			
8.0	R1	Intel Platform Support	Committed			
8.0	R1	Live view of network in Map, Tree, Lists	Committed			
8.0	R1	SANE integration	Committed			
8.0	R1	Inband/OutOfBand Mgt Enhancements	Committed			
8.0	R1	Increase Tasks supported by Task Mgr	Committed			
8.0	R1	Security Enhancements	Committed			
8.0	R1	Dynamic Capability Support for One Build	Committed			
8.0	R1	Wavetracker support via SAM-O	Committed			
8.0	R1	Service CAC	Committed			

8.0	R1	IPv6	Committed			
8.0	R1	IPMP between SAM Components	Committed			
8.0	R1	Policy Distribution Scaling	Committed			
8.0	R1	Licensing	Committed			
8.0	R1	Select Multiple NE Types for a Script	Committed			
8.0	R3	8.0 R3 Licensing	Committed			
8.0	R3	Service Within a Service (DCR 582118)	Committed			
8.0	R3	LTE User Bearer Management	Committed			
8.0	R5	Parameter Guide Automation	Candidate			
8.0	R5	DB Support	Committed			
8.0	R5	8.0 R5 Licensing	Committed			
8.0	R5	Capture final AA stats	Committed			
8.0	R5	Tunnel Selection Enhancements	Committed			
8.0	R5	VPRN Audit/Topology on RT	Committed			
8.0	R5	Service Enhancements	Committed			
8.0	R5	Switch Mode of policy from list	Committed			
8.0	R5	Configurable labels for LLDP	Committed			
8.0	R5	5620 SAM 3GPP Compliance for SAM-O	Committed		LA3.0	
8.0	R5	Detaching Windows	Committed			
8.0	R5	Scalability Testing	Committed			
8.0	R5	Ethernet CAC	Committed			
8.0	R5	Resource Overlay and Plug-In	Committed			
8.0	R5	Format and Range Support for Policies	Committed			
8.0	R5	Tools CLI commands available for SAM	Committed			
8.0	R5	Channelized card configuraton	Committed			
8.0	R5	Per node Trap rate reporting and trap log view	Committed			
8.0	R5	Computed Values for Stats	Committed			
8.0	R5	LLDP support (GNE)	Committed			
8.0	R5	Hardening of GNE profile	Committed			
8.0	R5	Alarm History Improvements - add missing fields	Committed			
8.0	R5	Alarm History Improvements - open from current ala	Committed			
8.0	R5	Alarm Correlation - auto manual clearing of correl	Committed			
8.0	R5	Auto-Provision Enhancements	Committed			
8.0	R5	XML U/S - Modify Templates incl def queries / curr	Committed			

8.0	R5	XML U/S - Unify S/T GUI Builder	Committed			
8.0	R5	SAP Copy/Move Support for ATM	Committed			
8.0	R5	Service && Composite-Service OAM	Committed			
8.0	R5	History of NE backups on File System	Committed			
8.0	R5	Global AA Definitions	Committed			
8.0	R5	OAM Results on Map (including CFM)	Committed			
8.0	R5	SAN Storage	Committed			
8.0	R7	Application Assurance Data via SAM-O	Candidate			
8.0	R7	3GPP compliance for SAM-O step 2 (CNBI R3.0.1)	Candidate		LA3.0	
8.0	R7	enabler for NPO QoS Alerts display as alarms in SAM	Candidate		LA3.0	
8.0	R7	Convert key generation to ASLM for SAM	Candidate			
8.0	R7	8.0 R7 Licensing	Candidate			
8.0	R7	Timezone Management	Candidate		LA3.0	

Table 1: 8.0 Feature Planning

1. SCALABILITY

Scalability Targets

Aggressive network growth targets, entry of 5620 SAM into new markets, latency, and increased product functionality and usage are driving capacity requirements upward. Changes to the 5620 SAM scaling philosophy and practices are required to be successful. All architectural scale and performance changes shall be put into the product in the R1 release and early in the development cycle.

Table 2 below shows the scale achieved in 5620 SAM release 8.0 R5. The target scale listed in 8.0R7 column are planned targets for 5620 SAM release 8.0 R7. Commitment to achieve increased targets is made once tested.

Note that:

- These limits require particular hardware specifications and specific deployment architectures.
- Scale limits for network elements including GNEs, 7705s, and 7210s assume a maximum sustained trap rate of 80 traps/second.

The following table represents the scalability targets for Release 8.0.

Criteria	8.0 R5	8.0 R7
Maximum network elements (excluding GNE)	7,000	12,000
Maximum number of GNEs (assumes 10 interfaces per)	15,000	18,000
Combined GNE/network elements (assumes 10	18,000/3000	18,000/3000

interfaces per GNE max)		
Combined network elements/GNE (assumes 10 interfaces per GNE max)	12,000/2,000	12,000/2,000
Maximum number of managed MDAs containing:	20,000	25,000
Max 7250 network elements	2,500 (= 5,000 MDAs)	2,500 (= 5,000 MDAs)
Max 7705 network elements	12,000 (= 7,000 MDAs)	12,000 (= 7,000 MDAs)
Max 9500 network elements	5,000 (=5,000 MDAs)	5,000 (=5,000 MDAs)
Max 7210 network elements	5,000 (= 5,000 MDAs)	5,000 (= 5,000 MDAs)
Max OMNISwitch 6000 series (1 MDA equivalent to 1 chassis)	5,000	5,000
Max OMNISwitch 9000 series (1 MDA equivalent to 1 NI)	1,000	1,000
Maximum number of SAPs	5,000,000	6,000,000
Maximum number of Services	1,000,000	2 Million
Maximum number of LSPs	50,000	50,000
Concurrent Clients		
Max OSS Clients [HTTP, JMS1]	30	30
Max GUI Clients	150	150
OAM Tests (10 minute interval)		
Standard Tests (not simultaneous with Lightweight)	6,000	6,000
Lightweight Tests (not simultaneous with Standard)		
Accounting Based Tests	50,000	50,000
Statistics (15 minute interval)		
Accounting Statistics	10,000,000	10,000,000
Performance Statistics	500,000	500,000
Combined Accounting/Performance	10,000,000/500,000	10,000,000/500,000
Alarms		
Outstanding Alarms	50,000	50,000
Alarm History (assumes 50,000 alarms per day)	One month	One month

Table 2: Scaling Commitments & Targets for Release 8.0

Performance Targets

The following table represents the performance targets for 5620 SAM R8.0. Factors that may result in fluctuation of these targets include:

- SAM Server and SAM Database hardware platforms (faster platforms switch faster)

- Network Activity
- User/OSS Activity
- DB activity (i.e. database backups)
- Network size
- Latency

Performance Item Description	8.0 Performance Targets
5620 SAM Client GUI Performance	
Time to launch a 5620 SAM Client GUI	~30 seconds
Time to launch a 5620 SAM Client GUI configuration form	~2 seconds
Time to save a 5620 SAM Client GUI configuration form	~2 seconds
5620 SAM Server Performance	
Time to restart the 5620 SAM Server when managing the maximum number of devices	~10 minutes
Estimated time to resynchronize one new router in domain	<20 minutes (subject to size of new router)
SAM DB Backup (without stats)	Up to 60 minutes (subject to network size)
SAM DB Restore	~45 minutes
SAM Server activity switch	<10 minutes
SAM DB switchover (by invoking through the GUI)	<10 minutes
SAM DB failover (manually invoked)	<20 minutes until complete recovery, including SAM Server restart
SAM DB failover (automatic)	<20 minutes until complete recovery, including SAM Server restart
Recovery of standby SAM Database after failover (This assumes a workstation is available and properly configured before the recovery begins)	< 75 minutes
5620 SAM-O Performance	
Number of services created per day by an OSS workflow for VLL Service type	Up to 25K per day (24 hours)
Average time to create 1 VLL service	~3.0 s
Average time to create 1 VPLS service (3 sites, 1 SAP/site)	~4.5 s
Average time to create 1 VPLS service (6 sites, 1 SAP/site, 30 circuits fully meshed)	~10 s
Average time to configure 100 VPLS Service on 3 Sites using one SAP	~16 min
Average time to add 1 IES interface to an existing service	~1.5 s
Average time to create 1 static route on a 7750 SR	~0.6 s
Average time to create 1 MAC ACL filter	~0.8 s
Average time to create 1 GRE SDP	~0.75 s
Average time to create 1 MPLS SDP	~1.0 s
Average time to create 1 MPLS path	~0.8 s
Upgrade Performance	

SAM Client Upgrade	<10 Minutes
SAM Complex Upgrade (Server, Database, Auxiliaries)	<6 Hours
SAM Upgrade Maximum Visibility Outage with SAM Redundant system	<15 minutes

Table 3: R8.0 Performance Targets

2. PLATFORM

A full description of the platform requirements is provided in the *Alcatel-Lucent 5620 SAM R 8.0 Planning Guide*. The following platforms can be present in a 5620 SAM 8.0 deployment:

- 5620 SAM GUI Client or Client Delegate
- 5620 SAM Server (which performs statistics collection)
- 5620 SAM Database
- 5620 SAM Statistics Collectors (Optional)
- 5620 SAM Redundancy (Optional)
- 5620 SAM CallTrace Auxiliary (Optional)

The 5620 SAM Database can be collocated with the SAM Server platform (for Solaris) or Standalone on a separate platform (for Solaris or Windows).

The 5620 SAM Clients run the GUI software that allows operators to use the system.

The 5620 SAM Server platform houses the software that interfaces with clients (GUI and OSS), mediates with the network, and accesses the database.

The 5620 SAM provides the option of redundancy, with the ability to configure a second system, or set of systems, to protect the 5620 SAM Server and 5620 SAM Database for redundancy.

The 5620 SAM also offers the option to install a 5620 SAM Auxiliary Statistics Collector platform, which will retrieve the statistics from the managed network and provide for higher statistics collection capabilities. A second 5620 SAM Auxiliary Statistics Collector platform can be installed to provide redundancy of the statistics collection capability. The supported deployment configurations are described in the *Alcatel-Lucent 5620 SAM 8.0 Planning Guide*.

5620 SAM Component	Windows	Sun SPARC	Sun x86 AMD 64-bit	Sun x86 Intel 64-bit
5620 SAM Server	Networks not exceeding 50 MDAs/3 clients or 30 MDAs/5 clients. Only recommended for lab use.	Networks not exceeding release 7.0 limits.	Supported	Adding Support in 8.0
5620 SAM Database			Supported	Adding Support in 8.0
Collocated Server & Database	Not Supported	Networks not exceeding 675 MDAs, 1000	Networks not exceeding 675 MDAs, 1000	Adding Support in 8.0. Network size restrictions

		GNEs, 5 clients, 1000 elemental STM tests every 10 minutes, 50,000 performance/ 100,000 accounting statistics records every 15 minutes	GNEs, 5 clients, 1000 elemental STM tests every 10 minutes, 50,000 performance/ 100,000 accounting statistics records every 15 minutes	will apply.
5620 SAM Client	Supported	Supported	Supported	Adding Support in 8.0
5620 SAM Client Delegate	Not Supported		Supported	Adding Support in 8.0
5620 SAM Auxiliary (Stats Collector)	Not Supported	Supported	Supported	Adding Support in 8.0
5620 SAM Redundancy (Server & Database)	Not Supported	Supported	Supported	Adding Support in 8.0

Table 4: Platform Support

Support x86 Intel 64-bit (NEBS compliant)

5620 SAM is supported on a platform that is NEBS compliant for customers requiring this. The platform chosen is a Sun x86 Intel 64-bit (x4170). Toward Q3/Q4 of 2010 testing is planned on a similar spec HP platform.

IPMP Support

In SAM 7.0R4 IPMP is supported between the SAM servers and the network elements. In SAM 8.0R1 5620 SAM supports IPMP between all SAM components for additional resiliency. This includes all GUI and OSS clients. Note that 5620 SAM will only support IPMP in the "active-standby" mode. IPMP load sharing is not supported.

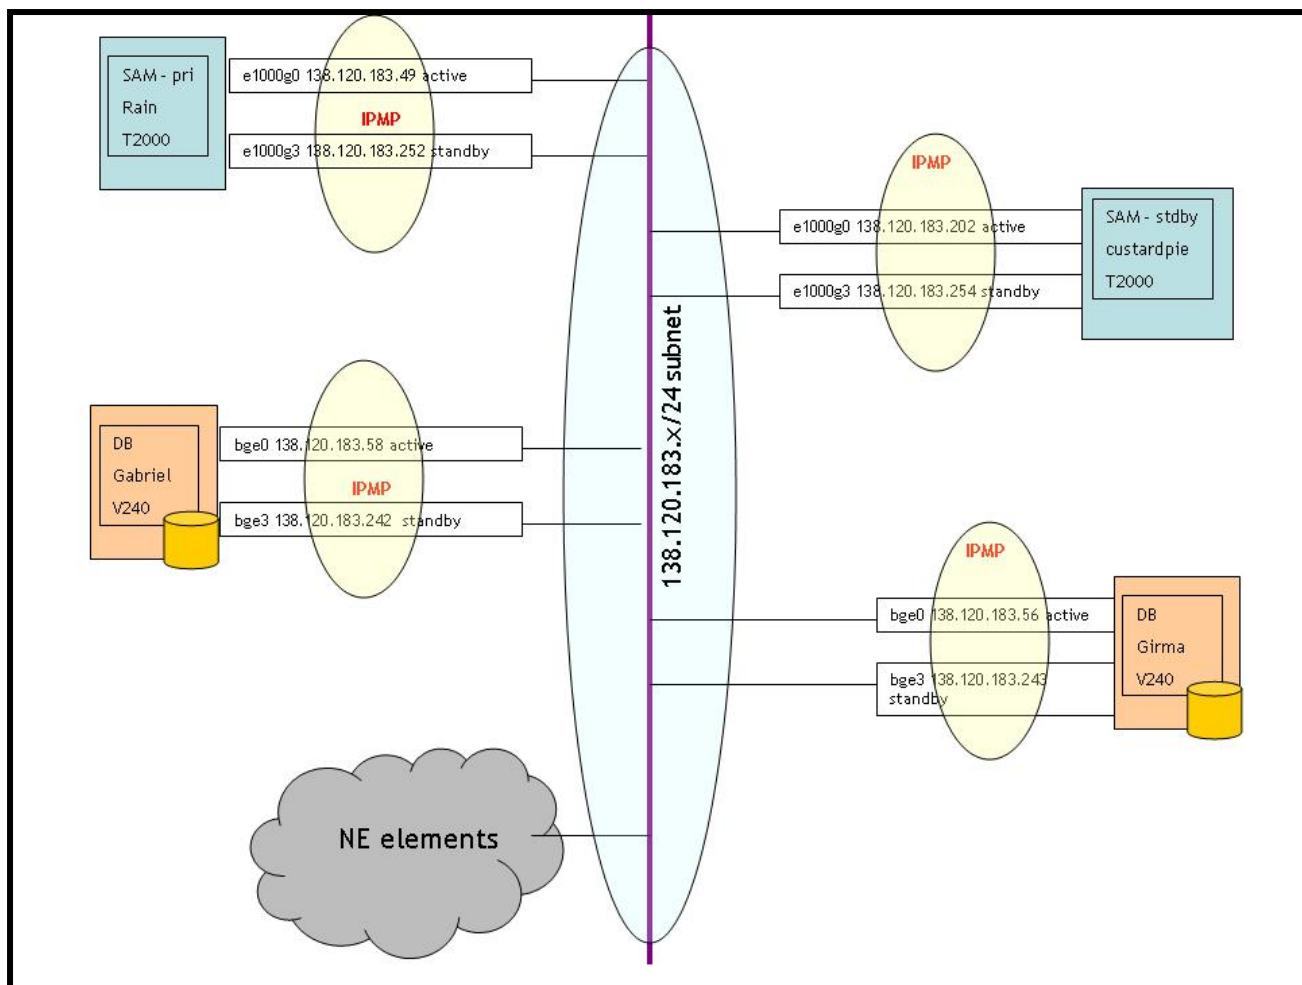


Figure 1: IPMP Support

Inband & OutOfBand Enhancements

Several inconsistencies have been seen between management IP addresses that may be found in different areas of SAM (Network Element, Node Discovery Controller, Management Ping, Discovery Manager, etc). Release 8.0 provides appropriate mediation handling of those IP addresses at discovery time. In previous releases of 5620 SAM, it is inaccurately assumed that the Primary IP Address represents the OOB IP address while the Secondary IP Address represents the IB IP address.

The naming of the GUI labels and the attribute names are changed from Primary IP Address and Secondary IP Address to OOB (Out Of Band) IP Address and IB (In Band) IP addresses respectively.

The GUI has two GUI sections under Management group in the NE configuration form - one section has the label "OOB Management", containing one field "Management Address", and the other has the label "IB Management", containing two fields "System Address" (mandatory) and an optional "L3 Management Interface".

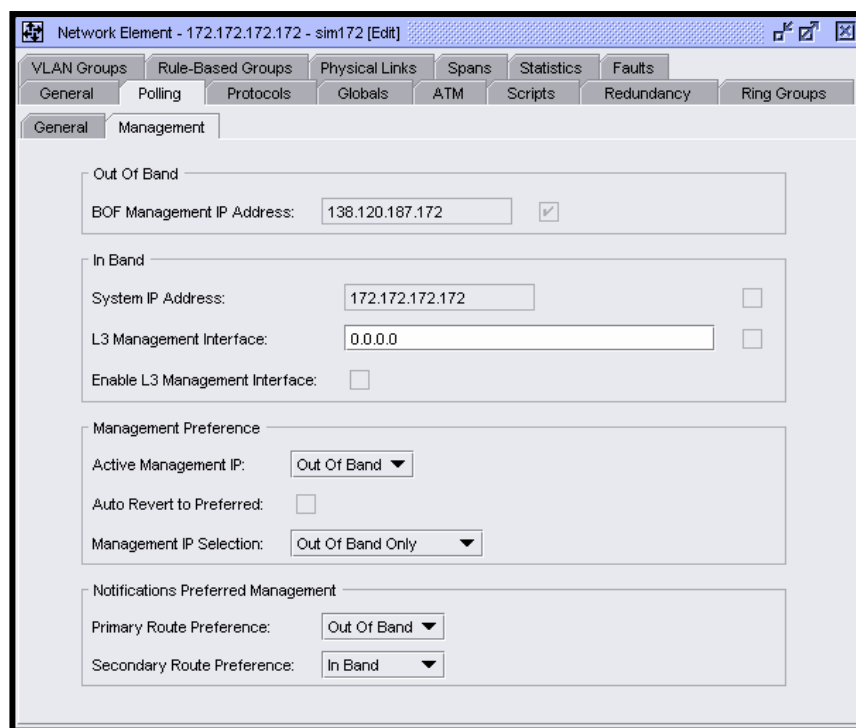


Figure 2: Management Window

Each IP address has a status associated with it.

The OOB IP address is always set to the IP associated with the NE management port (configured through BOF) if one exists and it is retrieved by SAM from the node. If the IP address associated with the NE management port is not configured, or it does not exist, then the returned MIB value is null and the OOB IP address is empty (null/empty in the database, empty for users).

The OOB IP address is not editable for the time being.

If the Management Ethernet port is not configured for the managed node then the status is set to: 'not configured', if the management port is not available for a node the status is set to 'not supported'.

There are two IP addresses defined as IB:

- IB System IP (inBandSystemAddress) is set to the NE System IP address and is not modifiable.
- L3 Management Interface address is set by default to the discovery IP address the operator inputs to manage the NE inband or NE System IP if the discovery IP address is either OOB IP Address or the NE System IP. The "L3 Management Interface" is only filled in if it is different from both the Management and System Addresses (configurable).
- This L3 Management Interface (inBandL3ManagementIf) is editable to allow the operator to change at any time the Network Interface Address used to manage the NE. Note that if the operator changes the Network Interface Address afterwards then the initial value for the discovery IP Address held by this variable is lost.

5620 SAM uses a third address to complete the IB management addressing. This third address is called the network interface address and is set to the discovery IP address if this address is different than OOB address and system address. This address is writeable in situations where an operator knows that there is a third IP address to manage the node but SAM was unable to detect.

Attribute	Modifiable	Explanation
ipAddress	Deprecated	Currently this attribute refers to Discovery IP address. The attribute will be deprecated.
systemAddress	No	This attribute will not change and it will hold the value of the NE System IP used in every FDN
inBandSystemAddress	No	New attribute to hold the value of the system address used for inband management. This could be different from the systemAddress above which is used in FDNs.
outOfBandAddress	No	New attribute to hold the value of the NE Management Ethernet Port. This will replace ipAddress.
inBandL3ManagementIf	Yes	This is a new attribute which refers to the IB L3 Management Interface address.

Table 5: SAM NE Attributes

Support for preferred management address SNMP table for SR

A new feature was added to the SR family of nodes in 7.0 R1 that allows the user to select which management interface to use when sending SNMP events to the outside world (SNMP Traps).

```
A:sim130>config>log# route-preference
```

```
- no route-preference
- route-preference primary {inband|outband} secondary {inband|outband|none}
<primary-preference> : [inband|outband]
<secondary-preference> : [inband|outband|none]
```

Shortly thereafter, the option of configuring this via SNMP has been added to the TIMETRA-LOG-MIB file.

```
tmnxEventPrimaryRoutePref OBJECT-TYPE
```

```
SYNTAX      INTEGER {
                inband (1),
                outband (2)
            }
```

```
tmnxEventSecondaryRoutePref OBJECT-TYPE
```

```
SYNTAX      INTEGER {
                inband (1),
                outband (2),
                none (3)
            }
```

5620 SAM now supports this enhancement. A new group 'Notifications Preferred Management Interface' is added to the Polling tab with the following items:

- 'Primary' with 2 check boxes: inband and outband
- 'Secondary' with 3 check boxes: inband, outband and none

The check boxes are predefined based on the values derived from the NE but the operator has the ability to re-set these values.

As per the 7750 SR User Guide, if the route preference is set to outband, then the source IP address of the SNMP trap is the NE Management port. For inband route preference, the SNMP trap source IP address is the system IP address of the managed 7750 SR.

If 5620 SAM preferred NE management is different from the notification preferred management, a warning alarm will be issued to notify the user of the differences.

This feature is only supported starting in SAM release 8.0 R1 for SR 7750 7.0 R1 nodes or higher. If a node is upgraded from 6.x to 7.0+, this functionality is automatically available. Similarly, if a 7.0 node is managed by a 7.0 5620 SAM and then the 5620 SAM is upgraded to 8.0, the feature is available on that same 7.0 node.

Preferred management address after un-manage

By default, when performing an Unmanage followed by a Manage of a node that has been originally discovered IB, the 5620 SAM always tries to connect using the OOB address first and when no responses are being received from the node, there is a switch to the IB address (Primary). The same behaviour exists for OOB discovered nodes where the user later switches to the IB (Secondary) IP address for management. Following unmanage/manage processes, the information on the preferred management address is lost.

Two new attributes are added to network discovery control: "Last Active Management IP" to hold the value of the last active management IP address to use when re-managing a node as this information is lost during the unmanage. The second attribute is a flag "Use original discovery IP" which gives the operator the choice to select the original discovery IP address used in case of re-managing the node. Re-managing attempts to use the 'Last Active Management IP' unless the flag "Use original discovery IP" is set.

Rename Primary/Secondary address fields to OOB/IB

Currently, it is assumed that IB is Secondary and OOB is Primary. This is not always true. With the 8.0 changes, Primary and Secondary addresses do not exist and the IB/OOB IP Addresses are always consistent:

- OOB IP Address is set to the NE Management Port IP if one exists and is configured, otherwise this is empty.
- IB System IP is always set to the NE System IP. L3 Management Interface is set to either the NE system IP if the operator selects to discover the NE through the System IP or to a L3 Management Interface address of the operator's choice.

Management Ping versus IB and OOB

Currently in SAM for the case where an operator enters the Network Interface in the Discovery Manager to manage a new node, we are faced with the situation where both the Primary and Secondary IP addresses are In Band addresses. However, the Management Ping feature blindly assumes that the Primary Address is of type OOB while SAM knows that it is an IB address; this is shown in the Ping Display Manager window.

With the 8.0 changes, the Ping Destination Tab has the OOB/IB ping destination selection only if OOB/IB IP addresses are available. If an OOB IP Address exists, then in the Management Ping

Destination tab the IP Address is set to the NE Management Port. If an IB Network Interface IP Address is set then Management Ping Destination 'IP Address' will be set to the appropriate IB Network Interface IP address.

IPv6

Enhancements are added in SAM Release 8.0 R1 to allow management of the routers over IPv6. SAM shall continue to support management of routers over IPv4 as before.

The scope of changes includes aspects like discovery, re-sync, deployment, trap handling over IPv6.

The scope of this feature is limited to interaction between the SAM servers and the IPv6 capable routers. The following are the current list of routers that can be managed over IPv6 by SAM:

- 7750 6.1/6.0/7.0
- 7450 7.0
- 7710 6.1/6.0/7.0

It does not address interaction between the SAM servers or with third-party servers that SAM interacts with.

The scope of the IPv6 management support is limited to the following:

- SNMP Mediation for OOB and IB
- Telnet for CLI cut-through
- FTP for Software Upgrades, Software Backups & Accounting Stats
- STM, File Browsing, Node ICMP Ping and Secure HTTP for OSS

The 5620 SAM also provides mixed IP protocol support where some of the routers in the network are managed over IPv4 while others are managed over IPv6.

For management of a router over IPv6, the IPv6 capable router should have IPv6 addresses configured for its management port (OOB) or interfaces used for management (IB), and the SAM servers should also have IPv6 addresses.

Configuring the address

The value of the IPv6 address is populated during installation. The installer prompts the user for the IPv6 address of the SAM server. The IPv6 address is optional.

The IPv6 address of the SAM server is kept in the nms-server.xml config file along with where the current IPv4 address of the SAM server is kept. A new entry added to the server config file to hold the IPv6 address of the SAM server looks like the following:

```
<snmp
    ip="20.1.1.1"
    ipv6="2aa2::55:221:28ff:fe46:ff4b"
    natEnabled="false"
    port="162"
    trapLogId="98" />
```

The router's BOF must be configured correctly. The IPv6 addresses must be added for the BOF, as well as the system interface /L3 Management interface

A:itb_john12# show bof

=====

BOF (Memory)

=====

primary-
image ftp://*:*@20.1.1.1/extra2/images/7750/8_0/B1-6/i386-both.tim

primary
config ftp://*:*@20.1.1.1/extra2/cfgs/john_b/itb_john12/7750/8.0/itb_john12_7750_8_0_IPv6.cfg

address	20.1.169.234/24 active
address	2AA2::1:0:0:A01:A9EA/64 active
static-route	20.1.1.0/24 next-hop 20.1.169.254
static-route	2AA2::55:0:0:0:0/64 next-hop 2AA2::1:0:0:0:FF
autonegotiate	
duplex	full
speed	100
wait	4
persist	on
no li-local-save	
no li-separate	
console-speed	115200

Figure 3: Sample BOF

```
. #-----  
. echo "Router (Network Side) Configuration"  
. #-----  
. router  
.     interface "IPv6-test"  
.         port 1/1/1:111  
.         ipv6  
.         address 4CC4::1/64  
.     exit  
. exit  
. interface "system"  
.     address 30.1.169.234/32  
.     ipv6  
.     address 24AA:C8:710:FFFF::1E01:A9EA/128  
.     exit  
. exit
```

Figure 4: Router Configuration

Discovery

SAM will prevent the creation of an IPv6 discovery rule if the IPv6 address is not configured on the SAM server.

An IPv4 address is also needed for the SAM server at this point in time. It is needed for the case where SAM is managing some routers over IPv4 and some over IPv6 and more so because the communication between the SAM servers is still over IPv4.

There is currently a restriction in the SR routers that support IPv6 management where they require an IPv4 management address to be configured for both OOB and IB for traps to be sent.

Furthermore, SAM will not discover a router if an IPv4 system interface is not configured on the router

SAM still identifies IPv6 managed routers by their IPv4 system interface, and therefore an IPv4 system interface will need to be configured for SAM discovery.

SAM will raise a critical alarm if the SAM user manages a router over IPv6 which does not have the IPv4 address for the management port for the OOB case OR either an IPv4 address on the management port or the system interface for the IB case.

This alarm is also raised if this IPv4 address is removed after discovery from the router.

SAM supports discovery of routers in a sub-net by providing the ability to specify an IPv6 address and prefix length, in a manner similar to what it supports today for IPv4 discovery.

SAM validates the IPv6 address entered in the discovery rule elements as it does with IPv4 addresses today.

- For example multicast, loopback addresses cannot be used in creating discovery rule elements.

When specifying the sub-net for IPv6 discovery the minimum permissible value that can be specified for the prefix is 120

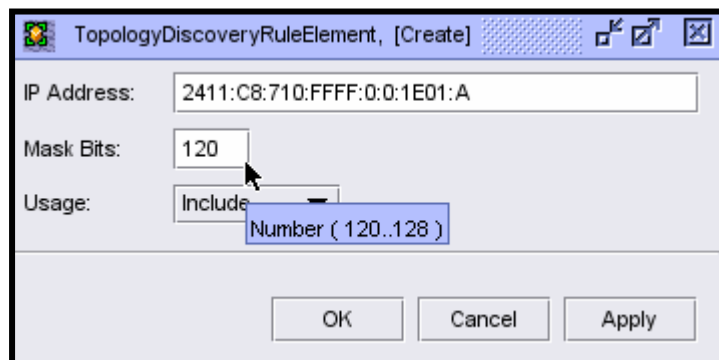


Figure 5: Discovery Rule Element

Determining the address used

During discovery, if SAM discovers that there are both an IPv4 and IPv6 address configured on the router for the system interface, SAM will use the preference specified in the discovery rule to pick the address for identifying the router in SAM, whether it is IPv6 or IPv4.

Once the decision is made on the address to use to identify the router it can NOT be reverted as this IP address is embedded in the FDNs of all equipment objects and can not be changed once discovered. For this reason SAM will raise a critical alarm if the address it used for identifying the router during initial discovery gets removed from the router.

IPv6 System ID/FDN support is not available. Once a router is discovered using IPv6, the system management address/System ID of the router will not be set to be an IPv6 address. SAM will still use the IPv4 system address for the Site ID and FDNs. The OSS will still use the IPv4 system address.

Configuring IPv6 Trap Destinations

If the router is discovered using IPv6, SAM will configure the trap destination on the router to send traps to the IPv6 address of the SAM server.

When configuring SNMP trap destinations in the router for IPv4 management, SAM currently uses the IP address of the SAM server followed by the port number (<ip_address>:<port>) to form the name of the trap destination which is the key of the trap destination table in the router.

There is a restriction in the router that limits the maximum length of the trap destination name string to 28 characters. This length is not sufficient to fit even the IPv6 address of the SAM server.

Since SAM cannot fit the whole IPv6 address in the string, SAM uses the IPv4 address of the SAM server with an extra suffix at the end to form the name of the snmp trap destination for the IPv6 case. There is no change to the format of the trap destination names for the IPv4 case.

Trap Target

This is the name field which is essentially a description/label of the trap-target, and does not imply any limitations on the IPv6 address. For SAM, it was decided that IPv6 trap targets will simply append '-v6' to the management platform's IPv4 address to indicate that this is the IPv6 trap-target (example below).

```
A:itb_john12>config>log>snmp-trap-group# info
-----
          description "5620sam"
          trap-target "20.1.1.1-v6:162" address
2AA2::55:221:28FF:FE46:FF4B snmpv2c notify-community
"RTBqatrap98"
          trap-target "20.1.1.2:162" address 20.1.1.2
snmpv2c notify-community "RTBqatrap98"
```

Figure 6: trap-target

Limitations

- The IPv6 capable SR routers currently do not support SSH over IPv6.
- The SR routers that support IPv6 also need an IPv4 address configured on the management port for it to send traps over IPv6 for OOB management.
- For IB management an IPv4 address is needed either on the management port or the system interface for the router to send traps over IPv6.
- There are some limitations in the SR routers which will limit its functionality if there is no IPv4 address configured in the System/Network Interface.
 - Features like MC APS, MC LAG, MC Ring, MC Sync and a few others expect a peer IPv4 address for it to communicate.

Future Evolution (R5 Candidate)

If the router supports both IPv4 and IPv6 management SAM shall provide the ability to switch the management protocol on the fly once the router is discovered using one protocol version.

SAM Database and Server Security Enhancements

New for installation is ./Oracle10g_PreInstall.sh script prompts for a password for the oracle user.

Only the oracle user now has rwx privileges in the oracle home

DBSNMP account has been dropped as it can be exploited by knowledgeable attackers. When a SAM database is created, the default Oracle user accounts SYS, SYSTEM, SYSMAN, DBSNMP, OUTLN, DIP, MGMT_VIEW are created in addition to the SAM user accounts. These accounts are secured by locking the accounts.

Privileges have been revoked from the PUBLIC role. The EXECUTE permission is removed from RBMS_EXPORT_EXTENSION for PUBLIC, DBMS_LOB for PUBLIC, DBMS_OBFUSCATION_TOOLKIT for PUBLIC.

Additional auditing log files have been added to assist with the increase in audits for monitoring db administrative changes. These are located in: <installdir>/5620sam/oracle10r2/rdbms/audit

Stronger user password verification now exists. The following rules for SYS and SYSTEM passwords are applied:

- The minimum length is 1. The maximum length is 30.
- The password must be different from the user name.
- The password must be different from the reverse of the user name.
- The password only contains letters and digits.
- The first character of the password must be a letter

The following Security Requirement PTSs have been implemented in SAM 8.0:

- PTS 560692 - An attacker could use the credentials to impersonate the "admin" user.
- PTS 560695 - An attacker could read files, including sensitive configuration information
- PTS 560697 - An attacker can obtain the layout of the JBoss directory structure by sending invalid requests.
- PTS 560699 - The Uninstaller files shall not have world writable permission
- PTS 560703 - The 5620 SAM password shall not visible for ps command

Network Element Backups & MME Perf Mgmt Files

NE configs are backed up into the database since SAM 5.0. Customers requested the ability to save to the file system.

SAM 8.0 adds the ability to automatically save a copy of the latest NE config on the filesystem. For a redundant SAM setup, the file will be 'rsync' to the redundant SAM and configured as an option in the installer:

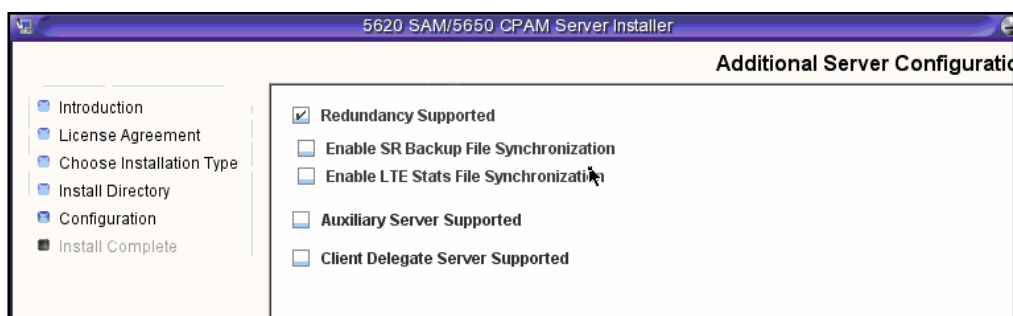


Figure 7: Installer Options

An entry can be changed in the nms-server.xml file if the feature needs to be enabled after installation. The 'latest' config file retrieved from the NE is left on the filesystem (at the indicated location).

```
<srBackup
  srBackupDirectory="<File Dir Location>"
  srBackupSyncEnabled="false"
  srBackupSyncInterval="30" />
```

Rsync needs to be enabled for this to work:

```
enableDbBackupRsync="true"
```

LTE statistics follow the same redundancy for performance management (PM) files from the 9471 MME. These are ftp'd from the MME and stored on the active SAM server file directory.

Enabling the LTE file sync option will synchronize the PM files to the standby. Future other LTE nodes may use a similar PM collection mechanism which is why the option is named as generic LTE and not refer to the 9471 MME directly.

One Build Support - Network Element Upgrades

This feature is a change to node software upgrades. There is now a common software image for the 7x50 and 7710 nodes. It is possible to upgrade the 7450 ESS, 7750 SR, and 7710 SR using one software image.

The Product Name field displays Alcatel-SR/ESS -7XXX to indicate that the software image is a common software image.

All NEs will use the same image.

If the user selects this image and proceeds with the "Upgrade Sites" operation, the Popup list will have all 7750, 7450, 7710s in the network, since this image can be used for any of them.

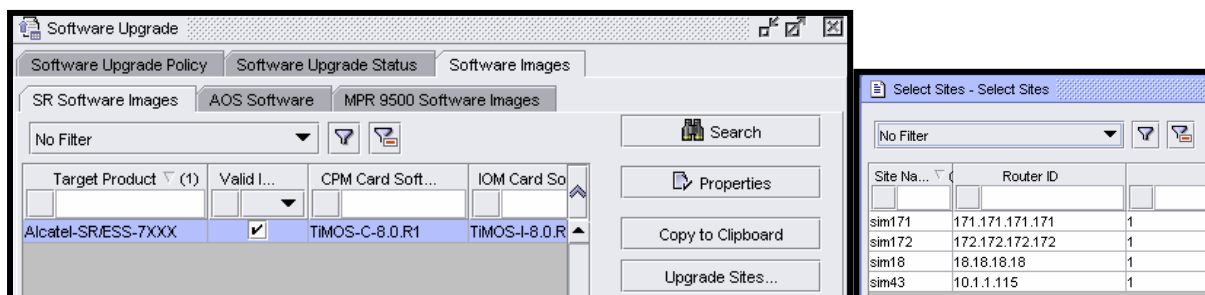


Figure 8: Common Software Image

GUI Framework Enhancements

Reduce JMS events processed by SAM GUI client

Issues have been found in the past caused by flooding the SAM GUI client with JMS events (Attribute Value Change, Create and Delete events). To reduce the event rate, the SAM Server will only send certain events to those SAM GUI clients that are interested in receiving these events.

In previous SAM releases most of the JMS events were sent to each SAM GUI client regardless if the client was interested in receiving the message or not. In this release the SAM GUI client will have to register to receive event notifications. SAM gains performance improvements because JMS will have to process and send less events (significantly decreases bandwidth and memory requirements).

The GUI tells the server by registering/unregistering for events for given FDNs (Fully Distinguished Names) or MO (Managed Object) class names. The server maintains a map of FDN/class names and sends events to GUIs that requested them.

The GUI has a registry of event listeners. Listeners are internal components of GUI e.g. forms, tables, tree nodes, etc.

No manual configuration is required, this is only between the SAM GUI and the SAM server. Other clients (OSS) do not require registration (they work as before)

The GUI now has a counter for those registrations. Whenever the first listener registers for a given FDN or MO class name, it dispatches the registration request to the server. Also, when the last listener unregisters for that FDN or MO class name the framework will dispatch the unregistration request to the server.

GUI clients need events for objects that are being edited, or objects that are listed in a table. For example if the "Services" manager frame is open the client will register for service events, etc. Once the form or table is closed the client does not need to listen to these events.

It is not possible to view the server-side registry.

"Reset" button on Configuration Forms

The "Reset" button on Configuration Forms is removed in this release. The alternative to the Reset button is to Cancel and re-open the configuration form.

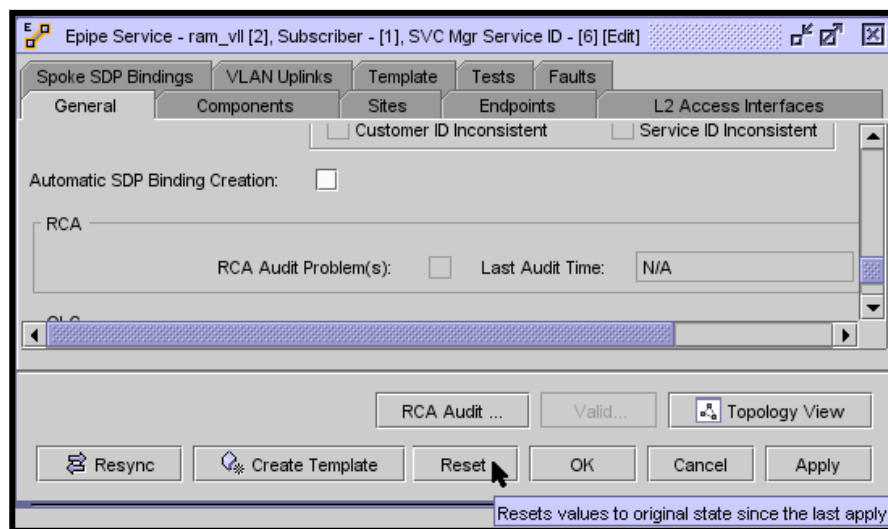


Figure 9: Reset Button

Last Search Time in Table Views

A label displaying the last search time is displayed above the Search button of lists that do not update their contents automatically. The search time is the client time when the Search button was clicked.

This label is displayed on the following tables:

- Browser View
- Tables in configuration forms containing objects indirectly related



Figure 10: Last Search

Changes in Layout of Configuration Forms

Groups can be displayed side-to-side. Example: Administration -> System Information Form. The first diagram shows the 7.0 form. The second diagram shows the same for in 8.0.

A screenshot of the 'System Information' form in Release 7.0. The form has a title bar and three tabs: 'General', 'Auxiliary Servers', and 'Faults'. The 'General' tab is selected. The form contains several sections with input fields and checkboxes. The 'Domain Name' field is set to '5620sam'. The 'Redundancy Enabled' checkbox is checked. The 'Primary Server' section has fields for 'IP Address' (138.120.187.11), 'Host Name' (dino), and 'Status' (Up). The 'Standby Server' section has fields for 'IP Address' (138.120.187.213), 'Host Name' (pool), and 'Status' (Up). The 'Primary Database Server' section has fields for 'Database Name' (samdb), 'Instance Name' (samdb1), 'IP Address' (138.120.187.11), and 'Host Name' (dino). The 'Standby Database Server' section has fields for 'Instance Name' (samdb2), 'IP Address' (138.120.187.213), and 'Host Name' (pool). The 'Redundancy Database State' section has fields for 'Switchover State' (Success) and 'Last Attempted Switchover Time' (2009-11-27 15:34:41).

Figure 11: Rel 7.0 Sys Info

Domain Name: 5620sam

Redundancy Enabled: ☒

Realignment Enabled: ☒

Realignment Status: Unknown

Primary Server

Host Name: 138.120.142.99 [Properties](#)

Preferred DB: N/A

Status: Up

Standby Server

Host Name:

Status: Not Applicable

Primary Database Server

Instance Name: cmettier

IP Address: 138.120.200.223

Host Name: 138.120.200.223

Standby Database Server

Instance Name:

IP Address:

Host Name:

Redundancy Database State

Switchover State: Not Attempted

Last Attempted Switchover Time:

Failover State: Not Attempted

Last Attempted Failover Time:

Standby Re-instantiation State: Not Attempted

Last Attempted Standby Re-instantiation Time:

Figure 12: Rel 8.0 Sys Info

Map Enhancements

Map Filter Functionality

The map filtering functionality introduced in Flat View topology maps in a previous release of SAM has been integrated into all other maps. This allows the user to filter the current view based on a newly defined or selected saved filter.




Figure 13: Filter Controls

The filter will apply when navigating between groups where as only objects that match the filter will be displayed.

The Filter combo box when not expanded indicates the current filtered state with either:

- No Filter for no filter applied,
- Filter Applied if an unsaved filter is applied, or
- The name of a named filter that is applied.

When the Filter combo box is expanded the following entries are available providing control of the filtered state:

- No Filter - represents No Filter applied and clears the applied filter if selected
- Filter Applied - represents an unsaved configured filter and applies it if selected

- <saved filter entry> - individual entries for all available saved filters and applies the filter if selected

When a filter is applied to the map only the objects that are returned by the filter are displayed with the following exceptions:

- Edges will automatically display their endpoints even if a Vertex filter would have excluded them.
- If no Edge filter is included with a Vertex filter only the edges that are between the included Vertex will be displayed, all others will be excluded.

Make Map from Group

The Map Group selector tree has been enhanced to include the following options in the right click context menu for all groups:

- Make Flat Map from Group - creates a new Flat Map for the selected group expanding all contained groups.
- Make Map from Group - creates a new Map for the selected group

If selected group is the same group as the current group and a filter is applied the filter will also be applied to the new Map.

Show Only Selected Option

All maps have been enhanced to include a new Show Only Selected display functionality. This functionality appears as an option in the right click context menu if one or more Vertex are selected on the map.

Selection of this option results in all other Vertex being hidden and any links not connected between selected Vertex also being hidden.

Once this functionality is applied to a map a new Cancel Show Only option appears in right click context maps for Vertex and the map background which turns off the functionality redisplaying all hidden objects.

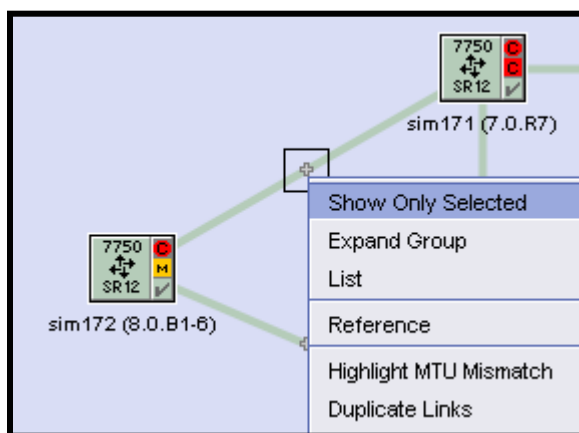


Figure 14: Show only selected

Info Table Usability

The Global Info Table functionality has been enhanced to include the following new display options:

- Mouse Over Display
- Header Display

This functionality only exists for Global Info Tables.

Show on Mouse Over option

The Show on Mouse Over option controls display of the Global Info Tables, if the option is off all Info Tables will be displayed, however if the option is turned on the Info Table will only be displayed if the mouse hovers over a map object for a period of time equivalent to tooltip time.

This option is OFF by default.

The following shows the Mouse Over display option in the Global Info Table pulldown menu:




Figure 15: Show On Mouse-over

Show Header option

The Show Header option toggles display of Info Table headers and attribute names. If this option is off then headers and attribute names are not displayed in the table. Info Tables also support a mouse over override of this option, if the option is off and the mouse hovers over an info table for a period of time equivalent to tooltip time, the headers and attributes will be displayed.

This option is ON by default.

The following shows the Show Header option in the Global Info Table pulldown menu:
The following shows an Info Table with Show Header option ON:



Figure 16: Show Header Option - ON

The following shows an Info Table with Show Header option OFF:

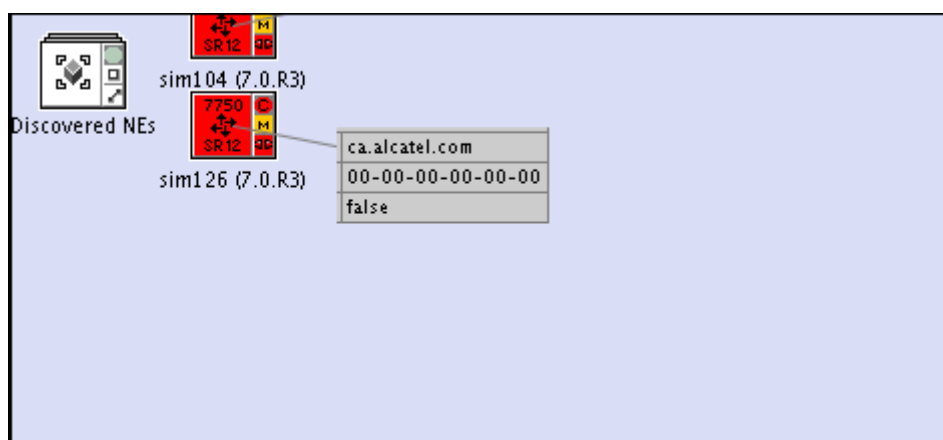


Figure 17: Show Header Option - OFF

500 Elements per Group

The default and maximum number of elements per group has been increased to 500 elements.

Icon Legend

All maps now have an icon legend.

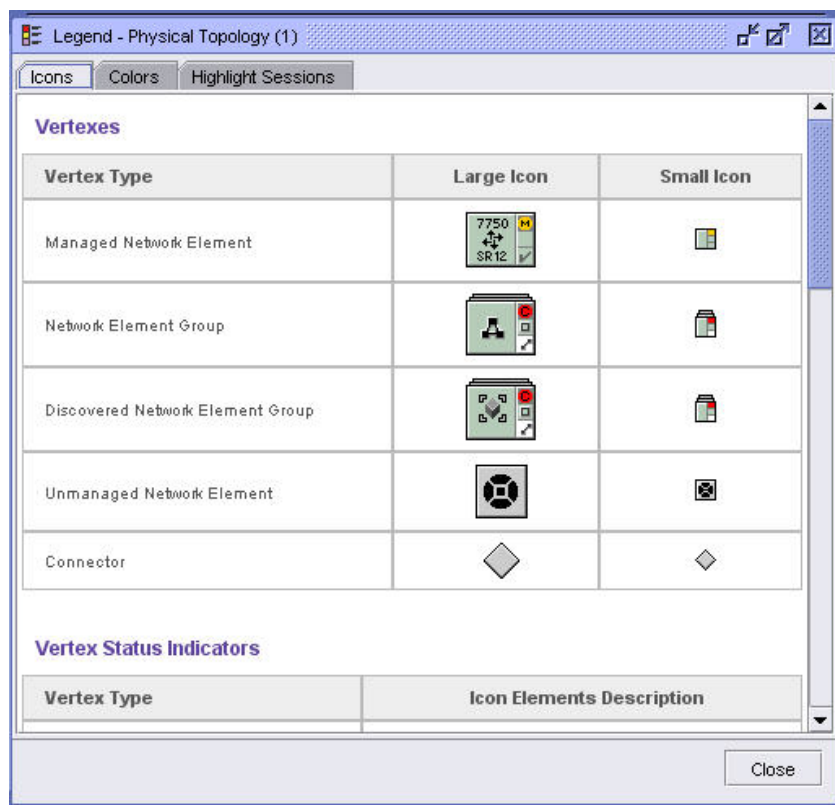


Figure 18: Icon Legend

Find Vertex/Edge Enhancements

The Find Vertex/Edge functionality has been enhanced with the following:

- Label change to: TBD
- List column ordering has been changed to move the following columns to the left making them more accessible: OLC State, Alarm Status, Aggregated Alarm Status and Descriptor Version
- Find list is auto-populating on display
- A Navigate button has been added that will pan the map to the selected object keeping the Find window open. This allows the user to quickly navigate between multiple objects in the find list.

Find by Label functionality

All maps have been enhanced with a Search field that will search through the vertex and edges in the current view and select the ones that have the search string in their label(s). The search takes place once a string is entered in the search field and Enter is hit. All resulting objects will be selected and the map pans to the first one returned.




Figure 19: Search By Label

The search value does not support wild cards.

The search value can be cleared by double clicking on it and hitting the Backspace or Delete key. This results in the entire value being selected and cleared which is standard text field functionality.

Show On Map

A “Show on Map” button has been added that will pan the map to the selected object keeping the Find window open. This allows the user to quickly navigate between multiple objects in the find list.

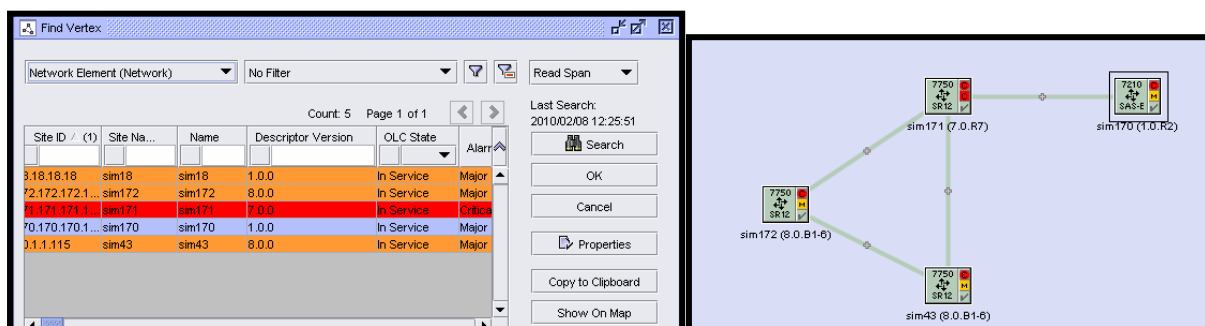


Figure 20: Show On Map

Logical Groups

A new group “Logical Groups” in the Equipment Navigation tree contains all the logical groups.

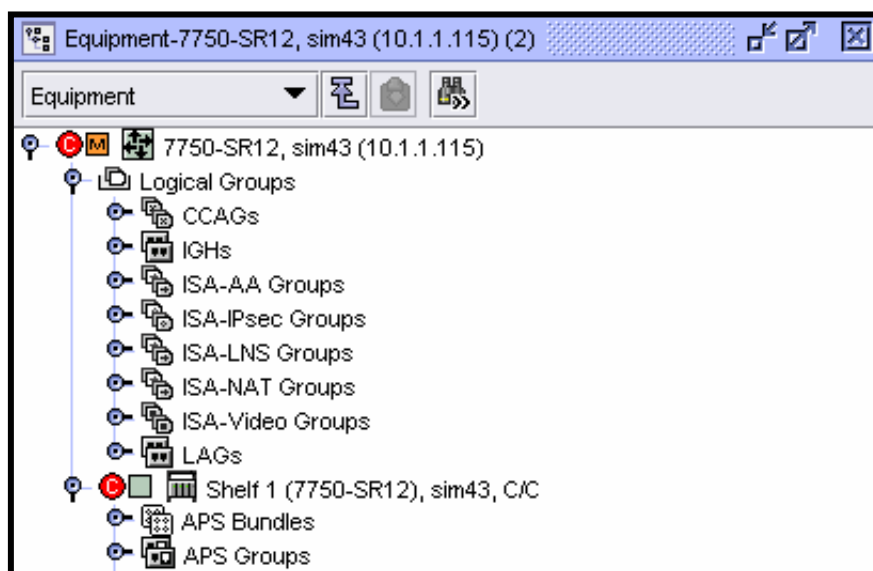


Figure 21: Logical Groups

Keyboard controls

The following keyboard shortcuts have been added:

- **CTRL A Select All** — All maps now support CTRL A keyboard shortcut to select all Vertex on the map.
- **Shift Pan** — All maps now provide the ability to perform a pan by holding down the shift key. Holding down the shift key results in the map temporarily switching over to pan mode allowing panning when both shift key and mouse left button are held during mouse movement.
- **Next/Previous Selected** — All maps now support pan to Next/Previous selected functionality. The F3 key is used to pan to Next selected object. The Ctrl-F3 key is used to pan to previous selected object.
- **Centre On Map** — Pressing the F3 key while selecting a node moves the node to the centre of the map.

Configurable labels for LLDP

Two fields namely Name and Description have been added to the DiscoveredPhysicalLink Object.

Name is read-only and will be auto populated. Examples are as follows:

- For SR's connecting back to back and both nodes are managed by SAM
`<sr-node name>:<port-name>--<sr-node name>:<port-name> (sim20:1/1/1--sim30:2/2/2)`
- For GNE device connecting to SR will be like
`<sr-node name>:<port-name> -- <GNE-node name> (sim20:1/1/1--LE-311)`
- If one end is unmanaged then name will be reflected like
`<sr-node name>:<port-name>--unmanaged-ne (sim20:1/1/1-un-managed ne)`

Description field data can be given by the user.

Span of Control/Scope of Command

Span on Services & Customers

Span is a grouping of the following types of SAM objects:

- Equipment (Topology) Group
- Network Elements (routers)
- Scripts (CLI, XMLAPI, and Auto-Config Profiles)
- Bulk Operations
- Test Suites
- Ring Groups
- VLAN Groups
- Service Templates

- Tunnel Templates

New in release 8.0 5620 SAM will offer services and customers as span objects.

- Services
- Customers

By default, the 5620 SAM includes the following 8 spans for each of the span object types.

Span ID	Span Name
1	Default Topology Group Span
2	Default Router Span
3	Default Script Span
4	Default Test Suite Span
5	Default Group Span
6	Default Bulk Operation Span
7	Default Service Span
8	Default Customer Span

Table 6: Default Spans

New customers and services will be put in the *Default Customer and Service Spans*.

When a span object is assigned to a specific user span, it remains in the corresponding *Default-xx Span*.

Users are not allowed to edit the contents of the default spans. Users can make a copy of any default span, and edit the contents of the copy when needed. The copy will NOT be updated when the corresponding default span is updated.

During an upgrade from to 8.0, each Span of Control profile that does not have all of the *Default-xx* spans assigned, will have the missing *Default-xx* spans assigned as read spans. This will ensure that existing users will not loose any privileges that they already have.

Services

Newly created services are automatically added to the *Default Service Span*. A user can only change a service that is in their write span.

In order to add a customer or components to a service, the user must also have access to the Customer and the Network Elements.

For a service component to be in a user's write span, the Customer, Service and Network Elements must each be in at least one of the user's assigned write spans.

- For example, an SDP Binding under a service is in the user's write span if the user has access to the Service, the Customer, the "from" node and the "to" node.

Span Types

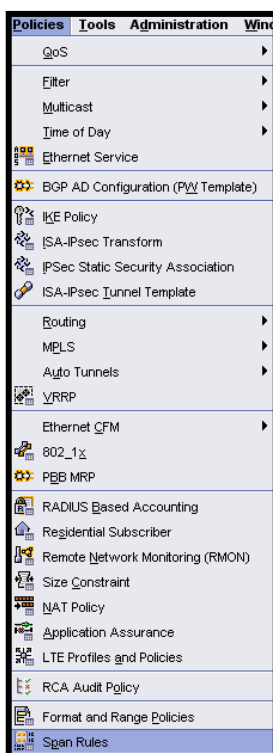
Four different types of spans can be chosen to add to a Span of Control Profile:

- View Access (objects contained in the span can be seen by the user, except in circumstances where the Scope of Command has specifically denied read privileges).
- Edit Access (objects contained in the span will adhere to the user's Scope of Command)
- Blocked Edit (blocks the user's Scope of Command for objects contained in the span)
- Blocked View (blocks the user from seeing the objects contained in the span)

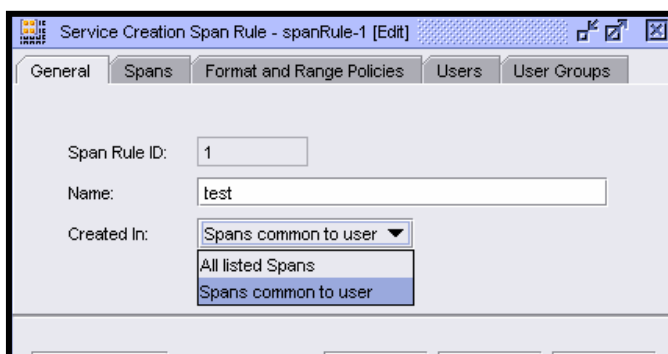
Span Rule

Procedure to configure the Span Rule

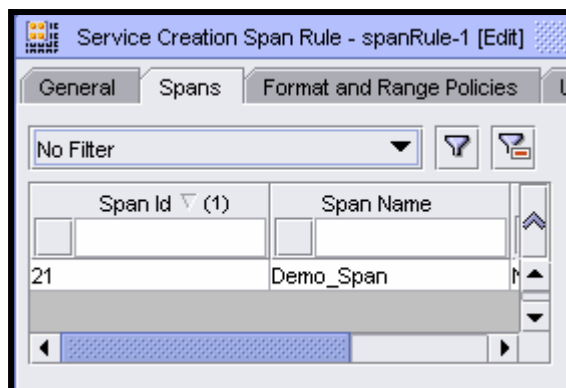
- Select the "Span Rules" from the Policies menu.



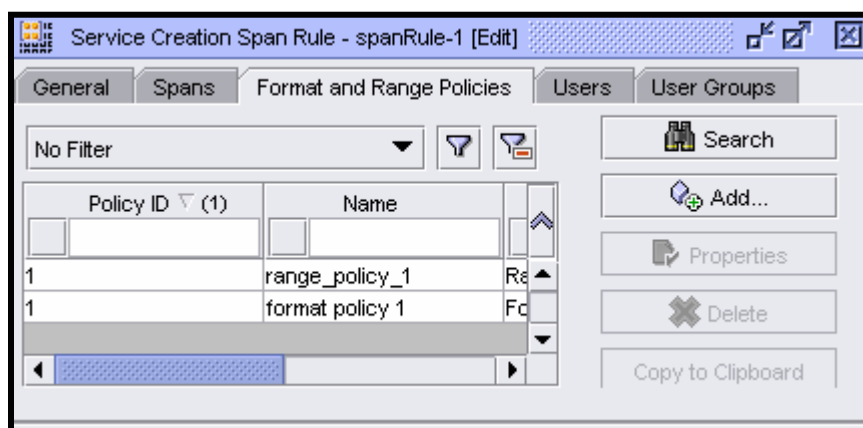
- Select the "Created In" option.



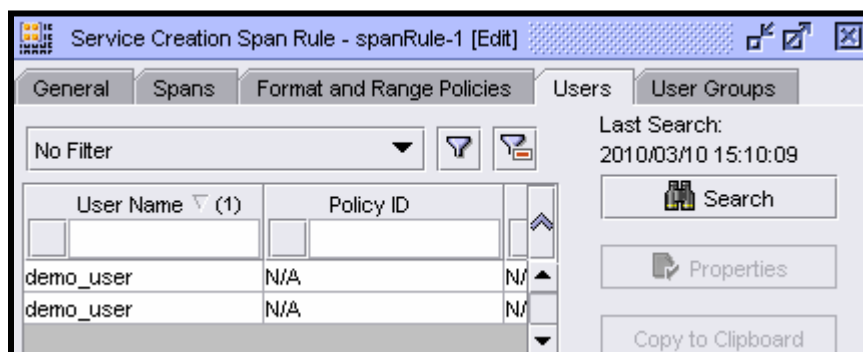
- Add the required span.



- Add the required format and range policies.



- The affected users and groups will be listed under the appropriate tabs.



Creating Services

When creating a service from the Create->Service GUI menu, the Span Policy associated with the user's Format/Range policy dictates which span(s) the service will be automatically added to, in addition to the default.

Customers

Newly created customers are automatically added to the *Default Customer Span*. A user can only change a customer that has edit access type.

In order to configure a customer on a node, a user must also have access to the Network Element.

The main use case for this feature comes from a customer whose operations team has certain customer prime-ships. Therefore they prefer that the operator only see objects that pertain to the customer they support.

One to many customers can be specified in a customer span. This span would get added to a Span Profile (which could include other spans...i.e. ones for network elements, ones for test suites, ones for services (when we support that), etc). The span profile gets applied to a User Group

Restricting Read Only Access

Automatic span filtering can only be turned off with the “Apply User Span of Control” user preference.

Starting in Release 8.0 R1, the Equipment, OSPF, ISIS, Routing, and Ring Group trees automatically filter the network elements by the user’s span when the User Preference for span is also checked. Automatic span filtering can only be turned off with the “Apply User Span of Control” user preference.

Topology Groups, Routers and Ring Groups that are not in the user’s span will be hidden.

Menu for Scope of Command & Span of Control

The following windows have been enhanced such that the create button now has a submenu. Span of Control Tab. Clicking on the “Create...” button now opens a small sub-menu beside it with the options “Span” and “Profile”.

- Selecting the “Profile” create option opens a new Span of Control profile configuration form.
- Selecting the “Span” create option opens a new Span of Control span configuration form.

Scope of Command Tab. Clicking on the “Create...” button now opens a small sub-menu beside it with the options “Role” and “Profile”.

- Selecting the “Role” create option opens a new Scope of Command role configuration form.
- Selecting the “Profile” create option opens a new Scope of Command profile configuration form.

System Administration

Task Manager Enhancements

The task manager was new in release 7.0. It is available on the 5620 SAM GUI and focuses on tracking status of write operations to the database and network deployments. Release 8.0 extends it in the following ways:

A new Tree tab is available to list the tasks and sub task hierarchy.

Bulk Change requests are a special case. The task manager tracks as a single batch entry in the manager and each individual request is as a child of that batch.

Now the bulk cancel operation is also tracked. This was not tracked as this was a read only operation before.

The following tasks are now supported:

- The OSS find and findToFile methods
- Resyncs
- Executes
- Tracking of scheduled tasks

Task Manager Reports

- Dump tasks in memory to file using OSS findToFile.
- Purged tasks are written to disk in the log directory in XML format. This can be appended to the OSS findToFile dump to get the full history of all tasks tracked since server startup.
- Tracks OSS findToFile and use of taskDescription as filter.
- Tracks OSS find and use of taskDescription as filter.

Task Manager reports are now available in the log directory (default is <install_dir>/5620sam/server/nms/log)

Policy Infrastructure

Explicit Policy distribution scalability

Prior to Release 8.0, the policy distribution is limited to the number of sites based on the number of items within the global policy and the policyDistributionMaxObjectLimit configured in nms-server.xml. This limitation is removed in SAM8.0.

A new distribute method is available for both GUI and OSS which has the same behavior as the global policy releasing in the following aspects:

1. The sites are grouped into multiple deployers to allow scaling.
2. The configuration parameter “policyDistributionMaxObjectsPerDeployer” defined in nms-server and number of items within the global policy is used for limiting the number of sites to be distributed per deployer.
3. If one site fails the distribution, the rest of sites continue to be distributed instead of failing the distribution for all sites.
4. The JMS message is sent to the client to list the sites failing the distribution instead of throwing the exception. As OSS clients already have to listen to JMS events for release failures, the new distribution method does not return information regarding success or failures.
5. The distribution can be followed through the Task Manager. Each distribution operation/request has a corresponding entry in the Task Manager.

Global policy created by the first local policy discovery

In releases prior to 8.0, when the first local policy is discovered, the global policy is created based on the first local policy configuration if it doesn't exist. The local policy configuration may be partially copied into the global policy because of the limitation of the discovery process.

In SAM 8.0, the global policy should be created “identical” (depending on applicability of properties and items) to the first discovered local policy if it doesn't exist in SAM.

Because the local policy might not be complete (e.g. entries/children of the policy have not been discovered yet) SAM will create an initially incomplete global policy. Two new properties will be introduced to show the different states/initialization phases of a global policy:

- *Creation State*: either “initializing” or “created”.
- *Created From*: either the siteld of the local policy if the global policy is or will be created from a local policy due to discovery OR the string “Sam” if the global policy was created through SAM (GUI/OSS).

Both properties are read-only and applicable to global policies, only.

When the initially incomplete global policy is created, “*Creation State*” is set to “initializing” and “*Created From*” is set to the siteld of the local policy that initiated the creation. When the router from which the local policy originated is completely discovered, the global policy is updated to match the local policy and the “*Creation State*” of the policy is set to “created”. If the operator (using GUI or OSS) modifies the global policy while its *CreationState* is “initializing” (before it is completed), the behavior is as follows:

- If discovery/full resync of the router identified by *Created From* is ongoing modification will be rejected (GUI/OSS).
- If discovery/full resync of the router identified by *Created From* is finished, modification using OSS will proceed as normal, while the GUI will present a warning and allow cancellation of the modification. If the modification is not cancelled, *Creation State* will be changed to “created” and the global policy will not be updated to match the local policy (at a later time).

The operator can use the *Creation State* property to search for incomplete policies in the event of discovery problems (e.g. SAM crash, switchover, parts of network not reachable), as this property will be visible in tables and policy config forms.

A side-effect of this way to create a complete global policy is that OSS clients will observe additional update events (due to the update of the creation state).

Note on task priorities: The policy framework will give priority to tasks performed as part of discovery over tasks that are performed as part of release/distribution.

Note on global policy creation behavior when local policy is created on a managed node through CLI (script) or other SNMP manager: when SAM is notified of a new local policy, SAM will resync the policy and the policy's potential children. But depending on the number of MIB entries involved, network latency and how fast the CLI (script) or SNMP manager executes the creation, the result may be partial, only. To correct this situation, the operator has to sync the global policy with the local policy. It should be noted that this behavior/procedure is not different from SAM pre 8.0.

Policy performance improvements

Multi-threading is introduced for the policy distribution, global policy releasing, and global policy creation by the first local policy discovery to improve the policy performance.

Task de-duplication is performed in the release case. This means, that as long as a release of a specific policy is not executed, but queued, and a subsequent release is requested, only one release be performed by SAM (the second release would not change the configuration of a node, as the first release already picks up the changes introduced in between the two releases. Release failures to specific nodes can be tracked through both release requests.

Additional performance gain will be achieved by reducing duplicate and unnecessary validation at the framework and application level, and by extending the policy framework's rollback capabilities on pre-commit validation errors.

Switch Mode of policy from list

A new button was added to the policy management window to allow the user to switch the policy to "Released" mode from the management window without opening the policy properties.

VI for CLI

In 8.0 users are able to use VI in CLI sessions. This allows the use of VI within SAM NE sessions (telnet and SSH). The user is able to use VI in a SAM spawned CLI session the same way as in any CLI session. All special characters are supported for cursor movement, pasting, deleting, etc.

Multiple NE Types for a Script

A user can now select more than one NE type for a script.

CLI Script, [Create]

General Versions Targets Spans

Script ID: 0 ☒ A

Name: demo_script

Description: demo_script

Type:

Use Latest Version: ☐

State: Enable ▼

Continue On Command Failure: ☒

Content Type: CLI ▼

Mode: Draft ▼

NE Types

NE Types ▼ (1)

Alcatel-SR-7710

Alcatel-SR-7750

Figure 22: Multiple NE Types for Scripts

Tools

SAP Copy/Move Support for ATM

SAPs created on a Port Termination with encap type ATM can be copied/moved to another Port Termination having the same encap type. This is an extension to SAP copy/move feature in SAM for ethernet ports.

Tools CLI commands available for SAM

The re-evaluation of PW template allows a user to apply changes to PWs using the template on existing BGP/ BGP-AD VPLS(s). From SAM, the function is available at three different levels: per PW, per template per router, and per template per service.

Scripts

XML U/S - Modify Templates incl def queries / curr

The 5620 SAM continues to extend the powerful template system for services and bulk configuration.

The 5620 SAM significantly extends service management by offering modify templates within all service template flows extending templates to the full service life cycle of add, drop and modification. A template *designer* can define a template as a modify class within template and UI builder tools. Modify templates, when executed by a template *user*, will query the related service objects so that the *user* sees current settings. For example, an iES SAP level template that allows users to change the QoS policy will display the fact that QoS policy 77 is currently applied on the SAP. Template *designers* have access to the full suite of UI builder functions including the ability to place fields as read only, read / write, set defaults, and so on.

The 5620 SAM significantly extends the user interface by unifying the bulk and service template interfaces. This gives service templates access to the powerful 'build up' tools offered in bulk configuration templates. Likewise bulk templates have access to some of the simplified finding and filtering features of the service template UI.

3. EQUIPMENT MANAGEMENT

7450/7750 IOM3/IMM Mixed-Mode Chassis Support

A number of forces are driving towards the introduction of additional routing into the 7450, including:

- Existing L2 TPSDA customers need path to IPv6 subscriber management
- Customers desire to offer some routing services - but not enough to warrant a 7750
- Many TPSDA deployments are Routed CO which currently requires the 7750
 - This becomes an issue if the customer does not need an entire chassis worth of bandwidth to handle the TPSDA traffic

The most immediate driver is the need to support IPv6 subscriber management functionality in networks where the 7450 is currently providing the subscriber management function. The SR and ESS product lines currently support 2 major IPv4 models, a bridged subscriber management model supported on the ESS and SR and a routed CO model supported only on the SR line.

Going forward for IPv6 it is preferred that only a routed CO model is supported that is currently only supported on the 7750 product lines. To allow introduction of these capabilities into the 7450 platform, 7750 IOM3-XP's and associated MDAs and or 7750 IMM's may be placed in the 7450 platform. Existing 7450 customers can add new 7750 IOM3-XP's and MDAs or IMM's into existing 7450 chassis and enable IPv6 subscriber management on the 7750 MDA ports. Ideally not just IPv6 subscriber management would be supported but all 7750 functionality, including IPv6 and multicast routing, full BGP support, VPRN services and on the associated 7750 IOM3-XP's and MDAs, full 7750 scale for these services.

A secondary benefit and driver of this feature is to allow existing 7450 customers to begin to expand the available services by adding L3 services into 7450 chassis as needed.

This feature is supported on 7450 ESS product family in R8.0. The following are to be supported in a 7450 in both legacy (current 7450 mode) and in mixed mode.

- 7750 IOM3-XP's
- IMMs
 - 48xGE (SFP and Copper)
 - 4x10GE
 - 5x10GE IMMs
 - 8x10GE IMMs
 - In 8.0R4 we should also support the 10x10GE IMM
- MDAs
 - Priority 1: 7750 MDA-XP's (20xGE MDA-XP, 2 or 4 x 10GE MDA-XP), current 7750 Ethernet MDAs
 - Priority 2: current 7750 POS MDAs
 - Post 8.0: 7750 ASAP MDAs, 7750 CES MDAs, 7750 ATM MDAs

The feature includes:

- Framework changes for dynamic capabilities (chassis+card)
- Mixed mode support
- Checks and validation (like ESS-1/6 not supported)
- 7750 IOM/MDA on 7450 but working as 7450
- 7750 IOM/MDA on 7450 operating in network mode
 - IPV6
 - BGP-All families including 2547
 - PIM/IGMP/MLD/MSDP
 - IPV4/IPV6-Routed CO
 - IPV4/IPV6 L3 Spoke interfaces
 - P2MP
- Software upgrades (CPM/AA/IPSec etc.)
- IMM support

SAM allows any number of 7750 IOM3s with 7750 MDAs or IMMs to be installed in a 7450. Access ports on these IOMs will provide 7750 services (see below).

- Mixed mode : Allows 7750 IOMs and IMMs to function in a 7450 and also allows 7750 functionality on those IOM3-XP's/IMMs
- Legacy mode: Allows 7750 IOMs and IMMs to function in a 7450 but only 7450 functionality will be supported. So basically the 7750 IOM3 & IMMs will act as 7450 cards.

SAM provides an option to enable and disable Mixed mode.

All network interfaces carrying 7750 associated services must also be located on a 7750 IOM3-XP or IMM card in mixed mode

Other functions include:

- IPv6, IPv6 Submgmt
 - Base IPv6 routing and associated protocols, including, IPv6 IGP support, BGP and 7750 IOM3-XP scale should be supported.
- VPRNs

- Access ports on the 7750 IOMs will provide access to VPRN services. This will provide an upgrade path from L2 to L3 services in a limited scope to customer that already deployed 7450s.
- 7750 IOM3-XP scale should be supported.
- Full scale IP (Internet routing), including full scale BGP support
 - Full IES services including BGP and full 7750 route table, IP interface and BGP scale (IOM3-XP) should be supported.
 - The 7450 will retain it's default address family support in mixed mode. As a result once changed to mixed mode, the administrator must configure the additional address families they wish to advertise via BGP. Likewise, they must remove those address families if changing out of mixed-mode.
- IP multicast routing
 - IPv4 - Support for IGMP and PIM protocols as well as 7750 IOM3-XP scale
 - IPv6 - MLD, PIMv6 and 7750 IOM3-XP scale
 - P2MP support for multicast
- Frame Relay, ATM and CES VLLs (Priority 2)
 - Long term L2 VLL services not supported on the 7450 should also be supported on the 7750 IOM3-XP and IMMs in a 7450 chassis.

In General, all service and protocols scale should match that of a 7750 with similar types of IOMs when the chassis is in mixed mode.

- If the 7450 has a mixture of IOM3s/IMMs and IOM1 then the scale should match 7750 mode B scale limits
- If the 7450 has all IOM3/IMMs then the scale should match 7750 mode D scale limits

Note: Since the 7450 line does not have IOM2s, mode C is not applicable

No licensing changes exist as a result of this feature. A 7750 MDA is charged as a 7750 MDA regardless of the chassis it is put in. A 7450 MDA is charged as a 7450 MDA regardless of the chassis it is put in.

New 7xx0 MDA support

The following MDAs are supported in 5620 SAM:

- 7750 SR 48-PT GE MDA-XP - TX
- 7450 SR 48- PT GE MDA-XP - TX

SAM 8.0 supports the configuration of 48-port TX MDA-XP on 7x50.

The configuration of the ports and the MDA will be similar to the existing MDA's.

The statistics support is similar to the ones in the existing Ethernet MDA

This card allows provisioning of 48 physical Ethernet ports via (8) mini-rj21 connectors that each deploy (6) ports.

The 48-port MDA-XP delivers full XPL+ 25Gbps capacity. This should be full line rate (imix) for up to (25) of the (48) available ports (flexibly used, as opposed to a specific 25 if at all possible).

The 48-port GE TX MDA-XP is supported on both the 7750 SR and 7450 ESS platforms, in the 12-slot, 7-slot, and 6-slot chassis's, in 5620 SAM Release 8.0.

The 48-port GE TX MDA-XP delivers on the full breadth of R8.0 SW features at DR4 and offers all the functionality of today's MDA-XP's.

Hybrid Port Support

A hybrid port allows the user to configure service SAP and IP network interface on virtual interfaces of the same physical port. A network interface on a hybrid port performs IP routing, LSR and LER functions exactly like the existing network interface on a network port. This feature is supported on SR/ESS 7x50 product family in Release 8.0.

The feature provides the following main functions:

- Hybrid Port Mode Configuration
- SAP Capability
- Network IP Interface Capability
- Hardware Support
- QoS Requirements
- Service destination based shaping (To be delivered post 8.0R1)

Business and Mobile Backhauling

The first application of the hybrid port is in the use of a single aggregation network to backhaul both mobile and business services as shown in Figure 23. The mobile services extend a PW from the CPE, e.g., 7705 SAR, all the way to the 7x50 Metro Edge Node (MEN) over an aggregated VLAN sub-interface. Each PW is then switched to another PW towards the 7x50 Mobile POP Node. The business services are presented to the 7x50 MEN as individual VLAN SAPs.

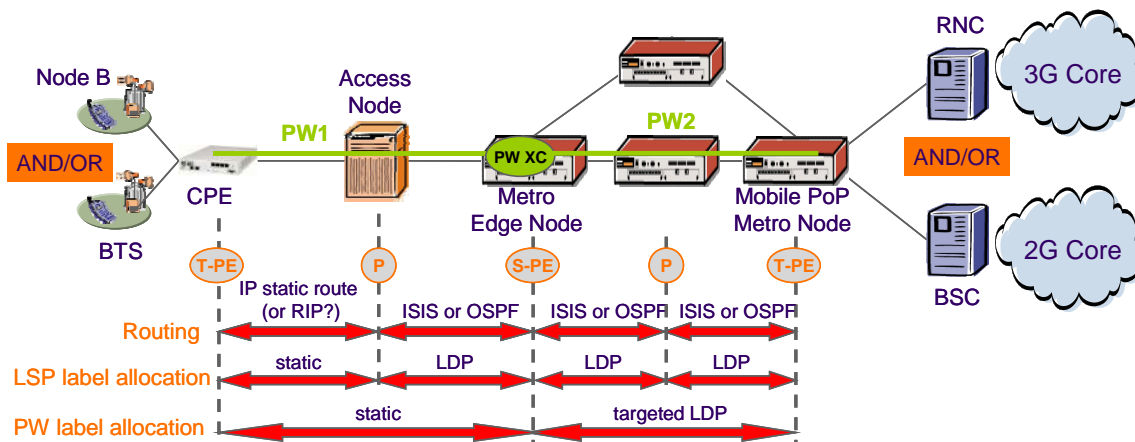


Figure 23: Business and Mobile Backhauling Application

The customer would like to use a single GigE uplink to forward packets of both services from the DSLAM to the 7x50 MEN. In other words, the 7x50 MEN will be presented with both a regular VLAN SAP and a VLAN network interface on the same GigE port. In addition, the customer requires that each network interface on a hybrid port provide per forwarding class queues. The following are the requirements in these types of applications which are not satisfied with the existing 7x50 network interface feature:

- Provide per VLAN sub-interface per forwarding class queues. This feature is provided by the Queue Group capability introduced in R7.0 [Queue-Group].

- Scale the number of IP network interfaces, including scaling IGP adjacencies and LDP adjacencies, to 2K. Each CPE at a Node B/BTS site requires the configuration of an IP network interface on the 7x50 in the hub or central site. When redundant PWs are used, there will be two interfaces to each CPE. The 7x50 currently supports a maximum of 255 network interfaces on regular network ports.
- Combine regular VLAN SAP with VLAN network interface on the same Ethernet port.
- Allow the ability to use the implicit NULL label to allow for the encapsulation using the PW label only (dry-martini PW) in the case of single-hop LSPs between the CPE and the 7x50 MEN. In the current implementation of a network interface, the 7x50 can send a labeled packet with the implicit NULL but cannot be configured to receive it.

Although the above application makes use of an LER capability at the 7x50 MEN, there are other applications where the 7x50 is an LSR and switches an LSP from a network interface on a hybrid port to a network interface on another hybrid port or on a regular network port.

Service Destination Based Shaping

In the second application, a network operator wants to shape or rate-limit traffic on per service and per destination basis. In the specific setup shown in Figure 24, this means shaping traffic on a per VPLS/per PW basis.




Figure 24 Service Destination based Shaping

Note: This requirement applies to network interfaces on a regular network port and on a hybrid port. It is captured here to make sure hybrid ports are designed to allow for this capability.

Hybrid Port Mode Configuration

The user can configure the hybrid mode only on Ethernet ports.

The default mode of an Ethernet port remains network. A port configured in mode hybrid allows both ingress and egress contexts to be configured concurrently. It can have only two encapsulation types: dot1q and QinQ.

SAP Capability

When the port is configured to the dot1q encapsulation, the user configures a SAP inside a service simply by providing the SAP ID which must include the port-id value of the hybrid mode port and an unused VLAN tag value. The format is *<port-id>:qtag1*. A SAP of format *<port-id>.** will continue to be supported.

The 4096 VLAN tag space on the port is shared among VLAN SAPs and VLAN network IP interfaces.

When the port is configured to the QinQ encapsulation, the user configures a SAP inside a service simply by providing the SAP ID which must include the port-id value of the hybrid mode port and the outer and inner VLAN tag values. The format is *<port-id>:qtag1.qtag2*. A SAP of format *<port-id>:qtag1.** will continue to be supported.

The outer VLAN tag value must not have been used to create an IP network interface on this port. In addition, the *qtag1.qtag2* value combination must not have been used by another SAP on this port.

All other features and options which are currently supported on a service Ethernet SAP continue to be supported.

Network IP Interface Capability

When the port is configured to the dot1q encapsulation, the user configures a network IP interface under *config>router>interface>port* by providing the port name which consists of the port-id of the hybrid mode port and an unused VLAN tag value. The format is *<port-id>:qtag1*. The user must explicitly enter a valid value for *qtag1*. The *<port-id>.** value is not supported on a network IP interface.

The 4096 VLAN tag space on the port is shared among VLAN SAPs and VLAN network IP interfaces.

When the port is configured to the QinQ encapsulation, the user configures a network IP interface under *config>router>interface>port* by providing the port name which consists of the port-id of the hybrid mode port and a VLAN tag value. The format is *<port-id>:qtag1.**.

The VLAN tag value must not have been used to create another IP network interface on this port. In addition, the VLAN tag value must not have been used as the outer VLAN in the SAP-id by another SAP on this port.

All other features and options which are currently supported on a network IP interface continue to be supported.

Hardware Support

Hybrid is supported on:

- IOM-3/IMM in SR7/SR12 with all chassis modes. It is not supported on the SR1 platform.
- IOM-3 in ESS6, ESS7, and ESS12 with all chassis modes. It is not supported on the ESS1 platform.
- 40G-7710.
- All Ethernet MDA and CMA ports. It is not supported on the HS-MDA (Q-MDA) and the VSM MDA.

QoS Requirements

This feature introduces changes to the IOM buffer pool configuration for an Ethernet MDA.

Each port configured in hybrid mode must support SAP queuing model and network interface queuing model concurrently. One way of achieving this is to have each hybrid port draw buffers from separate access and network buffer pools. To be more specific, each time the user configures a VLAN SAP or a VLAN network interface, the hybrid port must be able to create the corresponding queues by drawing from the following default buffer pools:

- One access ingress buffer pool (dedicated).
- One access egress buffer pool (dedicated).
- One network egress buffer pool (dedicated).
- One network ingress buffer pool (shared MDA-wide)

Each default buffer pool is allocated at the time the first queue needs to be created in the context represented by the pool, e.g., access egress.

Each of the hybrid port's access and network ports belonging to the hybrid port named pools is further split among these named pools using the access-allocation-weight and network-allocation-weight parameter values defined within the named pool policy applied to the hybrid port. These parameters are available to end user in GUI Client Ethernet port properties, only for ports with Hybrid Port set. Note: No new pool policy is introduced for Hybrid Port. It uses the existing named pool policy for network/access in SAM.

SAP-ingress and SAP-egress QoS policies continue to apply to ingress and egress of a SAP within a hybrid port. They manage queues within the ingress and egress access pools of a hybrid port.

Network QoS policies continue to apply to the ingress and egress of network interfaces within a hybrid port. They thus manage queues within the ingress and egress network pools of a hybrid port.

Scheduler policies continue to apply to ingress and egress of a SAP within a hybrid port. They thus manage scheduling of queues within the ingress and egress access pools of a hybrid port.

The port-scheduler policies continue to apply to egress SAP and to egress network interface within a hybrid port.

The queue group feature [Queue-Group] must be supported on both ingress and egress VLAN SAP and egress VLAN network interface.

M1-10GB DWDM TUN with DWDM Wavelength Selection

The DWDM channel can be configured on the port. The port can be configured to any one of the possible 89 different wavelengths. The wavelength can also be modified without resetting the MDA.

Support for OUT changes

The MDA supports OTU-2, a line rate of approximately 10.7 Gbits/s. The following parameters can be configured for OTU-2.

- Generic Forward Error Correction (G-FEC): This is used to extend non-regenerated optical transport distances.

- Enhanced Forward Error Correction (E-FEC): This option is used for long haul application support.
- Signal Failure and Signal Degrade thresholds and methods can be configured.
- Trail Trace Identifier for tx/Rx for Section Monitoring

A number of new properties are added to the existing OTU tab to view/ modify the new changes

pm-tti - Configure the Path Monitoring Tail Trace Identifier parameters

psi-payload - Configure the Payload Structure Identifier Payload parameters

psi-tti - Configure the Payload Structure Identifier Tail Trace Identifier parameters

Support for expected sm-tti changes.

Async Mapping - Enable/Disable OTU asynchronous mapping; synchronous when disabled

The following clocking modes are supported

- Support for Local (internal) and line (network) loop-back.
- Support for Local (free run) or loop-timing (recovered from the network)
- Support for system-synced timing on the optical link.

The following alarms are supported

1	loc	Rx PLL Loss of lock
2	los	Loss of signal transitions on the data
3	lof	Loss of OTU framing
4	lom	Loss of Multi-frame
5	otu-ais	OTU Alarm Indication Signal (all 1s, overwrites all OTU overhead, even framing bytes)
6	otu-ber-sf	SM Signal Fail (based on BPI8)
7	otu-ber-sd	SM Signal Degrade (based on BPI8)
8	otu-bdi	SM Backward defect indication
9	otu-tim	SM Trace Id Mismatch
10	otu-iae	SM Incoming Alignment Error
11	otu-biae	SM Backward Incoming Alignment Error
12	fec-sf	Signal Fail (based on FEC corrected bits)
13	fec-sd	Signal Degrade (based on FEC corrected bits)
14	fec-fail	FEC Mode mismatch (EFEC-GFEC) Uncorrectable rate (>10E-2)
15	fec-uncorr	One or More Uncorrectable FEC errors
16	odu-ais	ODU Alarm Indication Signal
17	odu-oci	ODU Open connection Indication
18	odu-lck	ODU Locked
19	odu-bdi	PM Backward Defect indication
20	odu-tim	PM Trace Id Mismatch
21	opu-tim	OPU PSI Trace Mismatch

Synchronization Enhancements on the SR & ESS

The synchronization capabilities involve enhancements to the existing infrastructure to support the deployment of the SR/ESS operating with a co-located BITS/SSU. The specific features being developed include:

- use of both BITS input ports by the active SF/CPM
- enabling of the BITS output ports
- selection of timing reference based on QL and priority
- support for SSM on E1 BITS interfaces

ESMC (Ethernet Synchronization Message Channel) Configuration on an Ethernet Port

Synchronization Status messages have been defined for various transport formats and are placed in prescribed overhead bytes for SONET & SDH signals and also in bit-oriented messages within the data link for DS1 (ESF) and E1 signals, for interaction with office clocks, i.e., BITS or SSUs. For Synchronous Ethernet interfaces, there is no equivalent fixed locations to convey SSMs, so the QL is transported using Ethernet frames. Specifically, the message channel is an Ethernet protocol based on an IEEE Organization Specific Slow Protocol (OSSP). The channel is called the Ethernet Synchronization Message Channel (ESMC).

Both the MDA must be configured for Synchronous Ethernet and the port must have SSM enabled for the ESMC packets to be generated and the received packets to be processed.

The following attributes are modeled:

network-type - decides encoding of synchronous status messages, i.e., whether to use SDH or SONET set of values. Configuring the network-type is only applicable to SyncE ports. It is not configurable on SONET/SDH ports

tx-dus - forces the QL value transmitted out the SSM channel of the SONET/SDH port or the Synchronous Ethernet port to be set to QL-DUS/QL-DNU. This capability is provided to block the use of the interface from the SR/ESS for timing purposes.

System Clock Configuration

In 8.0 configuration of BITS output ports on the SR/ESS are enabled. The following new attributes are modeled:

output - to configure and enable/disable the external BITS timing reference output to the SR/ESS.

line-length - configures the line-length parameter of the BITS output, This is the distance in feet between the network element and the office clock (BITS/SSU). There are 2 possible BITS-out interfaces, one for each CPM. They are configured together, but they are displayed separately in the show command. This command is only applicable when the interface-type is DS1.

ssm-bit - configures which Sa-bit to use for conveying SSM information when the interface-type is E1.

ql-override - configures the QL value to be used for the reference for SETS input selection. This value overrides any value received by that reference's SSM process.

input - to enable/disable the external BITS timing reference inputs to the SR/ESS. There are 2 possible BITS-in interfaces, one for each CPM. They are configured together

ql-selection - When enabled the selection of system timing reference and BITS output timing reference takes into account quality level. This attribute turns-on or turns-off SSM encoding as a means of timing reference selection.

Enhanced BITS in port redundancy

On 7750/7450 systems with redundant CPMs, the system will have two BITS input ports (one per CPM). These BITS input ports provide redundant synchronization inputs from an external BITS/SSU. On systems with cross coupled timing modules, the active CPM shall be capable of using either BITS input port for its synchronization. The following systems support cross coupled timing modules:

- 7750 SR12, SR7
- 7450 ESS12, ESS7

No SAM changes were made for the above feature, it is test only.

MultiService ISA support in the IOM3 for Video Services

In R7.0, the MultiService ISA (previously Video ISA) is supported in the iom-20g-b and the iom2-20g for video services.

In R8.0, MultiService ISA is supported on IOM3 also.

Prioritization Mechanism for RET vs. FCC

In R7.0, RET and FCC requests are processed with the same priority. Since RET generally has a more direct impact on a subscriber's "quality of experience"

In R8.0, this feature provides a mechanism to reserve an explicit amount of egress bandwidth for RET for all the ISAs within an Video Group. If the amount of egress bandwidth is less than the reserved amount, FCC requests are discarded and only RET requests processed. The bandwidth will need to be dynamically adjusted per ISA within the Video Group if ISAs become operational/non-operational within the group.

The following attribute is introduced.

resv-ret - Configure reserve RET bandwidth

ECMP fate sharing

This feature is supported on all 7x50 platforms including 7710.

A new group is introduced in Equipment tree “IGHs”. This holds all instances of ECMP fate sharing group [named interface-group-handler on node cli]. A new class is added to support ECMP fate sharing group, named “IGH”.

The GUI implementation is similar to that provided by “LAG” functionality.

E-LMI Protocol

The Ethernet Local Management Interface (E-LMI) protocol is used for enabling the CE to request and receive status and service attributes information about the Ethernet Virtual Connection (EVC) that it is getting from the PE.

E-LMI supported on all SR and ESS platforms for Ethernet access ports with dot1q encapsulation type.

A number of new attributes are modeled under Ethernet equipment:

- mode - Configures E-LMI UNI mode.
- t391 - Configures the Polling Timer for UNI-C.
- t392 - Configures the Polling Verification Timer for E-LMI
- n393 - Configures the Status Counter for E-LMI.

Increase the LAG limit to 16

5620 SAM Release 8.0 increases the maximum number of ports in a single LAG to sixteen.

This feature is supported on IOM3 and IMM line cards only with chassis Mode D.

If a system with a configuration of more than eight ports in a LAG is executed in a chassis that is not running in Mode D then only the first eight ports should be accepted in the LAG and other ports should be rejected with an appropriate error message

Enhancements to LAG Hashing for consistent per service forwarding

5620 SAM Release 8.0 provides a consistent forwarding path for all the frames belonging to an Ethernet service - ELINE/Epipe or ELAN/VPLS for both regular and PBB services. This feature provides this capability using ISID based hashing for PBB based services and Service ID for regular Epipe and VPLS.

The following new attribute is modeled.

per-service-hashing - enables consistent per-service hashing for Ethernet services over LAG or over Ethernet tunnels (eth-tunnel) using load-sharing protection-type. Specifically it enables the new hashing procedures for Epipe, VPLS, regular or PBB services.

SAP and Service Configuration Enhancement for Frame Relay

Under the “Frame Relay” tab, there is an additional field in the FRF.12 group:

interleave: This option provides a new mode of operation for fragmentation in the transmit direction of a FR SAP. It allows for interleaving of high-priority frames and fragments of low-priority frames. When this option is enabled, only frames of the FR SAP non-expedited forwarding class queues are

subject to fragmentation. The frames of the FR SAP expedited queues are interleaved, with no fragmentation header, among the fragmented frames. In effect, this provides a behavior like in MLPPP Link Fragment Interleaving (LFI). The receive direction of the FR SAP operates as described and supports both modes of operation concurrently, i.e., with and without fragment interleaving. By default value is set to “no interleave”.

Multi-Chassis PW Endpoint Support for VPLS

This feature provides the following functions:

- MC Endpoint configuration under MC Peer (Node Redundancy)
- MC Endpoint configuration under VPLS Endpoint
- Binding MC Endpoint with Spoke SDP
- Map Display for MC Endpoint peers and the associated Spoke SDPs

MC PW Endpoint for VPLS is an enhancement to existing Endpoint support (7x50 6.0) for redundant Spoke-SDP to VPLS. This enhancement provides a way to group the multiple spoke SDPs associated with the 2 MC endpoint peers to ensure that only one of them is active at any point of time, therefore to provide a new way to avoid the traffic loops among the VPLS instances.

The following picture depicts a customer network with Metro-mesh and Core mesh. Pseudo-wires are used to interconnect the PEn to the PE-c pair. In order to avoid loops, the customer requests that just one of the links (the bold, black one) will be actively used to interconnect different VPLS Meshes.




Figure 25: Customer Network with Metro Mesh and Core Mesh

In order to solve the above problem, the node 7.0 introduces the concept of MC PW Endpoint to provide multi-chassis resiliency at the endpoint side. The new construct, referred to as Multi-chassis PW Endpoint for VPLS (MC-PW-EP-VPLS) provides a way to group multiple PWs distributed between 2 or multiple chassis endpoints that are configured in the same VPLS.

The figure shows the solution to Multi-domain VPLS resiliency with MC Endpoints. A pair of endpoints are configured on PE1 and PE1' as peers (MC endpoints), each of them are associated with two spoke SDPs to PEC1 and PEC2. Based on the information received by the peer shelf and the local configuration, PE1 and PE1' will make a decision on which PW will become active.

The MC-EP solution is built with the following main components:

- PW Data plane - represented by four PWs per VPLS that connect the MC-EP pair to the other gateways. Only one of the PWs is activated based on the selection algorithm described below. The rest of them are blocked.
- Multi-chassis protocol is used to provide the following functions:
 - decide which of the MC-EP peers will be the master or slave
 - synchronize the PW information between the MC-EP peers.
 - although it relies on centralized BFD for fast fault detection, MC-EP protocol also contains its own keep-alive mechanism as a fallback (detects if MC-EP peer is in shutdown mode/miss-configured - not detectable by bfd)..
- T-LDP signaling - used to inform the other gateway pair (not necessarily running MC-EP) about the PW choices (A/S)

This new feature is supported in regular VPLS and in BVPLS context. Both GRE and MPLS (LDP and RSVP-TE) are supported for underlying tunnels associated with SDPs.




Figure 26: Multi-domain VPLS Resiliency with MC Endpoints

MC Endpoint support in Node Redundancy management

MC Endpoint configuration is supported under Node Redundancy management. It is configurable under the MC Endpoint Group object, which contains two matching MC endpoints under the two matched MC peers. Each peer contains attributes that are shared among all the underlying objects, e.g. MC LAG, MC Ring, etc. The two matched MC peers are managed via an MC peer group in SAM.

The MC endpoint group is managed as a child of MC peer group and can only be created under an MC Peer Group object. It also can be listed from Node Redundancy Manager but without creation functionality.




Figure 27: Associated MC Peer Group View

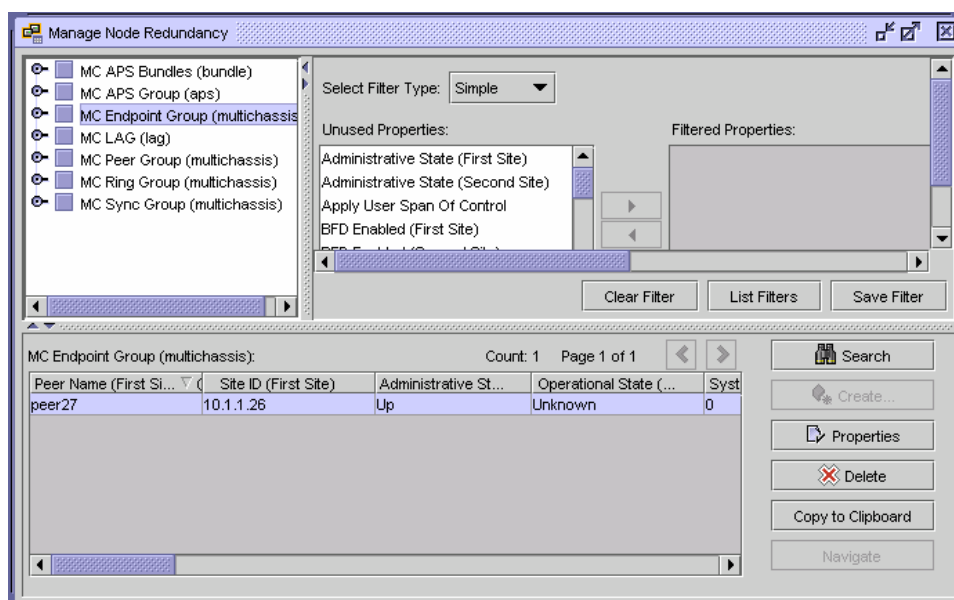


Figure 28: MC Endpoint Group in Node Redundancy Mgr

Each MC Peer Group managed by SAM contains two MC peers, each one is configured as the peer of the other one.

MC Endpoint support on VPLS Site

Maximum number of Multi-Chassis Endpoints per VPLS instance is 10

In SR 6.0, the maximum number of endpoint per VPLS instance is 2, now with multi-chassis endpoint support, this number is increased to 10.

SAM 6.0 supports Endpoint (single chassis) under VPLS site. The endpoint object is used to support the PW redundancy. Maximum two PWs can be configured under one endpoint in a VPLS site. Only one of the PWs is active and used for passing the traffic. If the active PW goes down, the backup PW will become active.

Support for MC Endpoint is added under VPLS site. All the attributes from the single chassis endpoint are also applicable to multi-chassis endpoint. For the endpoints created in SR 6.0, they are referred to as single chassis endpoint. Starting from SR 7.0, the user can create an endpoint as a single chassis or a multi-chassis endpoint on a 7.0 SR node (VPLS site). The endpoint creation form has a new attribute "Endpoint Type", with two options single chassis or multi-chassis.

Figure 29: VPLS Endpoint Configuration

Binding MC Endpoint with Spoke SDP

Maximum 2 Spoke SDPs per multi-chassis endpoint.

This restriction is similar to the restriction in 6.0, maximum 2 spoke SDPs per (single chassis) endpoint.

The user associates the spoke SDP with endpoint in the way that it is same as before. If this endpoint is a multi-chassis VPLS endpoint, the spoke SDP is bound to a multi-chassis VPLS endpoint. Otherwise, it is bound to a single chassis endpoint.

Alarm Support

The alarm IncorrectEndPointPeerConfig will be set on an MO of class MultiChassisEndpoint parented by Peer when there is not a match for that MO.

The alarm McPeerEndpointDown will be raised when MultiChassisEndpoint operational flag is not connected and the administrative state is up.

McPeerEPBfdSessionDown when operational state of the BFD session between the multi-chassis endpoint and its peer is changed to 'down'

Statistics Support

The new stats for the MC Endpoint will be collected. See the 5620 SAM Statistics Guide for a list.

MC Global stats:

The MC EndPoint Global Stats data represents the statistics data for all the MC Endpoint objects on this node. This stats tab is added under the network element form stats tab.

7750 SR c12 Enhancements

Support for the “7750 SR-c12” hardware was added in 5620 SAM release 7.0. The following additions have been made with this area in release 8.0.

CPM Filter and CPMQ

CFM filters and CPMQ are supported. The maximum number of queues that can be used for CPM filter and CPMQ is 2000. The capability on 7750 SR-c12 matches the one on 7750 SR7/12. Both Ipv4 and Ipv6 are supported.

Timestamp up to 128 bytes into the packet

7750 SR-c12 is able to timestamp up to 128 bytes in a packet. There are no changes in SAM as this is the internal variable on the Node.

Scalability Requirements

7750 SR-c12 has 40+Gbps throughput which is more than triple the throughput of the original 7710 SR-c12 system. This requires that the control plane also scales up to match the increased hw throughput. The new 8 core CPU based CFM is capable of increased control plane scalability. In release 7.0, the scalability limits of the 7750 SR-c12 was kept at the same level as that of the original 7710.

The following are the scalability improvements in R8.0:

Description	7750 SR-c12
Max SAPs per systems	64K*
Max IP interfaces per system	4K
SDP Spoke Termination	4K
Max Mac FIB	256K

Max ARP Entries	256K
Max VLLs	20K
Max VPLS instances	3K
Max BGP Peers	1K
Max VPN labels	256K -> 800K (for Cisco interop)
Max VPRNs	4K
max subscribers	64K*
DHCP sessions	64K*
PPPoE sessions	64K*
L2TP tunnels	1K*
L2TP sessions	64K*

OMNI Support

The following tables show functionalities as per OmniSwitch family of products supported by 5620 SAM releases.

OmniSwitch Product	SAM 6.0 R1	SAM 6.0 R3	SAM 6.1 R1	SAM 7.0 R1	SAM 7.0 R4	SAM 8.0 R1	SAM 8.0 R3 (candidate)
OS 6850 6.3.1	/	/	/	/	/	/	/
OS 6400 6.3.3	-	-	/	/	/	/	/
OS 6855 6.3.2	-	-	/	/	/	/	/
OS 9000 6.3.1 R2	-	-	-	/	/	/	/
OS 6850, 6400, 6855 & 9000 6.3.4 R1					/	/	/
OS6250, 6.6.1 (Metro & SME)						/	/
OS 9000E 6.4.2 R1						/	/
OS 6850, 6400, 6855 & 9000 6.4.2 R1							/
OS 6855 U24X, 6.4.2							/
OS6250, 6.6.2 (Metro & SME)							/

Table 7: OMNI Support

Functionality	SAM 6.0 R1	SAM 6.0 R3	SAM 6.1 R1	SAM 7.0 R1	SAM 7.0 R3	SAM 7.0 R4	SAM 8.0 R1	SAM 8.0 R3 (candidate)
OmniSwitch Equipment Management	/	/	/	/	/	/	/	/
Stack Configuration ¹	/	/	/	/	/	/	/	/

Ethernet Port Configuration	/	/	/	/	/	/	/	/
VLAN Service	/	/	/	/	/	/	/	/
QoS Management	/	/	/	/	/	/	/	/
NE Maintenance	/	/	/	/	/	/	/	/
AAA Security	/	/	/	/	/	/	/	/
Port Security	/	/	/	/	/	/	/	/
Protocols - Static Routing, IPv4 Multicasting (switching & routing)	/	/	/	/	/	/	/	/
IGMP Snooping	-	/	/	/	/	/	/	/
Notifications - Traps & Alarms	/	/	/	/	/	/	/	/
Ethernet Interface Statistics	/	/	/	/	/	/	/	/
OSSI	/	/	/	/	/	/	/	/
OAM - ICMP Ping & Trace	/	/	/	/	/	/	/	/
Full Spanning Tree management	/	/	/	/	/	/	/	/
UDP Relay/DHCP Snooping ³	-	-	-	/	/	/	/	/
Switch Health Monitoring	-	-	-	/	/	/	/	/
Ethernet OAM - Connectivity Fault Management	-	-	-	/	/	/	/	/
LAG - Link Aggregation Group	-	-	-	-	/	/	/	/
LLDP - Link Layer Discovery Protocol	-	-	-	/	/	/	/	/
Service Templates	-	-	-	/	/	/	/	/
Scheduling	-	-	-	/	/	/	/	/
Routing Protocols: OSPF, RIP, PIM ²							/	/
MVRF-Multiple Virtual Routing & Forwarding ²							/	/
MPLS, LDP support ²							/	/
VPLS support ²							/	/

Table 8: Supported OMNI Features

¹Stack Configuration is not supported for the OS 6855 6.3.2 and OS 9000/9000E

² Applicable to AOS9000E nodes in 8.0 R1/R3.

³ UDP Relay/DHCP Snooping is not supported on AOS9000E nodes in 8.0 R1.

The 5620 SAM Release 8.0 supports the following OS6250 Chassis types:

Chassis Type	Description
OmniSwitch OS6250-8M	8 copper GigE Ports, 2 Fiber/Copper GigE combo ports, and 2 fiber ports for stacking.
OmniSwitch OS6250-24M	24 copper GigE Ports, 2 Fiber/Copper GigE combo ports, and 2 fiber ports for stacking.
OmniSwitch OS6250-24MD	24 copper GigE Ports, 2 Fiber/Copper GigE combo ports, and 2 fiber ports for stacking. internal DC power supply
OmniSwitch OS6250-24	24 copper GigE Ports, 2 Fiber/Copper GigE combo ports, and 2 fiber ports for stacking.
OmniSwitch OS6250-P24	24 copper PoE Ports, 2 Fiber/Copper GigE combo ports, and 2 fiber ports for stacking.

Table 9: Supported OS6250 Chassis Types

The 5620 SAM Release 8.0 supports the following OS9000E Chassis types:

Chassis Type	Description
OmniSwitch OS9700E	The OmniSwitch 9700E is a high performance switch offering eight slots for Gigabit Ethernet and/or 10-gigabit Ethernet Network Interface (NI) modules. Additional two slots are reserved for primary and redundant Chassis Management Modules (CMMs). The OmniSwitch 9700E supports a maximum of three power supplies.
OmniSwitch OS9800E	The OmniSwitch 9800E is a high performance switch offering 16 slots for Gigabit Ethernet and/or 10-Gigabit Ethernet Network Interface (NI) modules. An additional two slots are reserved for primary and redundant Chassis Management Modules (CMMs). The OmniSwitch 9800E supports a maximum of four power supplies.

Table 10: Supported 9000E Chassis Types

The OS9000E series share a family of common Network interfaces for 10 Gigabit Ethernet and Gigabit Ethernet connectivity:

NI Module	Description
OS9-GNI-C24E	Network Interface with 24 Ports 10/100/1000 with RJ-45 support
OS9-GNI-U24E	Network Interface with 24 Ports 1000Base-X with SFP/MiniGBIC support
OS9-XNI-U2E	Network Interface with 2 Ports Unpopulated 10Gigabit Ethernet with XFP support

Table 11: Supported OS9000E NI Modules




Figure 30: OMNISwitch OS9800E

In 8.0 R1, OSPFv2, RIP, and PIM protocol support is introduced only for AOS9000E nodes.

OSPF

Loading OSPFv2

OSPFv2 should be loaded separately for each VRF instance on the node.

The OSPFv2 can be activated on AOS nodes (OS9000E in 8.0 R1) from 5620 SAM. -Routing instance properties window->Protocols->Enable Ospf2.

It can not be unloaded from SAM as it has to be done manually on switch by editing boot.cfg file.

CLI:

```
->ip load ospf
-> vrf rtr1
```

Enable/Disable OSPFv2 instance

The OSPFv2 protocol instance can be enable/disable from OSPFv2 protocol properties window by setting administrative state up/down.

CLI:

```
->ip ospf status enable
->ip ospf status disable
```

Creation/Deletion of an OSPF Area

An OSPF Area can be created/deleted from Routing Tree ->OSPFv2 Instance -> Create Area. Also from the properties window of OSPFv2 Instance -> Area Site Tab -> Add/Delete

An OSPF Area Id and Type can be specified at the time of creation. An Area Range can be created from Area Range Tab in Area Site properties.

CLI:

```
->ip ospf area 1.1.1.1
->ip ospf area 0.0.0.0 //Backbone Area
->ip ospf area 1.1.1.1 type stub //Stub Area
->ip ospf area 1.1.1.1 summary enable
->no ip ospf area 1.1.1.1
->ip ospf area 1.1.1.1 default-metric 0 cost 50
->ip ospf area 1.1.1.1 default-metric 0 type type1
->ip ospf area 1.1.1.1 range summary 192.5.40.1 255.255.255.0 effect noMatching
```

Creation/Deletion of an OSPF Interfaces

Once areas have been established, interfaces need to be created and assigned to areas. From SAM, interface has to be directly created under Area. The OSPF interfaces not assigned to any area will not be shown in SAM.

Area Site ->Interfaces ->Add/Delete

OR

From Routing Instance tree -> OSPFv2 ->Area . -> create Interface

Interface parameters can be configured in Protocol Properties Tab of interface Properties window. Interface Authentication can be configured from interface properties ->Authentication Tab.

CLI:

```
->ip ospf interface vlan-213
->ip ospf interface vlan-213 area 1.1.1.1
->ip ospf interface vlan-213 status enable
-> ip ospf interface vlan-213 auth-type simple
->ip ospf interface vlan-213 auth-key test
-> ip ospf interface vlan-213 auth-type md5
->ip ospf interface vlan-213 md5 7
-> ip ospf interface vlan-213 md5 7 key "test"
->ip ospf interface vlan-213 md5 7 disable
->ip ospf interface vlan-213 auth-type none
-> ip ospf interface vlan-213 dead-interval 50 cost 100
-> ip ospf interface vlan-213 poll-interval 25 priority 100 retrans-interval 10
-> ip ospf interface vlan-213 hello-interval 5000
```

Creation/Deletion of Virtual Links

A virtual link is a link between two backbones through a transit area.

Area Site ->OSPF Links->Virtual Links ->Add/Delete

OR

From Routing Instance tree -> OSPFv2 ->Area .. -> create Virtual Link

Specify neighbor router id and transit area id, configure protocol properties like intervals and Authentication on virtual link.

CLI:

```
->ip ospf virtual-link 0.0.0.1 2.2.2.2
```

Creation/Deletion of Static Neighbors

It is possible to configure neighbors statically on Non Broadcast Multi Access (NBMA), point-to-point, and point-to-multipoint networks.

Routing Tree ->OSPFv2 Instance -> Create Static Neighbor
OR
OSPFv2 Instance Properties -> Static Neighbor Tab->Add/Delete
Specify IP address of Neighbor and DR Eligibility for static neighbor.

CLI:
-> ip ospf interface vlan-213 type non-broadcast
-> ip ospf neighbor 1.1.1.8 eligible

Listing of Virtual Neighbors

The Virtual neighbors associated with virtual link can be seen under virtual link properties.

CLI:
->show ip ospf virtual-link
-> show ip ospf virtual-neighbor 0.0.0.0 10.0.0.1

Creation/Deletion of Directly Attached Hosts

The specified host must be directly attached to the router. ToS routing is the ability to make a forwarding decision based on a destination address and a desired Quality of Service (QoS). ToS routing allows link selection based on QoS when more than one path exists between a source and a destination. A metric value is the cost of all the hops necessary for a packet to reach its destination. Routers use the metric to determine the best possible path

OSPFv2 Instance Properties -> Host Tab -> Add/Delete.

CLI:
-> ip ospf host 172.22.2.115 tos 1 metric 10
-> no ip ospf host 172.22.2.115 tos 1

Listing of the OSPF object in the Routing Instance tree

If OSPFv2 is enabled, it will be shown under Routing instance tree. Same is applicable for VRF instances.

Graceful Restart configuration

Graceful restart can be configured from the OSPFv2 instance properties.

CLI:
->ip ospf restart-support planned-unplanned

Note: The route distribution between different protocols will not be supported through SAM.

RIP

RIP protocol support from 5620 SAM 8.0R1 is done only for OS9000E nodes.

Loading RIP

The user can load RIP from SAM by selecting RIP in: Routing Window->Routing Instance Properties->Protocols tab.

The RIP protocol can only be unloaded from node by taking out any occurrences of RIP in boot.cfg file.

CLI:

->ip load rip.

Enable/Disable RIP

The administrative state of the RIP can be changed from the SAM from: Routing Window->Routing Instance->Rip Properties->General Window tab.

CLI:

->ip rip status enable.

->ip rip status disable

RIP General Configuration

The RIP Site General window consists of global configuration parameters for RIP as listed in the Table 4.8.6.2.2.

CLI:

->ip rip host-route

->ip rip route-tag 0

->ip rip update-interval 45

->ip rip invalid-timer 270

->ip rip garbage-timer 180

->ip rip holddown-timer 10

Creation/Deletion RIP interfaces

The user has to first create the Routing Instance interface in order to create RIP Interface on the node. The RIP Interface can be created from: Routing Window->Routing Instance->RIP Properties->Interface Tab->Create.

The Routing Interface can be assigned to RIP Interface from RIP Interface General Tab. The authentication is provided for every RIP interface.

CLI:

->ip rip interface rip-1 status enable

->ip rip interface rip-1 metric 2

->ip rip interface rip-1 send-version v1

->ip rip interface rip-1 recv-version both

->ip rip interface rip-1 auth-type none

->ip rip interface rip-1 auth-key nms

Note: RIP instance can be enabled for multiple routing instance (VRF).

MPLS/LDP

The MPLS protocol is a licensed application and is restricted only to a licensed user. MPLS license has to be installed manually on OS9000E node. Please refer section “Installing the MPLS Software License” in OmniSwitch AOS Release 6 Network Configuration Guide.

The MPLS and LDP instance is automatically loaded when MPLS license is installed on the OS9000E node. Loading/Unloading MPLS/LDP protocol instances are not supported from SAM.

The license information will be shown in Network Element properties->Licenses Tab.

Enable/Disable MPLS/LDP Instance

The MPLS protocol instance can be enable/disable from MPLS protocol properties window by setting administrative state up/down.

The LDP protocol instance can be enable/disable from LDP protocol properties window by setting administrative state up/down.

CLI:

```
->configure router mpls shutdown
->configure router mpls no shutdown
->configure router ldp shutdown
->configure router ldp no shutdown
```

LDP Configuration

From Routing Tree view, open the LDP instance properties to configure LDP general properties, interfaces, graceful restart, targeted session and targeted peer.

The Graceful Restart can be configured under Common Tab.

The global interface parameters can be configured under Interface Properties Tab.

The Interface can be added to LDP protocol from Interface Tab.

The Targeted peer parameters can be configured under Targeted Peer properties Tab.

The Targeted Peer tab will list all the targeted peers available.

CLI:

```
->configure router ldp interface-parameters interface vlan-40
->configure router ldp interface-parameters interface vlan-40 hello 50 10
->configure router ldp interface-parameters hello 40 2
->configure router ldp targeted-session hello 20 2
->configure router ldp interface-parameters transport-address interface
```

Static LSP Configuration

A Static LSP (tunnel) is a user-defined path of Label Switching Routers (LSRs). Configuration of label mappings and MPLS actions is required on each router that will participate in the static LSP. Signaling protocols, such as LDP, are not required and there are no dependencies on the IGP topology or local forwarding table.

Static LSP can be created from Manage->Static LSPs->Create.

The user can provision source, destination, egress label and next-hop to create static lsp.

Static hops can be created from the Static hops tab for LSP.

CLI:

```
->configure router mpls static-lsp to-R3
->configure router mpls static-lsp to-R3 to 10.10.10.3
->configure router mpls static-lsp to-R3 push 777 next-hop 192.168.10.1
->configure router mpls interface vlan-1 label-map 666
->configure router mpls interface vlan-1 label-map 666 pop
->configure router mpls interface vlan-1 label-map 666 no shutdown
```

MPLS Interfaces

An MPLS interface is required on each switch (ingress, transit, and egress) that will participate in the static LSP tunnel.

To configure interfaces for MPLS protocol, open MPLS instance properties and add interfaces to interfaces tab. User can also create MPLS interface directly from MPLS instance by create action. Label map can be created in “Static Label Map” Tab under MPLS interface properties. User can also configure Static Fast Reroute Feature available only on AOS9000E by selecting Label Action as “Swap/Protect-Swap”.

CLI:

```
->configure router mpls interface mpls-vlan-10
->configure router mpls interface mpls-vlan-10 label-map 777 swap 888 next-hop
192.168.30.1
```

Note: MPLS and LDP configuration is applicable only under default Routing instance. MPLS and LDP configuration is supported only on AOS9000E nodes in 5620 SAM 8.0 R1.

VRF Instances

Starting in 5620 SAM 8.0R2, creation and deletion of VRF instance is supported from Network Tree view.

- OmniSwitch 9800E
 - Bridge Instance -1
 - Routing Instance - default
 - Routing Instance - vrf1
 - Routing Instance - vrf2

The same level of configuration is supported on each VRF instance which is supported on default routing instance (Routing instance -1).

- IP Routing and Network Interface
- Static Routes

In addition:

- Full Routing protocol configuration - OSPFv2, RIP
- QoS Policy support for vrf instance. (New parameter is added to QoS Condition form to support VRF instance name)
- Full Routing protocol configuration -PIM under each VRF instance in 5620 SAM 8.0 R3.
- UDP/DHCP Relay on VRF instance in 5620 SAM 8.0 R3.

Following VRF-related configuration will not be supported in 5620 SAM 8.0:

- AAA Radius Server Configuration on VRF instance
- Routing protocol BGPv4
- UDP/DHCP Relay on VRF instance

OMNI Scalability

The maximum number of OS6850 nodes supported by a SAM is 5000.
The maximum number of OS6855 nodes supported by a SAM is 5000.

The maximum number of OS6400 nodes supported by a SAM is 5000.
The maximum number of OS6250 nodes supported by a SAM is 5000.
The maximum number of OS9800 nodes supported by a SAM is 1000.
The maximum number of OS9700 nodes supported by a SAM is 1000.
The maximum number of OS9600 nodes supported by a SAM is 1000.
The maximum number of OS9700E nodes supported by a SAM is 1000.
The maximum number of OS9800E nodes supported by a SAM is 1000.

7210 SAS Support

5620 SAM Release 8.0 provides support for the 7210 SAS MX node.

7210 SAS M devices are service provider managed customer premise or last mile aggregation equipment that provides rich Ethernet services to Business or residential customers. Most of the customer premise and last mile aggregation devices that exist in the market today are either cheap plain vanilla switches or very expensive high end switch routers. 7210 SAS M devices provides a feature rich but low cost carrier grade platform that enables service providers to offer multiple services over Ethernet typically to a single customer to or to a small number of customers. These devices extend the feature rich Ethernet services to the customer edge enabling the service providers to meet the end-to-end SLAs and OAM capabilities. These devices enable carriers to provide end to end Ethernet access services in a cost effective way. They help carriers in cutting down on OPEX through comprehensive set of end-to-end OAM tools and seamless integration with Alcatel-Lucent SAM 5620.

7210 SAS MX w/fixed 10G port support is being targeted at customers who require more than a few gigabits (i.e. about 5-6Gbps) of bandwidth to backhaul in the metro access and aggregation network. Customers, who have larger data volumes, want to move from use of 1G ring or multi-homed topologies to 10G topologies. Use of 10G links in these scenarios, is more cost effective than laying out multiple 1G links and aggregating them.

The card will already be equipped in the 7210 SAS M nodes with all 24 ports default to Network mode with null encapType. It would have 24 Ethernet ports. All features and functions support by the 60 port 10/100 Ethernet Port MDA in 7450 5.4 are supported by this MDA in 7210 SAS-M 1.1R3 with the following exceptions.

Unsupported Ethernet Functions for SAS-M and SAS-MX

- QoS Pool Tab not supported
- If ModeType is Access, only Null and Dot1Q Encap are supported.
- QinQ Encap type not supported

A 7210-M node running a 2.0 version or later load will support enabling/disabling LACP Tunnel on the Ethernet port. This feature is inline with the SR node.

The 7210 SAS MX supports a 26 port Ethernet Daughter Card. This daughter card type shall be displayed in the Supported Daughter Card Types list as:

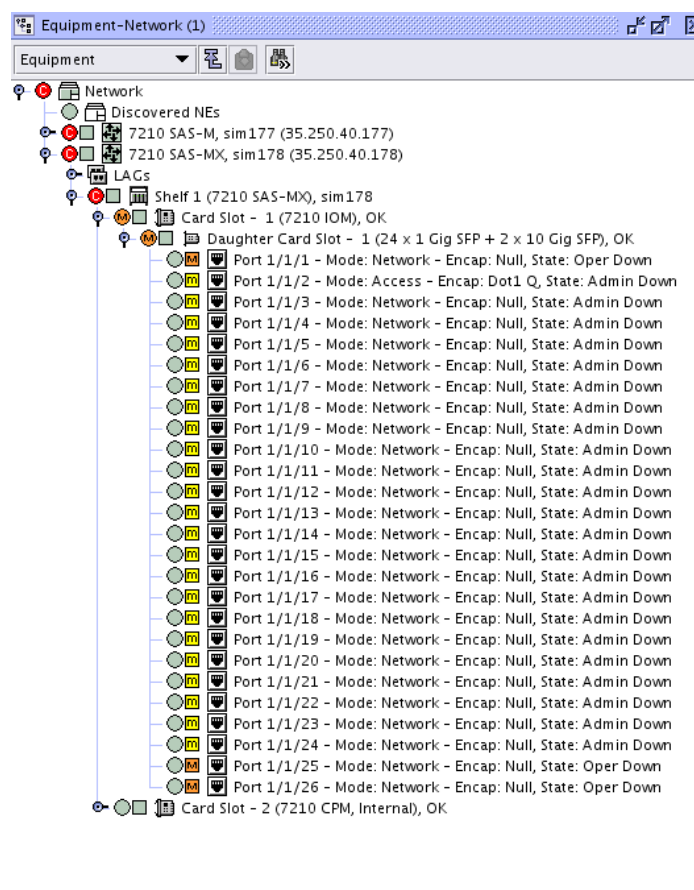


Figure 31: 7210 SAS-MX MDA Support

All features and functions support by the 24 port 10/100 Ethernet Port MDA in 7210 SAS-M are supported by this MDA.

SFP Diagnostics and Monitoring

The following parameters are available in the Port details:

SFP and XFP status: This indicates operational status of the inserted SFP/XFP. This is a status indicator and is non-editable.

DDM Event Suppression: This parameter is used to enable/disable the event reporting related to digital diagnostics monitoring from the node.

Since the digital diagnostics is supported only on the SFP enabled ports, the above parameters will not be displayed in the detail view of the ports which are not SFP enabled. For example the ports on the CES MDA are not SFP capable and so these parameters will not be displayed there.

A new tab “DDM” is provided in the port details which display the statistics information as received from the node. A snapshot of the Port screen is provided in the figure below:

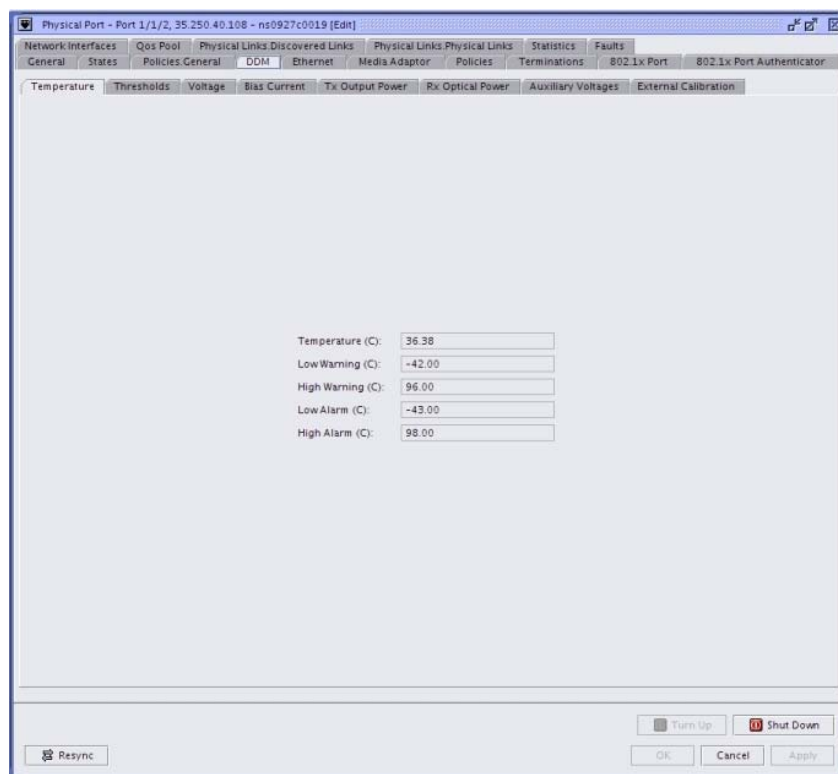


Figure 32: DDM Configuration

9500 Support

The 9500 is a family of products that provide Ethernet based backhaul of traffic via line of site microwave transmission. The 9500 family consists of the following:

- 9500 MSS: the 9500 Microwave Service Switch (MSS) provides a number of data control plane functions including radio management, access interfaces and termination, service type recognition, service classification, statistical multiplexing and service tunnel (pseudo-wires) for the transport of traffic across the network. This platform will ultimately be available in 3 different chassis sizes: MSS-8, MSS-4 and MSS-1.
- 9500 MXC: a microwave dish (Outdoor Unit - ODU) that connects to a 9500 Microwave Service Switch (MSS) via a RF modem located in the MSS. This unit will have both ETSI and ANSI versions.
- 9500 MPT: a microwave dish that connects to a 9500 MSS via a GigE interface located on the control card of the 9500 MSS. The MPT may also be deployed in standalone fashion i.e. directly connected to a 7705 SAR via GigE. The MPT supports many of the same service oriented function that are found in the MSS. The MPT will be available in ETSI/ANSI variants as well as versions of varying capacity and reach.

The 9500 MSS and 9500 MPT are marketed together as a system which is called the 9500 Microwave Packet Radio (MPR).

One of the primary network applications for the 9500 MPR is the backhaul of voice and data traffic (2G, 3G and 4G) across a mobile carrier's Radio Access Network (RAN). In the scenarios in which SAM will be deployed with the 9500 there will also be a significant 7705/7x50 component to the network.

Network Architecture

Microwave networks are inherently point-point in nature. The reference topology for a 9500 MPR microwave based RAN is shown below.

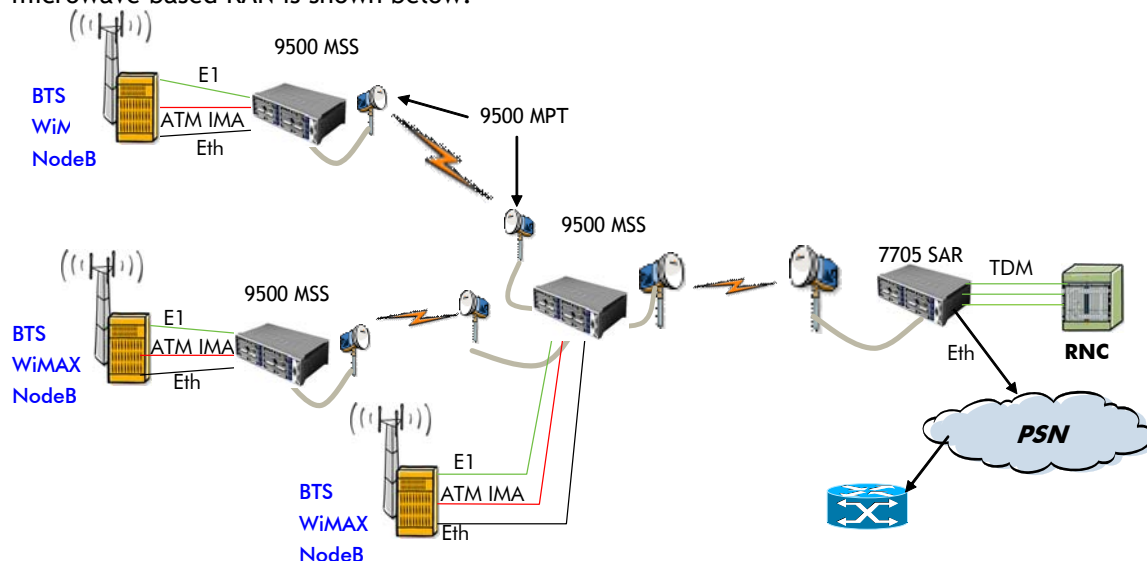


Figure 33: Reference Architecture for 9500

The physical topology is a simple ‘binary tree’ structure. This tree structure can be extended significantly in terms of its depth until the combined bandwidth requirements of all of the access (9500) nodes exceeds the capacity of the highest order radio link in the network.

In the access portion of the network the 9500 provides connectivity to the base stations using a variety of different L1/L2 technologies depending on the ‘mobile generation/technology’ and the BS vendor’s specific implementation.

Pseudo wires are used to create point-point service connections from the 9500 MSS at the BS and carry the traffic over the microwave based RAN to the 7705 SAR connected (directly or indirectly) to the MSO. The type of PW used depends on the connectivity between the BS and the 9500 MSS. Note that in the initial release (R1.0) of the 9500 MSS the ATM and TDM PW are based on MEF.8. Ethernet will be carried over a Ethernet VLAN (no PW). A subsequent release (R1.1) supports standards based TDM/ATM/Ethernet PWs. In both R1.0 and R1.1 of the 9500 the PWs will be statically setup. The L1/L2 technology used to connect the microwave network to the fixed portion of the network also varies between R1.0 and R1.1 of the 9500 MSS. In R1.0 the MEF.8 PW will be terminated on a 9500 MSS (not shown in Figure) adjacent to the 7705. The L1/L2 handoff to the 7705 will be the same as is used in the access to connect to the 9500 MSS for a specific traffic flow. In R1.1 the handoff between the 9500 MPT and the 7705 will be Ethernet. The service endpoints for the PWs will reside on the 7705.

UMTS Backhaul Architecture

Case 1 - TDM

In this scenario the BTS is connected to the 9500 via a TDM link (T1/E1). The traffic associated with the TDM link will be carried over the transport network on a single TDM PW. In R1.0 the TDM PW will be terminated on the 9500 MSS and the handoff will be via a TDM interface to the 7705. In R1.1 the TDM PW will terminate on the 7705 (The pattern of handoff/PW termination is the same for all of the following cases and will not be repeated).

Case 2 - 3G (HSDPA) Backhaul

In this scenario the Node B is connected to the 9500 via T1/E1 ATM (ASAP). The connection may be a VCC, VPC or n x T1/E1 IMA. The VCC, VCP or IMA connections are terminated on the 7705. The traffic is carried over the transport network to the RNC on a ATM PW - there is a 1-1 mapping between VCC, VPC or IMA and an ATM PW.

Case 3 - Ethernet Backhaul

In this scenario the WiMax BS is connected to the 9500 over Ethernet. The traffic may be tagged or untagged (aka raw mode). The 7705 will map traffic from individual VLANs (tagged mode) to a VLAN/Ethernet PW (1-1) or map all traffic (untagged) to a single VLAN/Ethernet PW.

CDMA Backhaul Architecture

Case 4 - 2G Backhaul

In this scenario the 1x BTS is connected to the 7705 via one or more TDM links (T1/E1). The link layer used is BS vendor dependant and is one of Frame Relay/Cisco HDLC/PPP. The 9500 will not terminate the layer 2 connection but map the T1/E1 into a TDM PW and backhaul the traffic over the transport network to the 7705.

Case 5 - 3G Backhaul over MLPPP

In this scenario the EVDO BS is connected to the 7705 over multiple T1/E1 links. The link layer in this case is MLPPP. The 9500 MSS will not terminate the MLPPP but map each T1/E1 link to a TDM PW and backhaul the traffic over the transport network to the 7750.

Case 6 - Ethernet Backhaul

This scenario is identical to that described above in case 3.

Transport Tunnel

The 9500 MSS/MPR requires support for a new type of PW transport tunnel - namely a VLAN tunnel (VLAN Path). These VLAN based tunnels are implemented on the 9500 given that there is no MPLS or GRE support on the 9500. There is a 1-n relationship between VLAN Path and VLAN Path Instances using that path.

7705 Support

7705 SAR R3.0 R1 functionality supported in 5620 SAM 8.0 R1

4 Port DS3/E3 ASAP Daughter Card

The 7705 SAR supports a 4 port DS3/E3 ASAP Daughter Card. This daughter card type shall be displayed in the Supported Daughter Card Types list as:

"4 x Channelized DS3/E3 ASAP"

Each DS3/E3 port on this daughter card can support only a DS3/E3 channel. This channel configuration is restricted as follows:

- ATM or PPP Auto encapsulation types for DS3
- PPP Auto encapsulation type for E3
- MTU for ATM encapsulation is 1524

- Clock Source default is Node Timed

The behavior of the DS3/E3 type property of the DS3/E3 port is as described for DS1/E1 ports on the 16 Port DS1/E1 ASAP Daughter Card.

The behavior of the encapsulation type property for this DS3/E3 channel is as described for DS0 Channel Groups on the 16 Port DS1/E1 ASAP Daughter Card. Channels on this daughter card can NOT be configured as Management IES SAPs.

8 Port Ethernet Daughter Card Enhancements

SSM

The 7705 SAR supports the enabling and disabling of Synchronous Status Messages for ports on the 8 Port Ethernet daughter cards.

Egress Scheduler Configuration

The 7705 SAR supports the configuration of the Egress Scheduler mode for network ports on the 8 Port Ethernet daughter cards. The Egress Scheduler mode can be configured for one of two values:

- Profile (default, and equivalent to operation of the scheduler prior to this release)
- Four Priority

This feature is not supported by the SR family.

16 Port DS1/E1 Daughter Card Members Per IMA/MLPPP Bundle

The 7705 SAR supports up to 16 members per MLPPP bundle on the 16 port DS1/E1 ASAP daughter card.

The 7705 SAR supports up to 16 members per IMA bundle on the 16 port DS1/E1 ASAP daughter card.

12 Port Serial Data Interface Daughter Card Port Display Enhancement

In Release 7.0 of 5620 SAM, the properties form of a port on the Serial Data Interface card has a Serial tab which contains the Type property. That tab has been removed, and the property has been moved to the General tab of the port properties form.

802.1ag/Y.1731 Ethernet OAM Enhancements

Existing Y1731 Tests From SR 7.0

The 7705 SAR supports EthTest, One Way Delay, and Two Way Delay tests in a manner that is consistent with Release 7.0.R1 of the SR family.

Additional MEP Counters

The 7705 SAR supports two counters under the Y1731 tab of the MEP properties form, for the number of LMR and DMR frames transmitted by the MEP. These counters are not supported by the SR family.

Single-Ended Loss Measurement Test

The 7705 SAR also supports the Single-ended Loss Measurement test. This test is not supported by the SR family. Support for this test in 5620 SAM is implemented in a manner that is consistent with existing Y.1731 tests.

The following results are reported by 5620 SAM for a Single-ended Loss Measurement:

- Duration: Test duration, in seconds.
- Received LMR/CCM Frames: Number of LMR frames received for this test.
- Tx Frames - Peer: Number of frames transmitted during this test by the far-end..
- Rx Frames - Local: Number of frames received during this test by the near-end
- Lost Frames - Local: Number of frames lost during this test at the near-end. This value is calculated by subtracting “Rx Frames - Local” from “Tx Frames - Peer”.
- Lost Frame Ratio - Local: Percentage of frames lost during this test at the near-end. This value is calculated by taking the ratio of “Lost Frames - Local” to “Tx Frames - Peer”.
- Tx Frames - Local: Number of frames transmitted during this test by the near-end.
- Rx Frames - Peer: Number of frames received during this test by the far-end.
- Lost Frames - Peer: Number of frames lost during this test at the far-end. This value is calculated by subtracting “Rx Frames - Peer” from “Tx Frames - Local”.
- Lost Frame Ratio - Peer: Percentage of frames lost during this test at the far-end. This value is calculated by taking the ratio of “Lost Frames - Peer” to “Tx Frames - Local”.

VLL Redundancy - Active/Standby

The 7705 SAR support for active/standby VLL redundancy is consistent with that for Release 8.0 of the SR family. Details of 5620 SAM support for this feature on 7705 SAR are included in the FS for this feature.

VLL Switching Site

The 7705 SAR supports VLL Switching sites.

The 7705 SAR allows Spoke SDP Bindings to be assigned to Endpoints on a VLL Switching site. The SR family does not allow this assignment. The restriction for SDP redundancy on 7705 SAR VLL Switching sites is that the Spoke SDP Binding toward the non-redundant destination must have T-LDP disabled. Note that auto-selection of service tunnels does NOT account for this restriction; consequently, an auto-selected tunnel for the non-redundant Spoke SDP Binding may have T-LDP enabled, and cause a server exception when Apply/Ok is selected. In this situation, it may be necessary for the 5620 SAM user to manually select a tunnel for the non-redundant Spoke SDP Binding.

VPRN

The 7705 SAR adds support for VPRN. This functionality is the same as that offered by Release 6.0 of the SR family, except that the 7705 SAR does NOT support:

- Certain properties and statistics under the DHCP tab of the L3 Access Interface
- Certain properties under the Routing tab of the Site
- BGP or OSPF under the Protocols tab of the Site

Furthermore, L3 Access Interfaces for VPRN services must be associated with IPCP DS0 Channel Groups or bundles on the 16 port DS1/E1 ASAP daughter card or ports on the 8 port Ethernet daughter card. If a daughter card has L3 Access Interfaces for a VPRN service, then the Access Ingress Fabric Profile policy associated with that daughter card must be configured for Aggregate mode.

Management IES

Management IES is NOT supported in 7705 SAR 3.0.R1.

Routing Policy

The 7705 SAR adds support for the Community, Damping, and AS Path policies, as supported by Release 6.0 of the SR family.

The 7705 SAR Routing Policy supports the same configuration for Default Action, Accept Action, From Criteria, and To Criteria as Release 6.0 of the SR family, except that the 7705 SAR does NOT support:

- Certain values of Protocol for To Criteria
- Certain values of Protocol for From Criteria
- Certain values of Family for From Criteria
- OSPF Instance ID for To Criteria
- OSPF Instance ID for From Criteria
- Multicast properties for From Criteria

IP Filter Policies

The 7705 SAR no longer has restrictions associated with the configuration of the 64 IP Filter entries per policy. This enhancement applies to both ACL IP Filters and CPM IP Filters.

Performance Statistics

Network Port Control Packet Performance Statistics

The 7705 SAR supports two sets of counters associated with network port control packets (ingress and egress).

Y1731 MEP LBR Performance Statistics

The 7705 SAR supports two counters for LBR frames transmitted (with and without TLV).

ACL IP Filter Entry Statistics

The 7705 SAR adds support for the Ingress Hit Byte Count within the Hit Count record for ACL IP Filter Entries.

Route Statistics

The 7705 SAR adds support for statistics related to route aggregation and BGP within the Route Stats record of VPRN sites and the Routing Instance.

Detailed Packet Discard Statistics Modification

The 7705 SAR adds support for inCsmQMediumPriDiscards, and obsoletes inCsmQFtpPriDiscards in the Detailed Packet Discard statistics record.

FIB Statistics

The 7705 SAR adds support for FIB statistics associated with BGP.

IP Service Tunnels

The 7705 SAR supports service tunnels with an Underlying Transport type of IPv4 for any VLL service. This value is not supported by the SR family. Service tunnels with IPv4 transport type must be manually created and bound to SDPs; rule-based auto-creation, service auto-bind, and tunnel auto-selection are NOT available for this tunnel type. Furthermore, the transport type of a service tunnel can only be set to IPv4 if the Source Node is 7705 SAR Release 3.0+, and the Destination Node is either 7705 SAR Release 3.0+, GNE, or an unmanaged IP address. Tunnels with IPv4 transport type are NOT supported by 5650 CPAM. As such, the Monitor and Navigate functions are not available for these tunnels¹.

DHCP Relay On Network Interfaces

The 7705 SAR supports DHCP Relay on a Network Interface. This is not supported by the SR family. The configuration of DHCP for Network Interfaces is identical to that for L3 Access Interfaces on the 7705 SAR, except that DHCP for Network Interfaces:

- Does not support Vendor Specific Options
- Supports enabling of “Copy Option 82 To Option 43”
- Supports additional circuit ID values: port ID, interface name

Routing Instance AS Number

The 7705 SAR supports an Autonomous System number under the Routing tab of the Routing Instance properties.

Route Aggregation

The 7705 SAR supports various properties associated with AS numbers under the Route Aggregation tab of the Routing Instance properties. The Route Aggregation supported by the 7705 SAR is therefore now identical to that supported by Release 6.0 of the SR family.

IP Addresses

The 7705 SAR supports the Broadcast Address Format property for VPRN L3 Access Interface Addresses. This property remains unsupported for Network Interface Addresses.

BGP

The 7705 SAR supports BGP. This functionality is the same as that offered by Release 6.0 of the SR family, except that the 7705 SAR does NOT support:

- Multihop
- MED
- Keychain
- Inter AS VPRN
- Confederations
- IGP
- Minimum TTL value

7705 SAR 3.0 R1 Functionality added in 5620 SAM 8.0 R3

CEM For 2 Port Channelized OC3/STM1 Daughter Card

The 7705 SAR supports CEM encapsulation for DS0 channel groups on the 2 Port Channelized OC3/STM1 daughter card.

OneToOne LSP Fast Reroute

The 7705 SAR supports the oneToOne type of LSP fast reroute. The configuration of LSP fast reroute type supported by the 7705 SAR is therefore now identical to that supported by Release 6.0 of the SR family.

7705 SAR 3.0 R1 Functionality NOT currently supported

The following features are supported in 7705 SAR 3.0 R1, but are NOT currently supported in 5620 SAM.

- **Timing Reference Quality Level** — The 7705 SAR supports timing reference quality level configuration for all three timing references: Reference One, Reference Two, and External. The 7705 SAR support for timing reference quality level is consistent with that for Release 8.0.R4 of the SR family. Support for this feature within 5620 SAM is planned for 5620 SAM 8.0 R5.
- **Y1731 Dual-Ended Loss Test** — The 7705 SAR supports the Dual-Ended Loss Test. The test is enabled as an option within the CC Test. Support for this feature within 5620 SAM is planned for 5620 SAM 8.0 R5.

GNE Support

Hardening of GNE profile

The GNE profile has been hardened in its ability to deal with CLI dialog that has different forms of end of line characters.

The GNE profile has been extended with the ability to specify a confirmation prompt as part of the CLI login process. After the login prompt and password the profile now allows specifying of a response string to a confirmation prompt.

The GNE profile has been extended with the ability to select one out a fixed set of 6 icons that should be displayed for devices discovered with the GNE profile.

LLDP support (GNE)

SAM now supports automated discovery of L2 links between GNE and managed routers that support link layer discovery (802.1ab). The support requires that the GNE has compatible implementation of 802.1ab and interworks with the managed node. SAM does not support the MIB on the GNE but uses the information on the managed node to figure out the physical link. As a consequence links between GNE are not supported.

4. ASSURANCE

Alarm Timestamps

This feature sets the firstTimeDetected and lastTimeDetected timestamps on the alarm when the trap from the node is received instead of when the alarm is created.

This improves the timestamp accuracy to show a time which is closer to when the event actually occurs.

In the past the timestamps could be off by several minutes, whereas now the timestamps should be much more accurate and only be a few seconds off.




Figure 34: Alarm Timestamps

Test Suite support for CFM tests

CFM Loopback, Link Trace, One-way and Two-way delay, and Ethernet tests are supported by STM test suite.

The user shall be able to create the Test Policy (TP) with the Entity Type: "Ethernet"

The TP Test Generation has an option to generate the Ethernet tests between every MEP. The Test Generation Strategy option is "Mesh MEPs".

The "Mesh MEPs" generation strategy finds all target MEPs in the listed MA entity list.

For each MA entity MEP it queries all MEPs in the same MA/MD and Level for their operational MAC addresses

CFM tests (defined in test definitions) for every operational MAC MA entity MEP are generated with a target MAC for all operational MEP MAC returned in the above query.

Future possibilities include a Mesh MEP and MIP option.

Updated MEPs /MAs should be automatically reflected in the generated tests list.

Test Policy, [Create]

General Test Definitions Usages

ID: 0 ☒ Auto-Assign ID

Name:

Description:

NE Schedulable: ☐

Testing Policy

Entity Type: Ethernet

Test Generation

Strategy: Mesh MEPs

Reset OK Cancel Apply

Figure 35: TP-General tab

The TP type “Ethernet” has the following test definition options: CFM Loopback, CFM Link Trace, CFM Eth Test, CFM Two Way Delay Test, CFM One Way Delay Test.

The Direction property is configurable on the test definition form.

The General tab of the Test Suite (TS) includes the “Ethernet” item under the “Entity Type” drop-down list.

If the “Ethernet” option is chosen then the “First Run”, “Last Run” tabs are not applicable.

Figure 36: TS-General tab

The user is able to select the MA by using the Tested Entities tab of the TS form:

Site ID	Maintenance Domain ID	Maintenance Association ID	Maintenance Association Name	Maintenance As
10:work:10.01.185.133	2	5	string	ee
10:work:10.01.185.134	2	5	string	ee
10:work:10.01.185.135	2	5	string	ee

Figure 37: TS-Tested Entities tab

Below is a list of other test suites that are extended to support CFM test generation.

- VLL (Ipipe, Epipe)
- VPLS
- VPRN

A Stop Test button is added to Test Suite to recursively call the child tests to stop execution. This is intended to help support the SAA continuous CFM tests.

The “Ethernet” Test Entity Generation shows “Test does not apply to non-Y.1731 MA's” generation log messages when Y.1731 Tests (Eth Test, OneWay & TwoWay Delay) are attempted to be generated on non-Y.1731 Maintenance Associations (Maintenance Association Name Type = “icc-based”)

OAM Results on Map (including CFM)

Service test result summary is now displayed on the service map if the test has been invoked from that same map. Only a few key parameters for a test is displayed. This info box is linked to the source object of the test.

As multiple results can be displayed on view, user can click on the "x" button to close the selected result. Also, there should be a pointer on that info box, which when clicked, STM shall bring up the usual detailed test result.

Ethernet SAP OAM (Mapping)

OAM mapping is a mechanism that enables a way of deploying OAM end-to-end in a network where different OAM tools are used in different segments. For instance, customer could have an EPIPE service which spans across the network using Ethernet access (CFM used for OAM), PW (T-LDP status signaling used for OAM), Ethernet access (E-LMI used for OAM). Another example is ipipe service, where one end is Ethernet, the other end is FR or ATM.

Fault propagation for a MEP can only be enabled when the MA that the MEP belongs to has remote MEP(s) configured as one of the following conditions:

- 23 No remote MEP configured
- 24 Only one remote MEP configured
- 25 Two remote MEPs configured and at least one of them is a local MEP

Fault propagation cannot be enabled for eth-tun control MEPs (MEPs configured under the eth-tun primary and protection paths).

Modify AIS selection

AIS level selection will allow the same level as the MA to be selectable. Currently only levels higher than the MD level are selectable.

Add MEP Fault Propagation

From the MEP Screen a new Fault Propagation selection is available (use-if-tlv or suspend-ccm) (New MEP Property TmnxDot1agCfmMepFaultPropagation.)

The value of tmnxDot1agCfmMepFaultPropagation specifies what action should be taken by the MEP if a fault is detected in the service containing the MEP.

Down MEP on Ipipe Ethernet SAP

SAM OAM has extended MEP support to include Down MEP SAPs on Ipipe Services

SAM shows the MEP tab on SAP/SdpBindings in Epipe, IES and VPRN services, and users are able to create MEPs via the standard Add button.

Epipe Service & 3.5 MIP for Epipe

SAM OAM shall extend OAM support to include Up MEPs and MIPs for SAP/SdpBindings on Epipe services. SAM currently supports only down MEPs on Epipe.

Enable MEP/MIP on Epipe Service

SAM will show the MEP & MIP tabs on SAP/SdpBindings in Epipe services, and users will be able to create Only Up MEPs via the standard Add button. MIP's will be created via the MipEnabled flag on the SAP/Sdp and the MAservice MHF-Creation default/explicit.

MEP, [Create]

General Remote MEP DB State

MEP

ID: 0 ☒ Auto-Assign ID

Maintenance Association

Maintenance Domain ID:

Maintenance Association ID:

Maintenance Site ID: 0.0.0.0

Maintenance Service ID: 0

Direction: Up

Administrative State: Down

CCM Messages Enabled: ☐

Priority Level For CCM Messages: 7

Low-priority Defect: macRemErrXcon

Mac Address: 00-00-00-00-00-00

Type: Regular

SAP, BINDING or PATH ENDPOINT: SAP

SAP/Binding

SAP

Name:

Port ID: 0

Encap Type: N/A

Service ID: 0

Service Name:

Operational State: Unknown

Administrative State: Up

Reset OK Cancel Apply

Auto MEP/MIP Creation

MEP Generation will function as currently implemented.

Service & Composite-Service OAM

Ethernet OAM MEG/MA and MEP can now be created from the service map. In addition, selected 1ag and Y.1731 tests can be created or executed on service map. The functions can be enabled with a check on the OAM config box at the bottom left of the service map.

Despite the name, this feature only support VLL and VPLS service. VPRN, VLAN, and composite service shall be addressed in a near future.

MAC address provisioning for MIP

SAM OAM shall extend MIP enabled flag on SAP/SdpBindings general tab to support a user configured MAC address.

Currently 7.0R4, MIPs are only enabled or disabled on SAP/ SdpBindings

Add MAC property to MIP

The MIP enabled flag on supported SAP/SdpBindings will also allow for a user entered MAC address (tmnxDot1agCfmSapMipSrcMacAddress).

Fast CCM timer for all MEPs

SAM OAM shall extend CCM transmit interval to allow fast CCM at 10ms and 100ms on SAP/SdpBindings where MEPs are provisioned.

Fast CCM must be configured for the lowest level MEP and one fast CCM per SAP/SdpBinding.

From the Maintenance Association screen fast CCM transmit interval selections will be available (10ms and 100ms), this was previously only applicable to Ethernet tunnel MEPs or 7705.

CCM Interval Enumeration:

Indicates the interval at which CCMs are sent by a MEP.

The possible CCM Interval values are:

intervalInvalid(0)	No CCMs are sent (disabled).
interval300Hz(1)	CCMs are sent every 3 1/3 milliseconds. (Not Used but in MIB)
interval10ms(2)	CCMs are sent every 10 milliseconds. (Fast MEP)
interval100ms(3)	CCMs are sent every 100 milliseconds. (Fast MEP)
interval1s(4)	CCMs are sent every 1 second.
interval10s(5)	CCMs are sent every 10 seconds.
interval1min(6)	CCMs are sent every minute.
interval10min(7)	CCMs are sent every 10 minutes.

Maintenance Association - sas:CTest-2:site-10.106.185.113 [Edit]

General | Template | Service | Local MEP | Remote MEP | MIP | Faults

Site

System ID (Loopback IP Address): 10.106.185.113 [Properties](#)

Management IP Address: 138.120.185.113

Maintenance Domain ID: 10

Maintenance Association ID: 3

Maintenance Association Name Type: string

Maintenance Association Name: ma1

CFM Test

Global ID: MD:hghghgh-MA:ma1 [Clear](#) [Properties](#)

Service ID: 13

CCM interval: 10 s

10 ms
100 ms
1 s
10 s
60 s
600 s

[Resync](#) [Reset](#) [OK](#) [Cancel](#) [Apply](#)

SAA Support for Ethernet CFM OAM tests including Y.1731

SAA functionalities are extended to support scheduling Ethernet CFM OAM tests. Test results are monitored using configured event thresholds, including jitter-event, latency-event, and loss-event. SNMP traps must be supported as well.

In order to support jitter and latency measurements, additional timestamping is available for non Y.1731 delay-measurement tests, to be specific, loopback and linktrace tests. An Organization-Specific TLV can be used on both sender and receiver nodes to carry the timestamp information.

CFM loopback, linktrace and DM SAA tests need to support send-count, interval, timeout, and FC. But existing CFM OAM commands will not be extended to support send-count and interval natively in 8.0.

The new tests supported are: *CFM Loopback*, *CFM Tracelink*, and *CFM Two-way Delay*.

Test Type	Thresholds Types
CFM Loopback	Jitter-in, Jitter-out, Jitter-rt, Latency-in, Latency-out, Latency-rt, Loss-in, Loss-out, Loss-rt
CFM LinkTrace	Jitter-in, Jitter-out, Jitter-rt, Latency-in, Latency-out, Latency-rt, Loss-in, Loss-out, Loss-rt
CFM Two Way Delay	Jitter-in, Jitter-out, Jitter-rt, Latency-in, Latency-out, Latency-rt, Loss-in, Loss-out, Loss-rt

Supported NeThresholds

Testsuites allow these Ethernet tests to be added.

SAA are extended to support Ethernet CFM OAM tests results. SAM parses Ethernet CFM OAM tests from the SAA Accounting file.

To specify a test as using the SAA Accounting files for its results, the test must be generated as part of a test suite and the test policy must specify “NE Schedulable” and “Accounting files”. If these conditions are met, then, on a supported node, the tests use the SAA Accounting file on the SR to store the results.

SAM (s)ftps the accounting file and parses it for the results it contains.

Continuous SAA CFM OAM Tests

SAM OAM extends CFM tests to support a continuous execution mode.

Currently, SAA test has a limited send-count, which means the test runs a certain number of probes then stops. 5620 SAM adds support for continuous SAA CFM tests, meaning the SAA test runs indefinitely until it is stopped by the operator.

The SAA test is still provisioned with a limited number of probes, because it allows easier statistics collection and output to the accounting file. The test however has a new flag, which is used to indicate that the test will run indefinitely by iteration. In every iteration of the test, SAA performs all the probes, collects the results and writes them to the accounting files. It is a requirement that the next iteration will run immediately after the last one finishes (response received or timeout). The time gap between test iterations should not be greater than the configured timeout value.

The three CFM tests involved in this upgrade are:

- CFM Loopback
- CFM Link Trace
- CFM Two Way Delay Test

CFM tests can identify a continuous run mode where the node will continuously run a test until a stop is requested, stops will use the current stop execution GUI button.

The following limits are based on the nodes and are enforced by SAM:

Item	Version 6.X	Version 7.X	Version 8.X
Max number of MDs per TiMOS system:	30	50	50
Maximum number of MEPS (local and remote) per MA:	64	64	64
Max number of MAs per TiMOS system:	1000*	2000	2000

*Note: For release 6.0, the node restricts this to 120 MAs per TiMOS. For future releases it is increased to 1000.

The following limits are based on the nodes and are not enforced by SAM:

Item	Version 6.X	Version 7.X	Version 8.X
Max number of MEP per TiMOS system:	1000	2000	5000
Max number of MEPS/MIPs per TiMOS system:	1000	2000	5000
Max number of MEPS (with CC enabled) per TiMOS system:	200	200	200

Max number of SAA tests: 50000

5. SERVICE & ROUTING ENHANCEMENTS

VLL Redundancy Enhancements

Network Redundancy

SAM currently colors VLL links based on the following operational and forwarding states.

Operational State	TX Active State	Overall State	Color
Link UP	Active	Active	Green
Link UP	Active Forced	Active	Green
Link UP	Backup	Standby	Purple
Link DOWN (Standby SAP on MC-LAG or MC-Ring)	N/A	Standby	Purple
Link DOWN	N/A		RED

Table 12: VLL Link States

Currently VLAN Up-Links are only created on SAPs using LAGs when the service is a B-VPLS/M-BVLPS. SAM shall extend this support to VLL services. The following figure depicts this scenario:

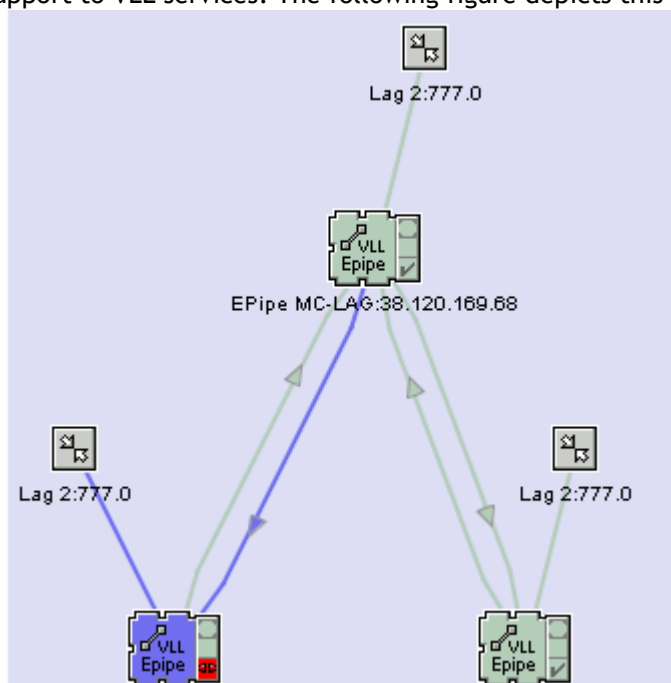


Figure 38: MC-Lag VLL Link State

SAM shall be enhanced to create VLAN Up-links on MC-LAG and SC-LAG SAPs and color the links based on the following operational and forwarding states.

Operational State	Overall State	Color
Link UP (Active SAP on MC-LAG)	Active	Green
Link DOWN (Standby SAP on MC-LAG)	Standby	Purple

Table 13: MC-Lag VLL Link States

Access Redundancy

SAM is enhanced to identify the 'active / standby' SAP. The enhancement also provides the navigation ability to the peer SAP from the map, component tree, and configuration form. A warning message is displayed in cases where the peer PE is not yet managed or the peer SAP is miss-configured.

SAM Alarms

The following two alarms are raised on VLL Access Redundancy mis-configuration:

- Un-managed Peer PE
 - Severity = Warning
 - Cleared: when the peer PE is managed
- Peer SAP Mis-Configured
 - Severity = Warning
 - Cleared: once the configuration is corrected

Segmented Service

This feature is part of the service management initiatives in SAM 80R5. This segmented service feature, in particular divides a service into one or more segments to facilitate service component provisioning/modification (80R5). For example, a PE can be added to one of the VPLS metro meshes of the service where the PWs can still be auto-created by SAM.

To maintain the current provisioning paradigm in both GUI (incl template) and OSS/I scenarios, segment is not required to be explicitly created. Segments will be auto discovered by SAM. Segments then can be highlighted on service maps where PW's of the same segment will have the same color.

In the next phase, SAM service manager will leverage segments to allow trouble shooting using OAM tests and service profile per segment. A service spanned multiple regions are also be managed by different operator. This feature shall allow span of control at the segment level.

Service Provisioning Enhancements

This features speeds up the GUI L2 service provisioning procedures with point and click actions using service map. Adding PE's, auto creation of PWs for VPLS and VLL in both direction, creating a new mesh or adding new PE into an existing mesh can all be done with the service map. Other components such as endpoint and SAP can also be created from the same view.

Tunnel Selection Enhancements

In 80R5, except for Ethernet CAC with PBB tunnel selection, service tunnel is basically SDP. Prior to R5, tunnel selection is based on tunnel type and the current load (number of bindings). In 80R5, two other factors are in the equation:

- Steering parameters - Included and excluded set of named parameters (0..31) are used in the same manner as in MPLS. Service tunnel is configured with a set of steering parameters. Service shall have a list of included and excluded parameters. A service uses a given tunnel if its steering parameter set is a subset of the service tunnel included set and does not overlap with the tunnel excluded list.
- Steering parameters are now part of tunnel (selection) profile which will include other attribute such as tunnel type. Tunnel template, BW, metric will be in future development.
- In short, a tunnel selection will be based on transport type, steering parameters, and load in R5.

Service Component Tree Navigation

The service component tree shall be enhanced to provide the ability to navigate to the following objects by opening a new window.

- Binding on the opposite direction for unidirectional SDP spoke/mesh binding. Right Click → <<Opposite>> Binding
- The link to B-VPLS for I-VPLS site. Right Click → B-VPLS
- The other SAP for the VLAN Up-Link. Right Click → <<Opposite>> SAP
- Redundant SAP. Right Click → Active/Standby SAP

Service Site Creation from Service Map

SAM is enhanced to provide the ability to create service sites (B-site, I-site, regular site) from the Service Topology Map. Select NEs, then Right Click → Create Site.

Service CAC

Service CAC is a SAM-only function and is introduced in 8.0R1. This function only applies to native Ethernet network using PBB. A flag in nms-server.xml must be checked to activate the function.

Highlights:

- Applicable for p2p service (multipoint to be considered in a future release)
- Applicable to PBB tunnel (SDP using GRE, LDP, RSVP and PBB tunnel over IP/MPLS will be considered in a future release)
- I-service will not be connected to the PBB tunnel (B-VPLS) if the B-VPLS has no sufficient BW.
- Link BW usage adjusted after service tunnel paths have been rerouted
- PBB tunnel total BW is calculated from the Ethernet link BW.
- Audit function is used to verify if a tunnel has enough BW to carry

The L2 PBB CAC function provides the following main functions:

- EVPL CAC with MSTP and G.8031 (unreserved BW) PBB Tunnel
- Topology change and BW usage recalculation
- CoS based CAC
- Modified BW requests while EVPL is in service

- System start-up and Audit
- Dual-homed EVPL CAC with MSTP-based PBB Tunnel
- Dual-homed EVPL CAC with G.8031 based PBB Tunnel
- EVPL Endpoints on closely located Access Switches.
- CAC on ES-AS link (provided by BAAIS)
- 3-pt Service CAC

The CAC algorithm is based on the concept that services ride on tunnels and tunnels use links. Even though the tunnels might not have any BW reserved (applicable to Ethernet layer and to many types of IP/MPLS based service tunnels), tunnel is used to define the set of links that the services are actually using.

Tunnel's BW availability is calculated at the service admission request time and is based on the links currently used by the tunnel's forwarding path. The booking of reserved BW by the service is done directly on the link (and directly on the tunnel for BW-reserved tunnel). The following picture summarizes the concept.

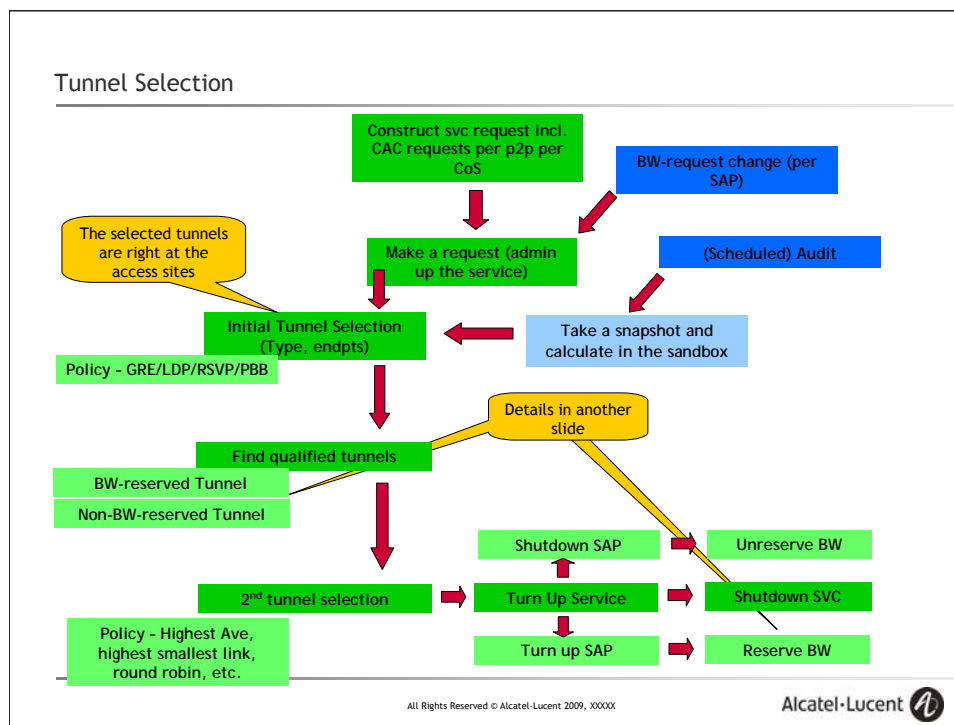


Figure 39: Tunnel Selection Flow

The link keeps track of the BW used by all of the services using that link. However, instead of having references directly to the services, the links keep track of BW usage by the tunnels. *The key reason is that a service can use different tunnels and the BW reserved on each tunnel can be different.*

Another reason for referencing the tunnels by the link is that monitoring the tunnels' active paths is simpler than monitoring individual service paths (not scale).

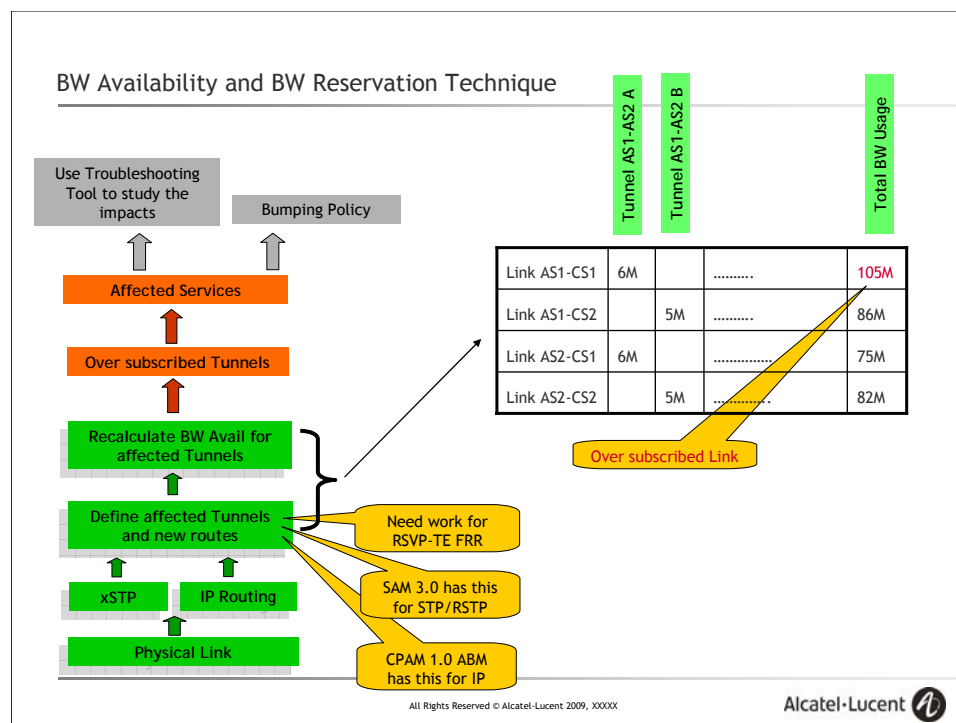


Figure 40: BW Reservation & Availability Calculation

Request for Admission Control

CAC is applied to the service tunnel and the access port. A service could have multiple SAPs and rides on multiple service tunnels. Thus, CAC function is invoked only when there is creation request for a service tunnel binding or a SAP, the default admin state of those two objects is "UP". If CAC verification fails, SAM will change the object's admin state to "DOWN". As one can always create a binding or a SAP via CLI, CAC for such case will be verified at the next Audit cycle.

The above reasoning works well for a unidirectional service tunnel such as SR/ESS SDP. SAM will have to enhance the rule to accommodate "multi-directional" service tunnel such as PBB.

5620 SAM allows the service creation in multiple steps. For example, an EPVL can have just one end (a service site, a SAP, and a binding to the service tunnel). Thus, the general rule is to invoke CAC (for a binding to a service tunnel or a SAP) whenever there is sufficient data entered. To invoke CAC you would need the Required BW and enough information to find the candidate tunnels (probably at least 2 endpoints). For SAM 8.0 an endpoint will be defined as a Service Site. Please see section 3.10 for some specific example and flow diagrams in which CAC will be invoked.

With p2p services, the BW request must be explicitly specified at the service level, i.e. SAM shall not derive BW request per CoS from SAP configuration at all.

CAC is now supported for SAP.

For future multipoint service, admission for that service will be examined when there are at least 2 endpoints defined. When more endpoints are added, admission verification will be done for the new endpoints, one by one.

Topology change and BW usage re-calculation

When there is a change in the physical topology (Link up/down), the topologies of some of the tunnels might be affected. The link utilizations for the down link as well the links are now used by the forwarding paths are changed, and are recalculated to provide a better picture of BW usage on those link.

If there is any link or tunnel is now overbooked, an alarm will be raised. The operator can then use the diagnostic tools (Service Test manager, Real time Stats plotter, and so on) to study the real impacts.

Possible events that could trigger this re-calculation:

- The effective BW being changed on a Physical Link.
- STP Change on a SAP. i.e. going from blocking to forwarding and vice versa.

Audit

The audit function is used to ensure that the links BW has been properly calculated. This will be invoked from the Service Manager. A full network audit could take some time and if during this time a service being created using CAC functionality will be prevented.

When requested it basically

- Resets (in the sandbox only) the available BW of all links
- Visits every tunnel in the network
- Adds up the BW requests per CoS for each and every EVPL currently using the tunnel. Note that rule X (if a service currently using the tunnel or not applies).
- Adjusts the available BW in each and every links currently used by the forwarding path of the tunnel.

Notes:

- It might take some time to complete an audit cycle. The topology of a tunnel that has been examined could be changed. In that case, the process shall be restarted.

Ethernet CAC

Three main enhancements will be made in R5:

- PBB tunnel itself can not be CAC'd. We are talking about multipoint service CAC'ing here. The PBB tunnel must use either mSTP or 8031. A CAC verification function shall be provided as the service can be created w/o SAM involvement. Also, actual used links (fwd links) are finalized by the network and based on the configured VlanLinks. It means the service can still be created and if the BW requirements are not met, an operation flag, CAC status, will be set to not sufficient BW.
- More accurate booking. In 80R1, BW used by a service will be booked (or subtracted from) all links used a given PBB tunnel. This is not completely optimized as in the case of MC-LAG as some of the VlanLinks are not carry the traffic even their states are forwarding.
- MC-LAG endpoint support. For dual homing scenario, p2p service endpoints could involve as many as 4 PE's. As mentioned in the above bullet, only links "connected" to the forwarding path of MC-LAG shall be booked. Also when MC-LAG switching happens, affected links' BW shall be recalculated.

IPVPN - VRF id based on strings rather than numbers

T

This feature allows service administrators to add an optional Service Name to services within the 7x50 and 7710 platforms.

On the node, services must be created using a Service ID. Once created, an optional Service Name can be configured on the service. In SAM, the name can be specified during Service Site creation.

Restrictions on the Service Name are as follows:

- Maximum character length of 64.
- Must not begin with a number. Special characters are accepted anywhere within the name.
- Must be unique across all configured services on that node

SAM already has a Name property on the Site level of each service that previously had no relationship to a property on the node. It will now be put to use by linking it in with the nodes new Service Name property. If the Name field is left blank the current behavior of SAM is to generate a default name (i.e. for a VPLS Site, the default name will be similar to "VPLS service-13 sim193 (38.120.200.193)"). This value will now be deployed to the node.

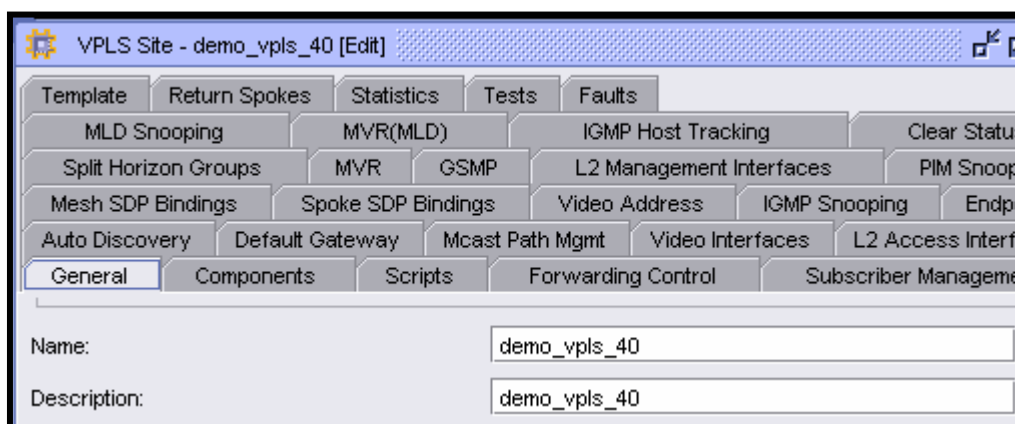


Figure 41: SAM Service Name

```

*A:sim172>config>service>vpls# info
-----
description "demo_vpls_40"
service-mtu 1300
stp
    shutdown
exit
service-name "demo_vpls_40"
sap 1/1/4:5 create
    collect-stats
exit
mesh-sdp 16:40 create
exit
mesh-sdp 18:40 create
exit
no shutdown
-----

```

Figure 42: CLI Service Name

There are a number of CLI command changes on the node (refer to nodal documentation for the complete list) which may now take a Service Name in place of the Service ID. Both of these cases have to be handled from within SAM. When performing one of these configuration changes on the node, an automatic translation of the Service Name to the Service ID is performed before sending the information to SAM. Likewise, SAM can only set the Service ID in the related MIB entries as there is no entry for the Service Name. This implies that in most cases, SAM should already have the capacity to handle these CLI commands without any changes. It is common practice in SAM to use a pointer to another service during these types of configurations, in which case there should be no change required in SAM. If SAM requires the user to manually enter the Service ID in a text field then the behavior will have to be changed to allow the user to specify a Service ID or a Name. If a Name is specified then we must match it to a Service ID before deploying to the node.

Given that we are using an old SAM property for a new purpose, there are upgrade considerations.

When upgrading a 7.0 node to an 8.0 node, if SAM has a Name configured for a service on the upgraded node then that value will be deployed to the node rather than syncing the blank value from the node, which would be the default behavior.

Diff-Serv Class type change during failures

The feature adds an option to configure a main Class Type (CT) and a backup CT for the primary path of a Diff-Serv TE LSP.

During a network failure event, there might not be enough network resources to signal all LSPs and deliver all in-contract and majority of out-of-contract traffic. In this case, a network operator can specify a back up TE Class Type to signal the LSP primary path when it fails and goes into retry. This will increase network efficiency and guarantee delivery of all in-contract traffic in a network failure event.

This feature is not supported on a P2MP LSP.

To provide support for this feature, the following displays have new attributes:

Manage > MPLS > Dynamic LSPs > Properties tab > Traffic Engineering and Protection group:

Main CT Retry Limit: This attribute configures the maximum number of retries the LSP primary path should be retried with the LSP Diff-Serv main Class Type.

When an unmapped LSP primary path goes into retry, it uses the main CT until the number of retries reaches the value of the new 'Main CT Retry Limit' parameter. If the path does not come up, it starts using the backup CT at that point in time. This attribute has no effect on an LSP primary path which retries due to a failure event. If the user entered a value of the 'Main CT Retry Limit' parameter that is greater than the value of the LSP Retry Limit, the number of retries will still stop when the LSP primary path reaches the value of the LSP retry limit. In other words, LSP retry limit represents the upper bound on the number of retries. This applies to both CSPF and non-CSPF LSPs.

Available values for this property are integers (0 to 10000) with a default value of 0 which means the LSP primary path will retry forever.

Manage > MPLS > Dynamic LSPs > LSP-Path Bindings tab > General tab > Traffic Engineering Properties group:

Diff-Serv Backup Class Type: This attribute enables and specifies which Diff-Serv backup Class Type will be used instead of the Diff-Serv main CT, to signal the LSP primary path when it fails and goes into retry.

When a LSP primary path retries due a failure, MPLS will retry a new path for the LSP using the main CT. If the first attempt fails, the head-end node performs subsequent retries using the backup CT. This applies to both CSPF and non-CSPF LSPs. When an unmapped LSP primary path goes into retry, it uses the main CT until the number of retries reaches the value of the new 'Main CT Retry Limit' parameter. If the path does come up, it starts using the backup CT.

Available values for this property are integers (-1 to 7) with a default value of -1 which indicates no backup class type has been configured for the LSP."

Diff-Serv Operational Class Type: This attribute specifies operational class type associated with the LSP. This property is read-only with type Integer.

Main Class Type Retries Remaining: This attribute specifies the number of remaining attempts the software will make before it starts using the backup class type for the LSP. This field will indicate "n/a" when the primary path comes up. While it is in retry, this counter will show 'infinite' if the Main CT Retry Limit parameter is left to the default value of 0. Otherwise, it will indicate the decreasing value as the number of retry attempts increases until it reaches zero at which point the path is retrying using the backup CT or exhausted the main LSP Retry Limit parameter. This property is read-only with type Integer.

Routing > RSVP > General tab > Diff-Serv TE group:

Diff Serv Model: A new option named "RussianDollModel" will be added to this property. Once selected, this option enables DiffServ TE with Russian Doll Model. The Russian Doll Model (RDM) LSP admission control policy allows bandwidth sharing across Class Types. It provides a hierarchical model by which the reserved bandwidth of a CT is the sum of the reserved bandwidths of the numerically equal and higher CTs.

For example if two Class Types are configured, CT0 will have a BC0=Maximum Reservable Link Bandwidth and CT1 a smaller percentage of BC0. The sum of reserved bandwidths of CT0 and CT1 should not exceed BC0.

The enabling or disabling of Diff-Serv TE on the system requires the RSVP and MPLS protocols to be shutdown.

BFD support of OSPF CE-PE adjacencies

The feature introduces BFD support to OSPF PE-CE interfaces. It enhances the fault detection for the PE-CE connection into a VPRN service.

Currently, when BFD is enabled for an OSPF interfaces within VPRN, SAM displays error message “BFD support is currently disabled on OSPF interfaces within VPRN”.

This message is removed for SR 8.0.

Spoke termination for IPv6 IES & 6VPE

The feature supports IPv6 spoke termination on IES and VPRN services. IPv6 on a spoke terminated interface has the same equivalent functions as IPv4 on spoke terminated interfaces in an IES or VPRN service.

Currently, SAM does not allow the bind of a Spoke SDP Binding to an IES/VPRN interface with IPv6 allowed. This restriction is removed for SR 8.0.

Support for MPLS hash label

The MPLS hash label allows LSR nodes in a network to load balance labeled packets in a much more granular fashion than allowed by simply hashing on the standard label stack. It will remove the need to have an LSR inspect the payload below the label stack to check for an IPv4 or IPv6 header. This feature applies to VLL, VPLS, IES, and VPRN services in R8.0.

The feature is supported on SR7/SR12 with chassis modes C and D. It is not supported on the SR1 platform. It is also supported on ESS chassis mode D and 7710 C4 and C12 chassis.

The MPLS hash label can be enabled or disabled for the SDP bindings on the following service types:

- VLL: Spoke SDP bindings for EPIPE, IPIPE and FPIPE. It is not supported for APIPE and CPIPE VLLs.
- VPLS: Spoke SDP bindings and Mesh SDP bindings on all VPLS sites (including B-VPLS and I-VPLS).
- VPRN: Interface terminated spoke SDP bindings.
- IES: Interface terminated spoke SDP bindings.

Figure 43: Configuration form of a spoke SDP binding

Hash Label can only be enabled on the SDP binding if the underlying transport of the service tunnel is MPLS type. An error message will be given if the underlying transport of the tunnel is GRE. It can also be configured on a VPRN site. When enabled on the site, the hash label will be included on packets forwarded on the following objects of the VPRN service:

- All RSVP or LDP LSPs to BGP next-hops when the service is configured in 'auto-bind [rsvp-te, ldp, mpls]' modes.
- All user specified SDPs.

It can only be enabled when the auto-bind of the VPRN site is set to 'LDP', 'MPLS' or 'RSVP'. A warning message will be given when a user tries to enable the hash label with auto-bind set to 'GRE' or 'None'.

The attribute on the site does not control the use of the hash label on an interface terminated spoke-SDPs (spoke interface). The hash label for the interface terminated spoke-SDP bindings can be enabled individually as stated above.

The screenshot displays the 'VPRN Site - vprn111 [Edit]' window. The top menu bar includes tabs such as IGMP Host Tracking, Clear Status, Route Aggregation, IPsec Security Policies, Template, Statistics, Tests, and Faults. Below this, a secondary menu bar lists various configuration categories like Source Addresses, GSMP, Local DHCP Servers, Self Generated Traffic, Mcast Path Mgmt Channels, and IP Mirror Interfaces. The main configuration area is divided into several sections: 'Network Element' with a System ID field (10.1.1.40) and a Select... button; 'Service' with Service ID (111) and Service Name (VPRN 111) fields; 'Customer' with Customer ID (1) and Customer Name (Default customer) fields. Below these are fields for Name (vprn111), Description (N/A), Administrative State (Up), Operational State (Up), Monitor Access Interface Operational State (unchecked), State Cause (SDP Binding(s) Down checked, Monitored Access(es) Down unchecked), Enable Hash Label (unchecked), GSMP Administrative State (enabled), and OLC State (Maintenance). At the bottom right, there are Turn Up and Shut Down buttons. At the bottom left, there are Copy..., Resync, and Create Template buttons. At the bottom right, there are OK, Cancel, and Apply buttons.

Figure 44: Configuring a VPRN site

Hash Label is also supported for PW-template. The new boolean attribute 'Hash Label' is added on the configuration of a PW-template policy.

This attribute is supported for the service templates that are associated with the all the correspondence service types.

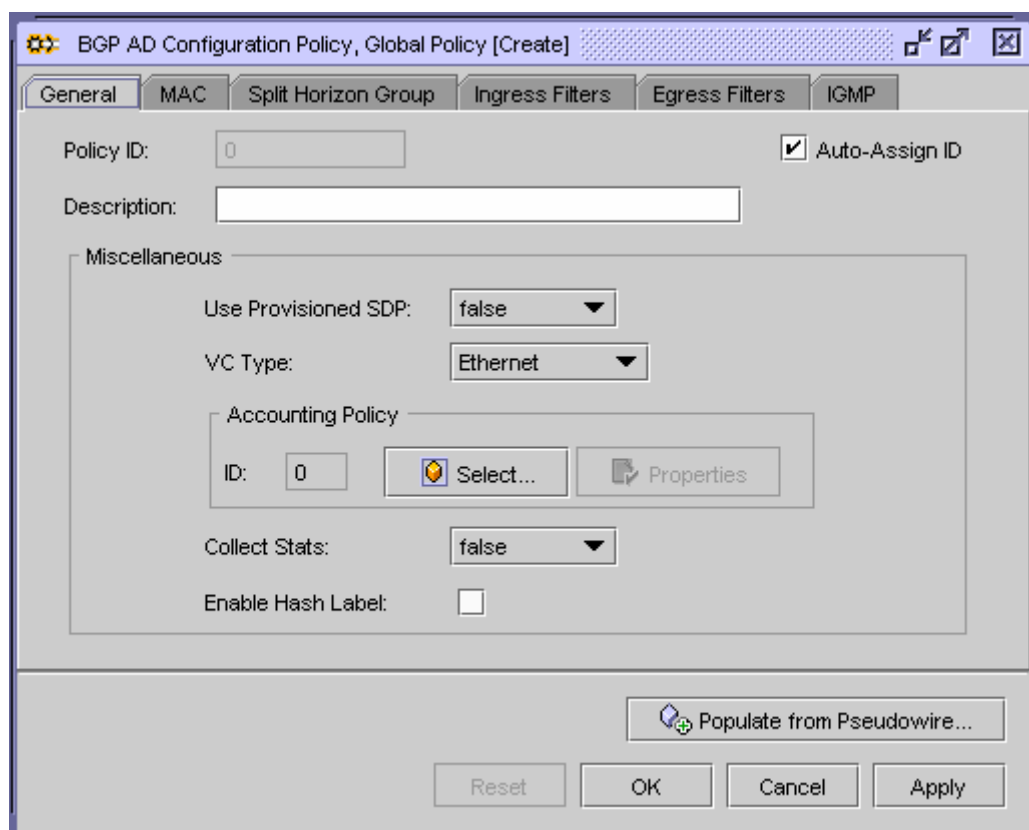


Figure 45: Creating a PW Template

RSVP LSP Primary and LDP LSP backup within a SDP

This feature allows the user to enable both RSVP and LDP on an SDP starting from 8.0 SR or ESS node. While both RSVP and LDP are enabled on an SDP, it can switch from using RSVP LSPs to an LDP LSP in case of failure of all RSVP LSP paths associated with this SDP.

A user has an option to set Mixed LSP Mode while creating/configuring an SDP from a 8.0 SR or ESS node (7750, 7450 and 7710). The following attributes are added to SDP in SAM 8.0 to support this feature.

Attribute	Access	Values	Description
Mixed LSP Mode	Read-Write	True (default) False	Specifies whether RSVP and LDP types of LSPs can be co-existed in this SDP
Revert Time	Read-Write	0-600 seconds, Infinite (never revert back) Default: 0	Specifies the time to wait before reverting back from LDP to a associated LSP when it is available. It is only applicable while Mixed LSP Mode is true.
LDP Active	Read	True/False (default)	Indicates whether the LDP is active on this SDP.

Note: LDP has an attribute “Tunnel Down Damp Time”, which is set to 3 seconds by default. This attribute specifies how long, in seconds an LDP waits before sending a tunnel down event to the route

table manager. When the LDP fails, the SDP will revert to the RSVP LSP only after the expiry of this timer. For an immediate switchover, this timer must be set to “0”.

The following figures shows the support for above attributes while creating/configuring an SDP from an 8.0 SR or ESS node.

- Creating an SDP from a 8.0 SR or ESS node:

While in step 4, while Underlying Transport is “MPLS”, the user can specify whether “Mixed LSP Mode” is true. If it is true, “LDP Enabled” will be set to true and the attribute “Revert Time” will be present for the user to specify the reverting time.

If “Mixed LSP Mode” is set to false, the behavior is the same as SAM 7.0.

- Configuring an SDP from a 8.0 SR or ESS node:

The user can modify the attribute “Mixed LSP Mode” and set it to true or false. If it is set to true, the attribute “Revert Time” will be present for the user to set the reverting time. If it is set to false, the attribute “Revert Time” will not be present.

If the user tries to set “Mixed LSP Mode” to false while there is/are RSVP LSPs associated with this SDP, SAM will reject the update with an error message “Can’t set Mixed LSP Mode to false while there is a RSVP LSP associated with this SDP”.

The screenshot shows a configuration window titled "IP/MPLS Service Tunnel (SDP) -". On the left is a "Steps" sidebar with 12 steps. Step 4, "Specify Transport", is highlighted. The main area is titled "Specify Transport" with the subtitle "specify transport for this Service Tunnel". It contains five configuration fields: "Underlying Transport:" set to "MPLS", "Mixed LSP Mode:" set to "true", "Ldp Enabled:" set to "true", "Revert Time (seconds):" set to "14", and "Signaling:" set to "TLDP". At the bottom are four buttons: "< Back", "Next >", "Finish", and "Cancel".

Figure 46: Creating an SDP from an 8.0 SR node

The screenshot displays the configuration interface for a tunnel in the Alcatel-Lucent 5620 Service Aware Manager. The window title is "Tunnel - 3, 10.1.1.26 [Edit]". The "General" tab is selected, showing the following configuration details:

- Identity:** Name: ttt1, ID: 3, Description: from-10.1.1.26-id-3, Underlying Transport: MPLS, PBB Ethernet Type: 0x88e7.
- MPLS Signaling:** Mixed LSP Mode: true, Enable LDP: true, Revert Time (seconds): 20, LDP Active: true.
- Source:** Source Site ID: 10.1.1.26, Source Site Name: sim26.
- Destination:** Destination Site ID: 10.1.1.27, Destination Site Name: sim27.
- States:** Administrative: Up, Operational: Up. State Cause checkboxes include OAM Validation Failed, Transport Tunnel Unstable, No System IP Address, Invalid Egress Interface, Keep-Alive Failure, Transport Tunnel Down, Signaling Session Down, and Tunnel Admin Down.
- MTU:** Administrative MTU: 0, Operational MTU: 1492, Advertised MTU Override: false.
- Class Forwarding:** Class Forwarding Capability: On, Administrative State: Up, Enforce Diff-Serv Lsp-Fc Map: Off.
- Metric:** Metric: 0.
- VLAN:** VLAN VC Ehtertype: 89024, No VLAN VC Ehtertype: checked.

At the bottom, there are buttons for "Navigate", "Turn Up", "Shut Down", "Valid", "Create/Edit Path Monitor", "Resync", "Create Template", "Reset", "OK", "Cancel", and "Apply".

Figure 47: Configuring an SDP from an 8.0 SR node

Due to the introduction of "Mixed LSP Mode" in SDP configuration, this new option is added anywhere the auto-tunnel creation is applicable.

1) For auto SDP binding creation in VLL, Mirror Service and VPLS, while 'Automatic SDP Binding Creation' ('Automatic Mesh SDP Binding Creation' for VPLS) is enabled, from drop-down list of 'Transport Type', a new option "Mixed LSP Mode" is available.

While "Mixed LSP Mode" is selected, SAM will create SDP bindings between sites by associating them with SDPs whose "Mixed LSP Mode" is true and operational up. If the attribute "Mixed LSP Mode" is not applicable to a site as a source of SDP Binding, or none of SDPs with "Mixed LSP Mode" is true, or none of them is operational up, SAM will choose operational RSVP SDPs first then operational LDP SDPs. If there are also no RSVP and LDP SDPs, no (Mesh) SDP Bindings will be created.

For "MPLS:RSVP-LSP" and "MPLS:LDP" options, SAM handles it same as before, i.e. choosing RSVP or LDP only SDPs, whose "Mixed LSP Mode" are false.

For "Any" option, SAM handles it same as before, i.e. choosing an SDP from all available ones (including ones with "Mixed LSP Mode" is true" now) on the node, this SDP should be operational up with least load factor; if none of them is operational up, it should be one with least load factor.

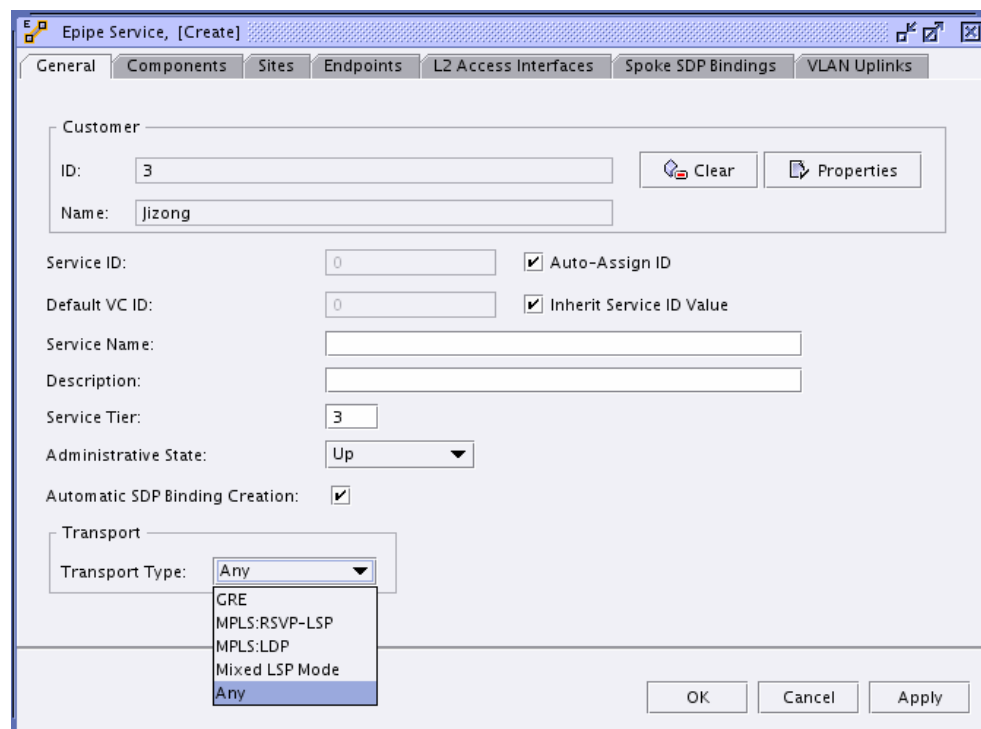


Figure 48: VLL creation with Automatic SDP Binding option

2) For auto tunnel selection in Spoke Connector while creating Composite Service, while 'Auto Select Tunnels' is enabled, from drop-down list of 'Transport Type', a new option "Mixed LSP Mode" is available.

While "Mixed LSP Mode" is selected, SAM will create SDP bindings between sites by associating them with SDPs whose "Mixed LSP Mode" is true and operational up. If the attribute "Mixed LSP Mode" is not applicable to a site as a source of SDP Binding, or none of SDPs with "Mixed LSP Mode" is true, or none of them is operational up, SAM will choose operational RSVP SDPs first then operational LDP SDPs. If there are also no RSVP and LDP SDPs, no (Mesh) SDP Bindings will be created.

For "MPLS:RSVP-LSP", "MPLS:LDP" and "Any" options, SAM handles it same as before, see description in section 1) for auto SDP binding creation in service.

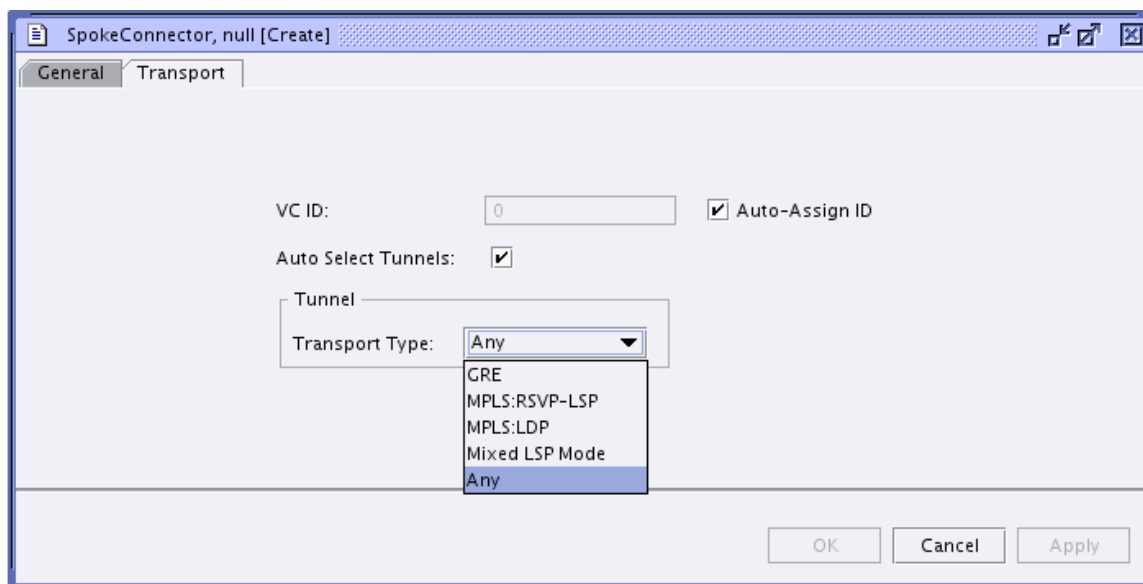


Figure 49: Spoke Connector with Auto Tunnel creation

3) For auto tunnel selection in Spoke SDP Binding creation, while 'Auto-Select Transport Tunnels' is enabled, from drop-down list of 'Tunnel Auto-Selection Transport Preference', a new option "Mixed LSP Mode" is available if it is applicable to the source site. Similar functionality is also available for creating Return SDP Binding.

While "Mixed LSP Mode" is selected, SAM will create SDP bindings between sites by associating them with SDPs whose "Mixed LSP Mode" is true and operational up. If none of SDPs with "Mixed LSP Mode" is true, or none of them is operational up, SAM will choose operational RSVP SDPs first then operational LDP SDPs. If there are also no RSVP and LDP SDPs, no (Mesh) SDP Bindings will be created.

For "MPLS:RSVP-LSP", "MPLS:LDP" and "Any" options, SAM handles it the same as before for auto SDP binding creation in service.

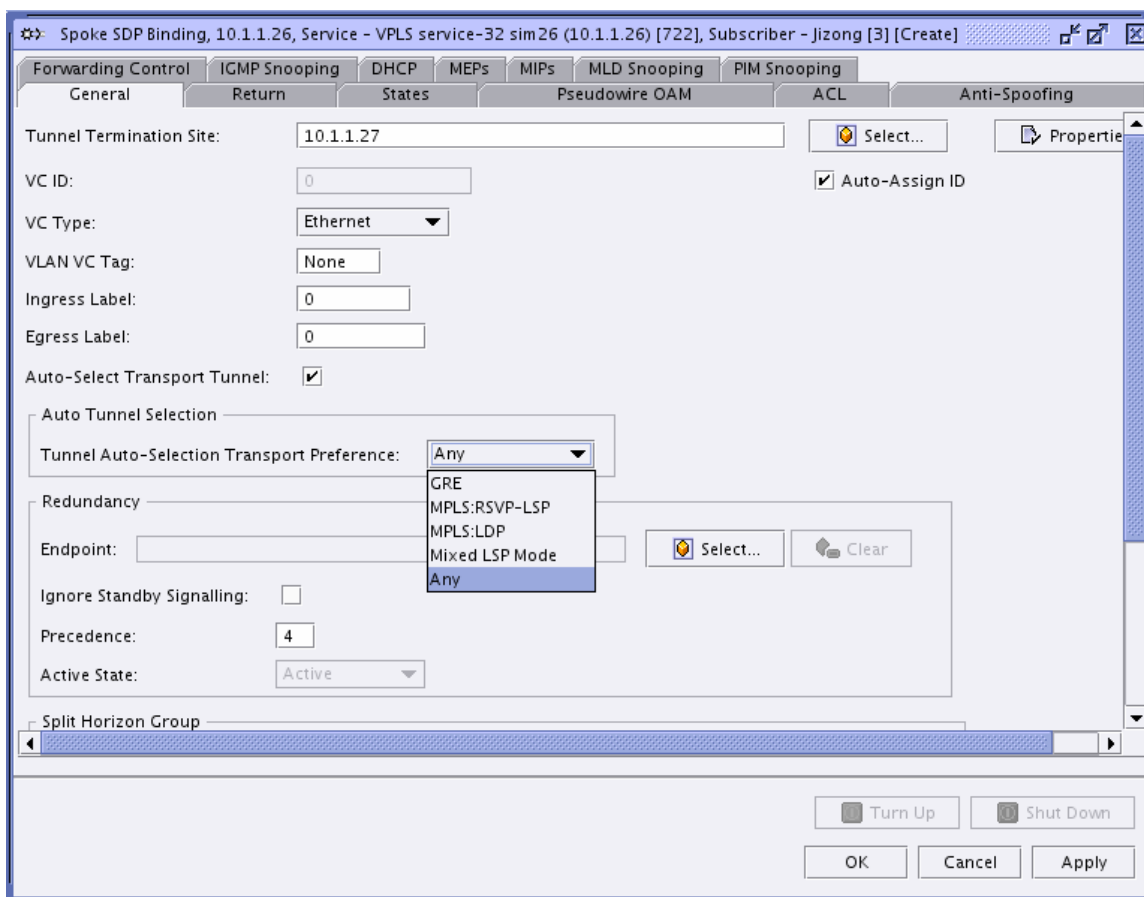


Figure 50: Spoke SDPBinding with Auto Tunnel creation

4) For Auto Tunnel Topology Rule (Mesh or Hub and Spoke) creation, while "Tunnel Type" is "SDP", a new option "Mixed LSP Mode" is available from the drop-down list of 'Underlying Transport'.

While "Mixed LSP Mode" is selected, SAM will create SDPs between sites with "Mixed LSP Mode" is true if it is applicable. Meanwhile, SAM will add LSPs to created SDPs as specified by the LSP templates, also, LDP will be enabled on SDPs.

If the attribute "Mixed LSP Mode" is not applicable to the source node, SAM will create the SDP with RSVP as underlying transport, plus adding LSPs to the created SDP as specified by the LSP templates. For " RSVP-LSP" and "MPLS:LDP" options, SAM handles it same as before, i.e. create SDPs with RSVP or LDP as underlying transport, "Mixed LSP Mode" should stay as false.

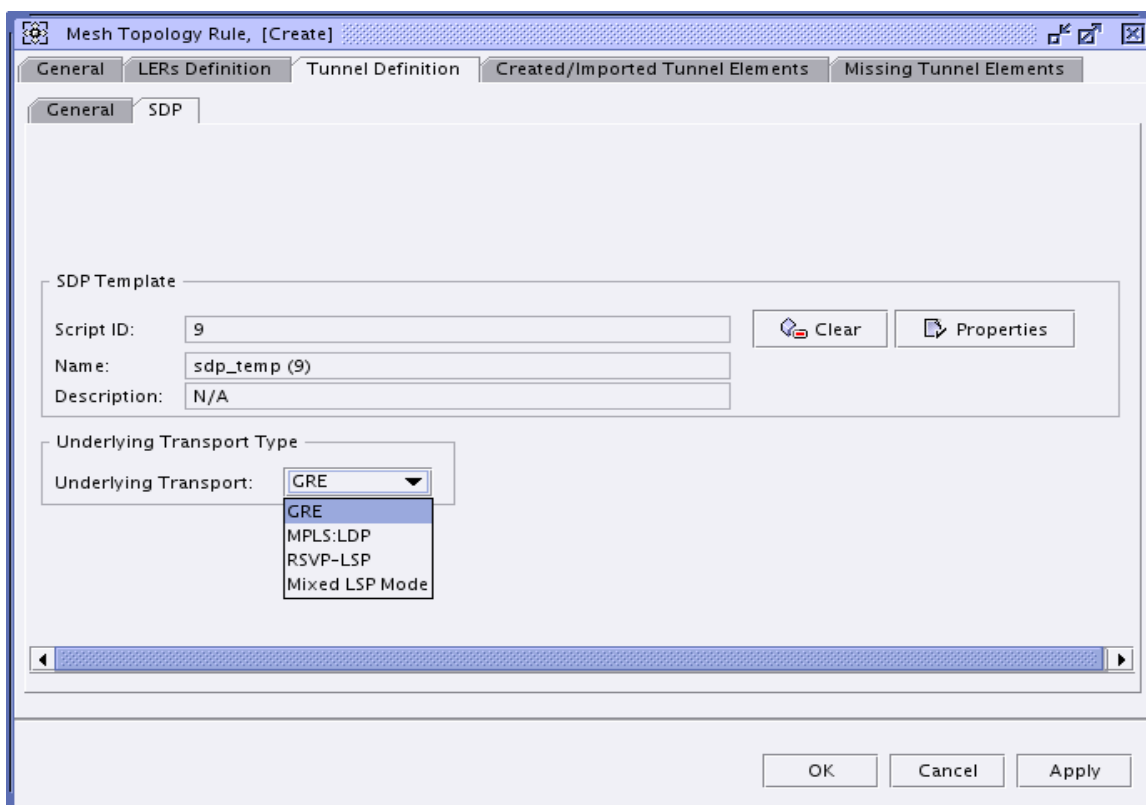


Figure 51: SDP creation with Auto Tunnel Topology Rule

5) For Template support, similar as above, a new option "Mixed LSP Mode" is available from the drop-down list of 'Transport Type' for service creation, spoke SDP creation, etc.

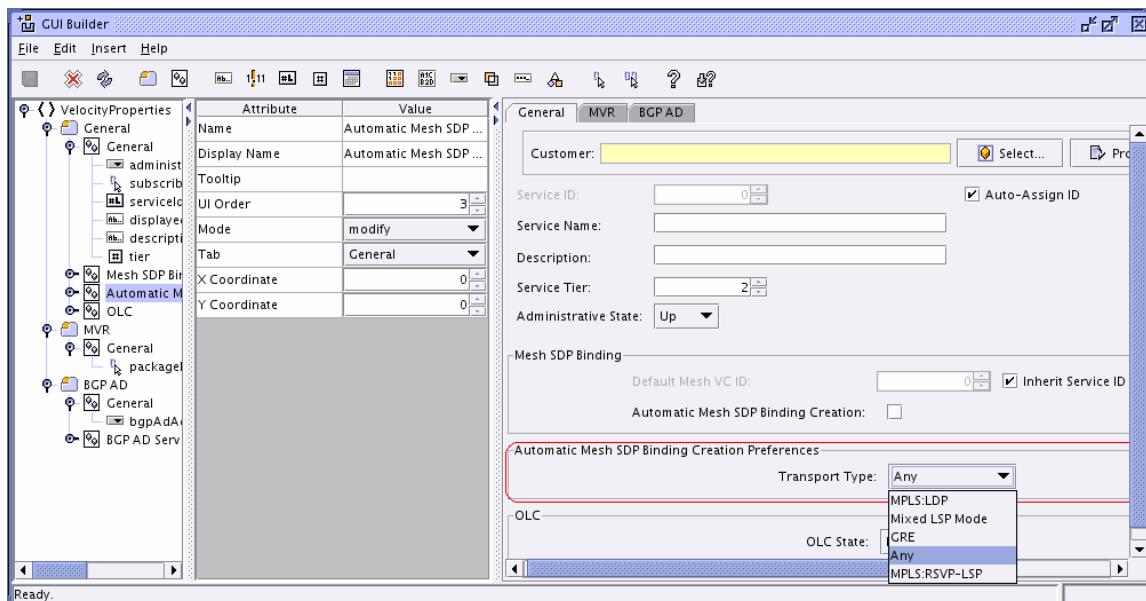


Figure 52: Auto Tunnel with "Mixed LSP Mode" in Template

Re-signalling of primary RSVP-TE paths

For an LSP starting from a 8.0 SR or ESS, make-before-break (MBB) procedures allow the operator to manually replace an existing MPLS path under the primary or secondary LSP path with a new path (must existing on the node).

If the MBB is successful, the new path will be used by the corresponding LSP path. If not, the LSP path will stay with its old MPLS path.

To support this feature in SAM, the user has an option to update the existing MPLS path under an LSP path, whose MBB is enabled and its source node is an 8.0 SR node, with another path that has been created on the node.

The following figure shows the new option provided by SAM. There is a new button “Update MPLS Path...” added in the bottom panel of LSP Path configuration form. If the user clicks this button, the SAM will launch a window for MPLS Path Selection.

A similar button is added to LSP configuration form (to right button panel in LSP path Tab). It will be enabled while a user selects a LSP Path, then the user can click this button to update the MPLS path associated with this LSP path.

The screenshot shows the 'LSP-Path Binding' configuration window for 'path27-26; id 1; Lsp lsp27-26; id 1 (from sim27 (10.1.1.27) to sim26 (10.1.1.26))'. The 'General' tab is selected, showing fields for LSP Name (lsp27-26), MPLS Path Name (path27-26), Type (primary), and Status (active). Below these are sections for States, Transport Properties, Traffic Engineering Properties, Make before Break, CSPF, and Administrative Groups. At the bottom, there is a row of buttons including 'Navigate', 'Update MPLS Path...', 'View MPLS Path...', 'Turn Up', 'Shut Down', 'Copy...', 'Resync', 'Create Template', 'Reset', 'OK', 'Cancel', and 'Apply'.

Figure 53: Configuring an LSP Path from a 8.0 SR node

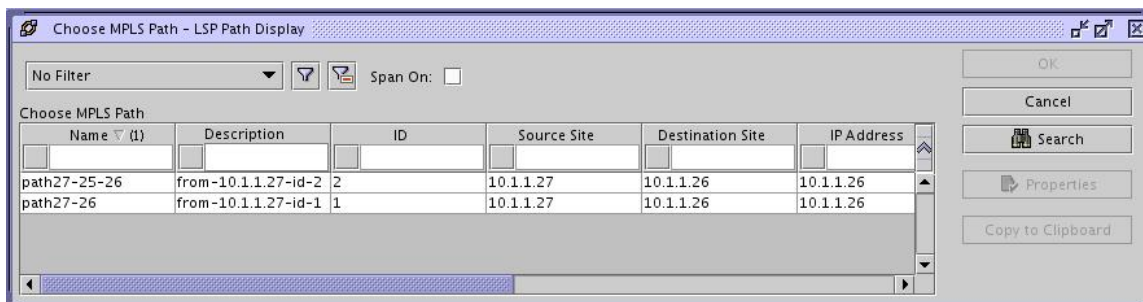


Figure 54: Selecting new MPLS Path window

Note: the MPLS paths listed in the window must not be used by the any LSP paths under the parent LSP.

After the user selects a new path from the list and clicks the “OK” button to close the MPLS Path list window,

SAM will immediately deploy this change to the node.

After deployed this change to the node, the SR node will send a trap to SAM to notify if the update is successful. If it is successful, the underlying MPLS path on LSP path will be updated with the one just selected; if not, there will be an alarm raised for this operation on LSP path object. The alarm will include the operation status (failed, aborted, ignored) and the reason (mbbRetryExceeded, lspPathGoingDown, startingHighPriMbb, restartingMbb, highPriMbbInProg)

Full IGP Shortcuts and Forwarding Adjacencies

This feature allows forwarding of packets to IGP learned routes using an RSVP-TE LSP.

The following are key attributes of this feature:

- Allow IGP-shortcuts to be enabled on RSVP LSP.
- Advertise the RSVP LSP with its admin/operational metric in IGP protocols such as IS-IS and OSPF, i.e., model the LSP as a link interface. This is referred to as OSPF or IS-IS forwarding Adjacency.
- Supporting IGP adjacencies across an RSVP LSP is not required.
- Support interoperability with Cisco ‘Auto-route Announce’ capability.

This feature allows customers to use shortcuts to engineer traffic traveling towards destination nodes that do not support MPLS LSPs. For example to traffic engineer IP traffic where the last router is not configured to support MPLS or to configure the rest of the IP routers to see the RSVP shortcut and include them in their SPF computation for reaching the non MPLS destinations. This requires LSP shortcut to be advertised in IGP as a link.

To support this feature, following displays have new attributes:

Routing > OSPFv2 Site > Behavior tab > OSPFv2 group

&&

Routing > ISIS Site > Behavior tab > General group:

RSVP Shortcut Enabled: This attribute specifies whether RSVP-TE shortcut for resolving IGP routes has been enabled or disabled for OSPFv2/IS-IS. This will instruct OSPFv2/IS-IS to include RSVP LSPs originating on this node and terminating on the router-id of a remote node as direct links with a metric equal to the operational metric provided by MPLS. This property will be presented as a check-box with a default of unchecked.

Routing > OSPFv2 Site > Behavior tab > OSPFv2 group

&&

Routing > ISIS Site > Behavior tab > General group:

Advertise Tunnel Links Enabled: This attribute specifies whether advertisement of LSP shortcuts into IGP has been enabled or disabled for OSPFv2/IS-IS.

This property will be presented as a check-box with a default of unchecked.

Manage > MPLS > Dynamic LSPs > Properties tab > Traffic Engineering and Protection group:

IGP Shortcut Enabled: This attribute specifies whether to exclude or include a RSVP LSP from being used as a shortcut while resolving IGP routes.

By default, all RSVP LSPs originating on a node that has “RSVP TE Shortcut” enabled, are included by OSPF and ISIS as direct links as long as the destination address of the LSP corresponds to router id of a remote node. RSVP LSPs with a destination address corresponding to an interface address of a remote node are automatically not considered by IS-IS or OSPF. The user can exclude a specific RSVP LSP from being used as a shortcut for resolving IGP routes by enabling this attribute.

This property is presented as a check-box with a default of checked (enabled).

Table below provides the outcome of configuration of the 'LDP Over RSVP Include' and 'RSVP Shortcut' or 'IGP Shortcut' attributes at both the IGP instance level and at the LSP level. Whenever both options are enabled, the RSVP Shortcut feature takes precedence and the RSVP LSP is used as a shortcut.

	Instance Shortcut-Enbl LDP/RSVP-Enbl	Instance Shortcut-Enbl LDP/RSVP-Dis	Instance Shortcut-Dis LDP/RSVP-Enbl	Instance Shortcut-Dis LDP/RSVP-Dis
Tunnel Shortcut-Enbl LDP/RSVP-Enbl	Shortcut (<i>Override case</i>)	Shortcut (<i>Override case</i>)	LDP/RSVP	None
Tunnel Shortcut-Enbl LDP/RSVP-Dis	Shortcut	Shortcut	None	None
Tunnel Shortcut-Dis LDP/RSVP-Enbl	LDP/RSVP	None	LDP/RSVP	None
Tunnel Shortcut-Dis LDP/RSVP-Dis	None	None	None	None

Table 14: LDP/RSVP Include & RSVP Shortcut Configuration Outcome

Enable/disable the no-propagate-ttl capability

The feature provides an option to specify if an LSP shortcut should operate in Uniform or Pipe mode. It allows the user to configure to hide or reveal the hops of their MPLS network when their customer packets are carried over an LSP shortcut.

The feature provides an option to enable or disable the propagation of the TTL from the customer IP packet into the MPLS packet. A TTL value of 255 will be inserted onto the pushed label stack when the propagation of TTL is disabled.

This feature is required on a P2P LDP/RSVP LSP shortcut used in static, BGP, or IGP route resolution.

The user can configure the behavior independently for local and transit IP packets for both LDP and MPLS LSP shortcuts. Two Boolean attributes will be introduced to LDP site and MPLS site for Network routing instance. They are “Shortcut Local TTL Propagate” and “Shortcut Transit TTL Propagate”. Both default values are ‘true’.

By default, the feature propagates the TTL from the header of transit and locally generated IP packets into the label stack of the resulting MPLS packets forwarded over the LSP shortcut. This is referred to as Uniform mode. When the TTL propagate is disabled, a TTL of 255 is programmed onto the pushed label stack. This is referred to as Pipe mode.

Figure 55: Configuration form of an LDP site

IS-IS and OSPF TE bandwidth updates triggered by threshold crossing events

Each LSP setup and removal triggers flooding of IGP TE-LSA update in the network creating a surge of TE updates in the network and leading to inefficient usage of signaling resources.

This feature reduces the surge by defining threshold levels of TE update per interface.

To support this feature threshold levels are defined based on a percentage of available bandwidth. By doing so, TE update frequency depends on bandwidth consumption instead of number of LSPs per interface which allows insignificant bandwidth updates to be ignored to reduce IGP TE updates.

Timer based triggered update is supported to allow forced update of bandwidth across nodes to sync the actual available/reserved bandwidth. This configuration is defined globally and applied to all RSVP interfaces. As an added enhancement, IGP TE update could be configured to be triggered if CAC failure is detected. In this case, TE-DB on ingress node is updated with the actual available bandwidth on CAC failure. A CAC failure triggered TE update overwrites a pending timer based IGP update.

To provide support of this feature the following displays have new attributes:

- **Routing > RSVP > General tab > Diff-Serv TE group:**
 1. **TE Threshold Update Enabled:** This attribute enables IGP TE update only based on bandwidth reservation thresholds per interface and blocks IGP TE update on bandwidth change for each reservation. Threshold levels are defined at global RSVP or per interface level.
 2. **TE Update On CAC Failure Enabled:** This attribute enables CAC failure triggered IGP update to allow IGP update with the actual available bandwidth on CAC failure. A CAC failure triggered TE update overwrites a pending timer based IGP update. This configuration is defined globally and applies to all RSVP interfaces. This attribute is available only when 'TE Threshold Update Enabled' is set.
 3. **TE Update Timer:** This attribute controls timer based IGP TE updates to allow forced update of bandwidth across nodes to sync the actual available/reserved bandwidth. This configuration is defined globally and applies to all RSVP interfaces. This attribute is available only when 'TE Threshold Update Enabled' is set.
- **Routing > RSVP > (new) IGP Update tab > (new) Up Threshold group:**

Threshold levels are supported for each direction (Up/Down). Threshold levels are for reserved bandwidth per interface meaning any reserved bandwidth change per interface is compared to the threshold levels and triggers an IGP TE update if the defined threshold level is crossed in either direction (LSP setup or teardown).

Threshold levels configured per node (at RSVP level) could be inherited by all configured RSVP interfaces. Available values for Thresholds are integers (-1 | 0 to 100 percent) where -1 indicates threshold level is disabled (not being used).

If user configures one or more thresholds to a non-default value, configured thresholds will be rearranged in ascending (for up thresholds) or descending (for down thresholds) order.

 - **Up Threshold 1:** This attribute configures specific threshold up level 1 per node and per interface. Default is 0.
 - **Up Threshold 2:** This attribute configures specific threshold up level 2 per node and per interface. Default is 15.

- **Up Threshold 3:** This attribute configures specific threshold up level 3 per node and per interface. Default is 30.
 - **Up Threshold 4:** This attribute configures specific threshold up level 4 per node and per interface. Default is 45.
 - **Up Threshold 5:** This attribute configures specific threshold up level 5 per node and per interface. Default is 60.
 - **Up Threshold 6:** This attribute configures specific threshold up level 6 per node and per interface. Default is 75.
 - **Up Threshold 7:** This attribute configures specific threshold up level 7 per node and per interface. Default is 80.
 - **Up Threshold 8:** This attribute configures specific threshold up level 8 per node and per interface. Default is 85.
 - **Up Threshold 9:** This attribute configures specific threshold up level 9 per node and per interface. Default is 90.
 - **Up Threshold 10:** This attribute configures specific threshold up level 10 per node and per interface. Default is 95.
 - **Up Threshold 11:** This attribute configures specific threshold up level 11 per node and per interface. Default is 96.
 - **Up Threshold 12:** This attribute configures specific threshold up level 12 per node and per interface. Default is 97.
 - **Up Threshold 13:** This attribute configures specific threshold up level 13 per node and per interface. Default is 98.
 - **Up Threshold 14:** This attribute configures specific threshold up level 14 per node and per interface. Default is 99.
 - **Up Threshold 15:** This attribute configures specific threshold up level 15 per node and per interface. Default is 100.
 - **Up Threshold 16:** This attribute configures specific threshold up level 16 per node and per interface. Default is -1.
 - **Reset to Default:** This button will reset all up thresholds to default values.
- **Routing > RSVP > (new) IGP Update tab > (new) Down Threshold group:**
Properties defined in this group behave the same as those defined in the 'Up Threshold' group. There are 16 'Down Threshold x' properties (where x ranges from 1 to 16 inclusive). Corresponding default values for levels 1 to 16 are 100, 99, 98, 97, 96, 95, 90, 85, 80, 75, 60, 45, 30, 15, 0, and -1 respectively. There is also a 'Reset to Default' button in this group which resets all down thresholds to their default values.
 - **Routing > RSVP > Interface > (new) IGP Update tab:**
This tab includes the 16 'Up Threshold' and the 16 'Down Threshold' attributes as defined in the global 'IGP Update' tab. These thresholds can inherit their value from the corresponding global RSVP level thresholds. Following additional properties are also added to this display:

1. **Update Pending:** This attribute indicates if the TE update which is to be sent to IGP on any bandwidth change is pending for this interface. Options are True/False. This property is read-only.
2. **Next Update:** This attribute indicates the time remaining in seconds before the next TE update would be sent to IGP for this interface. This property is applicable only when '*IGP Update pending*' is True. This property is read-only.

IMPLICIT NULL label option support on egress LER

The Implicit Null Label option allows a 7x50 egress LER to receive MPLS packets from the previous hop without the outer LSP label. This option is signaled by the egress LER to the previous hop during the LSP signaling with LDP or RSVP control protocols. In addition, the egress LER can be configured to receive MPLS packet with the Implicit Null label on a static LSP.

According to RFC 3032 "MPLS Label Stack Encoding", this is a label that an LSR may assign and distribute, but which never actually appears in the encapsulation. When an LSR would otherwise replace the label at the top of the stack with a new label, but the new label is "Implicit Null", the LSR will pop the stack instead of doing the replacement. Although this value may never appear in the encapsulation, it needs to be specified in the Label Distribution Protocol, so a value is reserved.

For LDP, RSVP, and RSVP Interfaces an "Enable Implicit Null Labeling" property will be added to SAM (as a checkbox in the GUI). SAM will be modified in the following forms:

- LDP properties window, under the Common tab. The default value will be false.
- RSVP properties window, under the Configuration group in the General tab. The default value will be false.
- RSVP Interface properties window, under the Protocol Properties tab. The RSVP Interface by default will inherit the value for Implicit Null Labeling from the parent RSVP site. This behavior can be changed by unchecking an "Inherit From RSVP" checkbox. Similar behavior is seen on the same form for the Diff-Serv TE properties.

The behavior for Static LSPs (and Static Label Maps) is slightly different. A similar "Enable Implicit Null Labeling" checkbox will be added to the GUI however the value of the checkbox is not deployed directly to the node as is the case for LDP/RSVP. Instead, the state of the checkbox will be used in determining the value deployed for the already existing Egress Label field (which according to RFC 3032, is a specific, predefined value of 3). The Egress Label field will be disabled when the checkbox is selected, enabled otherwise

The SAM GUI is modified (inclusion of checkbox which modifies Egress Label) in the following forms:

- Static LSP properties window, under the Source group. The default value is false. The Static LSP must be shutdown before modification on the Egress Label field (and thus the Implicit Null Label) may occur.
- MPLS Interface properties window, under the Static Label Map tab, on the creation form when Label Action "swap" is selected. The default value is false. The Static Label Map must be shutdown before modification on the Egress Label field (and this Implicit Null Label) may occur.

Currently the Egress Label field has a minimum accepted value of 16. The reason for this is that values below 16 are reserved (for such things as the Implicit Null label, which is represented by the value 3). The accepted range on this field will have to be changed to reflect these new lower values and some new validation will be put in place to block the still inappropriate values.

PBB (MMRP) Scalability for inter-domain services

The feature limits the scope of MMRP advertisements to a specific network domain using ISID-based filters for both MMRP control plane and BVPLS data plane.

The feature introduces a configurable option to instantiate a MMRP tree and related entry only when both a MMRP declaration and registration are received on a port.

A new policy is added in SAM. This policy is under Policies -> PBB MRP. MRP (Multiple Registration Protocol) is used in the text instead of MMRP (Multiple MAC Registration Protocol) to remain consistent with the CLI on the node and to remain consistent with the related properties on the B-VPLS Site, which appear in SAM under an 'MRP' tab.

Multiple 'Entries' can be created for a single policy.

Multiple ISID matching criteria can be specified for each Entry in an MRP Policy. The matching of ISID entries will be performed as follows:

The 'Low ISID' and 'High ISID' configure an ISID value or a range of ISID values to be matched by the MRP policy when looking at the related MMRP attributes (Group BMACs). The pbb-etype value for the related SAP (inherited from the ethernet port configuration) or for the related SDP binding (inherited from SDP configuration) will be used to identify the ISID tag

Multiple ISID ranges are allowed per entry. The following rules govern the usage of multiple ISID statements:

Overlapping values are allowed:

1. Low ISID: 1; High ISID: 10
2. Low ISID: 5; High ISID: 15
3. Low ISID: 16; High ISID: 16

The behavior on the node is to merge the overlapping ranges into a single range. For example, the above overlapping ranges would be merged by the node into:

- Low ISID: 1; High ISID: 16

There is no consistency check with the content of ISID statements from other entries. The entries will be evaluated in the order of their IDs and the first match will cause the implementation to execute the associated action for that entry and then to exit the MRP policy.

When a policy with overlapping ranges is distributed to the node, the node will merge any overlapping ranges. This would result in an inconsistency between the local policy and global policy. Therefore, when creating or modifying ISID ranges in a global policy, overlapping ranges will be merged in SAM to create a range that would match the result on a node. If the ISID range in a local policy is created or modified, no merging will be performed by SAM prior to deployment. Instead, the node will perform the merging and send a trap back to SAM, at which point SAM will resync the merged ranges.

To apply the MRP policy to a B-VPLS service, the following properties are added:

- Under Service -> B-Site -> B-L2 Access Interface -> Forwarding Control -> MRP
 - MRP Policy: A pointer to an MRP policy
- Under Service -> B-Site -> Mesh SDP Binding -> Forwarding Control -> MRP
 - MRP Policy: A pointer to an MRP policy

- Under Service -> B-Site -> Spoke SDP Binding -> Forwarding Control -> MRP
 - MRP Policy: A pointer to an MRP policy

The following properties are added to the ACL MAC Filter Policy under Policies -> Filter:

- MAC Filter Type: Specifies which type of entries this MAC Filter policy can contain. This property can only be changed if the filter is not applied and has no entries.
 - ISID: The only accepted match criteria for the filter entries are 'Low ISID' and 'High ISID'
 - Normal: All match criteria except 'Low ISID' and 'High ISID' are accepted. This will be the default value.

The following properties are added to the ACL MAC Filter Policy under the 'Match Criteria' group in Policies -> Filter -> Filter Entries -> Filter Properties:

- ISID: A SAM-only Boolean value that will specify whether or not to perform ISID matching. The default value will be 'false'.
- Low ISID: The lowest value of the 24 bit service instance identifier for this service that matches this entry. Can be equal to but not lower than the value of 'High ISID'. A value of -1 indicates no ISID matching will be performed.
 - Type: Long
 - Min: -1
 - Max: 16777215
 - Default: -1
- High ISID: The highest value of the 24 bit service instance identifier for this service that matches this entry. Can be equal to but not higher than the value of 'Low ISID'. A value of -1 indicates no ISID matching will be performed.
 - Type: Long
 - Min: -1
 - Max: 16777215
 - Default: -1

Unlike the MRP Policy ISID matching, only one ISID range can be specified for each Filter Entry. The 'MAC Filter Type' must be set to 'ISID' to set these properties.

Downstream on demand label for LDP (Tunnel only)

This feature adds support for Downstream on-Demand (DoD) label allocation as per RFC 5036 [rfc5036]. This feature can only be enabled on a link level LDP session and it will apply to prefix labels only, not service labels.

A new attribute "DoD Label Distribution" is added for LDP peer. The default is 'false'.

When this option is enabled, LDP will set the A-bit in the Label Initialization message when the LDP session to the peer is established. When both peers set the A-bit, they will both use the DoD label distribution method over the LDP session.

Figure 56: Creating an LDP Peer

PCP (dot1p) and DE bits transparency for PBB

This feature provides the ability to preserve the dot1p priority information of the incoming data and control traffic for PBB EPIPE and I-VPLS services. When the feature is enabled under PBB, an extra VLAN tag is added after the CMAC in the PBB header.

The following property is added in SAM under Service -> Epipe -> Site -> Backbone:

- **Force Q Tag Forwarding:** Specifies whether to enable addition of an IEEE 802.1q tag after the Customer MAC address when the PBB header is built, as it egresses the related B-VPLS service. This property can only be set on an IOM3 in chassis-mode 'D'. If the network element is not in chassis-mode 'D', this property will not be displayed.

The following property is added in SAM under Service -> VPLS -> I-Site -> Backbone and Service -> MVPLS -> I-Site -> Backbone:

- **Force Q Tag Forwarding:** Specifies whether to enable addition of an IEEE 802.1q tag after the Customer MAC address when the PBB header is built, as it egresses the related B-VPLS service. This property can only be set on an IOM3 in chassis-mode 'D'. If the network element is not in chassis-mode 'D', this property will not be displayed.

Since the 'Force Q Tag Forwarding' properties are only supported in chassis-mode 'D', the following validation is performed:

- 'Force Q Tag Forwarding' cannot be set when the node is not in chassis-mode 'D'
- The chassis-mode cannot be downgraded from 'D' if the 'Force Q Tag Forwarding' property is set on any of the node's (M)I-VPLS or EPIPE sites

When 'Force Q Tag Forwarding' is enabled, the I-VPLS or EPIPE service site MTU must be 22 bytes less than the mapping B-VPLS service (18 bytes + 4bytes VLAN header). The node performs this validation. Therefore, validation is added to check for this when any of the following occurs in SAM:

- The MTU on an (M)I-VPLS or EPIPE site is modified

- The MTU on a B-VPLS service site which is mapped to an (M)I-VPLS or EPIPE site with 'Force Q Tag Forwarding' enabled is modified
- A B-VPLS service site is mapped to an (M)I-VPLS or EPIPE service site which has 'Force Q Tag Forwarding' enabled
- The 'Force Q Tag Forwarding' property on an (M)I-VPLS or EPIPE site is enabled. No checking needs to be performed when the property is disabled because the validation when the property is disabled is less restrictive than when it is enabled

Once 'Force Q Tag Forwarding' is enabled in one (M)I-VPLS/PBB Epipe instance, it should be enabled in all of the related instances.

In SAM, if one site in an (M)I-VPLS/Epipe service has this property enabled, all other sites in the service should also have this property enabled. To handle an inconsistent 'Force Q Tag Forwarding' configuration, the following will be done in SAM:

A new operational flag 'Force Q Tag Forwarding Inconsistent' to indicate an inconsistent configuration will be added under the (M)I-VPLS and EPIPE service's 'State Cause' property. This flag will have no impact on the service's aggregated operational status. The modification of this flag will occur in the Service Model Post Processor. The post-processing task to update this property will be triggered after the 'Force Q Tag Forwarding' property is updated on a site in the (M)I-VPLS or EPIPE service, when a site is added to the (M)I-VPLS or EPIPE service and when a site is deleted from the (M)I-VPLS or EPIPE service.

A new alarm is introduced:

- Name: ForceQTagForwardingMisconfiguration
- Type: configurationAlarm
- Probable Cause: forceQTagForwardingInconsistent
- Default Severity: warning
- Implicitly Cleared: yes
- Raising condition: The 'Force Q Tag Forwarding Inconsistent' property transitions from 'false' to 'true'
- Clearing condition: The 'Force Q Tag Forwarding Inconsistent' property transitions from 'true' to 'false'

LDP Shortcut for IGP Route Resolution

This feature allows forwarding of IP packets to IGP learned routes using an LDP LSP.

When LDP shortcut is enabled globally (on default routing instance), IP packets forwarded over a network IP interface will be labelled with the label received from the next-hop for the route and corresponding to the FEC-prefix matching the destination address of the IP packet. In such a case, the routing table will have the shortcut next-hop as the best route. If such a LDP FEC does not exist, then the routing table will have the regular IP next-hop and regular IP forwarding will be performed on the packet.

This feature requires that an egress LER advertises and maintains a <label, FEC> binding for each IGP learned route. Also, it operates on the network interface participating in the IS-IS and OSPF routing protocols in R8.0. The resolution of BGP routes and static routes over an LDP shortcut is supported in the current release of TiMOS [IGP-Shortcut].

To support this feature in SAM, an option is added to enable LDP Shortcut globally under the routing instance for a SR/ESS 8.0 node (including 40G-7710).

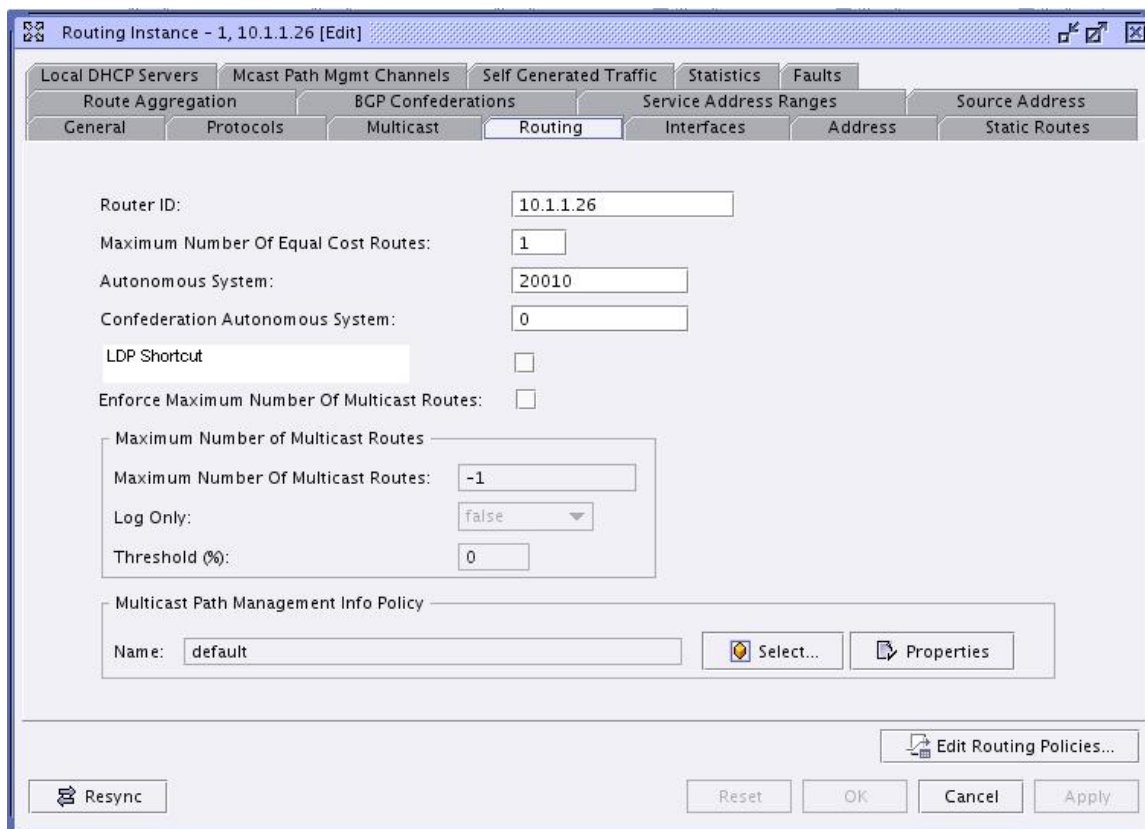


Figure 57: LDP Shortcut

GR helper for PE-CE protocols

The feature supports Graceful Restart functionalities for CE-PE routing protocols for 7750 and 7710.

Currently, GR helper function is only supported for routing protocols that are running in the default routing context. In SAM 8.0, the GR function for both BGP (all BGP site, group and peer) and OSPF is going to be supported on VPRN site as well, for both 7750 and 7710 R8.0 nodes. It is also supported by the service template. The following figures show the Graceful Restart helper for a BGP site and for an OSPF site on a VPRN routing instance.

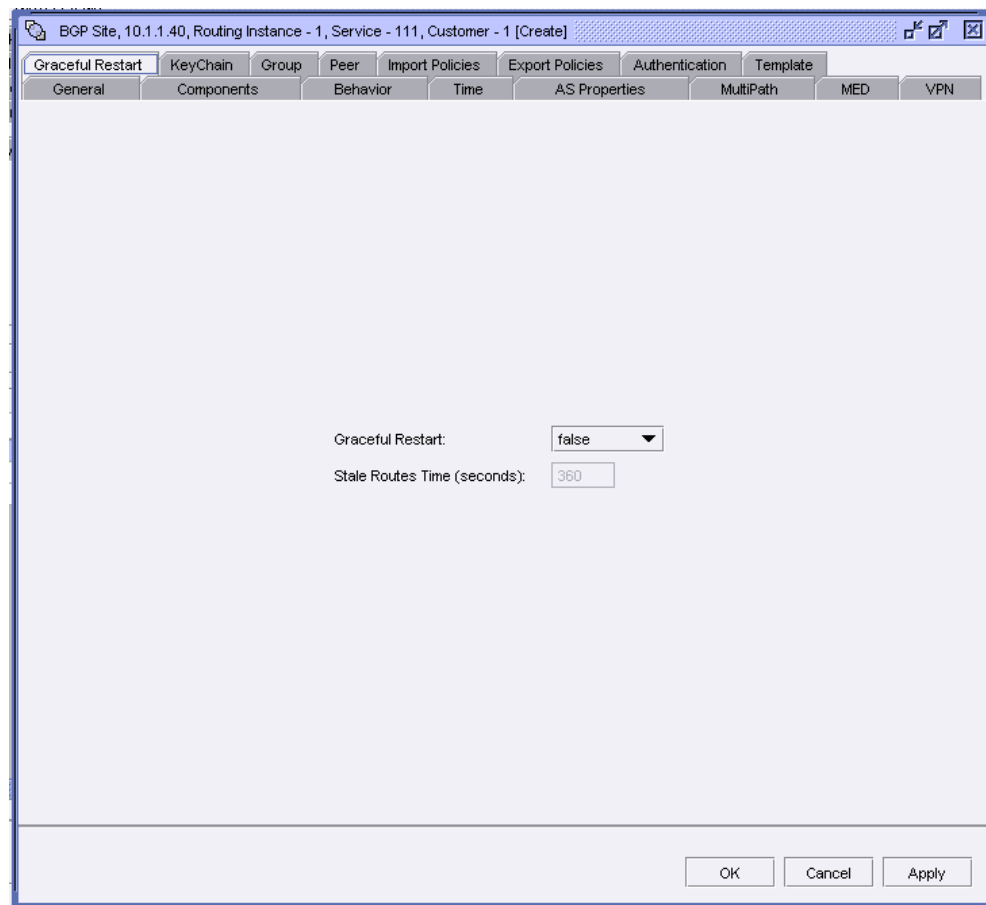


Figure 58: Graceful Restart configuration for an OSPF site on a VPRN site

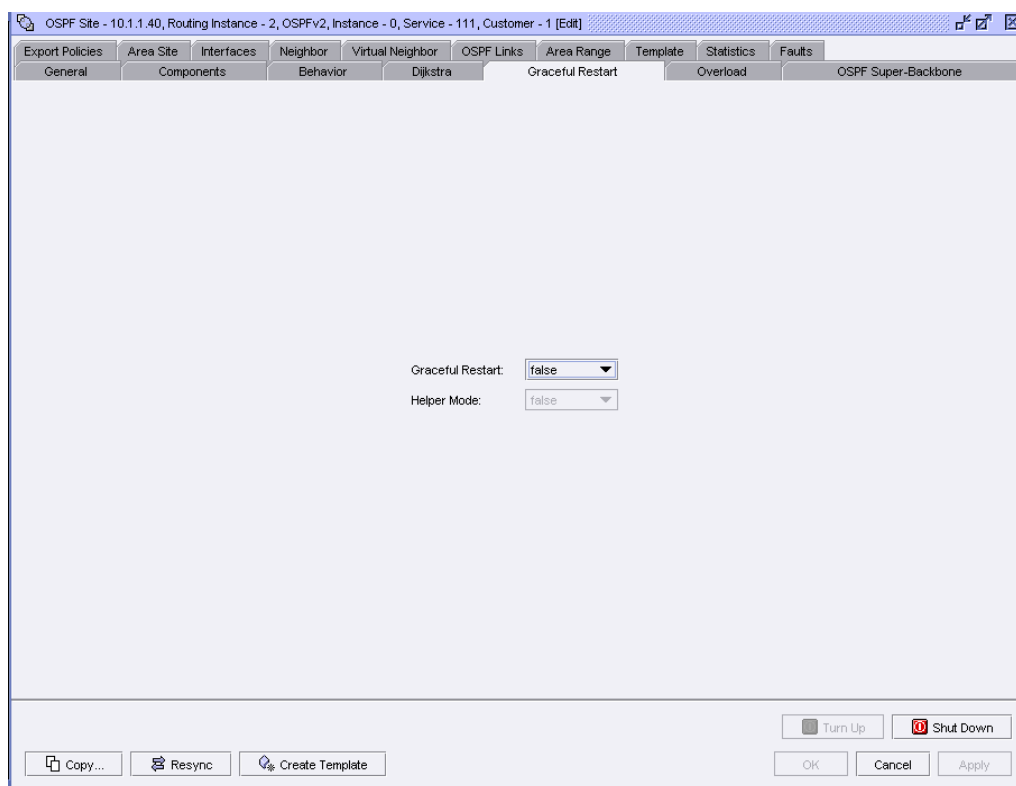


Figure 59: Graceful Restart configuration for an OSPF site on a VPRN site

Precedence support for LSP secondary paths

The feature is required to support the path revert behavior when both standby and non-standby secondary paths are configured. The preference is given to a standby secondary over non-standby secondary path with point to point LSP.

A new attribute '*Path Precedence*' is introduced on a standby LSP path. This value is defined to give priority to a specific standby path over other lower priority standby or non-standby secondary paths. The default value is 8 and the range is (1...255). Figure 13 gives an example of a standby LSP path configuration form with the '*Path precedence*' attribute.

LSP path wizard is updated to have this attributed configured during the creation time. The attribute is also added on the service template.

The screenshot shows the 'LSP-Path Binding' configuration window for 'path1, id 4; Lsp Isptest, id 2 (from sim40 (10.1.1.40) to sim41 (10.1.1.41) [Edit]'. The 'General' tab is active, displaying various configuration fields:

- General Fields:** LSP Name: Isptest, LSP ID: 2, MPLS Path Name: path1, MPLS Path ID: 4, Type: standby, Status: inactive.
- States:** Administrative: Up, Operational: Down, Failure Code: No Cspf Route Owner, Failed Site ID: 10.1.1.40, Bypass Tunnel Active: ☐.
- Transport Properties:** Maximum Transmitted Frame Size (MTU): 0, Operational Maximum Transmitted Frame Size (MTU): 0, Setup Priority: 7, Hold Priority: 0, Path Preference: 8.
- Traffic Engineering Properties:** Reserved Bandwidth (Mbps): 0, Operational Bandwidth (Mbps): 0, Hop Limit: 255, Record Actual Path: true, Record Label: record, Diff-Serv Class Type: 0, Operational Metric: 65535.
- Make before Break:** Make before Break: true, Resignal: N/A.

At the bottom, there are buttons for 'Copy...', 'Resync', 'Create Template', 'Navigate', 'Update MPLS Path...', 'View MPLS Path...', 'Turn Up', 'Shut Down', 'OK', 'Cancel', and 'Apply'.

Figure 60: Standby LSP path configuration

Multiple LDP LSR-IDs and LDP Instances

This feature provides the ability to configure and initiate multiple T-LDP sessions on the same system using different LDP LSR-ID, as well as the ability to use the LDP local interface address instead of the system address as the LSR-ID for the LDP sessions. In the current implementation, all T-LDP and LDP sessions must have the LSR-ID match the system interface address. This feature will continue to allow the use of the system interface but also any other loopback interface or local interface address on a per T-LDP/LDP session basis.

A new attribute '*Local LSR interface*' is introduced in a configuration of a LDP targeted peer. It defines the use of the address of a specific interface as the LSR-ID for the T-LDP session. The interface can be a regular interface or a loopback interface, including the system interface. By default, a t-LDP session uses the system interface. Figure 14 shows the properties configuration of a LDP targeted peer.

TargetedPeer - 10.1.1.41, Routing Instance - 1, 10.1.1.40 [Edit]

General Protocol Properties Statistics Faults

Keep-Alive Factor: 4 ☒ Inherit Value Keep-Alive Timeout (seconds): 40 ☒ Inherit Value

Hello Factor: 3 ☒ Inherit Value Hello Timeout (seconds): 45 ☒ Inherit Value

Passive Mode: false

Auto Created: true

Tunneling Enabled: ☐

Local LSR interface

Name:

Figure 61: Properties form of a LDP targeted peer

The user can also configure the use of the address of the local LDP interface as the LSR-ID to establish a link LDP adjacency and session which directly connected LDP peer. Figure 15 shows the new attribute 'Local LSR Type' defined on the LDP interface parameters. It can be either 'system' or 'interface'. The default is 'system'.

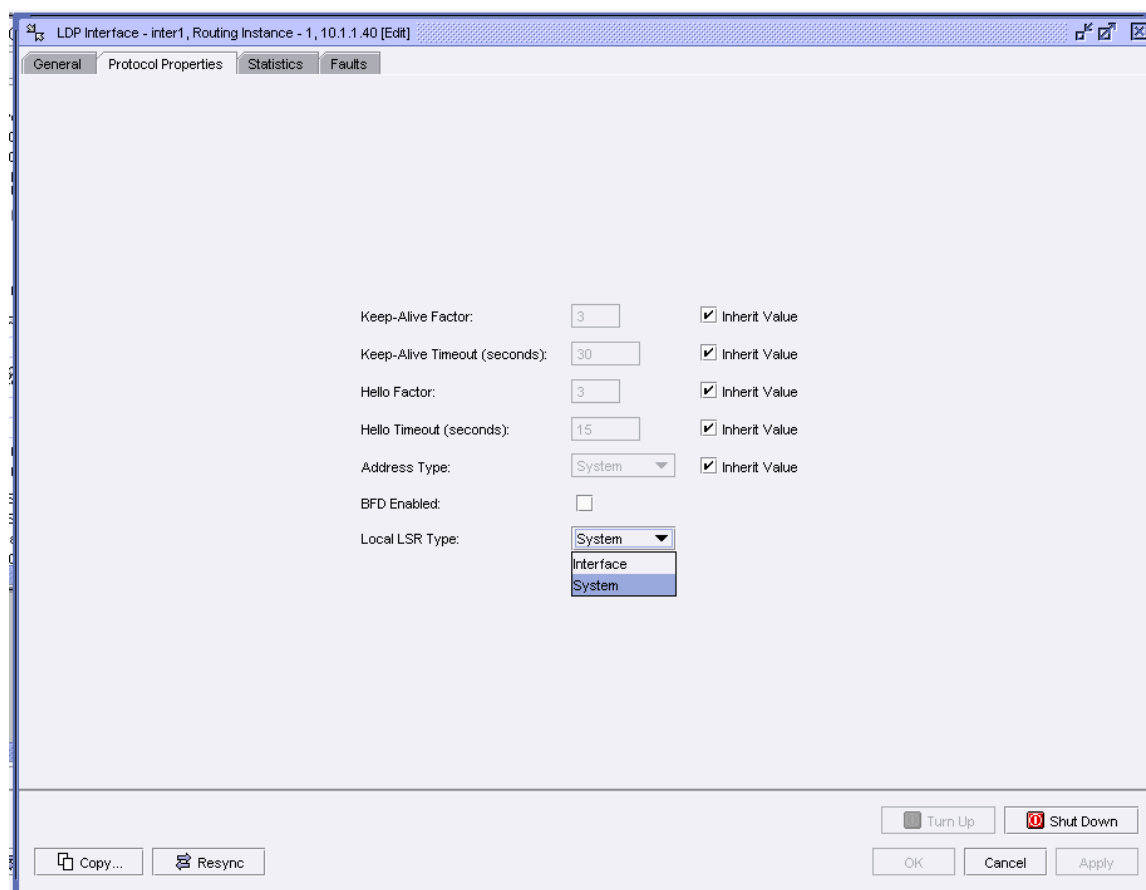


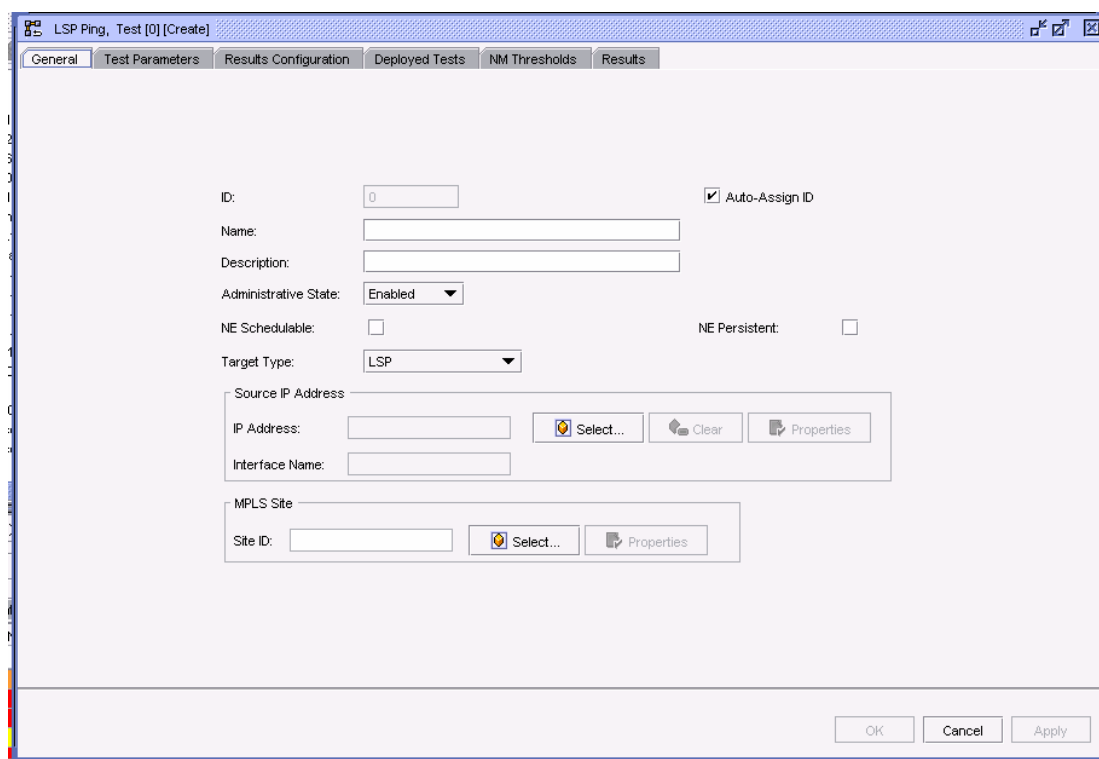
Figure 62: Configure LDP interface parameters

OAM Requirements

The feature also introduces some new OAM requirements. The OAM messages which operate over an LDP LSP and/or over a PW signaled by T-LDP and which require a response via the IP path must use a source IP address which is reachable within the same routing domain as that of the LSR-ID of the LDP or T-LDP session. In 8.0, the messages of the following OAM tools will use the T-LDP local LSR-ID as the source IP address:

sdp-mtu, sdp-ping, vccv-ping, vccv-trace, svc-ping, ldp-tree-trace

LSP Ping and trace messages for an LDP FEC MUST have the source address of the echo request message set to an address of the node sender of the message which must be reachable within the domain of the destination of the FEC. A new field '*Source IP address*' is added on the configuration of the LSP-Ping and LSP-Trace test (see Figure 16 and Figure 17).

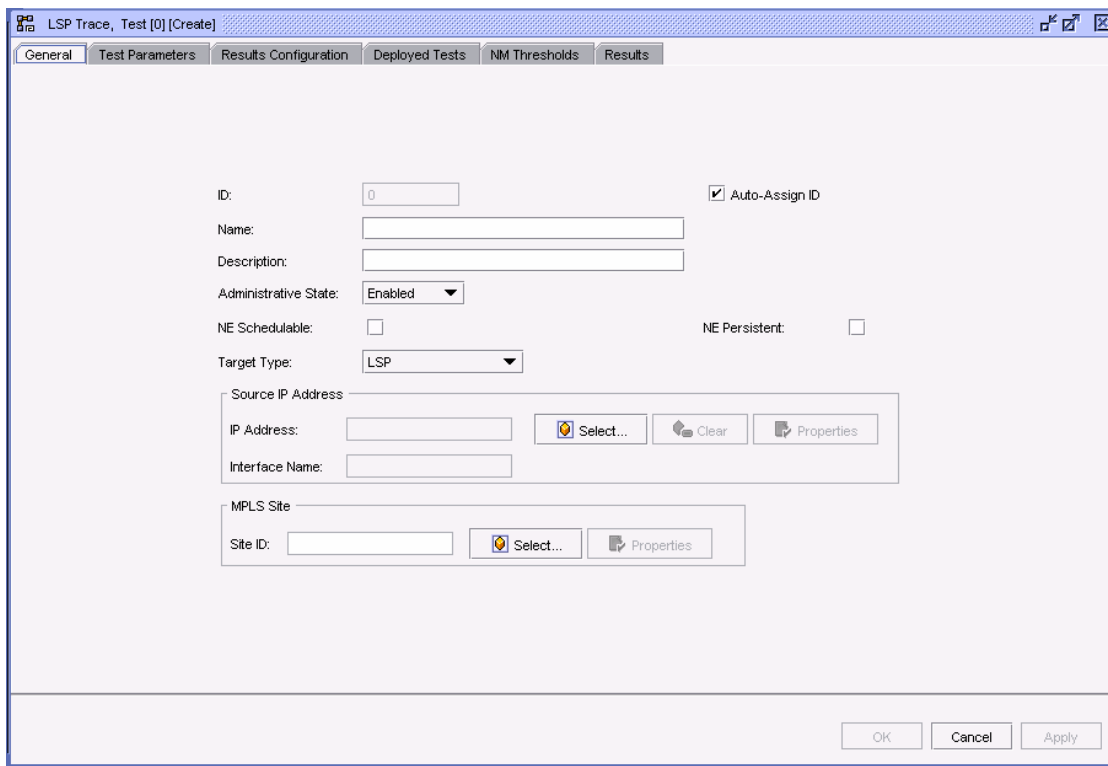


The screenshot shows the 'LSP Ping, Test [0] [Create]' dialog box. It has a tabbed interface with 'General' selected. The form contains the following fields and controls:

- ID:** A text box with '0' and a checked **Auto-Assign ID** checkbox.
- Name:** An empty text box.
- Description:** An empty text box.
- Administrative State:** A dropdown menu set to 'Enabled'.
- NE Schedulable:** An unchecked checkbox.
- NE Persistent:** An unchecked checkbox.
- Target Type:** A dropdown menu set to 'LSP'.
- Source IP Address:** A group box containing:
 - IP Address:** An empty text box, a 'Select...' button, a 'Clear' button, and a 'Properties' button.
 - Interface Name:** An empty text box.
- MPLS Site:** A group box containing:
 - Site ID:** An empty text box, a 'Select...' button, and a 'Properties' button.

At the bottom right are 'OK', 'Cancel', and 'Apply' buttons.

Figure 63: LSP-Ping test creation form



The screenshot shows the 'LSP Trace, Test [0] [Create]' dialog box. It has a tabbed interface with 'General' selected. The form contains the following fields and controls:

- ID:** A text box with '0' and a checked **Auto-Assign ID** checkbox.
- Name:** An empty text box.
- Description:** An empty text box.
- Administrative State:** A dropdown menu set to 'Enabled'.
- NE Schedulable:** An unchecked checkbox.
- NE Persistent:** An unchecked checkbox.
- Target Type:** A dropdown menu set to 'LSP'.
- Source IP Address:** A group box containing:
 - IP Address:** An empty text box, a 'Select...' button, a 'Clear' button, and a 'Properties' button.
 - Interface Name:** An empty text box.
- MPLS Site:** A group box containing:
 - Site ID:** An empty text box, a 'Select...' button, and a 'Properties' button.

At the bottom right are 'OK', 'Cancel', and 'Apply' buttons.

Figure 64: LSP-Trace test creation form

IP Pseudowire L3 Termination

The purpose of this feature is to allow termination of an Epipe/Lpipe into a L3 service, such as IES or VPRN. Spoke termination using Epipe is already supported in SAM. Spoke termination using Lpipe is now required for Release 8.0.

Previously, SAM allowed traffic to be sent across a spoke SDP binding from an Epipe to an IES/VPRN service. This feature will extend this functionality to allow traffic from an Lpipe in a similar fashion. To accommodate this, a property will be introduced in SAM on the spoke SDP binding creation form (when created from an IES/VPRN) that will allow the user to select which type of virtual circuit should be accepted.

- VC Type (drop-down box)
 - Type: Enumeration
 - Values:
 - ether (Epipe)
 - lpipe (Lpipe)
 - Default Value: Ether
 - Applicability:
 - Can only be set during creation of the spoke SDP binding
 - lpipe enumeration value only applicable to chassis mode C and D

The template framework will work correctly with this property without any modification.




Figure 65: Lpipe VLL spoke SDP binding termination into a VPRN service

In Figure 18, a spoke SDP binding from PE1 is terminated into a VPRN (could also directly connect to an IES service). A spoke SDP binding in this case would be created from an Lpipe on PE1 and terminate on a VPRN service on PE2. On the PE2 VPRN service we have to create a spoke SDP binding specifying the VC Type as 'lpipe'. If the PE1 VLL was Epipe then the VPRN service on PE2 would be required to use a VC Type of 'ether', which was the only supported type prior to 8.0. When this setup (as shown in Figure 18) is configured within SAM, a composite service is automatically created for the involved Epipe/Lpipe and VPRN/IES. This feature may require action on an RCA Audit. If the vc-type is specified as Epipe but an Lpipe is attempted to terminate into it we may make the user aware of this. More investigation required.

T-LDP status TLV (Active/Standby and Oper) support on VPRN

This feature enhances the fault propagation between PE devices in a spoke SDP termination between an Epipe/Ipipe and IES/VP RN service. It provides the ability to propagate faults in the SAP or PW to the PE where the IES/VP RN service is configured without withdrawing the PW labels. It also provides more efficient mechanisms for switching between active/standby spoke SDPs in a redundant configuration.

It is convenient to split the node specification into two discrete pieces.

- T-LDP Status Messaging on the IES/VP RN PE**
 In 8.0 the 7x50/7710 nodes have introduced support for T-LDP status messaging between PEs connected in a fashion portrayed in the *Figure 19*. This enhances the fault propagation between PE devices in a spoke termination on IES/VP RN service.

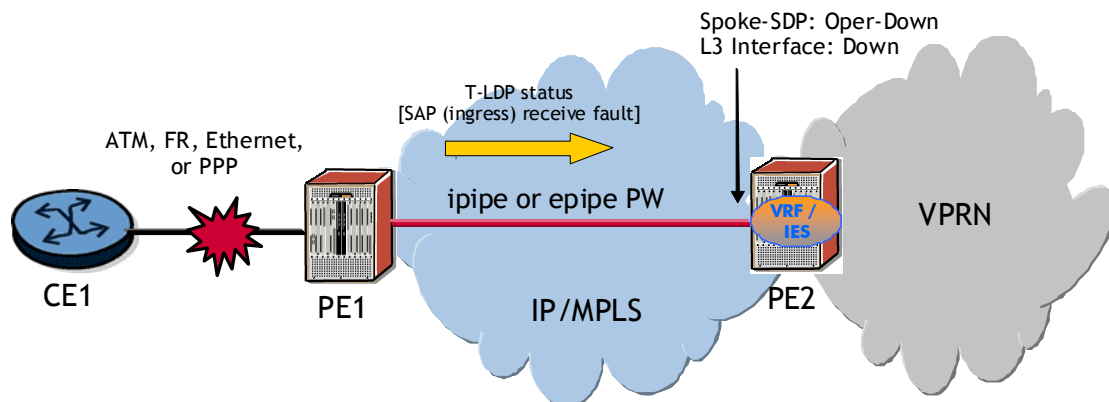


Figure 66: T-LDP Status Messaging

The PRD discusses all the possible messages that can be communicated between the PEs and what is the result of one PE receiving a specific message from the other. The fact that the node has sent/is sending a specific message is not published publicly. Looking at this from the 5620 SAM point of view, the result of any of this communication will (possibly) effect operational status of associated objects. The node will handle setting these states correctly and alerting SAM of the changes, meaning that within SAM everything will behave as it has in the past.

- Spoke SDP Redundancy into IES/VP RN**
 SAM already has a means of providing redundant spoke SDP bindings from an Epipe or Ipipe service, terminating into an IES/VP RN service. This feature provides a more efficient way (from the nodes perspective) for the VLL to switch between active and standby spoke SDP bindings.

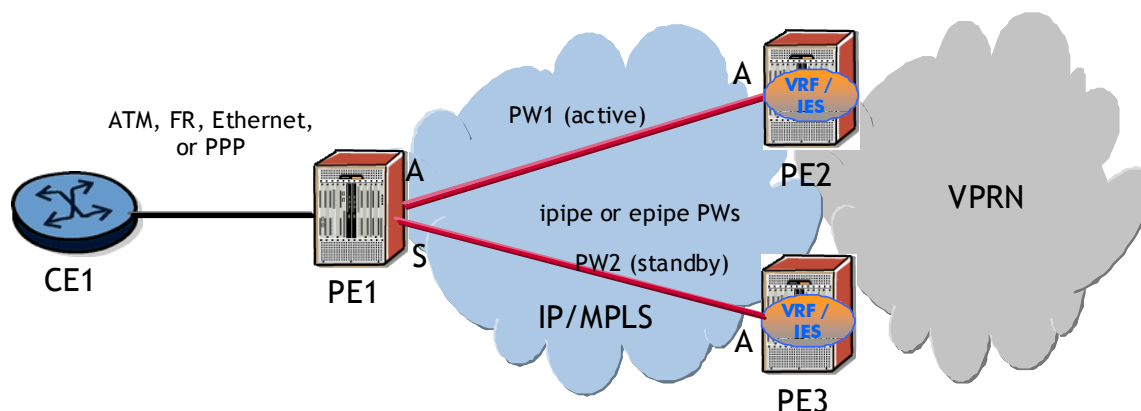


Figure 67: Redundant Spoke SDP Bindings between a VLL and VPRN

Consider the example configuration above. PE1 terminates two spoke SDPs into a VPRN. PE1 chooses to forward traffic on one of the spoke SDPs (active), while blocking traffic on the other(s). With the implementation of the status messaging for the first part of this feature, the node now has a more efficient way of switching between these active/standby SDP bindings when needed.

For backward compatibility reasons (in the figure, PE2 could be 8.0 while PE3 could be pre-8.0 and hence not have T-LDP status messaging), the node must be able to perform the switchover either the new way (T-LDP status messaging) or the old way (PW label withdrawal). To accommodate this, a new configurable property has been introduced on PE1 which indicates whether or not T-LDP status messaging should be used in a redundant setup.

To account for this in SAM, two changes are made. First, a new property is introduced to the Endpoint creation form. Details of the property are as follows:

- T-LDP Standby Signaling (checkbox)
 - Type: Boolean
 - Default Value: False
 - Applicability: Ipipe/Epipe VLL for 7x50, 7710, 7705

The template framework handles this new property.

The ability to enable T-LDP Standby Signaling on a VLL endpoint must be blocked under the following conditions:

- If vc-switching is enabled on the VLL site
- If Inter-Chassis Backup (ICB) is enabled on any spoke SDP bindings configured to use this endpoint
- If a SAP under this endpoint belongs to an MC-LAG/MC-APS

Secondly, when a PE that belongs to the VPRN/IES receives the T-LDP PW standby message the L3 Access Interface associated with that spoke SDP will go operationally down. To remain consistent in SAM, the spoke SDP using this interface will be made purple in the topology map, representing that this is the backup spoke SDP.

Option to Place IGP into Overload if Switch Fabric Fails

In a typical system in the event of a SFM failure, multicast traffic needs to be rerouted around the node. The defined failure scenarios include:

- There is only one SFM installed in the system

- One SFM (active or standby) failed in a dual SFM configuration
- In the process of ISSU
One solution is to use overload state in IGP to trigger the traffic reroute by setting the overload bit in IS-IS or setting the metric to maximum in OSPF. Since PIM uses IGP, a next-hop change in IGP will cause PIM to join the new path and destroy the old path, which effectively reroutes the multicast traffic. When the problem is resolved, the overload condition is cleared, which will cause the traffic to be routed back to the router.

To support this feature, following properties are added to Routing > Routing Instance Display > Routing Tab & Manage > Services > VPRN > Site > Routing/General Tab:

Single SFM Overload Admin State: This property specifies the administrative state of the IGP single SFM overload. When the value is equal to 'inService', IGP protocols (either IS-IS or OSPF) enter an overload state.

Hold-Off Time: This property is available only when property 'Single SFM Overload Admin State' is set to 'In Service'. It specifies the delay in seconds between the detection of the single-SFM condition and the IGP entering the overload state.

Following properties are added to Routing > Routing Instance Display > Routing Tab:

Single SFM State: This property is read-only and indicates the IGP single-SFM-overload state.

Single SFM Start Time: This property is read-only and it indicates the last time that the 'Single SFM State' had a transition from 'notApplicable' or 'normal' to 'overload'. If such a transition never occurred, this object contains a zero value.

Single SFM Interval: This property is read-only and indicates the duration of the most recent overload.

Load Sharing Multiple BGP Paths Even If AS-Path is Different

This feature allows the as-path comparison to be disabled on a per address family basis. This allows BGP routes which do not have equal AS paths to be considered equal and therefore load balanced across more BGP routes.

SAM currently supports property 'AS Path Ignore' in BGP for SR-OS which if enabled, disables the use of the AS path length in the BGP route selection process for all BGP address families. This feature will extend this functionality so that the system can be configured to ignore the AS Path length in the BGP route selection process only for certain address families.

To support this feature, following property is modified in Routing > BGP Site > VPN tab:

AS Path Ignore: This property is currently displayed under Behavior tab; as a part of this feature, it will be moved to the VPN tab and following options will be added to it. Each option will be presented in the format of a check-box with the default of unchecked. Once selected, the specified AS paths of incoming routes are not used in the route selection process for the given address families. By default, all options are un-selected:

- IPv4 - the AS-path length will be ignored for all IPv4 routes
- VPN-IPv4 - the AS-path length will be ignored for all IPv4 VPN routes
- IPv6 - the AS-path length will be ignored for all IPv6 routes
- VPN-IPv6 - the AS-path length will be ignored for all IPv6 VPN routes
- Mcast-IPv4 - the AS-path length will be ignored for all IPv4 multicast routes
- MVPN-IPv4 - the AS-path length will be ignored for all MVPN IPv4 multicast routes

- L2-VPN - the AS-path length will be ignored for all L2-VPN NLRIs

Following property is modified in Manage> Services > VPRN > BGP Site > VPN tab:

AS Path Ignore: property is currently displayed under Behavior tab; as a part of this feature, it will be moved to the VPN tab and following options will be added to it. Each option will be presented in the format of a check-box with the default of unchecked. Once selected, the specified AS paths of incoming routes are not used in the route selection process for the given address families. By default, all options are un-selected:

- IPv4 - the AS-path length will be ignored for all IPv4 routes
- IPv6 - the AS-path length will be ignored for all IPv6 routes

There are no new alarms for this feature. This feature does not impact Templates.

Since the type of this attribute has changed from Boolean to multi-select check-boxes, existing OSSI scripts should be modified to accommodate this change.

Support for RFC3107 BGP Label for L2 Services

To support Inter-AS VPLS/VLL service based on option C for 7x50 SR, a new SDP option based on BGP route tunnel will be added to SAM.




Figure 68: Inter-AS L2VPN

In figure 21, far-end PE loopback address is advertised in local AS via LDP or BGP-labeled routes (RFC3107) to provide inter-AS L2VPN services based on option C. MPLS SDP provides an option to use BGP route tunnels.

To provide support for this feature, MPLS SDP is enhanced to use BGP route tunnel to extend inter-AS support for L2VPN services. A single method of tunneling is allowed per SDP (i.e. LDP, RSVP-TE LSP or BGP route tunnel). BGP route tunnel method is excluded if multi-mode transport is enabled for an SDP.

Auto provisioned BGP route tunnel SDP is not supported in R8.0 (will be added in a later release). Only manually provisioned SDPs are supported in this release.

To support this feature, following properties are added to Manage > Service Tunnels > SDP (MPLS) > General tab > MPLS Signaling Group:

Enable BGP Tunnel: This property specifies whether the transport tunnel is BGP as opposed to LDP or RSVP signaled LSPs. This value cannot be set to 'true' if 'Enable LDP' is set to 'true' or if there is at least one RSVP or static LSP provisioned. This property applies to MPLS SDPs only

SDP Creation wizard will have this new attribute added to step 4 'Specify Transport'. When 'BGP Tunnel Enabled' is true, 'LDP Enabled' can not be set to true.

SDP binding is supported for VPLS/ MVPLS/ BVPLS/ MC-LAG/ VPRN/ IES and EPIPE services. Service creation window and SDP Binding creation window will be modified with this new attribute/option.

BGP Tunnel LSP ID: This property is read-only and indicates the ID for the BGP Tunnel LSP if 'BGP Tunnel Enabled' is set to 'true' and there is a valid LSP route to the SDP far-end IP address. Multiple BGP SDPs are assigned with the same ID as long as they are sharing the same LDP LSP. If multiple SDPs are available from the same PE to the same destination, the lower metric/preference one is chosen.

This feature has potential impact on CPAM Tunnel Management. Testing and potential additional coding is required to ensure there are no negative impacts.

PBB, RSVP LSP as transport within the AS, and LDP over RSVP as transport within the AS are not supported for 3107 BGP tunnel.

RSVP Shortcut for BGP Next-Hop Resolution

RSVP-TE LSP shortcut for BGP next-hop route resolution allows forwarding of IPv4 packets to routes resolved to a BGP next-hop using an RSVP-TE LSP. The RSVP LSP must have a destination address matching the /32 address of the BGP next-hop. When the RSVP-TE shortcut is enabled in BGP, a route resolved to a BGP next-hop will use the RSVP LSP towards the BGP next-hop if available in the tunnel table, if not, the regular IGP next-hop is used.

Support of Labeled IPv6 Routes (6PE) will be added in a later release. To provide support for this feature, existing 'IGP Shortcut' property on Routing > BGP Site> IGP Shortcut tab will be enhanced with new enumerations:

IGP Shortcut: Following new enumerations are added to this existing property:

- **MPLS:** This option instructs BGP to first attempt to resolve the BGP next-hop to an RSVP LSP. If no RSVP LSP exists or the existing ones are down, BGP automatically searches for the LDP LSP with a FEC prefix corresponding to the same /32 prefix in the tunnel table and resolves the BGP next-hop to it.
- **RSVP-TE:** This option means RSVP Traffic Engineering will be used to resolve paths to BGP next-hops.

PBB Ethernet Tunnel Enhancements

Release 7.0 provided support for Ethernet Tunnels. Some additions to this area are included in release 8.0.

When creating an Ethernet Tunnel a number of properties can be specified. Release 8.0 supports additional values and ranges as follows.

- Tunnel ID: Integer (min: 1, max: for pre-8.0 R1 nodes 64; for 8.0 R1+ nodes 1024)
- Protection Type: In 7.0R4, this can only be 'G8031 1:1'. As of 8.0 R1, can be either 'G8031 1:1' or 'Load Sharing'.
- Hold Time Up: The time, in deciseconds, used for the hold-timer for associated CC Session up event dampening (min: 0, max: 5000, default: 20). Applicable only to 8.0 R1+ nodes.
- Encapsulation Type: The encapsulation type of the ports that can be used as the member ports for the Path Endpoints. For pre-8.0 R1 nodes, can only be 'dot1q'. For 8.0 R1+ nodes, can be either 'dot1q' or 'qinq'.

When creating an Ethernet Tunnel Endpoints a number of properties can be specified. Release 8.0 supports additional values and ranges as follows.

- Tunnel Endpoint ID: Integer (min: 1, max: for pre-8.0 R1 nodes 64; for 8.0 R1+ nodes 1024)
- Hold Time Up: The time, in deciseconds, used for the hold-timer for associated CC Session up event dampening (min: 0, max: 5000, default: 20). Applicable only to 8.0 R1+ nodes.
- Encapsulation Type: The encapsulation type of the ports that can be used as the member ports for the Path Endpoints. For pre-8.0 R1 nodes, can only be 'dot1q'. For 8.0 R1+ nodes, can be either 'dot1q' or 'qinq'.

When creating a global Path, two local Path Endpoints must be created or added. Some changes to the values and ranges of some properties for 'Endpoints'.

- Path ID: An integer. For G.8031 1:1 protection, this can be either '1' or '2' in 7.0 R4. Beginning in 8.0 R1, this can be from 1 to 16.
- Control Tag (Outer Encapsulation Value): The control tag for the Path Endpoint. Once set, this property cannot be edited.
For 7.0 R4 network elements:
 - Min = 0, Max = 4094, Default = 0
 - A value of 0 specifies an invalid, unspecified control tag and will be mapped by SAM to -1 (0xffffffff) prior to deployment on the node.For 8.0 R1+ network elements:
 - Min = -1, Max = 4095 (4095=*) Default = -1
 - A value of -1 specifies an invalid, unspecified control tag.
- Inner Encapsulation Value: The inner encapsulation value for the control tag property for the Path Endpoint. Once set, this property cannot be edited. The value can be 0 to 4094 inclusive. This property will only apply to 8.0 R1+ nodes and only if the Ethernet Tunnel Endpoint's 'Encapsulation Type' is 'qinq'.

If the Path Endpoint already exists, it will not be necessary to set the above properties. Instead, the Path Endpoint can be selected as the 'Path Endpoint' property.

For 8.0 R1+ nodes, the Path Endpoint cannot be administratively enabled if:

- An operationally up same-fate SAP on the Ethernet Tunnel Endpoint does not have a tag configured for the Path Endpoint.
- The control tag or member port is not configured.

As of 8.0 R1, the following properties will be added to the 'General' tab of the Ethernet Tunnel Path Endpoint under the 'APS' group. These properties will not be available on creation and will only be applicable if the protection type is 'G8031 1:1'.

- Command Switch: The switch command to initiate on an Ethernet Tunnel APS group.
- Exercise Command Result: A read-only property that displays the result of the last issued 'Exercise' command.

Once these Ethernet Tunnels are created they can be selected as the terminating port when creating a B-SAP on a B-VPLS.

For 7.0R4 a Tunnel Endpoint can only be used by one B-SAP.

For Release 8.0, a Tunnel Endpoint can be used by:

- Epipe SAPs
 - egress (HSMDA), ingress (HSMDA), ring node configurations will be blocked
- Lpipe SAPs
 - egress and ingress HSMDA configurations will be blocked
- VPLS, i-VPLS, m-VPLS SAPs
 - Capture SAP, arp Host, arp reply agent, authentication, calling station id, dhcp, egress (HSMDA), host connectivity verify, igmp host tracking, ingress (HSMDA), msap, ppoe policy, static host, sub sla management, trigger packet configurations will be blocked
 - On the service, endpoint, gsmpp, host-connectivity verify, igmp host tracking, igmp snooping, interface, mac subnet length, mcr default gtw, radius discovery configurations will be blocked

As of 8.0 R1, fate-sharing is possible. The following properties will be added under the 'Port' tab of the SAP configuration form:

- A property 'Ethernet Tunnel Endpoint Control SAP' is added to specify whether or not this is the control SAP. This is a SAM-only property and is not deployed to the node.

For 7.0R4 network elements, the encapsulation for an Ethernet Tunnel Endpoint SAP is always 'null'.

For 8.0 R1+ nodes, the encapsulation is 'dot1q'.

In SAM 7.0R4, the tunnel can only use dot1q ports.

In SAM Release 8.0, the tunnel can use qinq ports. To support this, the 'Encapsulation Type' property is added to the Ethernet Tunnel Endpoint configuration.

'Create CFM Continuity Check'

The Continuity Check (CC) is an integral part of Eth-APS tunnel. SAM shall be able to accelerate the creation of (global and local) MAs, MEPs.

As of 8.0 R1, the Service Test Manager CC Test Creation form has an 'Ethernet Path' tab. This tab displays the global Ethernet Path that is associated with this test. The creation of this object automatically creates MEPs on both Path Endpoints.

Point to Multi Point LSP Enhancements

Point-to-multipoint (P2MP) MPLS LSP allows the source of multicast traffic to forward packets to one or many multicast receivers over a network without requiring a multicast protocol, such as PIM, to be configured in the network. A P2MP LSP tree is established in the control plane which path consists of a head-end node, one or many branch nodes, and the leaf nodes. Packets injected by the head-end node are replicated in the data plane at the branching nodes before they are delivered to the leaf nodes.

Tunnel Interface Changes in 8.0

Tunnel Interfaces have changed in SAM 8.0 to accommodate MIB changes. In SAM 8.0, Tunnel Interfaces are moved out of Pim and are associated with the Virtual Router Class. This applies to all NEs including 7.0. In SAM 8.0 the following is done to configure a Tunnel Interface: From the Routing Navigation Tree, select the Routing Instance. Click on the Tunnel Interface Tab. The Tunnel Interface is now configured as a child of the Routing Instance. It is no longer under the Pim Site.

SAM 8.0 Changes to IGMP Tunnel Interface

A validation has been added in SAM 8.0 for IGMP Tunnel Interfaces. In this case the LSP-name chosen for the IGMP tunnel interface must be associated with a pre-existing Virtual Router Tunnel Interface. In SAM 8.0, when the SAM user selects the LSP-name for the IGMP Tunnel Interface, a selection of Virtual Router Tunnel Interfaces with send-address 0.0.0.0 will be displayed. This differs from SAM 7.0, where the user will see a selection of P2MP LSPs when configuring the LSP Name.

Multicast Tree BFD

BFD on a multicast tree is similar to which over a point-to-point network except that leaf nodes are not allowed to communicate with the root node due to heavy overhead (by order of 1:N) inflicted on up stream routers. Hence in practice it becomes Unidirectional Failure Detection, or UFD.

Source Redundancy

During a network failure, independent of the location of failure, multicast receivers must be able to receive data stream with least period of interruption. Multicast data distribution over P2MP is able to achieve protection against failure in the core network of the service provider using fast reroute method. However fast reroute available on 7x50 SR is limited to link protection such that failures requiring node protection (e.g. head end node failure), or PE-CE link failures will have catastrophic consequences.

This feature introduces a 2-part solution to remedy this problem:

- Failure detection using BFD
- Standby P2MP LSP switch-over

The way BFD operates over a multicast tree is mentioned in previous section. When BFD is needed on a P2MP LSP, it can be enabled on PIM/IGMP tunnel interfaces on LERs. There are two types of protection BFD can provide and they target different types of failure:

- P2MP Ingress PE BFD for failure between an ingress LER (root node) and an egress LER (leaf node)
- P2MP Source BFD for failure between a multicast source and an egress LER (leaf node). This protection covers the path between multicast source and ingress LER.

Note that currently only RSVP P2MP is supported for this feature and LDP based P2MP would be supported in a future release.

Once BFD is operational, user can make configuration of a standby P2MP and its BFD session. This configuration has to be done on tunnel interfaces of egress LERs. BFD and standby P2MP LSP are closely associated together so that only a failure detected by BFD session will trigger a switchover to the designated LSP. Other failures will still proceed with fast reroute as resolution.

On a tunnel interface configuration form, a new “Source Redundancy” tab is added, it has two sub tabs: “General”, and Source Specific Parameters”. Under “General” tab, there are two groups: “BFD”, and “Standby Configuration”. Depending on whether the current tunnel interface is ingress or egress (specified on general properties), contents under “BFD” group are different, this is because of their different roles in a unidirectional BFD session.

Since a BFD session is initiated on the head-end root node, both transmit interval and multiplier can only be configured here on its ingress tunnel interface. Their values will then be distributed to leaf nodes and their egress tunnel interfaces. Session Id, however, must be configured explicitly on both ends with matching values; otherwise BFD session cannot be established. SAM shall ensure same session id is used on all LERs.

Likewise, hold timer is only available on an egress tunnel interface.

A failure on one single S2L would not result in a breakdown of entire P2MP multicast tree in most cases, A switchover to a standby LSP will also not swap the branches that are still operational. This means the decision of switchover to standby is made at leaf nodes. Hence, next group “Standby Configuration” is only available on an egress tunnel interface.

Under “Source Specific Redundancy” tab, users can configure multiple instances of source redundancy. Properties are similar to which under “General” tab and each represents an independent BFD session protection per source. The only difference is that on an egress tunnel interface, they share the same standby P2MP LSP with tunnel interface level protection. However, since each redundancy object contains a particular BFD session, source level redundancy is allowed to have its own standby BFD session over a common standby P2MP LSP.

There is a limitation on how many P2MP BFD sessions a node can have, which are 15 for tunnel interfaces and 35 source specific.

LDP for P2MP

LDP Support for P2MP can be enabled on LDP interfaces. Under “General” tab, a new checkbox “P2MP Enabled” is added below “Protocol” field. This field can only be seen/modified on a LDP interface but not on a targeted peer.

On LDP Routing Instance Configuration, “MBB Time” is added to the end of common tab. Taking value of range (0..10|3) seconds, it specifies the maximum time a P2MP transit node must wait before switching over to a new path if the new node does not send Make Before Break (MBB) Tag Length Value (TLV) to inform of the availability of the data plane.

Currently the only application of this feature is Global Tunnel Interface. Both PIM and IGMP are supported. Different from RSVP based P2MP LSP; P2MP Id is used as index for LDP based tunneling instead of LSP Id. It must be noted that BFD and Source Redundancy support that becomes available in Rel 8.0 is only applicable for RSVP based Tunnel Interfaces, not LDP ones.

Multicast Info Policy is also updated to support RSVP. While SAM still keeps reference to a P2MP LSP, it deploys P2MP Id to the node as identifier. There should be no change to the functionality otherwise.

P2MP-Lsp-Ping for RSVP P2MP LSP

The 7750 supports a new command `p2mp-lsp-ping`. The user specifies the P2MP LSP name and optionally the P2MP instance name. There is also an option to specify up to five addresses to limit the number of return packets. If no destination IP address is specified, then all possible S2L leaf nodes may send a response packet. The results are displayed for each destination IP address.

The parameters of the test are the following:

- 1) P2MP LSP-name
- 2) P2MP instance-name (optional)
- 3) S2L-dest-addr (optional up to 5 may be specified)

P2MP-Lsp-Trace for RSVP P2MP LSP

The P2MP LSP Trace is a new type of Trace test that allows the user to trace the path of a single S2L path of a P2MP LSP.

The parameters of the test are the following:

- 4) P2MP LSP-name
- 5) P2MP instance-name
- 6) S2L-dest-addr

The P2MP LSP trace probe results on all egress LER nodes eventually receiving the echo request message but only the traced egress LER node will reply to the last probe. The branch LSR or BUD LSR which as a downstream branch over which the traced egress LER is reachable will respond. The test result shows the Hops and Probes in the result display.

The P2MP LSP Ping and P2MP LSP Trace for RSVP P2MP LSP can be executed from within a Test Suite. Both tests can be generated within the Test Suite.

The Test Policy that has the Tested Entity Type set to MPLS shall allow the configuration of P2MP LSP Ping Definition and P2MP LSP Trace Definition.

The Test Suite with the Tested Entity Type set to MPLS LSP shall allow the selection of RSVP P2MP LSP within the Tested Entity list.

The auto-generation rules shall be updated to include the Test Suite type MPLS associated with a Test Policy with P2MP RSVP LSP Ping Definition or the P2MP RSVP LSP Trace Definition.

Hierarchical Policy Support

The feature provides the following main functions:

- Policer Control Policy support
- Shared Policer output queues support
- Queue Group Template and Port Queue Group enhancement
 - All override configuration is targeted for 8.0R4
- Access Ingress/Egress Policy enhancement
 - The ANCP policy enhancement is targeted for 8.0R4
- Subscriber Management enhancement
 - All override functionality is targeted for 8.0R4
- Policer Control Policy on Saps
 - All override functionality is targeted for 8.0R4

This feature is essentially the support of another set of policy type: The Policer Control Policy. The new policy types introduced with this feature are implemented using the existing policy framework and behave accordingly. The common SAM policy function is supported for this new policy, i.e. distribution, audit etc. The new policy will be applied to other SAM objects though assignment with pointers to policy or policy name. These assignments are listed in the policy form as binding tabs.

Configuration

The configuration of Policer Control hierarchies is performed as follows:

1. Configure Policy Control Policy, create Arbiters and configure the parent/child relationship to build a Policer Control hierarchy.
2. Create Policer inside the Ingress/Egress QoS Policy to configure the traffic control parameters. Assign parent arbiter to the policer so the traffic control is performed within the hierarchy configured in the Policer Control Policy.
3. Assign both Policer Control Policy to a SAP and/or a Subscriber. Once assigned, the node will try to find the arbiter assigned to the Policer from the Policer Control Policy. If find, the Policer be activated within the policer Control hierarchy configured in the Policer Control Policy.
4. Policer control hierarchies are now created on SAPs or on a subscriber context.

The following picture provides a summary of the related objects referred in this feature and there relationship (new objects are highlighted in yellow).

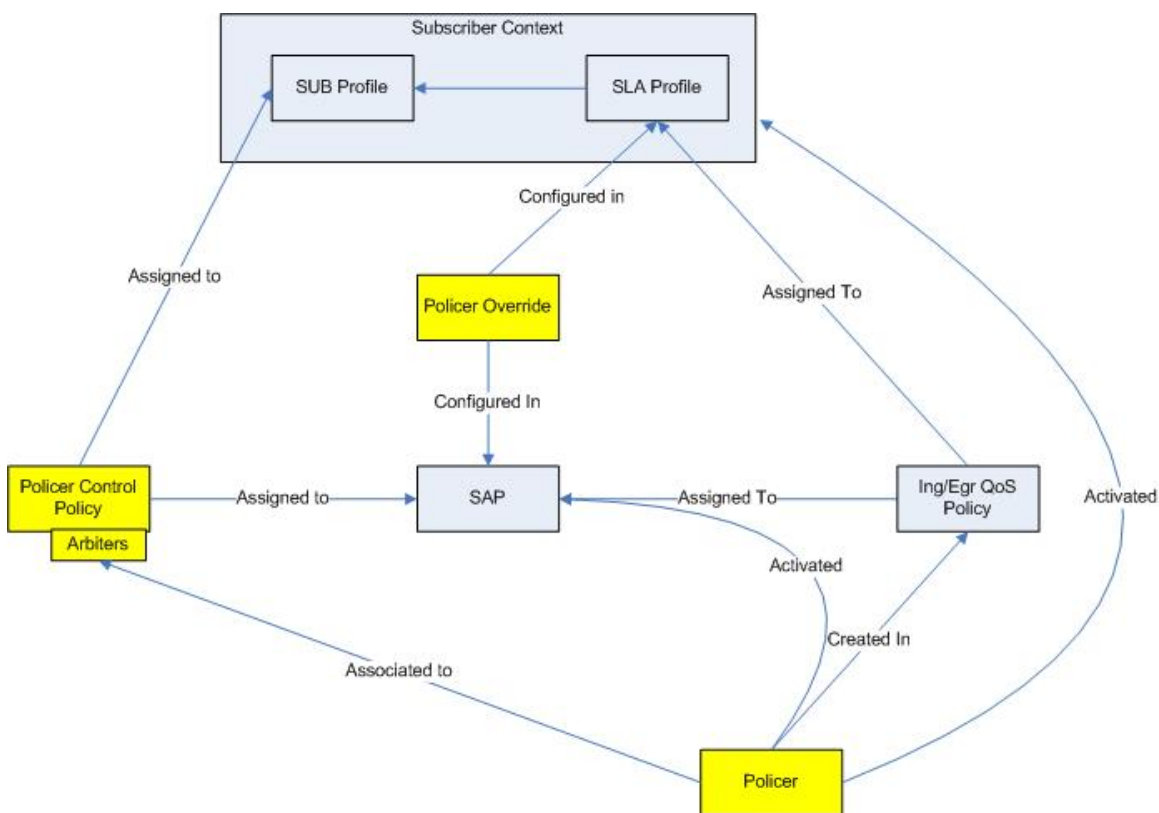


Figure 69: Policer Control Policy and related key objects

Note: Only key relations and objects are displayed.

Policer Control Hierarchies

Policer control hierarchies may be created on SAPs or on a subscriber context. To create a policer control hierarchy on an ingress or egress SAP context, a policer-control-policy must be applied to the SAP. Once applied, the system will create a parent policer that is bandwidth limited by the policy's max-rate value under the root arbiter. The root arbiter in the policy also provides the information used to determine the various priority level discard-unfair and discard-all thresholds. Besides the root arbiter, the policy may also contain user defined tiered arbiters that provide arbitrary bandwidth control for subsets of child policers that are either directly or indirectly parented by the arbiter.

A Policer can be created under QoS Policies. When the QoS policy containing the policer with a parent mapping to an arbiter name exists is applied to a SAP, the system will scan the available arbiters on the SAP (though the assignment of Policer Control Policy). If an arbiter exists with the appropriate name, the policer to arbiter association is created. If the specified arbiter does not exist either because a policer-control-policy is not currently applied to the SAP or the arbiter name does not exist within the applied policy, the policer is placed in an 'orphan' state. Orphan policers operate as if they are not parented and are not subject to any bandwidth constraints other than their own PIR. When a policer enters the orphan state, it is flagged as operationally degraded due to the fact that it is not operating as intended and a trap is generated. When a policer-control-policy is added to the SAP or the existing policy is modified, the SAP's policer's parenting configurations must be reevaluated. If an orphan policer becomes parented, the degraded flag should be cleared and a resulting trap should be generated.

For subscribers, the policer control hierarchy is created through the policer-control-policy applied to the sub-profile used by the subscriber. A unique policer control hierarchy is created for each subscriber associated with the sub-profile. The QoS policy containing the policer with the parenting command comes into play through the subscriber sla-profile which references the QoS policy. The combining of the sub-profile and the sla-profile at the subscriber level provides the system with the proper information to create the policer control hierarchy instance for the subscriber.

Policer Control Policy

The feature introduces the new GUI menu item "*Policer Control Policy*" under the "Policy->QoS" menu. The menu item gives access to the Policer Control Policy manager (a browser frame). Here the operator can perform all the actions - provided by the SAM policy framework - such as creating, deleting, and distributing policies.

As this new policy is a new type of QoS policy, the privilege required to configure (create/delete/modified) this policy are as following (any one of them)

- Admin
- slaMgm

The Policer Control Policy contains following Parameters

- Name: The unique name for this policy
- Description: A description for this policy, max length is 80.
- RootMaximumRate: The total maximum bandwidth limit for this policer (used by the root arbiter and its children).

- **MinThreshSeperation:** the minimum separation between any in-use discard thresholds. When more than one child policer is associated with a parent policer priority level

Arbiter Entry

Under the Policer control Policy, user can add/configure the Arbiters for specific Pocket Policing control at different tier. This can be done in the “Arbiter” tab.

There are two types of Arbiters: Tier 1 and Tier 2.

Following parameters are configurable for the Arbiters:

- **Name:** The name that identify this arbiter.
- **Description:** A description with max length of 80.
- **Rate:** the maximum bandwidth limit of this policer control policy arbiter for the given tier.
- **Parent:** The parent arbiter. Must be root for tier 1 arbiter. For tier 2, can be either root or a tier 1 arbiter of the same policy.

The maximum number of arbiters per policy is 32.

Priority Level Entry

Within each Policer Control Policy, users can configure the MBS contribution for 8 (1-8) different priority levels. The configuration of the 8 priority level is done in the “Priority Level” tab within the Policer Control Policy configure form. For each level, one can configure:

- **Cumulative MBS contribution:** the maximum amount of cumulative buffer space (in bytes) allowed for this level by this policer.
- **Fixed MBS contribution:** Is the cumulative buffer space for this level fixed. When this is set to be true for a priority level within the policy, the system will treat the defined mbs-contribution value as an explicit definition of the priority level's MBS.

Shared Policer Output Queue

To support the new Hierarchical Policing on the IOM3, on the SR node, a new *policer-output-queues* is automatically created and applied on each IOM that supports ingress policing. In SAM, this shared policer output queue is modeled the same as the existing shared queue. Under exist menu Policies -> QoS->SROS QoS ->Shared Queue), when the shared queue policy manage frame is opened, there will be a new entry called “Policer Output Queues”. The new entry is same as the old “default” shared queue, except it only have 16 queues instead of 32.

Queue Group Template and Port Queue Group Enhancements

Queue Group Template Policy Port Queue Group were introduced in SAM 7.0R4.

Queue Mbs unit changes

The Maximum Burse Size field have its unit changed from kbs to bps in following objects:

- Ingress Queue Group Template -> Queue
- Egress Queue Group Template -> Queue

- Ingress Port Queue Group -> Egress Queue override
- Egress Port Queue Group -> Egress Queue override
- Network Port Queue Group (Network Port)-> Network Egress Queue Override
- Access Ingress Policy -> Queue
- Access Egress Policy-> Queue
- L2/L3 Access Interface -> Override -> Access Ingress Queue
- L2/L3 Access Interface -> Override -> Access Egress Queue
- SLA Profile -> Override -> Override -> Access Egress Queue
- SLA Profile -> Override -> Override -> Access Ingress Queue

Note:

- OSSI can access both new and old field but can not modify both fields at the same time.
- On the GUI the new field mbs in bytes/s will always be used for view/configure the mbs rate. The value will be translated to Kbytes/s on nodes where only Kbytes/s is supported.

Forwarding Class for Egress Queue Group Template

A new tab Forwarding Classes is added for the Egress Queue Group Template configure form. Under this tab, user can configure the mapping of forwarding classes to the queues configure for this Egress Queue Group Template (Under the “Queues” tab).

Each entry has two parameters:

- Forwarding Class: Must be one of h1, ef, h2, nc, be, l1, af, l2
- Queue Id: Must be a queue ID that is configured under the “Queues” tab for the this Template.

Default Egress Queue Group Template

In release 8.0, the system maintains a special default egress queue group template (policer-output-queues) that is automatically applied to all Ethernet ports.

SAM will also create this default egress Queue group Template policy (With name “policer-output-queue”). The default Policy will have two queues created:

- Queue 1: Best Effect
- Queue 2: Expedite

All other parameters are default. This default Egress Queue Group Template Policy can not be deleted.

Note: There will be a upgrade issue if this on the a pre 8.0 node have a Egress Queue Group Template configured with this name (“policer-output-queues”). It is very involving to remove this policy automatically from SAM (if it does exist in pre 8.0SAM). A release note have to be issued to make sure the custom will remove any Egress Queue Group Template that have a name “policer-output-queues” before upgrade to 8.0.

Egress Queue

For the Queues configured under the Egress Queue Group Template Policy. User can now define the PIR/CIR either through setting the specific value as before (i.e. PIR/CIR in kbps) or define the PIR/CIR as a percentage of the egress port's line rate.

Following additional parameters are added under the Queue Configure form "PIR/CIR" tab:

- RateType: can be set to "Specific" or "percentage" When set to percentage, the rate controlled by the percentage (relative to the port rate) defined in AdminPirPercentage and AdminCirPercentage.
- AdminPirPercentage: the percentage of PIR to the egress port's line rate.
- AdminCirPercentage: the percentage of CIR to the egress port's line rate.

Note: These to percentage field will not be applicable when the "Specific" option is selected for the RateType.

Port Access Egress Queue Group Configuration

The following new attributes are added to the Port Access Egress Queue Group to define a set of string values optionally used by subscriber management to map subscriber and subscriber host's policed traffic to a specific egress port queue group:

- DestinationString: The DestinationString parameter is used to specify the subscriber destination string that must match for a subscriber host to be associated with the egress queue group for forwarding classes mapped to an egress policer
- OrganizationString: The optional OrganizationString parameter is used to specify the subscriber host organization string that must match for a subscriber host to be associated with the egress queue group for forwarding classes mapped to an egress policer

Under the "Override Egress Queues" tab of Access Egress Queue form, following new parameters were added to give the option of override the PIR/CIR as a percentage of the egress port's line rate:

- RateType: can be set to "Specific" or "percentage" When set to percentage, the rate controlled by the percentage (relative to the port rate) defined in AdminPirPercentage and AdminCirPercentage.
- AdminPirPercentage: the percentage of PIR to the egress port's line rate.
- AdminCirPercentage: the percentage of CIR to the egress port's line rate.

Note: These to percentage field will not be applicable when the "Specific" option is selected for the RateType.

MIB table: TIMETRA-PORT-MIBtPortAccEgrQOverEntry

Access Ingress/Egress Policy Enhancements

Policer

Users can create policers for the Access Ingress/Egress Policy. Each Policer contains the following parameters:

- Description: A description string, max length 80
- PirAdaption: The adaptation rule to be used while computing the operational PIR value
- CirAdaption: The adaptation rule to be used while computing the operational CIR value

- Parent: The parent arbiter (root arbiter or tiered arbiter). See Section 4.1 for more details.
- Level: The level of priority (used by the parent policer)
- Weight: The weight used by the parent policer
- adminPir: Administrative PIR
- AdminCir: Administrative CIR
- StatMode: Decide the Counter allocation when generate stats. Options are: no-stats, minimal | , offered-profile-no-cir, offered-priority-no-cir, offered-profile-cir, offered-priority-cir, offered-total-cir, offered-limited-profile-cir. For the detail allocation rules for each option, Please see the node PRD, for more details
- MBS: Max amount of buffer space in bytes allocated by this spice.
- highPriorityRate: The percent of allocated buffer space that is used exclusively by high priority traffic.
- CBS: Reserved buffer space by this policer
- packetByteOffset: modify the size of each packet handled by the policer by adding (positive number or subtracting (negative number) a number of bytes. Range: [-32..... 31]

Forwarding Class

For each Ingress/Egress Forwarding Class Configuration, user can configure additional parameters to map packets that match the forwarding class to the specified policer-id.

- UnicastPolicerId: Policer to map to for normal traffic
- MultiCastPolicerId: Policer to map to for multicast traffic. Only applicable for FC in ingress policy.
- Broadcast PolicerId: Policer to map to for broadcast traffic. Only applicable for FC in ingress policy.
- Unknown PolicerId: policer to map to for unknown traffic. Only applicable for FC in ingress policy.

Note: the policer must exist within the corresponding Ingress/Egress QoS policy

Subscriber Profile Enhancement

Applying Policer Control Policy

The policer control policy can be applied to subscriber profile and when applied, each subscriber can configure some override.

A new tab called “Policer Control” will be added to the subscriber profile configure form:

User will be able to configure following attributes under this new tab:

- Ingress policer control policy: pointer to the Ingress Policer Control Policy.
- Egress policer control policy: pointer to the Egress Policer Control Policy.

Configure Policer Control Override

Two new sub-tabs will be added under the override tab of the subscriber profile configure form: “Ingress Policer Control” and “Egress Policer Control”

Under these two new tabs, user can add override for the ingress/egress policer control:
Following parameters can be overridden:

- MaxRate:
- MinThreshSeperation

Within the policer Control Override, one can override the MBS for each priority level. These include:

- Cumulative MBS contribution: the maximum amount of cumulative buffer space (in bytes) allowed for this level by this policer.

SLA Profile Enhancement

Policer Override

User can override the ingress/egress policer on per SLA profile base for detailed traffic control per subscriber. The policer override can be configured under the new sub tabs “Ingress Policer” and “Egress Policer” under the existing “override” tab of the SLA profile configure form.

The following parameters can be overridden:

- StatMode: Decide the Counter allocation when generate stats. Options are: no-stats, minimal | , offered-profile-no-cir, offered-priority-no-cir, offered-profile-cir, offered-priority-cir, offered-total-cir, offered-limited-profile-cir. For the detail allocation rules for each option, Please see node PRD, for more details.
- adminPir: Administrative PIR
- AdminCir: Administrative CIR
- MBS: Max amount of buffer space in bytes allocated by this spice.
- highPriorityRate: The percent of allocated buffer space that is used exclusively by high priority traffic.
- CBS: Reserved buffer space by this policer
- packetByteOffset: modify the size of each packet handled by the policer by adding (positive number or subtracting (negative number) a number of bytes. Range: [-32..... 31]

Service Access Point

A policer-control-policy must be applied to the ingress or egress contexts of a SAP in order for the system to create an instance of the policy’s arbiters and parent policer. The parent policer is used to manage the child policers on the SAP into a single aggregate metering rate. The arbiters are used to create subsets of child policers and create an arbitrary bandwidth control element that can limit the amount of bandwidth allowed to the subset of child policers and how that bandwidth is distributed amongst the children.

After the policer is applied, policer-control-overrides may be defined that allow for SAP level manipulation of the policy’s parameters.

Assign Policer Control Policy

New parameters for applying Policer Control are introduced (under the QoS tab) for following access interface:

- VPLS : L2 AccessInterface, B-L2 Access Interface I-L2 AccessInterface
- IES: L3 AccessInterface, ServiceAccessPoint
- VPRN: L3 AccessInterface, ServiceAccessPoint

These Parameters are:

- IngressPolicerControlPolicy: Pointer to the Ingress Policer Control Policy.
- EgressPolicerControlPolicy: Pointer to the Egress Policer Control Policy.

Policer Control Override (8.0R4 candidate)

Two new tabs called “Ingress Policer Control” and “Egress Policer Control” will be introduced under the existing “Override” tab under the above object configuration form to allow user to configure the SAP level policer control override.

Following parameters can be overridden for the Policer Control Override:

- MaxRate
- MinThreshSeparation:

For each Policer Control Override, user can also override the priority level configuration.

A sub tab “Level” will be available for override following parameters:

- Cumulative MBS contribution: the maximum amount of cumulative buffer space (in bytes) allowed for this level by this policer.

Policer Override (8.0 R4 candidate)

Two new tabs called “Ingress Policer” and “Egress Policer” will be introduced under the existing tab “Override” of the L2/L3 Access Interface configuration form to allow users to configure the SAP level policer override

The following parameters can be overridden for the Ingress/Egress Policer Control Override:

- StatMode: Decide the Counter allocation when generate stats. Options are: no-stats, minimal|, offered-profile-no-cir, offered-priority-no-cir, offered-profile-cir, offered-priority-cir, offered-total-cir, offered-limited-profile-cir. For the detail allocation rules for each option, Please see node PRD, for more details
- adminPir: Administrative PIR
- AdminCir: Administrative CIR
- MBS: Max amount of buffer space in bytes allocated by this spice.
- highPriorityRate: The percent of allocated buffer space that is used exclusively by high priority traffic.
- CBS: Reserved buffer space by this policer
- packetByteOffset: modify the size of each packet handled by the policer by adding (positive number or subtracting (negative number) a number of bytes. Range: [-32..... 31]

Down MEP support for IES & VPRN Subscriber Group Interface Ethernet SAP

5620 SAM 8.0 supports

- VPRN and IES support in Service creation for Global Maintenance Associations
- Down MEP creation/discovery for VPRN/IES Subscriber Group Interface

VPRN and IES support in Service creation for Global Maintenance Associations

In SAM 8.0, users are able to create Down MEPs on relevant Ethernet VPRN and IES subscriber group interface SAPs when creating a new Service binding in a Global Maintenance Association and the new “MEP(s) Creation on Subscriber Group Interface SAPs” checkbox is selected. New Down MEPs will only be created for Ethernet SAPs on IES/VPRN subscriber group interfaces. For non-Ethernet SAPs or SAPs not on IES/VPRN subscriber group interfaces, the current behavior will not be affected. The Down MEPs will then be visible in the local MEP tab.

Down MEP creation/discovery for VPRN/IES Subscriber Group Interface

Users are able to create and list MEPS from the following properties display forms:

- VPRN Subscriber Group Interface Ethernet SAP
- IES Subscriber Group Interface Ethernet SAP

The behavior matches existing MEP properties form

5620 SAM supports creation of MEPs on VPRN/IES Ethernet SAPs for the following nodes releases:

- 7750 8.0
- 7450 8.0
- 7710 8.0

FC Mapping based on Exp Bits at VLL/VPLS SAP

The FC Mapping based on EXP Bits at VLL/VPLS SAP Support is available for 7x50 Ethernet MDA/IOMs, using the access ingress qos policy. This feature sets the forwarding class or sub-class enqueueing priority when a packet is marked with a MPLS EXP bits specified.

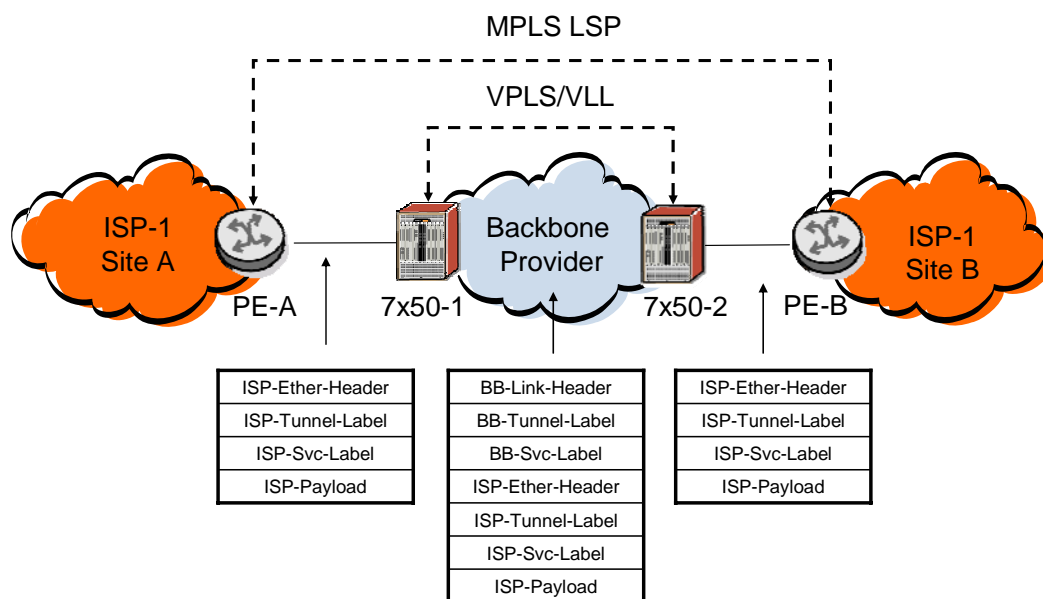


Figure 70: FC Mapping

Some backbone ISP want to provide VPLS/VLL to small ISPs as a site-to-site inter-connection service. Small ISP's router will connect to a 7x50's Ethernet L2 SAP, and the traffic will be encapsulated in a VLL/VPLS SDP. These small ISP's routers are PE routers typically. So, to provide appropriate QoS, 7x50s need to support a new classification option that based on received MPLS EXP bits.

This feature is supported as of 7750/7450/7710 8.0R1. It is supported on all Ethernet MDAs and all IOMs.

A new tab: `lsp-exp` will be supported on the access ingress qos policy. This tab is used for classification. Once the tab is selected, the "Search" and "Add" buttons will be enabled. The user can add `lsp-exp` rules. Adding a `lsp-exp` rule on the policy forces packets that match the MPLS LSP EXP specified to override the forwarding class and enqueueing priority based on the parameters included in the `lsp-exp` rule. When the forwarding class is not specified in the rule, a matching packet preserves (or inherits) the existing forwarding class derived from earlier matches in the classification hierarchy. When the enqueueing priority is not specified in the rule, a matching packet preserves (or inherits) the existing enqueueing priority derived from earlier matches in the classification hierarchy. When the user clicks the "Add" button, a new form appears with the following properties:

- *lsp-exp*: This value is a required parameter that specifies the unique MPLS LSP EXP value that will match the `lsp-exp` rule. A maximum of eight `lsp-exp` rules are allowed on a single policy. Valid values are from 0 to 7.

- *Forwarding Class*: The value must be one of the predefined forwarding classes in the SAM. Specifying the *Forwarding Class* is optional. When a packet matches the rule the forwarding class is only overridden when the *Forwarding Class* parameter is defined on the rule. If the packet matches and the *Forwarding Class* is not explicitly defined in the rule, the *Forwarding Class* is inherited based on previous rule matches. Valid values are: be, l2, af, l1, h2, ef, h1, nc.
- *Forwarding Sub Class*: The *Forwarding Sub Class* parameter is optional and used with the *Forwarding Class* parameter to define a preexisting sub-class. By Default *Forwarding Sub Class* is None and each *Forwarding Sub Class* must be explicitly defined. A *Forwarding Sub Class* can be 29 characters max.
- *Priority*: The *Priority* parameter is used to override the default enqueueing priority for all packets received on an access ingress SAP using this policy that match this rule. Specifying the priority is optional. When a packet matches the rule, the enqueueing priority is only overridden when the priority parameter is defined on the rule. If the packet matches and priority is not explicitly defined in the rule, the enqueueing priority is inherited based on previous rule matches. Valid values are: high, low and default (No override).
- *HSMD Counter Override*: The *HSMD Counter Override* is optional and only has significance on SAPs which are created on an HSMDA (a restriction not enforced by SAM). When specified, packets matching the MPLS EXP value will be mapped to the defined HSMDA exception counter-id for the packets queue group. The default behavior is to use the default counter on the queue group for the queue to which the packet is mapped. The specified *HSMD Counter Override* must be specified as an integer between 1 and 8.

Once the user adds a new rule and selects it, the “Properties” and “Delete” buttons will be enabled. This feature is applicable only to Ethernet L2 SAPs, but this restriction is not enforced by SAM.

PPPoE for Residential & Business Wholesale

The PPPoE for Residential and Business Wholesale provides the following main functions:

- Configuration and management of L2TP protocol and L2TP Site
- Configuration and management of L2TP Groups Configuration
- Configuration and management of L2TP Tunnel Configuration
- Management of L2TP tunnels and L2TP Tunnel Endpoints
- Management of L2TP Peers
- Management of L2TP Sessions
- Configuration and management LNS groups
- ESM enhancements for L2TP
- Performance statistics collection

L2TP it is a session layer protocol used to extend the PPP model by allowing Layer 2 and PPP endpoints to reside on different devices interconnected by a packet switched network.

For this reason L2TP has become a logical choice for wireline and wireless operators who wish to offer wholesale services to other service providers. Whereas the PPP session would normally exist between CPE (router or host) and the BRAS, L2TP extends the PPP sessions between the CPE and the PPP/L2TP termination point - known as the L2TP Network Server or LNS - via an intermediate L2TP Access Concentrator (LAC).

In release 8.0, to support L2TP feature, 5620 SAM allows configuration and management for the following new objects: L2TP Protocol Site, L2TP Groups, L2TP Tunnel, L2TP Peer, and L2TP Sessions. Enhancements to existing ESM functionality is required also to support this feature.

L2TP functionality is enabled by enabling the L2TP protocol on the Routing Instance. The Routing Instance is either the Base Routing Instance or a VPRN Router Instance Site.

Hardware Requirements

PPPoE for Residential and Business Wholesale feature is supported on release 7.0 and 8.0 on the following 7x50 platforms:

Platform	L2TP Support	Comments
7450 ESS12	No	Configured on “Mixed Mode”
7450 ESS7	No	Configured on “Mixed Mode”
7450 ESS6	No	Configured on “Mixed Mode”
7450 ESS1	No	Configured on “Mixed Mode”
7750 SR12C	Yes: LAC (??), Future:LNS	-
7750 SR12	Yes: LAC, LNS	LNS support only for release 8.0
7750 SR7	Yes: LAC, LNS	LNS support only for release 8.0
7750 SR1	Yes: LAC, LNS	LNS support only for release 8.0
7710 SR	Yes: LAC	-

Table 15: PPPoE for Residential & Business Wholesale

In order to enable L2TP protocol, the network element must be configured as chassis mode B,C or D. L2TP related functionality is supported only with Ethernet MDAs. IOM2 or IOM3 is required to configure Access Interfaces that are used for L2TP and IOM3 is required for network interfaces that are used for L2TP. IOM3-XP is required for configuration of the Broadband ISA MDA.

Typical L2TP Configuration

A typical L2TP configuration is applied to two network elements, of which one has LAC role and the other has LNS role. Only 7750 NEs release 8.0 can act as L2TP Network Servers. A network element can act as both LAC and LNS. L2TP protocol must be enabled for a Routing Instance.

At least an ISA-LNS Group must be configured for the NE with LNS role. For details on ISA-LNS Groups, see section 4.8.

L2TP Groups shall be configured on both NEs, but an LNS Group (with ‘Role’ attribute set to LNS) must be configured on one NE, and a LAC Group must be configured on the other NE.

L2TP Tunnel Configuration shall be created and applied to both NEs. A ‘Start’ operation shall be performed for the L2TP Tunnel configuration, on the NE that represents the LAC, in order to bring the tunnel to be operational. Alternatively, if no operational tunnel exists for the tunnel group, when an incoming L2TP session must be established on an L2TP Group, the tunnels configured for the group are automatically started.

If no group and tunnel configuration is present on the LAC NE, RADIUS authentication must return all attributes required for tunnel, when the session is authenticated.
A Tunnel Status shall be present on each NE representing the operational tunnel endpoint. In addition, 5620 SAM shall present the two tunnel statuses created on LNS NE and on LAC NE, as part of a global L2TP Tunnel object. L2TP Sessions shall be retrievable on LAC and LNS NEs, for the traffic generated by any PPPoE clients.

On an LNS NE, valid L2TP destinations, configured for L2TP tunnels, include: loopback interfaces in a VPRN service, loopback interfaces configured for the Base Routing Instance, or L3 Access Interfaces. The System Interface is not valid for L2TP Tunnel termination.

As existing functionality in 5620 SAM, PPPoE Sessions and hosts shall be retrievable, when created on LAC and LNS.

L2TP Protocol and L2TP Protocol Site

620 SAM 8.0 allows configuration of L2TP protocol, for a Routing Instance Site, on 7x50 nodes release 7.0 and 8.0. L2TP protocol can be enabled only for chassis mode C and D. The Routing Instance can either be the Base Router Routing Instance or a VPRN Router Instance Site.

When L2TP protocol is administratively enabled on a Routing Instance, an L2TP Protocol Site is automatically created for the Routing Instance. L2TP Protocol Site shall be listed on Manage Routing Instances Manager window. In addition to that, it is displayed on Routing Instances tree (for protocol sites of the base router) or on the VPRN Site configuration form, Protocols tab.

The following actions are supported for L2TP Protocol Site that is displayed under Routing Instances tree: Resync and Edit.

From the Groups Configuration tab, a 5620 SAM user shall create, delete, edit, list L2TP Group Configurations and perform operations on L2TP Groups.

From the Tunnels tab, a 5620 SAM user shall list and view L2TP Tunnel Endpoints and perform tunnel operations on selected tunnel.

From the Peers tab, a 5620 SAM user shall list, view L2TP Peers and perform operations on selected peer.

From the Sessions tab, a 5620 SAM user shall list L2TP Sessions, by retrieving them on-demand from the network and then view the L2TP Sessions.

L2TP Group Configuration

5620 SAM 8.0 allows configuration of L2TP Group Configurations on 7x50 NEs release 7.0 and 8.0. An L2TP Group Configuration represents the configuration for a group of L2TP Tunnels. Each tunnel can carry multiple L2TP sessions. If the tunnel carrying a session fails, another tunnel from the same tunnel group is used to carry the session.

A 5620 SAM user shall create, delete or edit an L2TP Group Configuration by selecting the Group Configuration tab on the L2TP Protocol Site configuration form or from the Manage Routing Instances tree.

L2TP Group Configuration form has the following tabs: General, Tunnel Configuration, Tunnels Status and Statistics. In addition to those tabs, LNS L2TP Group Configuration has the PPP tab.

A 5620 SAM user shall configurable PPP for an L2TP Group Configuration.

L2TP Group Operations

A 5620 SAM user shall be allowed to select an L2TP Group Configuration and perform a 'Drain' operation by clicking on the "Drain" button. A 'Stop Drain' operation is performed by selecting a L2TP Group Configuration and then clicking on "Stop Drain" button.

Drain operation forces all existing connections for the tunnels on selected L2TP tunnel group to time out. Until the operation is finished, no new connections are accepted.

On LNS L2TP Group Configurations a user is allowed to perform a 'Stop' operation, to stop the connection control on all tunnels from the group.

L2TP Tunnel Configuration Tab

A 5620 SAM User shall be allowed to configure multiple L2TP Tunnel Configurations for a L2TP Group Configuration. L2TP Group Configuration form has a tab named Tunnel Configuration, which allows creation, deletion, listing, editing and starting L2TP Tunnel Configuration.

A user is allowed to select a L2TP Tunnel Configuration that is listed on this tab and perform a "Start" operation by clicking on "Start" button.

Tunnels Status Tab

L2TP Group Configuration form shall have a tab, named Tunnels Status, to list operational information on all operational tunnels for the L2TP Group.

A user shall be allowed to select a L2TP Tunnel that is listed on this tab and perform a "Stop" operation by clicking on "Stop" button.

L2TP Tunnel Configuration

5620 SAM 8.0 allows configuration of L2TP Tunnel Configuration, on 7x50 nodes release 7.0 and 8.0.

An L2TP Tunnel is a connection that share common control channel, between a LAC and LNS. Within each L2TP Tunnel one or more L2TP Sessions exist. Each L2TP session transports PPP frames.

A 5620 SAM user can create an L2TP Tunnel Configuration, which is used as a template to create multiple operational tunnels. The endpoints of the operational tunnel are represented by a Tunnel Status object.

A user shall be allowed to create delete or edit L2TP Tunnel Configurations by selecting Tunnel Configuration tab on a L2TP Group Configuration form or by selecting an L2TP Group on the Manage Routing Instances tree and opening the drop-down menu.

The configuration form for an L2TP Tunnel Configuration, that is part of an L2TP Group configured for an LAC, has the following tabs: General, Tunnel Status.

In addition to those tabs, an L2TP Tunnel Configuration that is part of an LNS L2TP Group Configuration has the PPP tab with a number of attributes specific to LNS L2TP Tunnel Configurations.

L2TP Tunnel Status

A L2TP Tunnel Configuration does not have associated an operational status until the user performs a 'Start' operation for the tunnel or an incoming L2TP session establishes the tunnel. Depending on the tunnel configuration, multiple Tunnel Status instances can be associated with the tunnel configuration, since the tunnel configuration is used as a template for creation of multiple tunnels.

L2TP Tunnel Status information is available from the Tunnel Status tab on an L2TP Tunnel Configuration form.

L2TP Tunnel Operations

5620 SAM user shall be allowed to select a L2TP tunnel and perform the following operations:

- Start - operation attempts to start the control connection for the L2TP tunnel. Start operation is available only for non-operational tunnels.
- Stop - operation is available only for operational tunnels. Stop operation attempts to stop the control connection for the L2TP tunnel.
- Drain - operation forces all existing connections for the tunnel to time out. Until the operation is finished, no new connections are accepted. Drain operation is available only for operational tunnels.
- Stop Drain - Stop Drain operation is available only for active tunnels.

L2TP Tunnels

An Operational Tunnel object is a 5620 SAM representation of the two Tunnel Status objects that are created for the two endpoints of the L2TP tunnel.

An L2TP Tunnel Status is automatically created in the following scenarios:

- A “Start” operation is performed by the user on an existing L2TP Tunnel configuration
- An incoming L2TP session must be established using an existing group and tunnel configuration
- RADIUS authentication returns configuration for the tunnel when an incoming L2TP Session is authenticated

An L2TP Tunnel Configuration is configured as part of an L2TP Group Configuration on the LNS NE and LAC NE. After configuration, the L2TP Control messages must be started between the two endpoints in order to have the tunnel operational.

On the two NEs that are the endpoints of an operational tunnel, a Tunnel Status object is created that has the operational information about the tunnel. 5620 SAM correlates the information from the two endpoints and displays them as part of a Tunnel object. The properties form for Tunnel has two tabs: LNS Endpoint and LAC Endpoint to display the attributes of the two Tunnel Statuses and attributes of the Tunnel Configuration if tunnel was created from a Tunnel Configuration. The Tunnel Status attributes and Tunnel Configuration attributes are displayed under two sub-tabs.

Tunnels are listed from the ISA-LNS Manager window.

Purging L2TP Tunnels

Tunnel Status objects and Tunnel objects are not removed from 5620 SAM database when the tunnel is closed. When tunnel is closed its Operational Status attribute is set to ‘closed’. After a tunnel is closed, the associated Tunnel Status objects are stored for at least the interval specified by the attribute Destruct Time Out.

If one of the Tunnel Status objects is removed the Tunnel object is also deleted. The remaining Tunnel Status object is still listed on the Tunnel Status tab of the L2TP Protocol Site.

L2TP Peer

An L2TP Site has none or many L2TP Peers.

5620 SAM release 8.0 manages L2TP Peer information. L2TP Peer information shall be available from L2TP Protocol Site configuration form, Peers tab. On this tab shall be listed L2TP Peers for the L2TP Site. A user shall select a L2TP Peer from the list and open the properties form.

L2TP Peer Tunnels

5620 SAM shall list the L2TP Tunnels for a specific L2TP Peer, on the L2TP Peer properties form, on Tunnels tab.

By clicking on the 'Properties' button next to the Tunnel ID attribute, a user opens the properties form for the L2TP Tunnel Status object with assigned tunnel identifier, from the L2TP Peer.

L2TP Peer Operations

5620 SAM user shall be allowed to select an L2TP Peer and perform the following operations:

- Drain - by clicking on the "Drain" button. Drain operation forces a drain operation on all tunnels between current L2TP Site and a selected L2TP Peer.
- Stop Drain - by clicking on "Stop Drain" button.

L2TP Sessions

Due to the potential performance implications, L2TP Sessions are not fully managed by 5620 SAM 8.0. A 5620 SAM user shall be allowed to retrieve on-demand current L2TP Session information for a specific L2TP Session. Connection identifier for the L2TP Session must be known for the on-demand retrieval. L2TP Connection Identifier is an attribute of the PPPoE Session that is created on the LAC NE. L2TP Session retrieved on the LAC NE has as attribute Remote Connection Identifier, which is the local identifier for the L2TP Session on the LNS NE. The value of the Remote Connection Identifier can be used to retrieve on-demand information about the L2TP Session on the LNS NE.

L2TP Session Managed Routes

Due to the potential performance implications, L2TP Sessions Managed Routes are not managed by 5620 SAM 8.0.

L2TP ISA LNS Group

Broadband ISA is used for L2TP Network Server PPP Sessions termination. Sessions are distributed across a number of Broadband-ISA according to a load balancing algorithm.

ISA LNS Groups shall be configured for this reason. Each ISA LNS Group includes up to seven ISA Broadband MDAs. A maximum of four ISA-LNS Groups are allowed for one NE.

An ISA LNS Group is associated with specific L2TP inbound peers and groups and sessions traffic is automatically balanced across all available active ISA Broadband resource MDAs from the group.

ISA-LNS groups can be created, deleted or edited from the Manage Equipment tree window and from ISA-L2TP Manager window.

Alarms that are raised for ISA-LNS Groups are propagated to the Tunnel Status objects that are using the ISA-LNS Group.

A 5620 SAM user can configure a maximum of seven Group Members. For each Group Member, the user shall select a Broadband ISA MDA.

Broadband ISA MDAs can be configured only on IOM3-XP hardware on 7750 NEs.

Alarms rose for the ISA-LNS Groups or ISA-LNS Group Members are propagated on affected L2TP Tunnels.

ISA LNS Group Member Operations

5620 SAM user shall be allowed to select an ISA LNS Group Member and perform the following operations:

- Drain - by clicking on the “Drain” button. Drain operation forces a drain operation for all tunnels on selected ISA LNS MDA.
- Stop Drain - by clicking on “Stop Drain” button.

L2TP Manager

The ISA-L2TP Manager window is available from the Manage menu option and ISA sub menu option. The ISA-L2TP Manager is the main window for the management of ISA-LNS Groups related functions: listing, creation, deletion, drain operations on associated MDAs.

The ISA-L2TP Manager window allows management of L2TP Operational Tunnels. It provides listing and filtering functions for all operational tunnels managed by 5620 SAM.

A sub set of L2TP Operational Tunnels are listed on the L2TP Protocol Site, filtered for the routing instance of the L2TP protocol Site. In addition to the listing on the L2TP Protocol Site, the ISA-L2TP Manager offers a global view on all L2TP Operational Tunnels Managed by 5620 SAM.

Network Address Translation (NAT)

NAT is a popular tool for alleviating the IPv4 address exhaustion. It has become a standard, indispensable feature in routers for home and office Internet connections.

Network address translation involves re-writing the source and/or destination IP addresses and usually also the TCP/UDP port numbers of IP packets as they pass through the NAT. Checksums (both IP and TCP/UDP) must also be rewritten to take account of the changes.

The NAT feature on the SR provides the following capabilities:

- Ability to define a different NAT policy to different subscribers
- Support for external pools of up to /24 in prefix-length
- Ability to set per-subscriber maximums for number of flows
- Ability to set destination prefixes exempt from subscriber maximums. Exempt destinations use a reserved range of ports.
- Fixed port-range assignment - contiguous allocation to a single subscriber (ie, Sub1: 2000-2999, Sub2: 3000-3999, etc).
- Binding of NAT external interfaces to any VPRN/IES. The service with the external interfaces does not need to be the same service the ESM host is maintained in.
- Can operate in a traditional NAT mode (each host has a unique IPv4 address) or sub-aware NAT. IP addressing automatically preserved.
- IP destinations/ports may be excluded from the port limits by setting aside a reserved pool-space per subscriber.
- In traditional mode, unique source IP addresses (clients) may optionally cause the creation of a port-range reservation and log entry
- Special treatment for UDP-based DNS queries with shorter NAT timeouts

The Nat feature introduces the following new major objects:

Nat Policies, Nat Pools, L2 Aware Ip addresses, Nat Destinations, and ISA-NAT Groups.

Currently, only the 7750 supports NAT. 7710 will support NAT in the future.

The relation between these objects are as follows:

A subscriber profile can have a Nat Policy assigned to it. The Nat Policy references a NAT Pool. Nat configuration contains nat pools, L2 Aware Ip Addresses, and Nat Destinations. The Nat Pool has pointers to an ISA-NAT group, which contains the MDA's used by NAT.




Figure 71: Object Relations

NAT Policies

Each subscriber has a NAT-policy associated with it that indicates ALL IPv4 traffic will be passed to the NAT (note that the forwarding policy **MUST NOT** cause IPv6 traffic to pass through the NAT). The reason for having per-subscriber NAT policy instead of a route-lookup-to-NAT is because some subscribers are not expected to be subject to NAT.

NAT policies in SAM use the SAM policy framework. NAT policies are global policies in SAM. As with all policies, NAT policies are distributed to supported nodes and NAT policies are synchronized with a local policy on a node.

The NAT policy contains a pointer to the NAT pool to use with the policy. Other parameters in the policy are listed in the table below.

There is a new menu item on the policy menu: NAT Policies.

Nat Policy Display - David, Global Policy [Edit]

General Local Definitions Faults

Policy Configuration

Policy Scope: Global Policy

Configuration Mode: Draft

State: Initializing From: 10.101.185.135

Displayed Name: David

Description: N/A

NAT

NAT Pool

routerId: 1

Displayed Name: Roar

Filtering: Endpoint Independent

Port Reservation Count: 0

High Watermark: 0

Low Watermark: 0

Session

Session Limit: 65535

Reservation Count: 0

Session High Watermark: 0

Session Low Watermark: 0

Priority Session Forwarding Class Set: Default

To Tcp

to Tcp Establish: 7440

to Tcp Transfer: 240

to Tcp Syn: 15

to Tcp Time Wait: 0

To Ucp

to Udp: 300

to Udp Initial: 15

to Udp DNS: 15

To ICMP

to ICMP Query: 60

Table 16: Nat Policy parameters

Nat Configuration for IES and VPRN services

Nat configuration can be applied at a router level or for a specific VPRN service. Configuration applied at the router level applies to all IES services.

Configuration at the router level is added on the virtual router object.

Configuration applied at a VPRN service is added to the VPRN site for that service:

The Nat Configuration window allows the creation and addition of NAT pools, L2 Aware IPs, and NAT destinations.

NAT Pools

Network Address Translation translates outside IP Addresses and ports to another, recognized only by interior routers on a private network. In this way, NAT allows more devices to be connected to the network than the number of available external IP Addresses.

A NAT Pool contains the pool of external IP addresses and ports to be used for NAT. On the SR, a NAT Pool is associated with an ISA group.

The screenshot shows the 'Nat Pool Display - Roar [Edit]' window. The 'General' tab is selected. The 'Displayed Name' is 'Roar' and the 'Description' is 'N/A'. The 'NAT' section contains the following fields:

Field	Value
ISA Group	1
Nat Pool Type	l2Aware
Administrative State	Out Of Service
Port Reservation Type	blocks
Port Reservation Value	128
Port Reservation Allow Privileged	<input type="checkbox"/>
Block Usage	0
Block Usage High	<input type="checkbox"/>
High Watermark	0
Low Watermark	0

At the bottom of the window, there are buttons for 'Resync', 'OK', 'Cancel', and 'Apply'.

Table 17: Nat Pool Parameters

Nat Destinations

The Nat Destinations tab allows the user to configure NAT destination prefixes.

L2 Aware Ip Addresses

The L2 Aware IP tab allows the user to configure L2 Aware IP addresses for NAT.

ISA Groups and MDA's

A NAT ISA group is used to represent multiple hardware adaptors as a single entity, allowing for warm redundancy between multiple NAT ISA's.

MDAs are assigned to ISA Groups. An ISA group is associated with a NAT Pool, which contains the pool of outside IP addresses to be used with NAT.

ISA-NAT groups are in the equipment view. A new group can be created by right clicking and choosing "create."

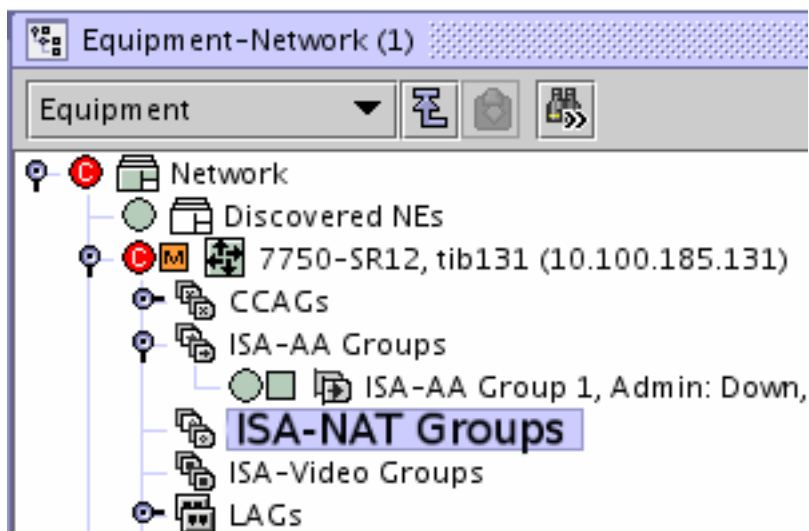


Figure 72: ISA-NAT Groups

Once an ISA-NAT group is defined, MDAs can be added to the group. MDAs can be added to the ISA-NAT group similar to the way MDAs are added to other ISA groups. (By right clicking on the group and choosing "create NAT-ISA group member".)

Assigning NAT Policies to Subscriber Policies

The way NAT is implemented on the SR requires that NAT Policies are assigned to subscriber profiles. The NAT Policy defines how NAT is used for that subscriber.

In SAM, the user will be able to assign a NAT policy to a subscriber.

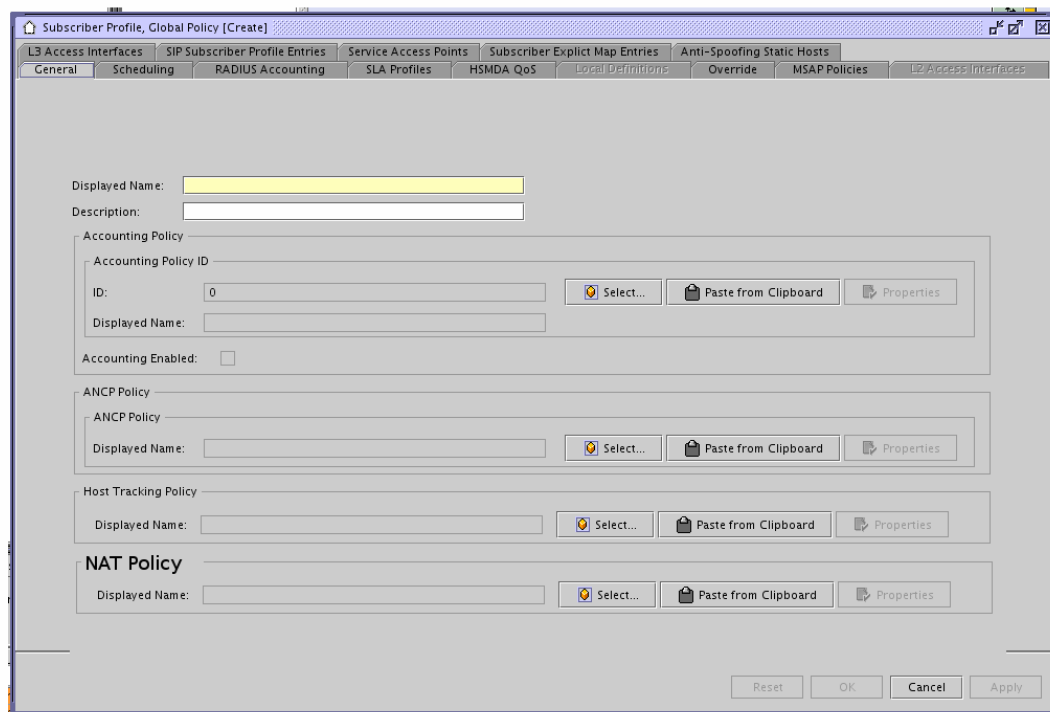


Figure 73: Subscriber Profiles

NAT and IP Filters

IP Filters will have a new action available: NAT. Filters are accessed from the policies/filter menu. In addition to the current actions: drop, forward, http redirect; the new action NAT will be available. When set to NAT, packets matching the filter entry are forwarded to the NAT function in the system.

Statistics

CPU and memory utilization statistics are available for the 7750. In addition, the following utilization statistics are available per ISA module:

ActiveInsideAddresses - Current count of Unique Inside Addresses / Subscribers (L2Aware) for a given NAT policy.

PeakInternalAddresses - High count of Unique Inside Addresses for a given NAT policy.

CreatedTCPSession - 32-bit counter incremented each time a TCP Session is created increment for a given NAT policy.

DestroyedTCPMappings - 32-bit counter incremented each time a TCP session is destroyed for a given NAT policy

CreatedUDPMappings - 32-bit counter incremented each time a UDP Session is created increment for a given NAT policy.

DestroyedUDPMappings - 32-bit counter incremented each time a UDP session is destroyed for a given NAT policy

CreatedQueryMappings - 32-bit counter incremented each time an ICMP query session is created increment for a given NAT policy

DestroyedQueryMappings - 32-bit counter incremented each time an ICMP query session is destroyed for a given NAT policy

On the SAM GUI, a statistics tab will be available on MDA objects. The stats will be displayed on this tab. The NAT Policy specific stats will be viewable on a per Policy basis.

Traps/Alarms

The NAT feature introduces the following new traps. An alarm will be raised when these traps are received.

tmnxNatIsaMdaSessionUsageHigh

"The *tmnxNatIsaMdaSessionUsageHigh* notification is sent when the session usage of a NAT ISA MDA reaches its high watermark ('true') or when it reaches its low watermark again ('false')."

An IsaMdaSessionUsageHigh alarm will be raised. This alarm will have a major severity and the alarm will be on the MDA, and also will be propagated to the NAT-ISA group, subscriber profiles, Nat Pools, and NAT Policies that are using that MDA.

tmnxNatPIBlockUsageHigh

"The *tmnxNatPIBlockUsageHigh* notification is sent when the block usage of a NAT address pool reaches its high watermark ('true') or when it reaches its low watermark again ('false')."

A NatPoolBlockUsageHigh alarm will be raised against the Nat Pool and propagated to any Nat Policies and Subscriber Profiles that use the Nat Pool.

tmnxNatPIBlockAllocationLsn

"The *tmnxNatPIBlockAllocationLsn* notification is sent when an outside IP address and a range of ports is allocated to a NAT host associated with a Large Scale NAT (LSN) pool, and when this allocation expires.

The allocated block is within the scope of the outside virtual router instance *tmnxNatNotifyVRtrID* and the outside IP address *tmnxNatNotifyAddr*; it starts with port *tmnxNatNotifyPort* and ends with port *tmnxNatNotifyPort2*.

The NAT host is identified with its inside virtual router instance *tmnxNatNotifyVRtrID2* and inside IP address *tmnxNatNotifyAddr2*.

When the block allocation is made, the value of the object *tmnxNatNotifyTruthValue* is 'true'; when the block allocation expires, it is 'false'."

A NatPoolBlockAllocation alarm will be raised against the Nat Pool. This alarm is informational only.

DOS Protection

In subscriber aggregation networks the 7750, 7450 and 7710 play an active role in several protocols. Subscribers either purposely or mistakenly may interfere with the operation of the node's processing capacity (for example excessive ARP handling) or other user's traffic.

Routing protocols such as OSPF and ISIS may also cause a threat as packets may be injected by customers (maliciously or by error) and may cause high overload of the CPM.

DoS protection is a significant point of concern for all service providers. This feature addresses DoS attacks when acting as a Subscriber Aggregation device and guarding against DOS attacks using unprovisioned protocols.

The following changes have been implemented with respect to Dos Protection policy in 5620 SAM 8.0.R1, in order to maintain compatibility with 7750, 7450, 7710 8.0.R1:

- Add a new "Out Profile Rate Limit" property to the global policy and to local policies on 8.0 NEs; this property corresponds to CLI command:

```
config
system
security
cpu-protection
[no] disable-protocol-protection
link-specific-rate rate
policy id-num
[no] per-source-rate per-source-rate
[no] out-profile-rate out-profile-rate
[no] overall-rate overall-rate
[no] alarm
```

- Create default global policies 254 and 255 on fresh install of or upgrade to 5620 SAM 8.0 R1 or later
- Global policies 254 and 255 cannot be deleted
- Local policies 254 and 255 cannot be deleted on 8.0 NEs; these policies can be deleted on pre-8.0 NEs
- Global policy 1 is modifiable
- Local policy 1 is modifiable/distributable or can be deleted on 8.0 NEs; this policy remains non-modifiable/non-distributable and can be deleted on pre-8.0 NEs
- Change the default Dos Protection policy from 1 to 254 (for Access) or 255 (for Network) on 8.0 NEs

6. TRIPLE PLAY MARKET

ESM Enhancements for L2TP

Enhancements are made to the following existing 5620 SAM objects: IES/VPN Group Interface, Subscriber Authentication Policy, RADIUS Accounting Policy, and Local User Database.

IES/VRPN Group Interface Configuration Enhancements

A 5620 SAM user shall configure an IES or VRPN group interface to be used to terminate L2TP Network Server PPP Sessions. IES and VRPN Group Interface configuration form, on General tab, has a new attribute.

An IES or VRPN Group Interface configuration form for Group Interfaces that are configured to terminate LNS PPP Sessions has a new tab (named LNS) where relevant attributes are available.

Local User Database Configuration Enhancements

5620 SAM 8.0 supports new configuration attributes on Local User Database, specific for L2TP.

Subscriber Hosts and PPPoE Sessions Enhancements

Subscriber Hosts information for those hosts that are present on LNS as the host associated with the L2TP Sessions, are retrievable on-demand as all existing types of subscriber hosts.

The type of the Origin attribute for the Subscriber Host is enhanced to indicate two new types: L2TP Session, IPCP Session.

L2TP Session type indicates that subscriber host is created as result of a PPPoE or L2TP (LNS) LAC session.

IPCP Session type indicates that subscriber host is created as result of a PPPoE or L2TP (LNS) IPCP session.

On the LNS NE, the subscriber hosts associated with the L2TP Session are created on an internal Service Access Point. The Service Access Point is created on a virtual port of the ISA-Broadband MDA. 5620 SAM does not manage Service Access Points with SAP Sub Type set to 'internal'.

PPPoE Session information for the PPPoE sessions that are tunneled to the LNS is retrievable on-demand as all existing types of PPPoE sessions.

The type of the Origin attribute for PPPoE sessions is enhanced to indicate one new type: L2TP. L2TP Type indicates that session is tunneled to an LNS NE via L2TP.

L2TP Performance Statistics

Performance statistics shall be collected for the following objects: L2TP Site, L2TP Group, L2TP Tunnel, L2TP Peer, and ISA-LNS MDA.

Database persistence for L2TP Sessions is not supported by 5620 SAM and then performance statistics collect

ARP Host

ARP host VPLS configuration

SAM supports the new dynamic ARP host type. The node is able to learn hosts of that type from the ARP requests exchanged (triggers). The characteristics of the ARP host type is similar to the DHCP host type in respect to scale and volatility.

On VPLS SAPs this functionality can be configured through the following properties:

- Administrative state: enabling or disabling the function
- Host limit: max. number of ARP hosts allowed on SAP

- Minimum authentication interval: specifies the minimum interval between two consecutive authentication attempts for the same ARP host

ARP host VPLS stats

For each VPLS SAP (where the ARP host function is enabled) SAM is able to collect statistics regarding the number of triggers and the number of created and deleted ARP hosts.

ARP host IES/VPRN Group interface configuration

Similar ARP host functionality as for a VPLS SAP is available on IES/VPRN Group interfaces. SAM will make this functionality configurable through the following properties:

- Administrative state: enabling or disabling the function
- Host limit: max. number of ARP hosts allowed on Group interface
- Host limit per SAP: max. number of ARP hosts allowed on a single Service Access Point of the Group interface
- Minimum authentication interval: specifies the minimum interval between two consecutive authentication attempts for the same ARP host

SAM will provide an operational value for the actual number of ARP hosts on the Group interface. The ARP host configuration is also available on subscriber interface on IES/VPRN if the interface is a retailer interface. The configurable properties are:

- Administrative state: enabling or disabling the function
- Host limit: max. number of ARP hosts

ARP hosts

As the ARP host is a new type of dynamic host, SAM will provide equivalent retrieval and display functionality as for DHCP based dynamic hosts. The host persistence (Fn19) will also include the ARP hosts.

ARP host managed routes

SAM will support the listing of managed routes for ARP hosts (IES/VPRN only). The routes are generated automatically by the system and are read-only.

IPSec VPNs

The IPSec application feature in SAM 8.0 supports

- IPSec Application Function
- IPSec Application Function Topology
- IPSec Session Management
- Alarm Management
- OAM
- Bidirectional Forward Detection (BFD) support

Possible IPSec Corporate Services that can be encountered in end-to-end scenarios are explained in the following subsections

In the Diagrams the following convention is used:

- Box represents a network element

- Circle represents a service site

Static IPsec Tunnel

In this scenario, one public L3 VPRN service is associated with its corresponding private VPRN service. The private VPRN service belongs to a larger private VPRN. The public service can be either a VPRN or an IES service. The private service can only be a VPRN service.




Figure 74: IPsec Static Tunnel

IPsec Tunnels on a Single Site

In the figure below, two public L3 VPRNs, VPRN 1 and VPRN 2 are connected to the private/secure service VPRN 3 through an IPsec gateway. The public services can be either VPRN or IES services. The private service can only be a VPRN service.

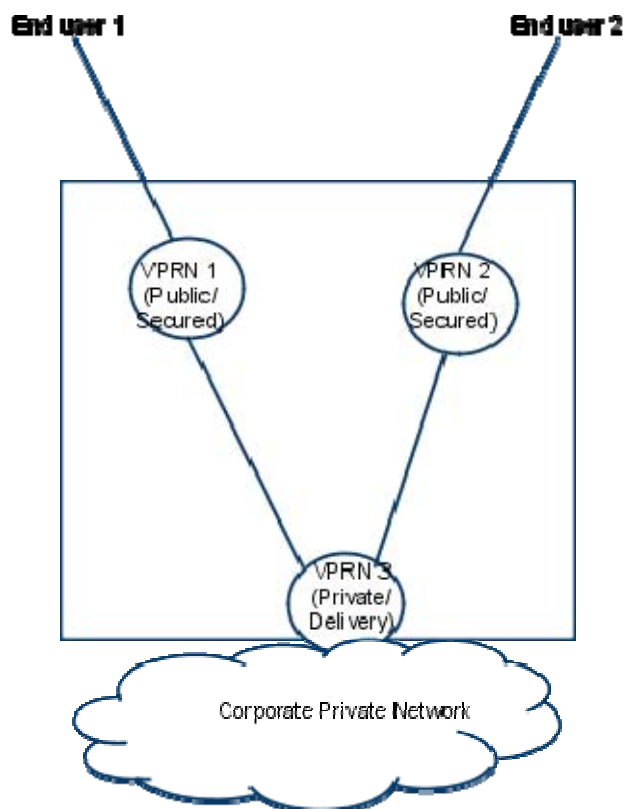


Figure 75: IPsec Tunnels on a Single Site

IPsec Tunnels on Multiple Sites

The following figure shows the scenario which is the extension of the scenario shown above, but across multiple sites. SiteA, SiteB and SiteC are part of the private service VPRN 4. On each of the sites, there is a secure IPsec tunnel between a public service and a private service. The public services can be either L3 VPRN or L3 IES services. The private service can only be a VPRN service




Figure 76: IPsec Tunnels on Multiple Sites

IPSec Application Function as part of the Corporate Network

In this application, the public IES service and the private VPRN service is connected to VPLS network through a CCAG or SCP. CCAG connect the private VPRN to the VPLS network. The scenario depicts the case where private VPRN is connected to the corporate network. In this example the VPLS is the corporate network.

IPSec Enhancements in SAM 8.0

IPSec Application Function

An IPSec infrastructure is the association between IPsec components (public and private services) to form a secured VPN. Together, this end-to-end solution is known as an IPSec Application Function (IAF). In SAM 7.0, only individual components of an IAF could be configured, but there was no end-end visibility or NMS specific value-adds (topology, OAM and trouble-shooting, impact analysis, correlation etc.) Also, in SAM 8.0 the configuration of the public and private service pairs for an IPSec VPN is much easier.

The configuration of the following components for the IPSec Application Function is the same as in SAM 7.0

- IPSec MDAs
- IPSec Groups
- IKE Policy
- IPSec Transform
- IPSec Tunnel Template

IPSec application function is launched from the main menu “Manage” → “IPSec Application Function”. This launches the IPSec Application Function Manager. The IPSec Application Function Manager allows management of secure IPSec Application Function instances (creation, listing, deletion etc.). The Topology view and the Topology flat view for each IPSec application function instance can be launched from the IPSec Application Function Manager.

The configuration of IPSec Application Function consists of the following steps:

Step 1: Establishing an instance of an IPSec Application Function. During creation, the IPSec Application Function ID, Name, and Description can be configured. “Aggregated Operational State”, and “Connection State” are applicable once the IPSec application function has been created. The properties for Step 1 are shown in table 1. The user also selects the “Corporate Service”. This shows how the private service is connected to the internal network. The Corporate service could be a VPLS, VLL, VLAN or a VPRN service. In most cases the Corporate Service and the private service are the same.

Step 2: Select or create the private service that will be used for the IPSec Application Function. A private service can only be a VPRN service. The steps to configure this include Select the service sites for the service. The user will have the option to either select the site that has connectivity to the “Corporate Service” as mentioned in Step 1 only if the service is different from the Corporate Service, or can select multiple sites from all the exiting sites available on that service. The user has to also select the following along with the site

- a. IPSec Interface on the private VPRN service

- b. Private IPsec sap for the terminating end point. The access point on the IPsec Interface.
- c. Outer encaps value for the private IPsec sap.

Step 3: Select or create a public service: The user is allowed to select one appropriate site from the list of sites that have been previously selected in Step 2. The public service is specific to the site selected. This restriction of the site selection to those selected in step 2 is based on the scenarios the IPsec Application Function will be deployed in the field. Step 3 establishes the relationship between the private service and the public services. On selection of the public service site, the following are to be configured

- a. L3 Interface on the public service
- b. Public IPsec sap for the terminating end point. The access point on the L3 Interface.
- c. Outer encaps value for the public IPsec sap.

The above scenario can be extended to show multiple public VPRN public services can be connected to different private services. For example, consider the configuration shown in Figure 3, and assume the Corporate Service is the private service.

There are 3 sites Site A, Site B and Site C.

- Site A includes : VPRN 4 and VPRN 3.
- Site B includes : VPRN 4 and IES 2.
- Site C includes : VPRN 4 and VPRN 1.

For the IPsec Application Function, we need to establish the following pairs.

- public service L3 VPRN 4 and private service VPRN 3 on Site A
- public service IES 2 and private service VPRN 3 on Site B
- public service L3 VPRN 1 and private service VPRN 3 on Site C

To create the IPsec Application function, the following are necessary for steps 2 and 3.

Step A: The user selects private VPRN 3 services on sites, A, B and C.

Note: In this step the same private service, VPRN 3 is selected on site A, site B, and site C.

Step B: In this Step, the user selects the following public services

- VPRN 3 on Site A
- IES 2 on site B
- VPRN 1 on site C.

Step 4: After selecting the private and the public services, the user configures the common parameters applicable to all services selected for the current IPsec application function instance of the IPsec Tunnel. The user will not have the option to configure each private service and public pair differently within an IPsec application function. If the user need different configuration for each pair, it will have to be done in a new IPsec Application function instance.

One of the properties that have to be selected for the IPsec tunnel is the tunnel type.

The tunnel type could be either Static IPsec Tunnel or Dynamic IPsec tunnel. For an IPsec Application, the tunnel type will be the same across all pairs of public and private VPRNs.

IPsec Application Function Discovery/Resync

On discovery, SAM will correlate the secure IPSec Application objects and create the IPSec Application function object.

Example: Consider the scenario in Figure 1, which shows public L3 services VPRN 1 and private service VPRN 2. Let's assume the order of discovery is VPRN 1, VPRN 2 and the tunnel type is Static

- When VPRN 1 is discovered, it is determined to be a private service or a public service. SAM has sufficient information looking at the L3 interface and the access point to determine it is a public service, but cannot associate it to a private service as VPRN 3 has not been discovered
- VPRN 2, a private service is discovered next. The information in the IPSec interface and the access point information specifies that it is a private service. When the IPSec static tunnels associated with the IPSec interfaces of VPRN 2 are discovered, the associated public services can be found if this corresponds to an existing IPSec Application Function object, then the existing object is updated, else a new object is created.

Consider the scenario where for the same example, the VPRN 2 is discovered first followed by VPRN 1.

- When VPRN 2 is discovered, SAM can figure out that it is a private service but cannot find the public service associated with this service. In this case the service is marked for post processing
- When VPRN 1 is discovered, the service does not have enough information to get the associated public service.
- On complete discovery of the node, the post processor selects VPRN 2 marked for post-processing. At this point it is able to associate the public L3 service VPRN 2. It checks to see if an IAF object exists and updates it, else it creates a new IAF object

This scenario described above is also applicable for dynamic tunnels.

IPSec Application Function Topology

The IPSec Application Function Topology view can be launched from “Manage” → “IPSec Application Function”. This launches the IPSec Application Manager. The IPSec Application Manager contains all the IPSec Application Functions that have been configured. Topology View or Topology Flat View can be launched by selecting an IPSec Application Function.

Topology View: This view shows the services involved in the IPSec Application Function. Figure 5 shows the topology view for the case where an IPSec Application Function has one public and one private VPRN. The topology view shows the IPSec private service(VPRN in this case, VPRN 21) and the IPSec public service(VPRN in this case, VPRN 8).The VPRN service properties form is launched by right-clicking the VPRN site icon.

If either the public or the private VPRN service icon is double-clicked, the detailed view of the VPRN site is shown along with the links and interfaces.

The relationship between a pair of public and private services can be viewed by double-clicking the link between the two VPRNs.

Figure 6 shows the topology view where multiple L3 public VPRNs are connected to a single private VPRN.




Figure 77: IPSec Topology View




Figure 78: Topology View of Multiple VPRNs connected to a single VPRN

- **Topology View Flat:** In this view both the IPsec L3 public service (VPRN 8) and IPsec private service (VPRN 21), along with the service sites, interfaces, and links between them are shown. The topology flat view shows the IPsec private and public services. In this example they are the VPRNs. Along with this, the L3 interface and the IPsec interface is shown.

The VPRN site properties form is shown by double-clicking the VPRN icon.

In the example below the properties form of the interface “L3-interface-1:ipsec-1:public:75” can be seen by double-clicking the interface icon or the link between the interface “L3-interface-1:ipsec-1:public:75” and “VPRN 8:172.22.24.118”.

The relationship between the private and public service, which includes the IPsec tunnel information can be viewed by double-clicking the link between the “VPRN 8:172.22.24.118” and “IPsec-interface-1:ipsec-1:private:55”.

The link between the VPRN’s is the IPsec tunnel. On double clicking the link the IPsec Tunnel configuration is shown. The status of the links will be the operational state of the IPsec tunnel.




Figure 79: IPsec Topology Flat View

Figure 80 shows the Topology flat view with multiple L3 VPRN’s being connected to a single private VPRN service.

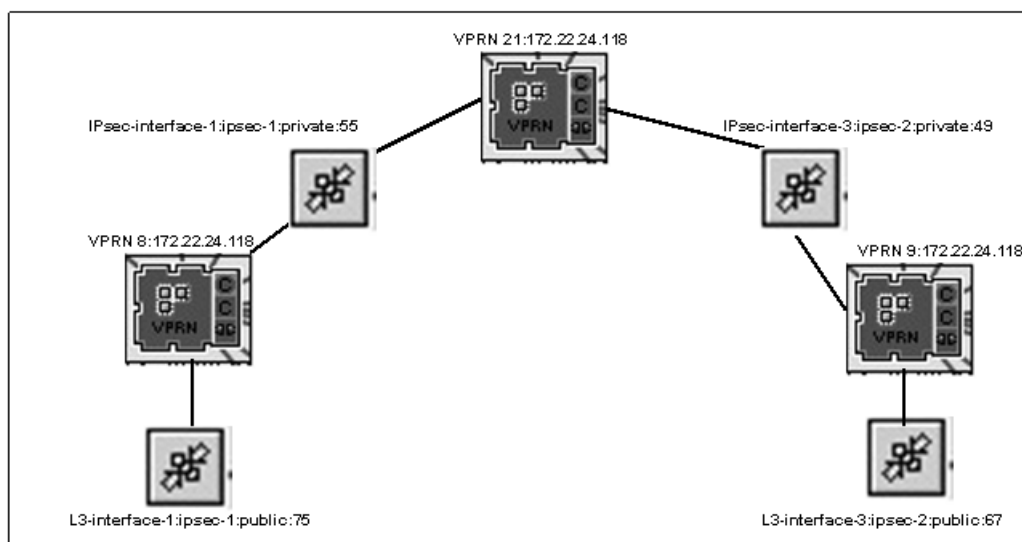


Figure 80: Multiple public VPRNs connected to single private VPRN

Alarm Correlation and OAM

IPSec dynamic tunnels and static tunnels state are propagated to all the IPSec objects containing IPSec tunnels.

For a private service, only static tunnels are applicable. An additional “State Cause” flag will be added to reflect the IPSec tunnel state on the following objects that contain the IPSec tunnel

- IPSec Interface
- VPRN Site
- VPRN Service

For Dynamic IPSec tunnels, an additional “State Cause” flag is added to reflect the IPSec tunnel state on the following objects that contain the IPSec tunnel.

- L3 Interface
- VPRN/IES Site
- VPRN/IES Service

In case of IAF the following states will be shown.

- Aggregated Operational State:
- Connection State:
- Operational Flag: A new flag “IPSec Tunnel State” will be added.

- Number of Services:
- Number of Connectors:

OAM will be supported using the STM (Service Test Manager) in SAM, in a similar fashion as other service tests. The purpose of these tests is to test the connection/communication between the public service site and the private service site. These tests are NOT designed to test the services contained in the public and private sites. The testing of services contained in the sites are provided in other OAM tests of the respective service site.

Below is list of tests that will be performed:

- Service Site ping
- ICMP Ping
- Trace route
- Additional tests such as L3 Service and Service Transport may be supported.

Service Test Management (STM)

Test Suites

SAM provides the ability to create a Test Suite for an IAF. The tests/policies added to the IPSec Application Function Test Suite will be applied/generated for all the public-private service pairs contained in the IAF, where applicable.

To create a Test Suite for an IAF, the user would perform the following steps:

- Open the Service Test Manager form
- Select Create
- Select Create TestSuite
- From the Entity drop-down, the user will have to option to select IPSec Application Function
- Generation of tests
- Execution of tests

Test Scheduling:

Test scheduling is supported in a similar manner as existing applications that use Test scheduling. A user is able to schedule a test by selecting the 'Schedule' button from the Test Suite form.

Test Policy

Where applicable, Test Policy is supported in a similar manner as existing applications that use Test policy.

IPSec Session Management

IPSec Session Management will support IPSec dynamic tunnels and IPSec Static Tunnels.

The information related to the remote user for dynamic tunnels and IPSec static tunnels for static sessions along with timestamps related to creation, deletion and the status is stored. For dynamic sessions, there is no guarantee that the user will use the same session.

Name	Type	Comments
Session Object	IPSec Static Tunnel or Remote User Tunnel	Dynamic Tunnel or Static Tunnel objects
Creation Time	Timestamp	Time the session was established
Last Changed	Timestamp	The time the session was last modified
State	Active/Inactive	Active if the IPSec session is established, and Inactive if the session has been deleted

Table 18: IPSec Session

IPSec sessions can be volatile, where the connections could be disconnected and connected repeatedly. This can result in creation and deletion of the IPSec sessions on SAM, which will degrade the performance in SAM. This is because SAM would have to process all the events related to creation and deletion of IPSec sessions. In order to address this volatility, the IPSec sessions are not deleted from SAM when the remote user disconnects the IPSec session. It is flagged as “Inactive”. Establishment of a session results in creation of an IPSec session if one does not exist. If an IPSec session exists, the state of the session is changed to “Active”.

SAM purges inactive records at a given threshold.

For better performance, the IPSec session objects are ONLY resynchronized once every twenty four hours, with no trap support. In other words IPSec sessions created between two consecutive twenty four hour resync periods will not be in SAM. Resyncs related to trap handlings are not supported. An option to allow manual resyncs is provided to make SAM database to reflect network state.

Size constraint is set to ensure that the record size does not exceed a particular threshold. If the threshold size is reached, older inactive sessions are deleted.

Bidirectional Forward Detection (BFD)

Bidirectional Forward Detection (BFD) is introduced in 8.0. BFD is used for IPsec tunnel failure detection.

Note: only BFD over static lan-to-lan tunnel is supported in 8.0.

To run BFD over IPsec tunnel like following, the following can be configured:

- Create a loopback interface in private VPRN service.
- Multiple multi-hop BFD sessions can be created on this interface; each session can be associated to one or more tunnels.(up to 500 tunnels)
- If one BFD session goes down, then associated tunnel will be torn down

Below are some configuration details for implementing BFD over IPsec tunnel

- Each tunnel can be associated to only one BFD session, however multiple tunnels can be associated to same BFD session.
- Only one BFD session is allowed between a given source/destination address pair.
- The BFD session should work in asynchronous mode.

- A BFD session can be associated to one or more tunnels but a tunnel can only be part of one BFD session.
- If one BFD session is associated to multiple tunnels, then the tunnel carrying the BFD traffic need to be brought up before any other tunnels.
- When system does NOT receive BFD packet from peer before detection time expires, or got signal down from remote peer, then the BFD session is considered down. System will also bring down the associated IPsec tunnels by:
 - Send delete payload message for all SAs of associated tunnels to remote peers
 - Remove all states or table entries of corresponding tunnels and SAs.

To configure the BFD:

1. On the private side, a user would designate a tunnel to be the BFD tunnel and enable BFD on that tunnel.
 - a. `config>service>vpn>ipsec-if>sap>tun>`
`[no]bfd-enable [service-id] interface interface-name dst-ip ip-addr`
`[no]bfd-designated`
2. On the public side, the user would configure the transmit interval
 - a. `config>service>vpn>if>`
`bfd transmit-interval [receive receive-interval] [multiplier multiplier] [echo-receive`
`echointerval]`

In SAM, the BFD information will be displayed on a 'BFD' tab in the IAF Display Manager.

Include DHCP-Options into RADIUS Authentication-Request

For RADIUS authentication, options in the DHCP-DISCOVER packets are used to identify a user. It shall be possible to insert all DHCP Options into a maximum of 2 RADIUS VSA attributes as part of the Authentication-Request message. Similarly, in the opposite direction, DHCP Options received in a RADIUS Authentication Response shall be extracted from the RADIUS VSA and mapped into DHCP Packets and forwarded to the client in the form of DHCP-Offer or DHCP-ACK packets.

In SAM Release 8.0, the ability to include the DHCP Options in RADIUS VSAs of Authentication Requests and Responses shall be supported as follows:

On the Subscriber Authentication Policy form, an extra checkbox labeled 'DCHP Options shall be added to the RADIUS Attributes.

7. APPLICATION ASSURANCE

ISA-AA Group CFLOWD implementation

Support for the cflowd implementation has been introduced in SAM 8.0.

Customers are looking to use off-board systems to use cflowd/Net-flow data to monitor and detect security threats. By monitoring sudden changes or deviations from expected trends, network attacks can be highlighted to the network operation groups or automated responses can be put in place to stop or mitigate such attacks. A CFLOWD Net-Flow record can contain a wide variety of information about the traffic in a given flow.

Although CFLOWD was supported on Node in SAM pre 8.0 release, in SAM 8.0 release we are supporting the implementation of CFLOWD with all the enhanced functionality provided by Node. In Application-Assurance policy CFLOWD can be attached.

Support for AA Group CFLOWD is introduced on Node in SAM 8.0 release. In SAM 8.0 release we are supporting the implementation of Volume and TCP performance parameters for AA Group CFLOWD. With Application Group / Application for an AA Group Policy is also enhanced in SAM 8.0 to support enabling of CFLOWD TCP Performance. It will be just an association with AA Group Policy.

The following properties are supported as a part of AA Group CFLOWD Volume and TCP performance reporting properties.

- Volume rate
- Volume admin status
- Performance sample rate
- Performance admin status
- CFLOWD admin status
- template-retransmit

As a part of Collector, the following properties are supported.

- Description
- Collector admin status
- Version (always it will be 10)

To create a Collector <ip-addr>:<port> is mandatory.

5620 SAM Release 8.0 supports the 7x50 features for CFLOWD implementation. It is enabled/disabled under the Network Element tab. Once the feature is enabled then multiple collectors for CFLOWD can be created.

If CFLOWD is disabled, then existing collectors are automatically deleted. Collectors are shown in another tab. Multiple collectors can be created under the CFLOWD collector tab.

It is possible to collect the statistics when CFLOWD is enabled for any Network Element. Currently version 5, 8 and 9 are supported for creating Collectors. So, if CFLOWD is enabled then all the statistics for packet sent, packet receive and packet errors are collected. If collectors are created, then statistics for respective collectors can also be collected.

5620 SAM Release 8.0 supports features for AA Group CFLOWD implementation in 7750, 7710, 7450 Mixed Mode. The volume and TCP performance related parameters can be configured as a part of AA Group CFLOWD and displayed in CFLOWD tab under the ISA-AA Group Display tab. We can configure collector objects which will be a combination of specific ip-address:port. These collector objects are shown in Collector tab which will be the child tab under CFLOWD.

Also the within an application group and applications it will possible to enable CFLOWD TCP performance. It is possible to enable CFLOWD TCP performance to multiple application group / application within a same AA Group Policy. These enabled application group and application will be shown in App-Grps and Application tab under CFLOWD tab.

To create AA Group CFLOWD, it is necessary to have an ISA-AA group should be configured on the node. For each ISA-AA group, one AA Group CFLOWD can be configured.

Statistics can be collected for Collectors and also for TCP performance. If there are collectors created for any ISA-AA Group CFLOWD then for all the collectors statistics data can be collected and plotted with existing plotter framework. If application and application group are associated for TCP performance enabling, the statistics data for TCP performance can be collected. The same can be plotted with plotter framework also.

Application Assurance Policy Enhancements

The Application Assurance Policy Enhancements feature is supported on all 7x50 platforms that provide Application Assurance capability. Thus Application Assurance Policy Enhancements feature is supported on the following 7x50 platforms:

System	Application Assurance Policy support
7450 ESS12	Yes
7450 ESS7	Yes
7450 ESS6	Yes
7450 ESS1	No
7750 SR12	Yes
7750 SR7	Yes
7750 SR1	No
7710 SR	No

Table 19: Application Assurance Policy Support

New Manage -> Application Assurance Window

A new Manage -> Application Assurance window is supported to list AA Groups, AA Partitions and AA Protocols created from SAM under entire network. The filter functionality is also provided to let user filter the AA Groups, AA partitions and AA Protocols as he/she wishes.

Any AA related objects which are not policy are added in this new window versus the exiting Policy -> Application Assurance window.

Multiple ISA-AA Groups and Partitions Support

The number of active ISA-AA groups is changed from 1 (pre R8.0) to 7 (R8.0) to allow AA resource partitioning/reservation for different types of AA service. The range of ISA-AA Group ID is 1 - 256.

Up to 128 partitions are allowed to be defined per AA Group for AA deployments with per VPN customization. The range of ISA-AA Group Partition ID is 1-65535.

A partition provides existing AA Group configuration but in a context of that partition only. This means application ID, policy and statistics configuration applies only to the given partition.

ISA-AA Group Configuration Enhancements

Under ISA-AA Groups tree node, SAM shall allow creation of more than 1 group up to 7.
For each ISA-AA group configuration, the new attribute called “Partitions” shall be added to ISA-AA group properties form.

Modification of “Partitions” attribute from ‘Enabled’ to ‘Disabled’ will be blocked if there are AA group/partition policies on the ISA-AA group. Pre-manual deletions required for AA group/partition policies and all their dependencies such as the application profile assignment on all SAPs.
This rule is also true when changing “Partitions” from ‘Disabled’ to ‘Enabled’.

ISA-AA Group Partition Configuration

The ISA-AA partition is configured under ISA-AA group property form.
The new tab “Partitions” is added to list and create ISA-AA group partition objects.

The bulk creation of ISA-AA partitions is provided by allowing the user to specify multiple parents and create ISA-AA partitions under different NEs and groups.

From new Manage -> Application Assurance window introduced, Select ISA-AA Group class, using filter to filter out the ISA-AA groups, then multi select ISA-AA group objects, press ‘Create Partitions’ button, the ISA-AA Partition configuration form opens.

Once the partition configuration done, Apply or OK will create AA partition objects under multi selected group objects.

The ISA-AA Group Partition contains the following configuration parameters:

In this window, the user is also able to list ISA-AA groups and partitions created under entire network. Filter also provided for filtering the groups as the user wishes. In this way, the user has very clear view of the ISA-AA groups under NEs.

AA Group Policy Management Enhancements

As the result of ISA-AA partition creation, the default AA group/partition policy will be created under this partition. The behavior is the same as the pre 8.0 release, that the AA group policy is auto created when the ISA-AA group is created.

The default AA group/partition policy is deleted upon the deletion of the ISA-AA group or partition. AA group/partition policy and its children policies shall have a new attribute called “ISA-AA Group Partition ID”. This attribute allows the association with an existing ISA-AA group partition on which the policy it’s children are applied.

AA Account Policy Enhancements in Release 8.0

The same as AA Group Policy, The same as AA Group Policy, new attributes “Group ID” and “Partition ID” shall be added to AA Accounting Policy to specify the stats collection is under this partition context. As the result of ISA-AA partition creation, the default AA accounting will be created under this partition. The behavior is different from the pre 8.0 releases, in that the AA accounting policy is auto created when the ISA-AA group/partition is created. The attribute Policy ID shall be used to associate the AA Accounting Policy with an existing ISA-AA group partition on which the policy and its children are applied. In other words, the attribute Policy ID not only identifies the AA Accounting Policy, but also

identifies the group/partition ID to which this policy applies. The default AA accounting policy is deleted upon the deletion of the ISA-AA group or partition.

AA Statistics/Debug Enhancements

Any statistics related objects such as AA SAP Summary, AA Subscriber Summary, AA Special Study SAP, and AA Special Study Subscribers currently under ISA-AA group property form are added to ISA-AA partition property form if the ISA-AA group is partition enabled.

The Statistics tabs are removed from AA SAP Summary, AA Subscriber Summary forms. Instead, AA SAP Summary Stats are displayed on the SAP and AA Subscriber Summary Stats shall be displayed on the Residential Subscriber Instance.

AA Protocol Stats are displayed under the isa-aa group/partition Statistics tab, instead of AA Protocol properties form.

MDA Slot Number assigned to an AA Subscriber/SAP is displayed on the AA Sub/SAP Summary form. In 8.0, MDA Slot Number assigned to an AA SUB/SAP may indicate Unassigned. This is a new status for 8.0 related to load balancing on a reboot. It is possible that aa-sub in the config file may not be able to fit after reboot, and are left in an UnAssigned state.

Group Id/ Partition Id fields are added to all AA Accounting Statistics records.

AA Policer Enhancements

A new attribute "ISA-AA Group ID" shall be added to AA Policer form as the ISA-AA group now supports up to 7 groups. Distribution of AA Policer will fail if the ISA-AA group ID it refers to does not exist on NE.

Object References and Filters to AA Group/Partition Policy Enhancements

Any existing pointer references to AA Group Policy and its children objects shall be modified to add partition context. Any filter which has static AA Group ID to 1 shall be modified to a specific group ID and partition ID.

Protocol Shutdown for New Signature Upgrade Support

This feature is to provide signature upgrades without a need for automatic policy change. All protocols introduced in R1 of a given release, are designated "Parent" signatures for a given release and cannot be disabled. All protocols introduced in post-R1 of a major release as part of any isa-aa.tim ISSU upgrade must by default be "shutdown". Operators must be allowed to enable and shutdown any post-R1 signatures introduced in R1 of a given major release. Enabling/disabling of a new protocol takes affect for new flows only.

A user with Application Assurance Management role is able to edit an Application Assurance Protocol from Application Assurance management window.
The bulk update of Application Assurance Protocols are allowed through the SAM Bulk Operations tool.

Application Assurance Protocols are moved to new Manager->Application Assurance window. Select Application Assurance Protocol class, new button called "Update Protocols" should be enabled, press it, the generic bulk operation window should be opened, from there, the user is able to bulk update the

protocols for shutdown or turn up the protocols. Application Assurance Protocol Stats are relocated under the Statistics tab on the ISA-AA Group/ Partition properties form.

Custom Pattern Based Protocol Support

This functionality is for operators to identify operator-specific or end-customer-specific (applicable to per VPN business deployment) custom-built applications that cannot be uniquely identified using ALU-provided global scope protocol signatures. Delivering on-demand signatures does not scale and may often be not acceptable (as a signature makes an application visible to other operators). Custom protocols and ALU-provided protocol are functionally equivalent.

Custom protocol is used in Application Filter in the same way as ALU-provided protocol in SAM.

Application Filter Expression Match Extensions

This feature is to provide greater flexibility in application definition.

The maximum Expression Index is increased from 3 to 4.

There are 4 new types for the Expression:

“SIP Media Type”, “Citrix Application”, “HTTP User Agent” and “H 323 Product ID”.

Capacity Based Load Balancing

Capacity-Cost based load balancing allows a cost to be assigned to diverted SAPs (via the app-profile) and this is then used for load-balancing SAPs between ISAs as well as for a threshold that notifies the user if/when capacity planning has been exceeded.

The load balancing decision is made based on the AA capacity cost of aa-sub SAPs and ESM subs. The capacity cost is configured against the app-profile. When assigning a new diverted aa-sub to an ISA, the ISA with the lowest cost (that also has sufficient resources) is chosen.

Attributes Load Balance Status, Number of Unassigned ESM Subscribers, Number of Unassigned SAP Subscribers and Number of Unassigned Spoke SDP Subscribers are not auto-updated. They are updated only on an ISA-AA Group resync or a full node resync.

Spoke-SDP aa-sub

Node Rel 8.0 adds the ability to divert spoke-SDPs to the AA-ISAs. App-profiles must be assignable to a spoke-SDP, for divert from these service types: VPLS, IES, VPRN, Epipe.

In SAM, Application Profile configuration parameter is added to a Spoke SDP Binding that is associated with a L3 Access Interface on IES and VPRN, and to a Spoke SDP Binding on VPLS and Epipe.

Spoke SDP AA Performance and Real-time Stats

SAM supports AA performance and real-time stats collection on the diverted Spoke SDP.

These are supported by logToFile for collection by RAM.

ISA-AA only Upgrades

In pre-8.0 SRs, ISA-AA Only Upgrades were tied to the CPM ISSU upgrades. SR 8.0 onwards, ISA-AA only upgrades are independent of CPM ISSU restrictions, which means that ISA-AA upgrades shall be allowed from any load to any other load, starting with R1, within the major release.

8. LTE

Note: See the 5620 SAM LTE Release Description for further details.

The 5620 SAM offers great ease and accuracy in managing the SR family of nodes (including SAR, SAS, and ESS nodes). With Alcatel-Lucent entering the 4G mobility market (through LTE) 5620 SAM supports this new facet of the 7750.

In its new form, the 7750 acts as a SGW or a PGW depending on the application given to the Groups managing the ISMMG Cards that have been inserted. The 5620 SAM will detect the mobility dedicated 7750s and will discover them as being a variant of the 7750 SR (this is based on the SysObjectID that is different on an SRMG node when compared to an SR node).

5620 SAM support includes the introduction for a new type of Service called the Mobile Service on which an operator can run diagnostics using existing 5620 SAM tools, the management of User Bearers and the EPS path discovery and the ability of relate to the underlying connectivity (routed network, physical links, ports) to offer a better view of the managed network and to quickly analyse the impact of certain events in the network (ie : port or card down, link broken, routing protocol errors, etc). Lastly, 5620 SAM will also offer the 5620 SAM operator a tool that will allow creation and AGW (SGW and PGW) in 10 mouse clicks; the AGW Creator Facilitator will guide the operators through every step of the way.

Key Elements

The following are represented in the 3GPP Architecture example shown below.

SGW: It routes and forwards user data packets, while also acting as the mobility anchor for the user plane during intereNodeB hand overs and as the anchor for mobility between LTE and other 3GPP technologies (terminating S4 (not shown in the illustration above) interface and relaying the traffic between 2G/3G systems and PGW). For idle state UEs, the SGW terminates the DL data path and triggers paging when DL data arrives for the UE. It manages and stores UE contexts, e.g. parameters of the IP bearer service, network internal routing information. It also performs replication of the user traffic in case of lawful interception.

PGW: This gateway provides connectivity from the UE to external packet data networks by being the point of exit and entry of traffic for the UE. A UE may have simultaneous connectivity with more than one PGW for accessing multiple PDNs. The PGW performs policy enforcement, packet filtering for each user, charging support, lawful Interception and packet screening. Another key role of the PGW is to act as the anchor for mobility between 3GPP and non3GPP technologies such as WiMAX and 3GPP2 (CDMA 1X and EvDO).

PCRF: The Policy and Charging Rule Function component enables operators to have rules based, real time dynamic control over bandwidth, charging and usage.

HSS: The Home Subscriber Server is an integrated network for telecommunications carriers that uses the IP protocol as its foundation for packetized voice, video and data.

MME: The Mobility Management Entity is the key control node for the LTE access network. It is responsible for idle mode UE tracking and paging procedure including retransmissions. It is involved in the bearer activation/deactivation process and is also responsible for choosing the SGW for a UE at the initial attach and at time of intraLTE

Hand-over involving Core Network node relocation. It is responsible for authenticating the user (by interacting with the HSS).

SGSN: The Serving GPRS Support Node (SGSN) is responsible for the delivery of data packets from and to the mobile stations within its geographical service area. Its tasks include packet routing and transfer, mobility management (attach/detach and location management), logical link management, and authentication and charging functions.

eNB: The so called “Evolved NodeB (eNodeB)” are new enhanced base stations as per 3GPP standards. This enhanced BTS provides the LTE air interface and performs radio resource management for the evolved access system.

The figure below shows a series of Control (EPS) Bearers that must be present to allow communication between different components of the LTE network.

S1u: Between SGW and eNB (GTP based)

S5 : SGW and PGW using GTPC, GTPU, or PMIPv6 (can be referred as S8)

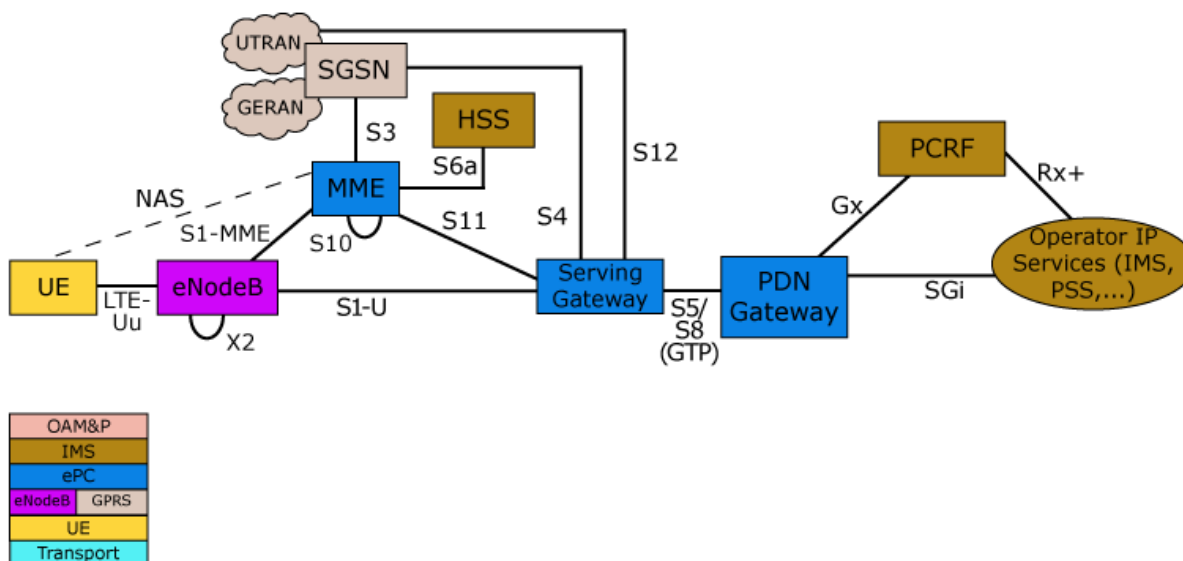
S11 : SGW and MME using GTPC

S12 : SGW and UTRAN

Gxc: SGW and PCRF

Gx: PGW and PCRF

S2a: PGW



Example of 3GPP Architecture

5620 SAM Functions

Here is a list of the functionalities that will have to be introduced in the 5620 with regards to the extended management of LTE Network Elements (primarily SGW and PGW nodes) and additional applications to augment the value of the 5620 SAM in this space.

- Mobile Service
- EPS Path Discovery
- AGW Configuration Facilitator
- Service Data Flow SDF
- Bearer Context Statistics
- Extension to Statistics Plotting
- Performance Management Job
- Control Plane Threshold Configuration
- PolicyBased Configuration Audit

9. DEPRECATIONS

Deprecations This Release

5620 SAM LSP MAP

The 5620 SAM caps all development work for the 'LSP Topology' map feature set in R7.0 and shall remove the feature from R8.0R(TBD). The 'LSP Topology' map is being removed from 5620 SAM because other features within 5620 SAM are far better suited to managing SROS IP/MPLS content complexity at scale - eg. LDP-over-RSVP or CoS tunnels - than current map technology can reasonably deliver. Moreover, more advanced features outside of the scope of 5620 SAM that are provided in 5650 CPAM dramatically exceed the utility of the old map in current releases. Therefore, resources are being concentrated on enhancing LSP functionality within in SAM rather than the legacy 'LSP Topology' feature. Some of the NSM portfolio features that deliver equivalent, and superior, functionality for LSP management are as follow:

- 5620 SAM's tunnel manager delivers dramatic improvements in R5.0 / R6.0 / R7.0 addressing many new SROS IP/MPLS features and required scale. The existing 'LSP Topology' feature simply could not provided practical usability at network scale exposing far too many edge-to-edge lines to be readily used. Moreover, the complexity of new features like switching or LDP-over-RSVP make existing challenges far harder to resolve in a map style interface.
- 5620 SAM significantly enhances the scalability and utility of tunnel feature sets by extending a new filtering and finding capability to the 5620 SAM tunnel manager. Next generation filtering and finding allows a user to rapidly find the tunnels he / she is interested in overcoming the 'ball of string' inevitably exposed in a map with hundreds or thousands of tunnels. Filtering and finding is far more effective at identifying LSPs of interest within the tunnel manager and is the preferred solution for this feature.
- 5620 SAM's tunnel automation feature delivered in R6.0 significantly improves life cycle maintenance by automating tunnel creation and change. This feature is further enhanced in R7.0 to support definition of tunnel of naming conventions through policy, new CSPF attributes, CoS tunnels, change management, and re signaling. Moreover, the automation features reduce the number of incompletely configured LSPs for mesh, ring or hub and spoke configurations making it easier to maintain LSPs. The legacy map is not suitable for such application as listing along with filtering are more appropriate when trouble shooting at this level.
- 5620 SAM's alarm management has been significantly enhanced in R7.0 to improve network side correlation again making the workflow at the alarm layer more efficient than map inspection.
- 5620 SAM's LSP manager will be further enhanced in R8.0 to leverage correlation to further simplify identification of incomplete meshes, failed tunnels or other common problems formerly available in the map. Moreover, we will capture many of the common 'LSP map' use cases such as find failed LSPs, find isolated sites without LSPs, and policy audit within the tunnel automation.
- 5650 CPAM R1.0 / R2.0 significantly supersedes all existing 'LSP Tunnel' capabilities while providing far easier work flows for visualization and trouble shooting. For example, users can easily place any LSP type and path on an actual IP topology in the network which is impossible with the present LSP map. The presence of check points, path & tunnel histories, and audit functions dramatically aid trouble shooting while being fully in sync with the 5620 SAM tunnel management features in R5.0 & R6.0.
- 5650 CPAM R3.0 will further this trend by leveraging IP level capabilities to support tunnel concatenation (LDP-over-RSVP), CoS tunnel placement and related features as correlated to the real time IP topology.

SAM-O Deprecations

5620 SAM Release 7.0 was the last release to support Java Release 5.0 for the 5620 SAM-O.

Deprecations in Future Releases

Dynamic Shelf Drawings

5620 SAM Release 6.x was the first release that implemented static shelf drawings for newly supported NEs. The trend toward static drawings will continue, and in a future release of the 5620 SAM, the Equipment Manager will be removed and all functionality provided by it will be moved to individual equipment properties forms.

SAM-O Deprecations

The following jars are deprecated and will be removed in a future release. samOss.jar should be used to connect to the 5620 SAM server:

samOssJBoss.jar

samOssAnyServer.jar

The SchemaChanges80.html file on the product DVD-ROM contains information on deprecated XML API content. The file is available from an installed 5620 SAM client in the
<installation_directory>→client→nms→distribution→User_Documentation→5620_SAM-O_documentation/XML_Reference

Platform Deprecations

5620 SAM Release 8.0 is the last release to support Sun Microsystems T-series servers hosting either a 5620 SAM server or database.

5620 SAM Release 8.0 is the last release to support two CPU core platforms, such as the v240/v245/v215/v210, for live deployments of the 5620 SAM server and database.

5620 SAM Release 9.0 is the last release to support Windows on the 5620 SAM server and database.

10. ORDERING INFORMATION

Sales Engineering Information

5620 SAM H/W PLATFORM sizing web tool [link](#).

5620 SAM DVD-ROM Request [link](#).

5620 SAM License Key [link](#).

5620 SAM Pricing Information [link](#).

- New part numbers will be added for new elements supported.
- Note that 7750 MDAs running in mixed mode on a 7450 will count as 7750 MDAs. 7450 MDAs running in mixed mode on a 7750 will count as 7450 MDAs.

Note that as of release 7.0 Integrated Service Adapters no longer require a Premium license.

Note that if a customer reaches the limit of licensed Premium MDAs but still has spare Standard MDAs and wishes to add another Premium MDA, two standard MDAs will be counted.

5620 SAM Technical Support [link](#).

5620 SAM Documentation [link](#).

Licensing Information

The 7705 SAR-8 requires two license pools, one for the 7705 SAR-8 chassis count and one for the 7705 SAR-8 daughter card count. Each discovered 7705 SAR-8 chassis shall consume one 7705 SAR-8 chassis license; and each daughter card within a discovered 7705 SAR-8 shall consume one 7705 SAR-8 daughter card license. If either license pool is negative, no additional 7705 SAR-8 nodes can be managed.

The 7705 SAR-F requires a license pool, separate from those used for the 7705 SAR-8. The 7705 SAR-F license pool shall be handled similarly to the 7705 SAR-8 chassis license pool.

7210 SAS M & 7210 SAS MX[ETR] is supported as node based license. User has to buy separate license for SAS-M and SAS MX (ETR). It is to be noted that the license for SAS-MX is also shared with SAS-MX ETR. So if user has 5 SAS-MX licenses, one can manage 3 SAS-MX and 2 SAS-MX ETR or 1 SAS-MX and 4 SAS-MX ETR.

Oracle Ordering Information

When purchasing the 5620 SAM application, sufficient licenses must be acquired for the utilization of the Oracle product, which is embedded into the 5620 SAM Database software. The licenses required to operate the Oracle product are computed based on the number of CPUs in the server workstation on which the 5620 SAM Database component will operate.

The table below provides examples of the type and possible number of CPUs that can be included in the servers that Alcatel-Lucent recommends for the operation of 5620 SAM.

The additional Oracle licensing requirements can be computed by looking at the number of CPUs that are included in the Server workstation on which the 5620 SAM Database component is installed. When a 5620 SAM installation is redundant, these Oracle licensing requirements must be doubled.

Sun Server	Number of CPUs	Additional Oracle Licensing Requirements
Single-core CPU Platforms (mono-core)		
SunFire v240 Server	Up to 2	None
SunFire v440 Server	Up to 4	If 2 CPUs, no additional licensing If 4 CPUs, additional licensing for 2 CPUs
SunFire v445 Server	Up to 4	If 2 CPUs, no additional licensing If 4 CPUs, additional licensing for 2 CPUs
Dual-core CPU Platforms		
SunFire v490 Server	2 or 4	If 2 CPUs, additional licensing for 1 CPUs (dual-core) If 4 CPUs, additional licensing for 3 CPUs (dual-core)
SunFire v890 Server	Up to 8	If 2 CPUs, additional licensing for 1 CPUs (dual-core) If 4 CPUs, additional licensing for 3 CPUs (dual-core) If 6 CPUs, additional licensing for 5 CPUs (dual-core) If 8 CPUs, additional licensing for 7 CPUs (dual-core)
Multi-core CPU Platforms (4-core or 8-core)		
SunFire T2000 Server (small 4-core)	1 CPU (4-core)	None
SunFire T2000 Server (large 8-core)	1 CPU (8-core)	None
AMD Platforms (Solaris x86)		
SunFire x4100	1 or 2 single or dual-core CPUs	No additional licensing requirements
SunFire x4200	1 or 2 single or dual-core CPUs	No additional licensing requirements
SunFire x4600	Up to 8 dual-core CPUs	If 2 dual-core CPUs, no additional licensing requirements If 4 dual-core CPUs, additional licensing for 2 CPUs (mono-core) If 6 dual-core CPUs, additional licensing for 4 CPUs (mono-core) If 8 dual-core CPUs, additional licensing for 6 CPUs (mono-core)

Table 20: Platform CPU and Oracle Equivalent

11. REFERENCES

The following documents have been referenced:

- ITU-T Rec. G.8261/Y.1361 (05/2006), Timing and synchronization aspects in packet networks, 2007/03/01.