EMC® Smarts®

Features from Version 9.2.x Applicable to Version 9.3.x SAM, IP, ESM, MPLS, VoIP, and NPM Managers

P/N 302-001-048

REV 01

EMC²

Copyright © 2014 EMC Corporation. All rights reserved. Published in the USA.

Published June, 2014

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

The information in this publication is provided as is. EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose. Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

EMC^2, EMC, and the EMC logo are registered trademarks or trademarks of EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners.

For the most up-to-date regulatory document for your product line, go to EMC Online Support (https://support.emc.com).

CONTENTS

Chapter 1 9.2.x Release Product Features and Technical Notes

Features from 9.2.x release products that apply to 9.3.x 6
Features in 9.2.2 release products 6
Technical notes from 9.2.x release products that apply to 9.3.x 6
Technical notes for 9.2.2 release products 6
Technical notes for 9.2.1 release products 7
Fixed problems from 9.2.x that apply to 9.3.x 8
Fixed Problems in 9.2.x 8

Chapter 2 Service Assurance Manager 9.2.x Features and Technical Notes

Service Assurance Manager features from 9.2.x that apply to 9.3.x 12
Features in SAM 9.2.2 12
-Features in SAM 9.2.1 12
Service Assurance Manager technical notes from 9.2.x that apply to 9.3.x ... 16
Technical notes for SAM 9.2.2 16
Technical notes for SAM 9.2.1 18
Service Assurance Manager fixed problems from 9.2.x that apply to 9.3.x ... 19
Fixed problems in Service Assurance Manager 9.2.x 19

Chapter 3 IP Manager 9.2.x Features and Technical Notes

IP Manager features from 9.2.x that apply to 9.3.x 24
Features in IP Manager 9.2.2 24
Features in IP Manager 9.2.1 38
IP Manager technical notes from 9.2.x that apply to 9.3.x 55
Technical notes for IP Manager 9.2.2 55
Technical notes for IP Manager 9.2.1 60
IP Manager fixed problems from 9.2.x that apply to 9.3.x 61
Fixed problems in IP Manager 9.2.x 61

Chapter 4 Server Manager 9.2.x Features and Technical Notes

Server Manager features from 9.2.x that apply to 9.3.x 66
Features in Server Manager 9.2.2 66
Features in Server Manager 9.2.1 66
Server Manager technical notes from 9.2.x that apply to 9.3.x 67
Technical notes for Server Manager 9.2.2 67
Technical notes for Server Manager 9.2.1 71
Server Manager fixed problems from 9.2.x that apply to 9.3.x 72
Fixed problems in Server Manager 9.2.x 72

Chapter 5 MPLS Manager 9.2.x Features and Technical Notes

MPLS Manager features from 9.2.x that apply to 9.3.x 74
Features in MPLS Manager 9.2.2 74
Features in MPLS Manager 9.2.1 74
MPLS Manager technical notes from 9.2.x that apply to 9.3.x 74
Technical notes for MPLS Manager 9.2.2 74
Technical notes for MPLS Manager 9.2.1 74

MPLS Manager fixed problems from 9.2.x that apply to 9.3.x..... 75
Fixed problems in MPLS Manager 9.2.x 75

Chapter 6 Network Protocol Manager 9.2.x Features and Technical Notes

Network Protocol Manager features from 9.2.x that apply to 9.3.x 78
Features in Network Protocol Manager 9.2.2..... 78
Features in Network Protocol Manager 9.2.1..... 78
Network Protocol Manager technical notes from 9.2.x that apply to 9.3.x..... 79
Technical notes for Network Protocol Manager 9.2.2 79
Technical notes for Network Protocol Manager 9.2.1 80
Network Protocol Manager fixed problems from 9.2.x that apply to 9.3.x 81
Fixed problems in Network Protocol Manager 9.2.x 81
Fixed problems in Network Protocol Manager 9.2.x patch documentation .. 81

Chapter 7 VoIP Availability Manager 9.2.x Features and Technical Notes

VoIP Availability Manager features from 9.2.x that apply to 9.3.x 84
Features in VoIP Availability Manager 9.2.2 84
Features in VoIP Availability Manager 9.2.1 86
VoIP Availability Manager technical notes from 9.2.x that apply to 9.3.x 88
Technical notes for VoIP Availability Manager 9.2.2..... 88
Technical notes for VoIP Availability Manager 9.2.1..... 88
VoIP Availability Manager fixed problems from 9.2.x that apply to 9.3.x..... 89
Fixed problems in VoIP Availability Manager 9.2.x..... 89

CHAPTER 1
9.2.x Release Product Features and Technical Notes

This chapter provides information for EMC Smarts Service Assurance Manager, EMC Smarts IP Manager, EMC Smarts MPLS Manager, EMC Smarts Server Manager, EMC Smarts Network Protocol Manager, and EMC Smarts VoIP Availability Manager. It includes the following sections:

- Features from 9.2.x release products that apply to 9.3.x 6
- Technical notes from 9.2.x release products that apply to 9.3.x 6
- Fixed problems from 9.2.x that apply to 9.3.x 8

Features from 9.2.x release products that apply to 9.3.x

- "Features in 9.2.2 release products" on page 6
- "Technical notes from 9.2.x release products that apply to 9.3.x" on page 6
- "Fixed Problems in 9.2.x" on page 8

Features in 9.2.2 release products

- "New behavior for the sm_migrate utility" on page 6

New behavior for the sm_migrate utility

The sm_migrate utility prompts you to copy the security configuration files including the runcmd_env.sh file. You can either choose to copy the files into the local directory of the new installation or configure the files later, manually. If you choose to copy and have changed the site secret in your previous installation, run the sm_rebond command to encrypt the files:

```
./sm_rebond --basedir=C:\InCharge\IP\smarts
```

After copying the files, the data in the old runcmd_env.sh file is appended to the new runcmd_env.sh file and the new Smarts 9.2.2 version data is commented.

That is, if you select to reuse the runcmd_env.sh file from your old installation, the sm_migrate utility migrates the old values and will also append a commented out section to the bottom of the runcmd_env.sh file with the default environment values in the new version for your evaluation.

Technical notes from 9.2.x release products that apply to 9.3.x

- "Technical notes from 9.2.x release products that apply to 9.3.x" on page 6
- "Technical notes for 9.2.1 release products" on page 7

Technical notes for 9.2.2 release products

Instructions for changing default Cryptographic settings

Changing Pseudo-random Number Generators (PRNG) algorithm

CQ 649798

The Smarts Global Console SSL implementation uses ECDRBG128 (Dual EC DRBG) as the default PRNG algorithm. Following NIST's decision to recommend against the use of Dual_EC_DRBG, it is recommended that you discontinue the use of Dual EC DRBG, and move to a different PRNG (HMACDRBG256) at this time.

Changing the use of the default PRNG algorithm

Add the following line in the java.security file located at BASEDIR/_jvm/lib/security/ directory:

```
com.rsa.crypto.default.random = HMACDRBG256
```

Apply this change to the following components:

- Smarts Global Console
- Any other custom programs written by you using the DMT Java remote API

The *EMC Smarts Version 9.2 Security Configuration Guide* is updated to provide instructions on how to change the default cryptographic settings in Smarts components as well as provide information on how to change the use of the default PRNG algorithm.

Replacing or importing TLS certificates

By default, a Smarts Domain Manager generates a self-signed TLS certificate that it will present to the incoming TLS connections. However, you can configure a certificate to use by performing the following step:

Place the certificate and key files in $SM_SITEMOD/conf/tls/$DOMAIN.crt and $SM_SITEMOD/conf/tls/$DOMAIN.key where $DOMAIN is the name of the domain used to register with the broker. The files must be in PKCS#8, PEM encoded form with no password.

Self-signed certificates generated by Smarts are stored in the process memory. To regenerate default self-signed TLS certificates, stop and restart the Smarts processes.

Technical notes for 9.2.1 release products

There are no technical notes that apply to all 9.2.1 release products.

Fixed problems from 9.2.x that apply to 9.3.x

"Fixed Problems in 9.2.x" on page 8

Fixed Problems in 9.2.x

Features introduced in the Fixed Problems tables from 9.2.x.release notes are included here.

CQ/SR/Issue number	Problem summary	Description of fix	Found in version
Fixed problems common for all 9.2.x products			
CQ 577981 CQ 577738	--env=SM_SITEMOD crashed the VoIP domain in Windows. When Smarts was installed as a Windows service, with SM_SITEMOD specified using the --env option, the SM_SITEMOD search path did not get computed correctly when the service started.	Apply the environment variable settings (which are encoded in a special command line argument) early from the runcmd.exe wrapper program, so that the values are used in any search path computations performed by runcmd. Changed Components: Executables: runcmd.exe (and hence all shipped executables in smarts/bin) Libraries: sm_sms	9.2.1
CQ 644152 CQ 604257 SR 55843484	The activation command did not set the Standby domain to the Active domain.	The failover server did not process the command to TakeOver a domain. When the shell script actions that are part of the failover solution are modified using sm_edit, these actions fail when the failover server runs them. The first one is due to a limitation on the failover server. Backups are currently run sequentially and there were too many domains attached to the failover manager which was resulting in too much time being spent on backups. The recommendation is to split the domains into multiple failover servers so that the time spent doing backups is optimized and the failover/takeover requests are handled quicker. For example, put SAM, OI, BIM, NPM into one failover server and IP, MPLS into another failover server. The second issue was due to a bug in the shell scripts. Locating files under SM_SITEMOD/actions were not handled properly. They were looking at SM_HOME/actions. This is now fixed. The scripts can be modified using sm_edit and the failover server processes the actions from the modified script.	9.2.0.1

CQ 644510 CQ 644456	When the failover server was restarted, it started running backups and operations could not be performed until the backups completed.	A parameter “AutoBackupsEnabled” is now provided in the failover.conf file. The “AutoBackupsEnabled” parameter controls backups that are run automatically by the Failover Domain Manager. You can change the value to FALSE in the user interface, only if you want to disable the backup tasks. **Note:** (1) Changing the value in the Failover.conf file will result in unsuccessful failovers. (2) When starting the failover server for the first time, you must start it with AutoBackupsEnabled = TRUE. On subsequent restart, it can be set to FALSE for no auto manual backup.	9.2
CQ 577477 SR 50887122	Server's memory usage increases when there is a rate function defined over a negative window.	A negative value converted to unsigned resulted in a big value for the window over which to calculate the rate. This resulted in a big amount of samples that were required to cover the window. So new samples kept being added to the history collection data structure, which was causing the increased memory usage. The fix was to limit the collection size. The limit was set to 1000 samples. The rate function in this situation will be suspended with error MR_NOT_ENOUGH_TIME.	9.2.1
CQ 597446 SR 49540808	A problem gets suspended with "DX-PROBLEMUNKNOWN-Problem unknown to correlator because it has no symptoms", even though it has symptoms, and consistencyUpdate has been run after topology changes.	The problem was observed when server was started with --subscribe option that starts internal subscriber on start up. The correlator's consistencyUpdate has not been called by the time subscriptions started, so the problems initially gets rejected with DX-PROBLEMUNKNOWN error because they have not been added to the codebook yet. Normally, these problems will be resubscribed in the course of consistencyUpdateCompletion by resubscribeRejections call, and the error will be cleared. However, it may happen that by the time resubscribeRejections runs some of the problems are still in pending state, since their subscritpion processing has not been finished. These problems will not get resubscribed, and the error will not clear until the next consistencyUpdate call. The fix was to make resubscribeRejections being called after all pending subscriptions are processed.	9.2.1
CQ 577631 CQ 577591	Made the TLS cipher suite preferences configurable.	Allowed SM_TLS_SUITE_LIST to be set to a colon-separated list of cipher suite specifier strings. The list of allowable options, including FIPS compatibility, is listed in the RSA BSAFE MES documentation. The client and server will choose the first item in the client's suite list that also appears in the server's list. In other words, client and server must agree on acceptable options, but the client's preference wins. Also, removed undocumented support for FIPS TLS Suite B. Changed Components: Libraries: sm_clsapi	9.2.1

CQ 577628 CQ 577588	Allow CI_FlowBufferHead_U default buffer size to be tuned through an environment variable.	The default buffer size for CI_FlowBUfferedHead_U affects performance, but is not configurable at runtime, so any investigation of performance changes or tuning requires custom-built libraries. Introduced an environment variable SM_FBH_BUFSIZE to allow adjusting the default. It currently defaults to 2048 (the old hardcoded default).	9.2.1
CQ 642779	sm_migrate.pl error message.	sm_migrate.pl must be invoked using sm_perl.	9.2.1
CQ 567978	Enable run-time control over locking in SM_RepositoryServer get methods.	Introduced method EnableTransaction to MR_LockInterface model class to set the mask defining which methods of SM_RepositoryServer should start repository transaction. Also added method GetTransactionMask, to get the current mask. By default, transaction is disabled in all methods. Added an option to selectively unblock readers in SM_RepositoryServer transactions, depending on the method initiating the transaction, a single-property get or a get of multiple properties. Set the default value of MR_UnblockReaders::allowed flag to disallow unblocking in all cases. **Note:** Do not change these settings unless advised by EMC support. Changed Components: Libraries: sm_repos Model file: MR_LockInterface.mdl	9.2.1

CHAPTER 2

Service Assurance Manager 9.2.x Features and Technical Notes

This chapter includes the following sections:

- Service Assurance Manager features from 9.2.x that apply to 9.3.x..... 12
- Service Assurance Manager technical notes from 9.2.x that apply to 9.3.x..... 16
- Service Assurance Manager fixed problems from 9.2.x that apply to 9.3.x 19

Service Assurance Manager features from 9.2.x that apply to 9.3.x

- "Features in SAM 9.2.2" on page 12
- "-Features in SAM 9.2.1" on page 12

Features in SAM 9.2.2

Support for Internet Explorer (IE) 10 and IE 11

The *EMC Smarts Support Matrix* provides a list of officially supported browsers for launching the EMC Smarts Global Console.

-Features in SAM 9.2.1

SAM release 9.2.1 includes important new features and changes to the existing features:

- "SAM heartbeat monitoring" on page 12
- "Improvements to SAM archival time" on page 15
- "Service Assurance Manager technical notes from 9.2.x that apply to 9.3.x" on page 16
- "Service Assurance Manager technical notes from 9.2.x that apply to 9.3.x" on page 16

SAM heartbeat monitoring

All Smarts servers maintain an uptime attribute through the SM_System::SM-System object, which is updated once every minute. This uptime attribute can be perceived as the heartbeat of a Smarts server. In this release, SAM server has been enhanced with the functionality to monitor the heartbeat of a Smarts server. SAM heartbeat monitoring is a feature available from SAM that provides the capability to monitor the health conditions of the heartbeats of its underlying domains. In addition to monitoring heartbeats of underlying domains, this new feature is also capable of detecting two abnormal conditions:

- burstyNotification - Extreme burst arrivals of events or notifications, which could overwhelm the system resulting in slow response time or even freeze.
- hyperNotifThreadHung - This use case applies only to SAM -> SAM communication in hierarchical SAM configuration and to SAM -> OI communication. The feature can detect hanging of the hyper-notify thread associated with the underlying domain, which may indicate the notification processing at the local SAM server being suspended.

The feature is facilitated by the introduction of the ICS_Heartbeat model. Once SAM heartbeat monitoring is enabled from top level SAM for underlying servers, an instance of ICS_Heartbeat is created correspondingly for this domain in SAM. ICS_Hearbeat is responsible for generating the heartbeat events as well as events indicating the two abnormal conditions stated above. In a hierarchical SAM configuration, ICS_Heartbeat monitoring needs to be enabled for every SAM and OI server.

Issues addressed

SAM heartbeat monitoring, once deployed, provides a mechanism that can quickly identify the nature of some performance issues that are encountered in the field, thereby reducing troubleshooting time or providing guidance to hardware upgrade.

For example, if SAM experiences slow response time or freeze, it can be quickly determined whether the problem is due to burst arrivals from a connected domain, or due to a slow network connection, or high CPU usage. If it is due to the burst of Notifications, burstyNotification alerts for the domains that have sent all the notifications are seen. If the issue is due to a slow network connection, heartbeatDelayed notifications are received.

Further, if the SAM domain stopped processing new notifications from one or more underlying domains, it can be quickly determined whether the problem is due to hanging threads (hyperNotifThreadHung) or due to a broken network connection (heartbeatDelayed).

Heartbeat monitoring events

ICS_Heartbeat model generates four types of events when the event conditions are met. All events are converted to SAM notifications, therefore are visible from a Console.

- heartbeatReceived - Indicates a heartbeat is received from the monitored domain. No action is needed, which indicates the system is operating normally.
- heartbeatDelayed - Indicates a heartbeat has missed its deadline (determined by a configurable threshold) of 180 seconds. No immediate action is needed. This event will be cleared when a new heartbeatReceived event has arrived. If heartbeatDelayed is not cleared within several minutes, it could indicate a connection problem with the domain server, or the domain server is overwhelmed by operations that may hold the repository for too long, such as reconfigure or repository save or a hook script with a bug, therefore delaying reporting its heartbeat. If there is a consistent heartbeatDelayed event, it may indicate a hardware upgrade opportunity as well if the problem is not caused by a buggy hook script. If heartbeatDelayed notification is persistent, run the sm_monitor tool in ‹BASEDIR›/smarts/bin directory against the domain for which the heartbeatDelayed notification was generated.
- burstyNotification - Indicates the condition that notification or event arrivals per minute exceeds a threshold of 3000 per minute. No immediate action is needed. This event indicates the reason/root cause behind the response time slowness. It can pin-point the source of the bursts. To recover from the slowness, temporarily disconnect the source or reduce the bursts by enabling filtering (such as trap filters). If the bursts are legitimate and persistent, it might be an opportunity to upgrade the hardware or rethink on the system architecture, for example, if your system is capable of handling over 3000 notifications per minute, increase the threshold to avoid false alerts. “Setting up SAM heartbeat monitoring” on page 14 provides more details.

- hyperNotifThreadHung - Indicates the condition that the hyper-notify thread stopped processing events for more than the time period determined by a threshold of 180 seconds. Take action to resolve this issue. If the heartbeatDelayed alert is generated at the same time, it could indicate a network connection issue or a freezing peer server. If the heartbeatReceived event is generated normally, it could indicate that the hyper-notify thread within the local domain has stopped processing notifications which may be due to either a failed hook script or a remote function call that has never returned (such as a prolonged DNS request). By taking a sequence of the stacktrace using the sm_monitor tool in ‹BASEDIR›/smarts/bin, the exact cause can be determined. The sm_monitor tool needs to run against the domain running the hyper-notif thread, for example, top level SAM or SAM subscribing to OI.

Setting up SAM heartbeat monitoring

SAM heartbeat monitoring is currently not enabled by default.

To enable heartbeat monitoring:

1. In a hierarchical SAM configuration, heartbeat monitoring needs to be enabled at every layer of the hierarchy. Enable heartbeat monitoring once the SAM servers are up and running and the underlying domains are configured.

   Use dmctl to initialize and set up heartbeat monitoring by typing the following command:

   ```
   dmctl> invoke ICS_HeartbeatMonitor::ICS-HeartbeatMonitor initialize
   ```

   Once this command is executed, a conf file named as heartbeat.conf is automatically created in the local/conf folder, which contains all the appropriate entries for that SAM domain.

2. Use sm_edit to open smarts/local/conf/ics/heartbeat.conf (or the smarts/local/conf/icoi/heartbeat.conf) file. Check the following ICS_Heartbeat attributes for each instance based on the monitoring requirements and data from the past. The set of attributes include:

   - heartbeatDelayThreshold - the maximum delay time (in seconds) of the monitored heartbeat before a heartbeatDelayed event is fired by the corresponding ICS_Heartbeat instance.
   - notificationThresholdInOneHeartbeat - when notifications which had arrived from the connected domain exceeds the threshold, a burstyNotification event is fired by the corresponding ICS_Heartbeat instance.
   - hyperNotifQuietTimeThreshold - when the time period has passed in which there is no notification received through hyperNotif, a hyperNotifThreadHung event is fired by the corresponding ICS_Heartbeat instance.

3. Close the heartbeat.conf file if the default settings meet the monitoring requirements. This completes the setup.

4. If there are changes to be made to heartbeat.conf, then again use dmctl to reconfigure heartbeat monitoring by typing the following command:

```
dmctl> invoke ICS_HeartbeatMonitor::ICS-HeartbeatMonitor reconfigure
```

Once the above steps are executed correctly, heartbeat events will be displayed in the Console. Attach a Console directly to SAM and verify the arrival of heartbeat events to determine whether heartbeat monitoring has been set up correctly.

Controlling SAM heartbeat monitoring

ICS_HeartbeatMonitor is a manager object for controlling this feature. This object provides the following functions to enable and disable SAM heartbeat Monitoring. All the functions can be executed through dmctl.

- enableMonitorAll - Enables all heartbeat monitoring (which is default after initialize() is called)
- disableMonitorAll - Disables all heartbeat monitoring
- enableMonitorForDomain (in string domainName) - Enables monitoring for a specific domain
- disableMonitorForDomain (in string domainName) - Disables monitoring for a specific domain
- enableNotifyHeartbeatAll - Allows heartbeatReceived to be reported (which is default after initialize() is called). This does not turn off monitoring.
- disableNotifyHeartbeatAll - Does not allow heartbeatReceived to be reported
- enableNotifyHeartbeatForDomain (in string domainName) - Allows heartbeatReceived to be reported only for a specific domain.
- disableNotifyHeartbeatForDomain (in string domainName) - Does not allow heartbeatReceived to be reported for a specific domain
- resetHeartbeatAll() - Resets heartbeat statistics maintained by all ICS_Heartbeat Instances.
- resetHeartbeatForDomain (in string domainName) - Resets heartbeat statistics for a specific domain.

Improvements to SAM archival time

This is a performance improvement of the existing SAM archive feature. In the earlier release, it used to take the SAM server a long time to archive notifications with many audit entries. In this release, the code has been optimized so that SAM does not take a long time to archive notifications, even though it contains a large audit log.

Service Assurance Manager technical notes from 9.2.x that apply to 9.3.x

- "Technical notes for SAM 9.2.2" on page 16
- "Technical notes for SAM 9.2.1" on page 18

Technical notes for SAM 9.2.2

- "Details on how KeepAlive works" on page 16
- "After a crash, the failover server does not restart the Failover System" on page 16
- "How to configure trapFlood traps and TrapdDiscardingTraps" on page 17
- "Timeout value is set by default for all underlying Domain Managers" on page 17
- "Resolving issues related to unusual growth in log file size" on page 18
- "Users with custom code that uses EMC Smarts API" on page 18

Details on how KeepAlive works

CQ 641853 / CQ 641846 / SR 52796808

The Perl script acts as a client that connects to the EMC Smarts Domain Manager. The Perl script and the Domain Manager communicate through a TCP connection.

The KeepAlive messages check the connection between the clients, including Perl clients, and the EMC Smarts Domain Managers. Under certain circumstances, the client's operating system might incorrectly insert the urgent data into the regular data stream, which corrupts the regular data stream and causes problems in the connection. This leads to an error on the client side.

For a detailed explanation on how KeepAlive messages work, and to disable the feature, refer to SM_DISABLE_KEEPALIVES in the "Environment variables for users" section of the *EMC Smarts Foundation 9.2 System Administration Guide*.

After a crash, the failover server does not restart the Failover System

CQ 649126

When a failover server crashes, the server restarts from the point before it shut down and does not restart the Failover System.

For example, INCHARGE-SA is started as primary server A and secondary server B on two different machines. When a failover occurs, server B becomes the primary server and server A becomes the secondary server. Now if the failover server crashes, then after a restart of the failover server, INCHARGE-SA server B will still start as the primary server, even though it is set up as the secondary server in the failover.conf file, that is, server B will not pick up the status from the failover.conf file.

The roles cannot be reversed through the console.

How to configure trapFlood traps and TrapdDiscardingTraps

CQ 642856 / CQ 642389

"Starting and stopping the SNMP Trap Adapter" section in the *EMC Smarts Service Assurance Manager Adapter Platform User Guide* provides information on how to use the sm_service utility to manually start and stop the SNMP Trap Adapter.

This section describes how to configure trapFlood traps and TrapdDiscardingTraps.

.1.3.6.1.4.1.733.10 6 1 and .1.3.6.1.4.1.733.10 6 2 are TrapdDiscardingTraps. The following traps are sent by the sm_trapd queue when it reaches the QUEUE_LIMIT_MEGS or the QUEUE_LIMIT_SECS limits:

- .1.3.6.1.4.1.733.10 6 1 is sent when the Trap receiver queue has reached its limit.
- .1.3.6.1.4.1.733.10 6 2 is sent when the Trap receiver is not discarding traps any more.
- .1.3.6.1.4.1.733.10 6 3 and .1.3.6.1.4.1.733.10 6 4 are TrapFlood. The following traps are sent by the sm_trapd queue when it determines that certain sources are flooding it with traps:
  - .1.3.6.1.4.1.733.10 6 3 is sent when the Device is flooding the trap receiver.
  - .1.3.6.1.4.1.733.10 6 4 is sent when the Device is not flooding the trap receiver any more.

These traps are introduced in SAM 9.x and later releases. The trap_mgr.conf file provides more information on these traps.

You can create the flood trap using NOTIF by importing the trap_mgr.conf or by creating the Event Configuration Information (ECI) similar to the definition in the trap_mgr.conf file. The *Notification Module User Guide* provides more information on how to create the ECI.

Timeout value is set by default for all underlying Domain Managers

CQ 576857

In the 9.1 release, a timeout value is set by default for Event driver, Notification driver, and Topology driver. A default timeout value of 600 seconds is enabled for all underlying Domain Managers connected to SAM. Previously, SAM would wait indefinitely for a response from the underlying Domain Managers. Now with the default timeout value, SAM will wait for 600 seconds for the underlying Domain Managers to respond after which, it will end the connection with the Domain Manager and start a new connection.

The timeout value is assigned to the variables, TimeOutBE and TimeOutTopo in the files – ics-event-driver.asl, ics-event-processing.asl, and ics-nl-processing.asl. TimeOutBE is a timeout set for events and notification driver timeouts, and TimeOutTopo is a timeout set for topology driver timeout.

If you would like to modify the timeout value, you need to edit the variables in the dxa conf file corresponding to the underlying Domain Manager.

Resolving issues related to unusual growth in log file size

CQ 636047

If a huge growth in the log file size is observed, for instance in the Broker log:

1. Re-install the Broker as service with environment setting SM_NO_LOG_CLIENTS=TRUE.
2. Restart the Broker service.

NOTICE

To disable the client logging feature subscription for all the servers or Domain Managers, set SM_NO_LOG_CLIENTS=TRUE in the runcmd_env.sh file, and then restart the Broker and Domain Managers.

Users with custom code that uses EMC Smarts API

If the code is based on Java, Perl, or C API, modify and retest the code according to the changes in this release. If a dynamic model is used, recompile the code.

The *EMC Smarts Dynamic Modeling Tutorial* explains the concepts and methods of dynamic modeling.

Technical notes for SAM 9.2.1

There are no technical notes for Service Assurance Manager 9.2.1.

Service Assurance Manager fixed problems from 9.2.x that apply to 9.3.x

"Fixed problems in Service Assurance Manager 9.2.x" on page 19

Fixed problems in Service Assurance Manager 9.2.x

Features introduced in the Fixed Problems tables from 9.2.x.release notes are included here.

CQ/SR/Issue number	Problem summary	Description of fix	Found in version
Fixed problems in SAM 9.2.x			
CQ 644260 CQ 577986 SR 56283120	Overloaded logs were added to the OI server. The log messages help debug if traps are not converted into notifications because of the unknownagent setting. Notif has different log levels that you can specify when configuring it.	In earlier releases, the log level was set to ALWAYS. From this release, the log level is changed to DEBUG. These messages are logged only when the log level is set to DEBUG. Set the log level during the start of the SAM/OI server and restart the SAM/OI server if you need to enable this debugging.	8.1.4.4
CQ 644369 SR 56401364 CQ 644761 CQ 644210	When multiple domains with different domain names were started from the same SAM installation, the heartbeat.conf file was overwritten and the first domains entries were erased from the configuration file.	Running instances of the SAM server are distinguished from their server names. Thus the heartbeat configuration file name is now suffixed with the SAM server name. So when multiple SAM server instances are started from a single installation, as long as the corresponding server names are different, which should always be true in practice, the heartbeat monitor module should be able to automatically pick the correct configuration file to import or export heartbeat monitor related configurations. The location of the heartbeat monitor configuration file is not changed.	9.2.1
CQ 644394 CQ 604782 SR 55693508	Aggregate/Presentation SAM WAS_ACTIVE logic did not work. The scenario was that the Presentation SAM was receiving the same notification from two Aggregate SAMs and then the first Aggregate SAM goes down. Because we are not notifying the notification received by the second Aggregate SAM, the notification in the Presentation SAM was marked as WAS_ACTIVE because the EventStateTable of the notification had only the first Aggregate SAM entry. Ideally the notification should have been ACTIVE since the second Aggregate SAM was still running.	The fix is to NOTIFY the notification even if we have received the same notification from the same source. By this we have the second Aggregate SAM entry in the EventStateTable and the above scenario will be handled with the notification still ACTIVE.	9.2.0.1
CQ 644813 CQ 605991 SR 56797280 CQ 606891 CQ 642703	Device information was not correctly propagated from the Domain Manager to SAM for the FibreChannelSwitch class. Prior to the 9.2.1 release, because IP and ESM were following different naming conventions for the FC switch, SAM was creating duplicate instances (one for instance from IP and one instance from ESM).	From the 9.2.1 release, although the instance names are different in IP and ESM, SAM uses IP address which is the only common attribute between IP and ESM to avoid creating duplicate instances. Instance in SAM is created with IP address irrespective of the naming convention followed in the underlying domain.	9.2.1

CQ 642855 CQ 642742 SR 54337508 CQ 642855	In the Smarts 9.2 Installation Guide -> Overview -> Product and version compatibility -> Table 1 - Interoperability chart for EMC Smarts Service Assurance Manager, the EMC Smarts VoIP Availability Manager had only 4.0 and 8.1.x listed.	EMC Smarts VoIP Availability Manager version 9.x should also be listed. The EMC Smarts Version 9.2.2 Installation Guide for SAM, IP, ESM, MPLS, VoIP, and NPM Managers has been updated to reflect this change.	9.2.0.1
CQ 644791 CQ 644725 CQ 649340 SR 57271350 CQ 609664 CQ 609489	Heartbeat events were not updated after SAM was restarted.	This issue is now fixed by starting the heartbeat monitoring whenever SAM/OI server is restarted. An editable flag AUTO_START_HEARTBEAT is provided in the *ics-init.asl* and *icoi-init.asl* files to enable/disable the heartbeat monitoring auto restart. By default this flag is set to TRUE.	9.2.1

CQ 577496 SR 51566768 CQ 649082	SNMP v3 credentials need to be transferred to OI from IP and be used as part of SNMPv3 trap processing and handling. Keying in credentials twice (once in IP and then again in OI) is an overhead for users.	SNMP v3 Credentials present in IP Domain Manager will be automatically transferred to the OI server from IP during topology synchronization. The credentials will be saved in the InChargeDomain Object that will be read by the Trap Adapter Process. The Trap Adapter Process imports credentials into the Local Credential database from the OI server at startup or on reload configuration. To reload configurations for Trap Adapter Process: ./sm_trapd -b ‹broker-name› -s ‹TrapAdapter-Name› --reloadConfig **Automatic Transfer of Credentials from IP to OI to Trap Adapter process :** Configuration steps: 1. imk.dat in the IP installation and the INCHARGE-OI install area (from where the Trap Adapter process is started), should contain the same pass-phrase. To ensure this, recreate or upgrade the imk.dat using the sm_rebond command, in both the IP and OI install locations. 2. Start the OI server. Automatic transfer of credentials from IP to OI is disabled by default. To enable, run the following dmctl command : ./dmctl -b ‹broker-name› -s ‹OI server-name› put ICF_TopologyManager::ICF-TopologyManager::importSNMPv3Credentials TRUE 3. Underly the IP server to OI server. Now any credentials configured in the IP Server will be transferred to the OI server. 4. Now start a Trap Adapter attaching to the OI server. Credentials will be read by the Trap Adapter Process from the OI server during startup and configuration reload. 5. SNMPv3 Credentials would be synced from IP to OI at the end of Topology Synchronization. **Note:** After performing the reloadConfig operation for the Trap Adapter, configured as a Receiver to read the changed or new SNMPv3 Credentials from the OI server, wait till the first new trap is processed by the Trap Adapter for existing devices, so that the TrapAdapter successfully reads the changed or new SNMPv3 credentials from the OI Server.	9.2.1

CQ 642445 CQ 639335 CQ 559952 CQ 639295	SAM servers crashed due to out of memory.	The previousEvent ‹string set› has been added only for the HyperNotification functionality (added from SAM 8.X onwards). In case of HyperNotificaton, all incoming notifications are verified against currentEvents and previousEvent list. If it does not exist in either of them, it treats this event as an invalid notification. We do not have any logic to cleanup this list (we move the notification from previousEvent list to currentEvents in case of renotify and reverse action is done in case of deleting an event). Now we clean the old cache data periodically. This value is defined by the new flag: SM_CACHE_CLEANUP_SECONDS. Default 259200 seconds, that is, 3 days.	9.2.1
CQ 642135 CQ 642134 CQ 598052 CQ 640941	Migration of dynamic model files.	The documentation related to the section “Migration of dynamic model files” in the Smarts Installation Guide should be updated with the following information: 1. Recompile the .mdl files using the SAM dynamic model compiler. 2. Copy all the .dat files generated, to the local/l10n directory. 3. Start the SAM server.	9.2.1

CHAPTER 3
IP Manager 9.2.x Features and Technical Notes

This chapter includes the following sections:

- IP Manager features from 9.2.x that apply to 9.3.x 24
- IP Manager technical notes from 9.2.x that apply to 9.3.x 55
- IP Manager fixed problems from 9.2.x that apply to 9.3.x 61

IP Manager features from 9.2.x that apply to 9.3.x

- "Features in IP Manager 9.2.2" on page 24
- "Features in IP Manager 9.2.1" on page 38

Features in IP Manager 9.2.2

- "Changes in discovery Rogue Access Points" on page 24
- "Changes in generating "CLI credentials not found" discovery error" on page 24
- "HP Blade System discovery" on page 24
- "Cisco UCS Enhancement" on page 33
- "Protect trap receiver against uncontrolled memory growth for IP" on page 35
- "Changes to oidInfo.conf file and its location after an upgrade" on page 36
- "Device type changes due to recertifications" on page 36
- "Support for Cisco IOS 15.0" on page 36

Changes in discovery Rogue Access Points

CQ 613263

There was no parameter in IP Domain Managers to control the discovery of Rogue Access Points. This resulted in discovery and alerts of unwanted Rogue Access Points on Cisco wireless controllers. The EnableRogueAPDiscovery parameter was introduced in the ‹BASEDIR›/IP/smarts/conf/discovery/tpmgr-param.conf file to control Rogue Access Points discovery. By default, the value is set to TRUE. Set it to FALSE if you would like to stop discovery of Rogue Access Points.

Changes in generating "CLI credentials not found" discovery error

CQ 614051

By default, the IP Domain Manager generates the following discovery error message for devices that are discovered via CLI:

```
CLI credentials not found. Please provide the same in Device Access
tab of Polling and Thresholds
```

In some cases, users are satisfied with the discovery results based on SNMP only and do not wish to provide CLI credentials. The DisableCLIError parameter was introduced in the ‹BASEDIR›/IP/smarts/conf/discovery/tpmgr-param.conf file to prevent unnecessary discovery error alerts. The default value is TRUE. Set the parameter to FALSE if you would like to stop generating these discovery errors.

HP Blade System discovery

IP has been enhanced to support discovery and monitoring of HP Blade System C7000. The Blade System is modeled as ComputeFabric, similar to Cisco UCS. New model classes like Blade and ComputeFabricManager are introduced.

The Blade System C7000 provides the power, cooling, and I/O infrastructure needed to support modular server, interconnect, and storage components. The HP Blade system comprises of the following:

- HP Onboard Administrator — The OA or Onboard Administrator is the management entity of the HP Blade System.

  NOTICE

  The HP Onboard Administrator module provides the MIB details to manage the Compute Fabric and Blade Enclosure Manager. Both the Compute Fabric and the Blade Enclosure Manager are created using the ComputeFabricManager name.

  The Onboard Administrator is located in the enclosure and performs the functions listed here. A second optional OA provides redundancy to these functions:

  - Component insertion and removal detection
  - Component identification and connectivity
  - Power and cooling management
  - Components control

- HP Blade System C7000 Chassis
- Blade Server

Both SNMP and CLI are used for discovery.

NOTICE

CLI credentials of the Onboard Administrator (Active) is a prerequisite for complete discovery. You can configure the CLI credentials of the Onboard Administrator in **Domain Manager Administrative Console** › **Polling and Threshold** › **Device Access** tab.

Components like ComputeFabric, BladeEnclosureManager, Blade, Fan, Power Supply, and ComputeFabricManager are discovered through SNMP polling of the Onboard Administrator.

CLI information is used to find enclosure components like VirtualConnect, Switches, and Hosts installed on Blades and build relationships with the ComputeFabric.

The Onboard Administrator's redundancy is supported by creating a redundancy group and defining related events like SwitchOver, SwitchOverFailed, AtRisk, ReducedRedundancy, and so on.

The discovery of the HP Blade system involves the mapping of the HP Blade System NetConnectors to the HP Blade System. In the current implementation, the following switches are considered as HP Blade System NetConnector.

- HP VC-Ethernet Module — System ObjectID .1.3.6.1.4.1.11.5.7.5.1
- HP VC-FC Module — System ObjectID .1.3.6.1.4.1.3873.1.7
- HP VC-FlexFabric Module — System ObjectID .1.3.6.1.4.1.11.5.7.5.1
- Cisco Catalyst Blade Switches (Ethernet module), Model: 3120X, 3120G, 3200 — System ObjectID .1.3.6.1.4.1.9.1.748 or .1.3.6.1.4.1.9.1.948

- Cisco MDS Switches — System ObjectID .1.3.6.1.4.1.9.12.3.1.3.606
- BrocadeSilkWormSwitch
  - System ObjectID — .1.3.6.1.4.1.1588.2.1.1.43
  - System ObjectID — .1.3.6.1.4.1.1588.2.1.1.72)

Discovery of Blade System c7000 Enclosure (ComputeFabric)

The Blade System c7000 Enclosure is represented by the class *ComputeFabric* in IP topology.

It is defined by the following relationships:

- *ComposedOf* relationship with Card
- *ManagedBy* relationship with ComputeFabricManager
- *PackagedCooledBy* relationship with Fan
- *PackagesSystem* relationship with BladeEnclosureManager, ComputeFabricManager, FibreChannelSwitch and Switch
- *PackageTempMeasuredBy* relationship with TemperatureSensor

To manage the HP Blade System C7000, add the active Onboard Administrator through the **Add Agent** option in **Domain Manager Administrator Console**. After adding the device name or IP address, IP Availability Manager discovers it.

NOTICE

IP Availability Manager discovers the Onboard Administrator as ComputeFabricManager.

NOTICE

Provide the active OA management IP address with the corresponding user credentials in the CLI Access Group, when discovering the HP Blade System C7000 equipment. IP requires a read only CLI access to OA and privileged access need not be provided.

Discovery of the HP Blade System C7000 Chassis (BladeEnclosureManager)

IP Availability Manager discovers and represents the HP Blade System C7000 Chassis as the *BladeEnclosureManager* in the topology.

A HP Blade System C7000 chassis can have up to 16 blades. These blades connect to Interconnects through the Midplane.

When you discover the HP Blade System C7000 chassis, IP discovers the following components present in the HP Blade System C7000 entity:

- Physical blades
- Power supply
- Fan
- Temperature Sensor

The class BladeEnclosureManager is defined by the following relationships:

- *ComposedOf* relationship with Blades, Fan, and PowerSupply.
- *SystemPackagedIn* relationship with ComputeFabric.

NOTICE

IP uses the following CLI commands during HP Blade discovery:
SHOW EBIPA
SHOW SERVER PORT MAP ALL

The status of Blades, Power Supply, and Fans are fetched from the Onboard Administrator through SNMP requests. Other components are monitored through the respective device SNMPAgents.

Discovery of HP OnBoard Administrator (ComputeFabricManager)

IP Availability Manager discovers the HP OnBoard Administrator as *ComputeFabricManager.* The *ComputeFabricManager* is defined by the following relationships:

- *ComposedOf* relationship with Interface, Memory and Processor
- *ConnectedVia* relationship with Cable and IPNetwork
- *HostsAccessPoints* relationship with IP and MAC
- *HostsServices* relationship with SNMPAgent
- *ListsOfHostsAccessPoints* relationship with IP and MAC
- *Manages* relationship with ComputeFabric
- *MemberOf* relationship with Partition
- *NeighboringSystems* relationship with Router and Switches
- *PartOf* relationship with ComputeFabricManagerRedundancyGroup
- *SystemPackagedIn* relationship with ComputeFabric

Discovery of Blade

A *Blade* represents a compact, self-contained server that contains core processing components, hard-drives, memory, network cards that fit into an enclosure with other blade servers. These do not have their own power supplies or cooling solutions (including fans) as these are supplied by their enclosure.

Blade is defined by *PartOf* relationship with BladeEnclosureManager.

Discovery of ComputeFabricManagerRedundancyGroup

The *ComputeFabricManagerRedundancyGroup* is a specialized redundancy group comprising of more than one ComputeFabricManagers.

The ComputeFabricManagerRedundancyGroup is defined by the *ComposedOf* relationship with ComputeFabricManager.

NOTICE

If credentials are not provided, the device is placed on the pending list with a discovery error.

Root-cause analysis for HP Blade System C7000

The root-cause problem diagnosed for Blade, BladeEnclosureManager, Switch and ComputeFabricManager are listed in Table 1 on page 28, Table 2 on page 28, Table 3 on page 28 and Table 4 on page 28 respectively.

Table 1 Root-cause problem diagnosed for Blade

Root-cause problem	Events on Blade	Events on Host	Impacted system
Blade::Down	• Unresponsive • MightBeDown • Down	• Host::Unresponsive • Host::MightBeDown	• Blade • Host (on which the Blade is installed)
BladeEnclosureManager::Down	• Unresponsive • MightBeDown	• Host::Unresponsive • Host::MightBeDown	• BladeEnclosureManager • Blades • Hosts installed on each blade

Table 2 Root-cause problem diagnosed for BladeEnclosureManager

Root-cause problem	Description	Events on Blade	Impacted systems
BladeEnclosureManager::Down	Indicates that all the blades included in the BEM are down.	• Unresponsive • MightBeDown	• BladeEnclosureManager • Blades present on it • Hosts installed on each blade.

Table 3 Root-cause problem diagnosed for Switch

Root-cause problem	Description	Events	Impact
Switch::Down	Indicates that the switch is down.	• Blade::Unresponsive • Host::Unresponsive	• Blade • Host

Table 4 Root-cause problem diagnosed for ComputeFabricManager

Root-cause problem	Description	Events	Impact
ComputeFabricManager::Down	Indicates that the Onboard Administrator discovered as ComputeFabricManager is down.	ComputeFabricManager::Unresponsive	ComputeFabricManager

Root Cause and impact scenarios for ComputeFabricManager

Let the two ComputeFabricManager in the RedundancyGroup be ComputeFabricManager1 and ComputeFabricManager2.

Scenario 1: ComputeFabricManager is NOT part of the RedundancyGroup

Table 5 on page 29 lists the root-cause and impact scenarios when the ComputeFabricManager is NOT part of the Redundancy Group.

Table 5 Root-cause and impact analysis for ComputeFabricManager

Scenarios					Root-cause and impact analysis		
Compute Fabric Manager1	**Compute Fabric Manager2**	**Blade**	**Blade Composed Of Host**	**Host**	**Problem**	**Events**	**Impacted classes**
Unresponsive	NA	Unresponsive	No	NA	Compute FabricManager 1::Down	Unresponsive	Blade
						Unresponsive	Compute FabricManager1
Unresponsive	NA	Unresponsive	Yes	Responsive	Compute FabricManager 1::Down	Unresponsive	Blade
						Unresponsive	Compute FabricManager1
Unresponsive	NA	Unresponsive	Yes	Unresponsive	Compute FabricManager 1::Down Blade::Down **Note:** This occurs when all the blades are not unresponsive.	Unresponsive	Compute FabricManager1
						Unresponsive	Blade
						Unresponsive	Host
Unresponsive	NA	Unresponsive	Yes	Unresponsive	Compute FabricManager 1::Down BladeEnclosureManager::Down **Note:** This occurs when all the blades are unresponsive.	Unresponsive	Blade
						Unresponsive	Compute FabricManager1
						Unresponsive	Host
						Unresponsive	Blade Enclosure Manager

Scenario 2: ComputeFabricManager is a part of the RedundancyGroup

Table 6 on page 30 lists the root-cause and impact analysis of the ComputeFabricManager, when it is part of the RedundancyGroup.

Table 6 Root-cause and impact analysis for ComputeFabricManager (page 1 of 2)

Scenarios					Root-cause and impact analysis		
Compute Fabric Manager1	Compute Fabric Manager2	Blade	Blade Composed Of Host	Host	Problem	Events	Impacted classes
Unresponsive	Unresponsive	Unresponsive	NO	NA	ComputeFabric Manager1::Down ComputeFabric Manager2::Down	Unresponsive	Blade ComputeFabric Manager1 ComputeFabric Manager2
Unresponsive	Unresponsive	Unresponsive	Yes	Responsive	ComputeFabric Manager1::Down ComputeFabric Manager2::Down	Unresponsive	Blade ComputeFabric Manager1 ComputeFabric Manager2
Unresponsive	Unresponsive	Unresponsive	Yes	Un responsive	ComputeFabric Manager1::Down ComputeFabric Manager2::Down Blade::Down **Note:** This occurs when all the blades are not unresponsive.	Unresponsive	Blade ComputeFabric Manager1 ComputeFabric Manager2 Host
Unresponsive	Unresponsive	Unresponsive	Yes	Un responsive	ComputeFabric Manager1::Down ComputeFabric Manager2::Down Blade::Down BladeEnclosure Manager::Down **Note:** This occurs when all the blades are unresponsive.	Unresponsive	Blade ComputeFabric Manager1 ComputeFabric Manager2 Host BladeEnclosureManager
Unresponsive	Responsive	Unresponsive	Yes	Un responsive	ComputeFabric Manager1::Down Blade::Down	Unresponsive	Blade ComputeFabric Manager1 Host

Table 6 Root-cause and impact analysis for ComputeFabricManager (page 2 of 2)

Scenarios					Root-cause and impact analysis		
Compute Fabric Manager1	Compute Fabric Manager2	Blade	Blade Composed Of Host	Host	Problem	Events	Impacted classes
Unresponsive	Responsive	Unresponsive	No	NA	ComputeFabric Manager1::Down Blade::Down	Unresponsive	Blade ComputeFabric Manager1
Unresponsive	Responsive	Unresponsive	No	NA	ComputeFabric Manager1::Down BladeEnclosure Manager::Down **Note:** This occurs when all the blades are unresponsive.	Unresponsive	Blade ComputeFabric Manager1 BladeEnclosureManager

Map representation

Figure 1 on page 31 depicts the HP Blade System C7000 support – Topology view.



Figure 1 HP Blade System C7000 support–Topology view

Figure 2 on page 32 depicts the HP Blade System C7000 support – Map view.



Figure 2 HP Blade System C7000 support – Map view.

Figure 3 on page 32 depicts the HP Blade System C7000 support – Map view (Expanded).



Figure 3 HP Blade System C7000 support – Map view (Expanded)

Cisco UCS Enhancement

Blade Support

In this release, support is added to discover and monitor Blades in the Cisco UCS. The hosts installed on Blades can be accessed through the ComposedOf relationship of the Blade.

The Blade Down event will be generated based on the status provided by the Cisco UCS agent and the status of the Host which is installed on the Blade.

Discovery of Blades on a BladeEnclosureManager

The Blade servers present on the BladeEnclosureManager are represented as the class *Blade* in the IP topology. A Blade is a compact, self contained server containing core processing components, hard-drives, memory, and network cards that fit into an enclosure with other Blade servers.

IP 9.2.2 supports the discovery and monitoring of the Blade present on a BladeEnclosureManager.

The Blade is defined by the following relationship:

- *PartOf* relationship with BladeEnclosureManager. The BladeEnclosureManager has a *ComposedOf* relationship with the Blade.

Table 7 on page 33 lists the MIB Object IDs used for discovery and monitoring of the Blade instance on a BladeEnclosureManager.

Table 7 MIB OIDS used for discovery and monitoring of Blade

Purpose	Attribute in IP	MIB attribute type	MIB Object ID
Discovery	Blade Instance	Octet String	cucsComputeBladeDn 1.3.6.1.4.1.9.9.719.1.9.2.1.2
	Display Name	Octet String	cucsComputeBladeName 1.3.6.1.4.1.9.9.719.1.9.2.1.2
	Model number	Octet String	cucsComputeBladeModel 1.3.6.1.4.1.9.9.719.1.9.2.1.32
	Serial number	Octet String	cucsComputeBladeSerial 1.3.6.1.4.1.9.9.719.1.9.2.1.47
	Vendor	Octet String	cucsComputeBladeVendor 1.3.6.1.4.1.9.9.719.1.9.2.1.53
Monitoring	Admin state	Integer	cucsComputeBladeAdminState 1.3.6.1.4.1.9.9.719.1.9.2.1.5
	Oper state	Integer	cucsComputeBladeOperState 1.3.6.1.4.1.9.9.719.1.9.2.1.42

Root-cause analysis for Blade on a BladeEnclosureManager

When a Blade fails, the host which resides on the Blade also goes down. The root-cause analysis for a Blade failure scenario is listed in Table 8 on page 34.

Table 8 Events generated for Blade

Problem	Events (symptom) of the problem	Description	Impacted system
Blade::Down	Host::Unresponsive	Indicates that the system is not responding.	• Blade • Host
	Blade::Unresponsive	Indicates that the system is not responding.	

Map representation

Figure 4 on page 34 depicts the Cisco UCS enhancement – Topology view.



Figure 4 Cisco UCS enhancement – Topology view

Figure 5 on page 35 depicts the Cisco UCS enhancement – Map view.



Figure 5 Cisco UCS enhancement – Map view

Protect trap receiver against uncontrolled memory growth for IP

CQ 644557 /CQ 644587

When the trap receiver is flooded with traps, the trap receiver would crash, citing out-of-memory issue. This issue can be prevented by setting the values for the QUEUE_LIMIT_MEGS and QUEUE_LIMIT_SECONDS parameters in the trapd.conf file to 200 MB and 480 seconds, respectively.

NOTICE

By default, QUEUE_LIMIT_MEGS and QUEUE_LIMIT_SECONDS are commented out in the trapd.conf file. If trap flooding occurs in your environment, it is recommended to have values set for the above attributes depending on your environment. Verify the Trap Queue size by running sm_monitor -queue and then fine tune the values based on your environment. By having values set for the above attributes, traps will be discarded in trap flooding situations.

Additionally, the following four forward statements have been added to the trapd.conf file:

FORWARD: * .1.3.6.1.4.1.733.10 6 1 host:port

FORWARD: * .1.3.6.1.4.1.733.10 6 2 host:port

FORWARD: * .1.3.6.1.4.1.733.10 6 3 host:port

FORWARD: * .1.3.6.1.4.1.733.10 6 4 host:port

Where:

- Host is the hostname of the machine or IP address where the trap adapter attached to the OI server is running.

- Port is the port number where the trap adapter attached to the OI server is listening.

Traps .1.3.6.1.4.1.733.10 6 1 and .1.3.6.1.4.1.733.10 6 2 are sent by the OI trap receiver where they are converted into notifications in OI and then propagate to the SAM Console when it reaches the thresholds set for the QUEUE_LIMIT_MEGS or the QUEUE_LIMIT_SECS parameters.

Traps .1.3.6.1.4.1.733.10 6 3 and .1.3.6.1.4.1.733.10 6 4 are sent by the sm_trapd queue when it determines that certain sources are flooding it with traps.

On reaching the limits set by these parameters, the sm_trapd queue starts discarding traps and a notification *trapdDiscardingTraps* is sent to the Service Assurance Manager. The alert is cleared once the incoming trap rate to sm_trapd reduces to a rate below which these traps can be processed by the queue.

Use the sm_edit utility to change the values for these parameters in the trapd.conf file, based on your requirement.

Changes to oidInfo.conf file and its location after an upgrade

CQ 641616 /CQ 644111

The OIdInfo tool is a troubleshooting tool that provides detailed information about MIBs used by the IP domain to discover and monitor devices in the topology. The *Oidinfo.conf* file is the configuration file for the OIDInfo tool. To get the SNMP details of the entity, specify the SYSTEMCLASS, SYSTEMINSTANCE, TYPE, ENTITY, and the INSTANCE of the entity.

In earlier releases, “Type” was set to Mandatory in the oidInfo.conf file. In this release, the “Type” is changed to Optional and the default value is changed to Monitoring.

CQ 649682

After an upgrade from IP 9.2.0.0 to IP 9.2.2, the oidInfo.conf file is available at conf/discovery/ and also at conf/oidInfo/. In earlier releases, the file available at conf/discovery/ was used but from the Smarts 9.2.2 release, the file at conf/oidInfo/ is used for running the OIdInfo tool.

Support for Cisco IOS 15.0

CQ 649503

IP 9.2.2 now supports IOS 15.0 and the IP topology shows instances of temperature sensor, power supply, and fan for devices running on IOS 15.0. IP Manager cannot create these instances if a device is running on IOS versions earlier than 15.0 because the required MIBs are not present in older IOS versions. For example, Temperature sensor, Power supply, and Fan instances may not be present for devices which ran on IOS 12.0.

Device type changes due to recertifications

CQ 649699 / CQ 607333 / SR 56713712

The following changes pertain to changes in device types because of recertifications:

For Cisco C2951 (.1.3.6.1.4.1.9.1.1043), the device type is changed from Switch to Router in IP 9.2.2. Due to this change, Ports are no longer discovered for the device.

The device type for Cisco ONS-15454 (.1.3.6.1.4.1.3607.1.20.10.70) is changed from Router to Switch in IP 9.2.2. Due to this change, MIB2 IfTable entries are discovered as Interfaces (if the underlying IP is present) and Ports (for those which do not have underlying IPs).

The device type for Alcatel 7700 (.1.3.6.1.4.1.6486.800.1.1.2.1.3.1.5) is changed from Bridge to Switch in IP 9.2.2. This change enables discovery and monitoring of cards, CPU, and memory.

Features in IP Manager 9.2.1

IP Manager release 9.2.1 includes important new features and changes to the existing features:

- "Virtual Router Redundancy Protocol (VRRP) support for Aruba wireless controllers" on page 38
- "WirelessAccessPoint support enhancement" on page 44
  - "WirelessRadioInterface discovery and monitoring"
  - "Option to delay WirelessAccessPoint::Down problem"
  - "MultipleWirelessAPDown event"
- "Authentication Server support enhancement" on page 47
- "Connectivity Polling for Wireless Devices" on page 49
- "Logic for processing RouteDistinguisher for Juniper devices changed" on page 50
- "Enhancement to FibreChannelSwitch discovery" on page 50
- "OID Info tool enhancements" on page 50

NOTICE

These features are available only for Aruba Wireless Controllers.

Virtual Router Redundancy Protocol (VRRP) support for Aruba wireless controllers

IP Availability Manager Version 9.2.1 supports discovery and monitoring of wireless controllers that are configured for redundancy. The IP topology class **WirelessControllerRedundancyGroup** has been introduced to support this functionality.

WirelessControllerRedundancyGroup is a collection of two or more WirelessControllers that are configured for redundancy to manage the WirelessAccessPoints discovered by the WirelessControllers.

NOTICE

VRRP support for wireless devices is applicable only for Aruba Wireless Controllers.

Aruba OID certified for VRRP

All IP 9.1 certified WirelessController supports VRRP. One prerequisite is that the VRRP data should be available through SNMP.

Aruba OID .1.3.6.1.4.1.14823.1.1.13 is recertified as WirelessController in the Smarts 9.2.1 release.

Enable VRRP trap processing for Aruba devices

CQ 577828 / CQ 643367 / CQ 568043

The wireless VRRP feature in IP uses a Trap sent by the Wireless controller to monitor the Switchover of the WirelessAccessPoints from one WirelessController to another.

A typical setup will have a Trap Exploder receiving the Traps from devices and then forwarding them to the IP Domain Manager. In order to forward the wireless VRRP traps received from devices to the IP Domain Manager:

Use the sm_edit utility to modify the trapd.conf file in the installation running the Trap Exploder:

```
#cd <BASEDIR>/smarts/bin
#./sm_edit conf/trapd/trapd.conf
```

Add the following line to trapd.conf:

```
# Traps required by InCharge IP Domain Manager for Wireless VRRP
   Monitoring
#
FORWARD: * .1.3.6.1.4.1.14823.2.2.1.1.100.1038 6 1  host:port
```

where host is the IP address of the IP Domain Manager and the Port is the port on which the IP Domain manager is listening for Traps. The default port is 9000.

NOTICE

The trap exploder has a predefined section in trapd.conf which contains a list of traps to be sent to IP domains. This list is expanded in the Smarts 9.2.2 release to contain VRRP traps.

Relationship of WirelessControllerRedundancyGroup

WirelessControllerRedundancyGroup is defined by:

- *ComposedOf* relationship with WirelessAccessPoints.
- *LayeredOver* relationship with WirelessControllers.

Attributes of WirelessControllerRedundancyGroup

Table 9 on page 39 lists the attributes of WirelessControllerRedundancyGroup.

Table 9 WirelessControllerRedundancyGroup attributes (page 1 of 2)

Attributes	Description
ControllersList	List of WirelessControllers which are configured in redundancy mode and are part of a WirelessControllerRedundancyGroup.
GroupID	Unique identifier for the WirelessControllerRedundancyGroup
SwitchOverTime	Time duration (in seconds) that is allowed for the switch over to stabilize. The status is checked after this time interval. The time interval starts after the SwitchOver Trap is received from the WirelessController. Default: 1800 seconds. **Note:** The SwitchoverTime can be configured in **Polling and Threshold** › **Threshold** › **WirelessControllerRedundancyGroup**.
LastChangedAt_seconds	Time (in seconds) when the attribute IsSwitchover last changed.

Table 9 WirelessControllerRedundancyGroup attributes (page 2 of 2)

Attributes	Description
IsSwitchover	Indicates that the switchover between WirelessControllers has taken place for the WirelessControllerRedundancyGroup.
NumOfWAPsUpBeforeSwitchover	Indicates the total number of WirelessAccessPoint objects that were UP before the Switchover Trap is received.
NumOfWAPsDownAfterSwitchover	Indicates the total number of WirelessAccessPoint objects that went down after the Switchover Trap is received.
TotalNumberOfWAPs	Indicates total number of WirelessAccessPoint objects in this WirelessControllerRedundancyGroup.

Events of WirelessControllerRedundancyGroup

Table 10 on page 40 lists the events that are generated for WirelessControllerRedundancyGroup.

Table 10 WirelessControllerRedundancyGroup events

Events	Description
SwitchOver	Indicates the occurence of switchover (VRRP). **Note:** The **SwitchOver** event is generated whenever a SwitchOver Trap is received from a WirelessController in IP Availability Manager. The event remains active for the time specified against the SwitchoverTime, and clears automatically after that. **Note:** The SwitchoverTime can be configured in Polling and Threshold.
SwitchOverFailed	Indicates the failure of a switch over. After receiving the VRRP SwitchOver trap from the Controller, the active wireless APs are identified. After the configured SwitchOverTime has elapsed, the status of Wireless APs will be reverified for the failures. If any one Wireless AP's status was Down which was UP before, the SwitchOverFailed event will be generated.
AtRisk	Indicates that the number of operational members in the redundancy group is less than or equal to the AtRiskThreshold value (set to 1 by default) but is not zero. **Note:** If the number of operational members is zero, an AllComponentsDown event occurs. For example, assuming AtRiskThreshold = 1, the state of a three-member redundancy group having one operational member (and two non-operational members) is AtRisk.
ReducedRedundancy	Indicates that the number of operational members (WirelessControllers) in the redundancy group is greater than the AtRiskThreshold value, but is less than the total number of members in the redundancy group. For example, and assuming AtRiskThreshold = 1, the state of a three-member redundancy group having two operational members (and one non operational member) is ReducedRedundancy.
AllComponentsDown	Indicates that none of the member objects (WirelessControllers) in the redundancy group are operational.

Threshold settings

Table 11 on page 41 lists the threshold settings to configure the attribute **SwitchOverTime**. This attribute defines the length of time (in seconds) that the Switchover is allowed to stabilize, and status checked after this time interval. The Time Interval starts after the SwitchOver Trap is received from the WirelessController.

You can set the threshold settings for this attribute in **Polling and Threshold** › **WirelessControllerRedundancyGroups** › **Settings**.

Table 11 WirelessControllerRedundancyGroup threshold settings

Threshold parameter	Allowed values	Description
SwitchOverTime	0 – *9999* Default: 1800 (in seconds)	This is the time interval (in seconds) after receiving the trap, when the FailOver status will be verified.

VRRP SwitchOver trap information

Trap Syntax :‹trap-oid generic specific› ‹Uptime› ‹Varbind - VRRPOperID› ‹type› ‹Value› ‹Varbind - VRRPMasterIP› ‹Type› ‹Value› ‹Varbind - VRRPOperstate› ‹type› ‹Value›

Example Trap varbind: .1.3.6.1.4.1.14823.2.2.1.1.100.1038 6 1 5000 .1.3.6.1.4.1.14823.2.2.1.1.100.100.19.0 i 1 .1.3.6.1.4.1.14823.2.2.1.1.100.100.20.0 a 10.31.9.229 .1.3.6.1.4.1.14823.2.2.1.1.100.100.21.0 i 3

Representation of WirelessControllerRedundancyGroup in Map Console

Figure 6 on page 42 depicts the map of a WirelessControllerRedundancyGroup.



Figure 6 Map representation of WirelessControllerRedundancyGroup

Table 12 on page 42 explains the icons used in the map and their relationships.

Table 12 Description of WirelessControllerRedundancyGroup map

Icon	Representation
	WirelessControllerRedundancyGroup
	WirelessController
	WirelessAccessPoints

Table 12 Description of WirelessControllerRedundancyGroup map

Icon	Representation
	ComposedOf relationship with WirelessAccessPoints
	LayeredOver relationship with WirelessController

Representation of WirelessControllerRedundancyGroup in Topology Browser

Figure 7 on page 43 shows the WirelessControllerRedundancyGroup viewed in Topology Browser.



Figure 7 WirelessControllerRedundancyGroup in Topology Browser view

WirelessAccessPoint support enhancement

The WirelessAccessPoint support has been enhanced with the following:

- Discovery and monitoring support for wireless radio
- Delay APDown event
- Root cause and impact analysis for multiple Access Point down

WirelessRadioInterface discovery and monitoring

The support for WirelessAccessPoint has been enhanced to discover and monitor wireless radio, which is a component of the Wireless Access Point. The wireless radio on the WirelessAccessPoint is discovered as **WirelessRadioInterface** by IP Availability Manager. The user is alerted when the number of wireless clients per radio interface on the WirelessAccessPoint exceed the threshold.

Relationship of WirelessRadioInterface

The WirelessRadioInterface is defined by *PartOf* relationship with the WirelessAccessPoint and the WirelessAccessPoint is defined by *ComposedOf* relationship with the WirelessRadioInterface.

Object IDs probed for discovery

Table 13 on page 44 lists the Object IDs that are probed for discovery of wireless radio.

Table 13 Object IDs probed

Object IDs	Use
.1.3.6.1.4.1.14823.2.2.1.5.2.1.5.1.3	Used for the creation of WirelessRadioInterface.
.1.3.6.1.4.1.14823.2.2.1.5.2.1.5.1.2	Used for retrieving the type of radio interface.
.1.3.6.1.4.1.14823.2.2.1.5.2.1.5.1.7	Used for retrieving the number of WirelessClients per radio.
.1.3.6.1.4.1.14823.2.2.1.5.2.1.5.1.6	Used for retrieving the current utilization of radio interface.

Attributes of WirelessRadioInterface

Table 14 on page 44 lists the attributes of the class WirelessRadioInterface.

Table 14 WirelessRadioInterface attributes

Attributes	Description
AttachedWirelessClientThreshold	Indicates the maximum number of WirelessClient that can be associated with this WirelessRadioInterface at a given time.
HighUtilizationThreshold	Upper threshold for interface utilization expressed as a percentage of the total bandwidth.
NumberOfWirelessClientsOnRadio	Indicates the number of WirelessClient currently associated with the radio.
CurrentUtilization	Indicates the current utilization of the WirelessRadioInterface as provided by its WirelessController.

Events for WirelessRadioInterface

Table 15 on page 45 lists the events that are generated for classWirelessRadioInterface.

Table 15 WirelessRadioInterface events

Events	Description
WirelessClientsThresholdReached	Indicates that the maximum number of WirelessClients that can be connected to the WirelessRadioInterface has reached.
HighUtilization	Indicates that the utilization of the WirelessRadioInterface is greater than the specified High Utilization Threshold.

Threshold settings of WirelessRadioInterface

Table 16 on page 45 lists the threshold parameters of the threshold group WirelessRadioInterface.

Table 16 WirelessRadioInterface threshold settings

Threshold parameters	Allowed values	Description
MaxAttachedWirelessClientThreshold	0 – *99999* Default: 30	Maximum number of WirelessClients that can be attached with the WirelessRadioInterface at a given point in time, which when exceeded will generate the **WirelessClientThresholdReached** event.
UtilizationThreshold	0 – 100	Upper threshold for radio interface utilization expressed as a percentage of the total bandwidth, which when exceeded will generate the **HighUtilization** event.

You can set the threshold parameters in **Polling and Threshold** console › **WirelessRadioInterface Group** › **WirelessRadioInterface** › **Settings** › **WirelessRadioInterface Threshold.**

Option to delay WirelessAccessPoint::Down problem

The AccessPoint::Down problem has been customized with the option to delay the generation of the event for a specified time. The time can be specified in **Polling and Threshold** console › **Threshold** › **System Resource Group** › **WirelessAccessPoint** › **MaxAPDownDelayTime**. The default value is set to 0, which means that when the access points go down, then the event for AccessPoint::Down problem will be generated instantly.

For example, if the **MaxAPDownDelayTime** is set to 1800 seconds, the event for AccessPoint::Down problem will be generated only after half an hour.

Table 17 on page 46 lists the new attributes of WirelessAccessPoints.

Table 17 New WirelessAccessPoints attributes

Attributes	Allowed values	Description
MaxAPDownDelayTime Threshold	0 – *99999* (in seconds)	Maximum threshold time for which Down problem will be supressed for the WirelessAccessPoint.
ExceededAPDownDelayTimeThreshold	TRUE, FALSE	A value of TRUE indicates that: • The Down problem has occurred, and • The time set against **MAXAPDownDelayTime** has elapsed.
LastChangedAt_seconds	Unsigned integer	Time when the adapter last changed state. This typically represents the time at which this application monitored the change in state. It does not necessarily represent the actual time that the state changed.

MultipleWirelessAPDown event

The event **MultipleWirelessAPDown** has been introduced to alert the user only when the configured number of WirelessAccessPoints go down for a specified time duration.

The configuration can be set in **Polling and Threshold** › **Threshold** › **System Resource Group** › **WirelessController**. The parameters are **MaxAPsDownPerController** and **MaxMultipleAPDownDelayTimeThreshold**. Table 20 on page 47 lists the details of these threshold parameters and settings.

Attributes of WirelessController

Table 18 on page 46 lists the new attributes of WirelessController.

Table 18 WirelessController attributes

Attributes	Description
NumberOfAPDownThreshold	Maximum number of AccessPoints that can be down simultaneously on a Controller at a given time.
NumberOfDownAPs	Number of AccessPoints that are down on the Controller at a given time.
IsMultipleAPDown	Indicates whether Multiple AP went down and crossed the threshold.
LastChangedAt_seconds	The time at which the adapter last changed state. This typically represents the time at which this application monitored the change in state. It does not necessarily represent the actual time that the state changed.

Events of WirelessController

Table 19 on page 46 lists the new event generated for WirelessController.

Table 19 WirelessController event

Event	Description
MultipleWirelessAPsDown	Indicates that the Number of APs Down on this Controller has exceeded the Threshold.

Threshold settings for WirelessController

Table 20 on page 47 lists the threshold settings for WirelessController. These threshold settings determine the generation of the **MultipleWirelessAPDown event.**

Table 20 WirelessController Thresholds settings

Threshold setting	Default Value	Description
MaxAPsDownPerController	100	The maximum number of APs down in a controller.
MaxMultipleAPDownDelayTime	720	Maximum length of time that the Max number of AccessPoints that can be down simultaneously on this controller

Authentication Server support enhancement

The Authentication Server support has been enhanced to alert the user about authentication failures and authentication timeouts. Two new events **FailureThresholdReached** and **TimeOutThresholdReached** have been introduced to support the functionality. Table 22 on page 48 lists the details of these events.

Attributes of AuthenticationServer

Table 21 on page 47 lists the new attributes of AuthenticationServer.

Table 21 Authentication Server attributes (page 1 of 2)

Attributes	Description
NumberOfAuthenticationFailures	Indicates the number of times the connection from the Client to the Authentication Server failed. **Note:** This attribute is available in instrumentation class of the AuthenticationServer.
NumberOfAuthenticationTimeouts	Indicates the number of times the timeouts occurred when the Client were trying to connect to the Authentication Server. **Note:** This attribute is available in instrumentation class of the AuthenticationServer.
NumberOfAuthenticationRequests	Indicates the number of times the authenticaiton requests were sent by the Clients to the Authentication Server. **Note:** This attribute is available in instrumentation class of the AuthenticationServer.
DeltaOfAuthenticationTimeouts	Indicates the delta of AutheticationTimeout value between two consecutive polling cycles. **Note:** This attribute is available in instrumentation class of the AuthenticationServer.
AuthenticationFailureThreshold	The upper threshold for authentication failures expressed as a percentage of the total number of authentication requests. This is the threshold for generating the event FailureThresholdReached event.

Table 21 Authentication Server attributes (page 2 of 2)

Attributes	Description
AuthenticationTimeOutThreshold	This serves as the criteria for generating TimeOutEvent. Maximum number of authentication Timeouts between two polling cycle, beyond which the TimeOutThreshold event will be generated.
FailedAuthenticationsCount	Lists the controllers that have crossed the failure threshold and their failed authentication count.
AuthenticationRequestsCount	Lists the controllers that have crossed the failure threshold and their usage count.
AuthenticationTimeOutCount	Lists the controllers that have crossed the timeout threshold and their Timeout values.
DeltaOfAuthenticationTimeOut	Lists the controllers that has crossed the timeout threshold and their TimeoutDelta values.
FailureReached	If TRUE, indicates that the criteria for generating FailureEvent has reached.
TimeOutReached	If TRUE, indicates that the criteria for generating Timeout event has reached.

Events of AuthenticationServer

Table 22 on page 48 lists the new events generated for AuthenticationServer.

Table 22 Authentication Server events

Events	Description
FailureThresholdReached	This event is generated when: • (NumberOfAuthenticationFailures › ((AuthenticationFailureThreshold*NumberOfAuthenticationRequests)/100)) and • NumberOfAuthenticationRequests › 50
TimeOutThresholdReached	This event is generated when the difference between two consecutive polling cycle is greater than zero.

Threshold settings of AuthenticationServer

You can set the threshold settings in **Polling and Threshold › Threshold › System Resource Group › AuthenticationServer › AuthenticationServer Thresholds.**

Table 23 on page 48 lists the threshold settings for AuthenticationServer.

Table 23 AuthenticationServer threshold settings

Threshold parameter	Default values	Description
AuthenticationFailureThresholds	0 – 100 Default: 50	Maximum percentage of the authentication failures, which when exceeded will generate the FailureThresholdReached event.
AuthenticationTimeOutThreshold	0 – *99999* Default: 0	Maximum number of authentication timeouts between two consecutive polling cycles, which when exceeded will generate the TimeOutThresholdReached event.

Connectivity Polling for Wireless Devices

The connectivity polling settings for polling groups have been segregated between wireless domain devices and other wired devices. Two new connectivity polling settings have been exclusively introduced for wireless devices:

- Connectivity Polling for Wireless Devices
- Connectivity Polling for Wireless Devices — External Poller

For all wireless devices such as WirelessControllers, AuthenticationServers, WirelessAccessPoints, WirelessClient and so on, add these settings in **Polling and Thresholds › Polling Groups › Connectivity Polling For Wireless Devices setting › Settings**.

The "Connectivity Polling for Wireless Devices" and "Connectivity Polling settings" for other devices are mutually exclusive. The Connectivity Settings do not instrument the wireless devices. Use the **Connectivity Polling settings for Wireless Devices** to instrument the wireless devices only.

NOTICE

If you are upgrading and your existing topology contains Wireless Devices, replace the current Connectivity Polling setting with the Connectivity Polling for Wireless Devices setting in the respective polling groups in the Polling and Threshold Console. This has to be done manually as IP Availability Manager version 9.2.1 does not overwrite the existing settings in the old RPS file. "CQ 603145 / CQ 603416 / CQ 643403" on page 49 provides instructions on how to run the ic_correctSettings.asl script to correct the settings for the default polling and threshold groups.

Wireless topology is not monitored after an upgrade

CQ 603145 / CQ 603416 / CQ 643403

When you upgrade from an older version to IP 9.2.1, Polling settings for the Wireless Devices is not updated and hence monitoring does not happen for these devices. The settings need to be manually updated for the Wireless groups. This is applicable only for Wireless groups including WirelessController, WirelessAccessPoint, WirelessClient, and AuthenticationServer.

A new polling setting named as "Connectivity Polling for Wireless Devices" is added in IP 9.2.1 which is applicable for Wireless Devices classes. This setting needs to be applied to all wireless devices in the polling groups to manage the wireless entities.

Remove the old Connectivity Polling setting and add the new "Connectivity Polling for Wireless Devices" for Wireless Polling groups. This can be done by running the script ic_correctSettings.asl on the upgraded server. The script changes the "Connectivity Polling" setting to "Wireless Connectivity Polling" setting for the WirelessController, WirelessAccessPoint, and WirelessClient. For the Authentication Server group, the settings need to be applied manually. This script is applicable for post IP 9.0 upgrade scenarios. For IP 9.0 and lower version upgrade/migrate scenarios, the settings need to be updated manually.

Script usage

```
sm_adapter -s ‹IP server› -DCORRECTSETTINGS=TRUE ../rules/utils/ic_correctSettings.asl
```

NOTICE

Post upgrade and migrate, the priority of the Wireless groups need to be moved above the "Other System" groups.

NOTICE

The ASL script which was used in the Smarts 9.2.1 release for enabling the Wireless Connectivity Polling setting for Wireless Devices in case of an upgrade, may no longer be required in the Smarts 9.2.2 release. It is recommended not to use ic_correctSettings.asl for the 9.2.2 release.

Logic for processing RouteDistinguisher for Juniper devices changed

CQ 567883/SR 53199934/CQ 642413

In the earlier release, RouteDistinguisher values were not calculated properly for devices with decimal and hexadecimal RouteDistinguisher values.

In this release, the logic in ic-juniper-ip-tag.asl is changed such that the RouteDistinguisher values are handled properly for devices having decimal values.

Enhancement to FibreChannelSwitch discovery

The FibreChannelSwitch is discovered by EMC Server Manager from ProSphere along with associated FC Ports, Port links, Fabrics, and SAN connectivity.

FibreChannel switches has been recertified as FibreChanelSwitch in IP. These instances are now displayed under FibreChannelSwitch class instead of Switch class.

The details of the implementation, and the contrast with the earlier implementation are listed in Table 24 on page 50.

Table 24 FibreChannelSwitch discovery

Current implementation	Previous implementation
Service Assurance Manager has only one instance of the FibreChannel switch under the topology class FibreChannelSwitch.	Service Assurance Manager had two instances of the FibreChannel switch under two different classes, namely FibreChannelSwitch and Switch.
Additional attributes define the FibreChannelSwitch, namely AgentAddressList and IPAddress	These attributes were absent in the earlier implementation.

OID Info tool enhancements

- "Running multiple instances of the OID Info tool --input (or -i) option" on page 51
- "Using OID Info tool with Dynamic Model Classes" on page 51
- "Processing specific classes" on page 52
- "Changed output format" on page 53
- "OID Info tool dialog box changes" on page 55

Running multiple instances of the OID Info tool --input (or -i) option

In the 9.2.1 release, multiple instances of the OIDInfo tool can be run using the --input (or -i) option. This option can be used to specify an alternate configuration file to be used for running the tool. If not specified, the tool looks at the default configuration file, that is, conf/oidinfo/oidInfo.conf. The alternate configuration file should adhere to the syntax of the default configuration file. This option is beneficial if multiple instances of the tool need to be run for debugging issues with instances on different devices.

An alternate configuration file can be specified as follows:

```
sm_perl sm_oidinfo.pl -server=INCHARGE-AM --compare
--input=conf/oidinfo/oidInfo-1.conf
```

The alternate configuration file can be created using the sm_edit command as follows:

```
sm_edit --create conf/oidinfo/oidInfo-1.conf
```

Using OID Info tool with Dynamic Model Classes

The OIDInfo tool by default traverses select relationships for devices to collect instrumentation objects. If the dynamic model is used to introduce new relationships, entities or instrumentation classes, the tool configuration will need to be modified to traverse the new entities. This can be achieved by editing the configuration file conf/oidinfo/ClassRelations.conf.

The entries under the section “Common Relationships” in the ClassRelations.conf are applied to all types of system objects. Few relationships specific to the IP domain are added under the section “Specific to IP”. Any custom entries that need to be made can be added after this section.

NOTICE

Any changes to this file will need a server restart.
Changing this configuration will impact both the command line and console invocation.
Entries in this file are validated when the server is started and invalid entries will be ignored.
Class names and relationship names are case sensitive and they need to be defined using the same names as defined by the model.
Relationships need to be defined one-way only, to avoid loops.

The configuration contains the following entries by default:

```
#
# Common Relationships
#
UnitaryComputerSystem    ComposedOf
UnitaryComputerSystem    HostsAccessPoints
UnitaryComputerSystem    HostsServices
UnitaryComputerSystem    SystemPackagedIn
#
# Specific to IP
#
WirelessController    Controls
WirelessController    AuthenticationServicedBy
WirelessController    Controls->WirelessAccessPoint->Hosts
WirelessAccessPoint   Hosts
```

Entries in this configuration file need to follow the following syntax:

Column 1 should be the class name and Column 2 should be the relationship name. It is mandatory to specify a relationship name when a class name is added. The entry is ignored if not specified. Relationships can be defined as single-level if entities are defined directly under it. If multiple relationships need to be traversed to get to the entity, then the relationship need to be defined as multi-level.

The syntax for single-level relationships is as follows:

```
<ClassName>        <RelationshipName>
```

For example, the configuration contains the entry “UnitaryComputerSystem ComposedOf”. This would collect all entities under the ComposedOf relationship for system objects. The class UnitaryComputerSystem is the base class for all other system classes such as Switch, Router, Host, and so on. So these entries need not be defined for each of these classes again.

The syntax for multi-level relationships is as follows:

```
<ClassName> <RelationshipName>-><ClassName>-><RelationshipName>
```

The relationship parameter specifies the sequence of traversal to get to the entity object. This needs to be defined as a sequence of ‹RelationshipName›-›‹ClassName›. The symbol “->” is used as a separator. This sequence can terminate at a relationship name or a class name. If the entry terminates at a relationship name then all objects under that relationship are picked up. If the entry terminates at a class name then only objects of that class are picked up from the preceding relationship name.

For example, to get to a WirelessClient object from a WirelessController, we will need to traverse the following path: Controls->WirelessAccessPoint->Hosts. As this entry terminates at the relationship “Hosts”, all objects under this relationship will be picked up. The relationship “InstrumentedBy” need not be added to single/multi-level definitions. This is automatically traversed by the tool.

Processing specific classes

The OIDInfo tool can be used to process specific entity instances under a device. This will be helpful if data for a specific instance only is needed and processing of all other instances can be skipped. An instance name pattern can be defined in oidInfo.conf to filter out specific instances. This needs to be defined as the last column.

For example, if we need data for only Interface1 on a Router, the entry needs to be made as follows:

Router dev-7604.emc.com Monitoring Interface *IF*/1

The last column defines a pattern to be used for filtering out instances. This pattern is applied to instrumentation objects of the entity. In the above example, the pattern *IF*/1 is applied to instrumentation objects of Interface and the number 1 specifies the instance number of the object.

The following patterns can be applied for entities in IP:

- Interface - *IF*/‹number›
- Port - *PORT*/‹number›
- Card - *CARD*/‹number›

- Memory - *MEM*/‹number›
- Processor - *PSR*/‹number›
- PowerSupply - *PWR*/‹number›
- TemperatureSensor - *TEMP*/‹number›
- VoltageSensor - *VOLT*/‹number›
- Fan - *FAN*/‹number›
- Disk - *DISK*/‹number›
- FileSystem - *FS*/‹number›

If an entity is instrumented by multiple instrumentation classes, then the pattern can also be set on the instrumentation class name.

For example, interface objects are typically instrumented by multiple instrumentation objects of different type for availability and performance. For instance, Interface_Performance_CiscoRouter, NetworkAdapter_PerformanceData_MIB2_ifMIB, Interface_Fault_MIB2. It is possible to filter out only the fault related data as follows:

Router dev-7604.emc.com Monitoring Interface *Interface_Fault_MIB2*/1

Changed output format

The output format of the OIDInfo tool is changed to use the csv format. The tool generates the following files for every entry defined in oidInfo.conf.

- An XML file, which contains a snapshot of the configuration used to generate the output.
- A CSV file, which contains the actual output of the tool.
- A log file, which contains details of execution for debugging any tool related issues.

Each of these files is named by appending a timestamp to indicate the time of execution.

For example, if the conf file has the following entry,

Router dev-mplsp8.emc.com Monitoring Card|Processor|Memory

the tool will generate the following files under logs/oidInfoLogs:

Router_dev-mplsp8.lss.emc.com_05-15-2013-102832.csv

Router_dev-mplsp8.lss.emc.com_05-15-2013-102832.xml

sm_oidInfo-05-15-2013-102832.log

The csv file would have values as shown below:

Type: Indicates the type of the entry. Possible values are Discovery/Monitoring.

Component: For Monitoring entry types, it indicates the class name of the device entity. For Discovery entry types, it indicates the discovery probe name.

Target: For Monitoring entry types, it indicates the instrumentation class name for the Component. For Discovery entry types, it indicates the discovery driver name.

OID: Indicates the SNMP OID associated with the Target.

Index: Indicates the OID index used to run SNMP queries and would also indicate the instance number of the object in the server topology. Supported for entry type Monitoring only.

SNMPValue: Indicates the value retrieved via SNMP from the device. This column is displayed if the option --display (or -d) or --compare (or -c) or --poll (or -p) is used. Supported for entry type Monitoring only.

TopoValue: Indicates the value of the parameter in the server topology. This column will be displayed if the option --compare (or -c) is used. Supported for entry type Monitoring only.

TimeStamp: Indicates the time at which the data for the entry was collected.

Description: Indicates the textual description for the parameter. For Monitoring entry types, this gives the parameter name in the instrumentation class for the OID. For Discovery entry types, this gives the ASL file name associated with the discovery driver and an optional description for the discovery probe.

The XML file will have entries as follows:

```
<Topology>
<Device SysOID=".1.3.6.1.4.1.2636.1.1.1.2.18" Name="dev-mplsp8.emc.com">
<Entity type="Monitoring" component="Card" target="Card_Fault_Juniper">
<oid name=".1.3.6.1.4.1.2636.3.1.15.1.8" desc="jnxFruState"/>
</Entity>
<Entity type="Monitoring" component="Card" target="Card_Fault_Juniper_Redundancy">
<oid name=".1.3.6.1.4.1.2636.3.1.14.1.7" desc="jnxRedundancyState"/>
<oid name=".1.3.6.1.4.1.2636.3.1.15.1.8" desc="jnxFruState"/>
<oid name=".1.3.6.1.4.1.2636.3.1.13.1.6" desc="jnxOperatingState"/>
</Entity>
<Entity type="Monitoring" component="Memory"
   target="Memory_Performance_JuniperRouter">
<oid name=".1.3.6.1.4.1.2636.3.1.13.1.11" desc="jnxOperatingBuffer"/>
<oid name=".1.3.6.1.4.1.2636.3.1.13.1.15" desc="jnxOperatingMemory"/>
<oid name=".1.3.6.1.4.1.2636.3.1.13.1.12" desc="jnxOperatingHeap"/>
</Entity>
<Entity type="Monitoring" component="Processor"
   target="CPU_Performance_JuniperRouter">
<oid name=".1.3.6.1.4.1.2636.3.1.13.1.8" desc="jnxOperatingCPU"/>
</Entity>
</Device>
</Topology>
```

The ‹Device› tag gives the SystemObjectID and the Name of the specified device in the server topology.

The ‹Entity› tag provides configuration of discovery and monitoring entries used to generate the output. The attribute type gives the type of the entry, that is, Discovery/Monitoring. The attribute component gives the discovery probe name or the entity class name. This can be directly mapped to the Component column in the csv file. The attribute target gives the instrumentation class name or the discovery driver name. This can be mapped to the Target column in the csv file.

The number of entity tag entries can vary depending on the filtering defined in oidInfo.conf. Here the xml file contains entries for only monitoring of Card, Memory and Processor because the following pattern is used.

Router dev-mplsp8.emc.com Monitoring Card|Processor|Memory

Each entity tag can contain multiple ‹oid› tags. This gives the OIDs that are associated with a particular entity. If the entity has no OIDs, then there will be no ‹oid› tags under that entity. The attribute name gives the SNMP OID. This can be mapped to the OID column in the csv. The attribute description gives a textual description of the OID. This can be mapped to the Description column in the csv.

OID Info tool dialog box changes

The OIDInfo dialog in the SAM Admin Console is modified to address the following issues:

- Window resizing: The window can now be maximized or resized.
- Grid Scrolling: The grid view can now be scrolled horizontally and the columns can be resized.

An additional option of exporting the grid to a file (“Save Table”) is provided. The grid can now be saved by specifying a separator for the columns.

IP Manager technical notes from 9.2.x that apply to 9.3.x

- “Technical notes for IP Manager 9.2.2” on page 55
- “Technical notes for IP Manager 9.2.1” on page 60

Technical notes for IP Manager 9.2.2

QUEUE_LIMIT_MEGS and QUEUE_LIMIT_SECONDS are commented out in trapd.conf in IP

CQ 609017

By default, QUEUE_LIMIT_MEGS and QUEUE_LIMIT_SECONDS are commented out in the trapd.conf file. If trap flooding occurs in your environment, it is recommended to have values set for the above attributes depending on your environment.

Verify the Trap Queue size by running *sm_monitor -queue* and then fine tune the values based on your environment .

By having values set for the above attributes, traps will be discarded in trap flooding situations.

Exclude list functionality in the TCAT tool

CQ 567524

Certain objects or relationships can be excluded from comparison in the TCAT tool. This exclude functionality can be achieved by commenting out the unwanted elements in the TopoRulesDiff.txt file by adding “#” in front of that element.

#OBJ|Cable

#OBJ|TrunkCable

OBJ|Firewall

OBJ|IP

#OBJ|NetworkConnection

OBJ|LoadBalancer

#OBJ|AccessPoint

The TCAT tool will not process those entities or elements that are commented during comparison.

DisableAutomaticSettingsUpgrade in the discovery.conf file

CQ 649104

During server start in 9.2.2 with old repos (7.x onwards), the priority and settings of any new “Polling Groups” and “System Resource Group” are taken care automatically. You need to only modify the priority of your defined groups. This automatic setting correction will be done only once on an old RPS during first time start.

If you do not want to use this feature, you can disable it by setting the value of “DisableAutomaticSettingsUpgrade” to TRUE in the discovery.conf file. If this feature is disabled, you may not see some polling and threshold settings in the server Polling and Thresholds window by default. To get these settings, load the respective settings.conf manually. These conf files are available in the installation directory path: $SM_SITEMOD/conf/‹icf|devstat|health›/.

NOTICE

During server restart with the old repos, if “ComputeFabric” comes below “Switches” in “Polling Groups” and “System Resource Group”, then manually move the “ComputeFabric” up before “Switches” to apply the proper settings. Do not modify the default value of “DisableAutomaticSettingsUpgrade”, which is FALSE. It is recommended not to use “ic_correctSettings.asl”, as it is no more valid from the 9.2.2 release.

Different number of power supplies discovered on the same device in IP 9.2 and IP 9.2.1

CQ 577867

The device with sysOID 1.3.6.1.4.1.14823.1.1.13 is recertified as WirelessController in IP 9.2.1. Hence while running the TCAT tool, explanation for this entity that is newly discovered in IP 9.2.1 appears as “*In 9.2.1.0 (for the related Object): Originally Recertified in 9.2.1. Please refer to the Certification Matrix. Impacted Classes : Interface Port SNMPAgent IP Processor Memory Chassis*”. For already existing entities, it appears as “*Certified in an earlier release*”. Extra power supplies will get created because of this recertification. No functionality is impacted due to this behavior.

Ports and interfaces may have 0.0 as the current utilization value

CQ 636366

In earlier releases, when ports and interfaces did not have instrumentation or were not managed, the current utilization was grayed out. In IP 9.1, when Ports and interfaces do not have instrumentation, the IP server displays the currentUtilization value as 0 for those ports and interfaces.

Issues related to unusual growth in log file size

CQ 636047

If you see a huge growth in the log file size, for instance in the Broker log:

1. Re-install the Broker as service with environment setting SM_NO_LOG_CLIENTS=TRUE.

2. Restart the Broker service.

NOTICE

To disable the client logging feature subscription for all the servers or Domain Managers, put the entry in the runcmd_env.sh file, and then restart the Broker and Domain Managers.

Users with custom code that uses EMC Smarts API

If the code is based on Java, Perl, or C API, modify and retest the code according to the changes in this release. If a dynamic model is used, recompile the code.

The *EMC Smarts Dynamic Modeling Tutorial* explains the concepts and methods of dynamic modeling.

Fewer interfaces may be discovered on Hosts

CQ 574723 / CQ 574196 / CQ 542498

In IP 8.1.1, all interfaces defined in ifMIB were discovered if NIC Teaming was disabled on a Compaq device having Sys OID .1.3.6.1.4.1.311.1.1.3.1.*.

In IP 9.0 and later versions, if NIC teaming is disabled, we discover only those interfaces that have IP addresses configured on them.

SHM monitors IP standby server using the old port number after a restart

CQ 572156 / CQ 564417

Issue

If you create a standby server without registering with the Broker and mention the standby server details in the /conf/serverhealth/stdby_servers.conf file, and then perform a *Discover All* in Server Health Monitor (SHM), the standby server gets monitored in SHM.

If the standby server is now shut down and restarted, and if the new Port Number and Process ID is mentioned in the /conf/serverhealth/stdby_servers.conf file, the old standby server appears in the Pending Devices list after a *Discover All* is performed in SHM.

Resolution

When you restart the standby server, ensure that you use the *- -port* option and provide the same port number used earlier before the restart. This will prevent the earlier or the old standby server to get listed in the Pending Devices list during discovery.

Less number of processors and more interfaces may be discovered

CQ 541555 / CQ 545618 / SR39030228

Issue

Less number of processors and more interfaces might be discovered.

Resolution

In earlier releases, processors were created by polling the values from hrDeviceDescr and the processor utilization was taken from the UCD-Mib for all the processors created. Thus CPUs were created for Hosts in huge numbers. Now processors are created from the Host-resources-MIB, based only on the polled values of hrProcessorLoad. If hrProcessorLoad does not populate a value, then the processor(s) are created based on the values from ssCpuIdle(UCD-Mib). So if no data is populated for hrProcessorLoad (Host-Resources-Mib) but ssCpuIdle(UCD-Mib) has an entry, only a single CPU will be created for many hosts.

More interfaces seen may be because of the NIC teaming feature, which is enabled by default. However, if many Host SNMP agents are slow, on querying the NIC teaming OIDs, the agents may stop responding which may result in a partial discovery. If you are not using the NIC teaming feature, disable it by setting *EnableTeamedNICDiscovery* to FALSE in the tpmgr-param.conf file and then verify.

Further, ASNMP for specific devices may be timing out due to which the devices are getting added to the pending list. The devices will either go to the pending list or it will get discovered partially with a message saying "SNMP request timed out or had an error". To resolve this issue, reduce the "maxOIDsPerPacketForASNMP" value to less than 19 or disable ASNMP for the devices by setting the *DoNotUseASNMPPattern* flag in the tpmgr-param.conf file to TRUE. Then restart the server and rediscover the topology.

Creation of VirtualBridges and Bridges relationships

CQ 573309 / CQ 573334

During discovery, if a switch port is associated with a MAC address that matches the VirtualMACPattern configured in the file tpmgr_param.conf, a VirtualBridges relationship is created for the switch port. The VirtualBridges relation is also formed for Systemless MAC irrespective of the VirtualMACPattern.

NOTICE

Systemless MAC cannot have bridges relationship. The default VirtualMACPattern has the VMware specific MACs, VRRP, and HSRP MACs.

Limitation

By removing the VirtualMACPattern from the tpmgr-param.conf file, the Bridges relationship gets established. On the contrary, when the VirtualMACPattern is added again in the configuration file, the VirtualBridges relationship does get established, but the existing Bridges relationship is not removed.

Resolution

To avoid having both Bridges and VirtualBridges relationships, delete the device from the topology and discover the device again. This will enable the creation of VirtualBridges relationship only.

Incorrect layering of ATM interfaces may result in erroneous notifications

CQ 545619 / CQ 502966 / SR37491304

Issue

User had the custom setting to explicitly manage the AAL5/ATM subinterfaces. Because of this, the RCA was not proper for a specific fault.

Resolution

As per design, IP unmanages certain *AAL5/ATM subinterfaces* interfaces in ic-post-port-manager.asl. This is done to produce proper RCA.

In this case, these unmanaged interfaces were managed explicitly by the user script. This resulted in a wrong RCA of *Interface::Down* for the AAL5 interface that is explicitly managed by the user when the actual problem was connected Router DOWN.

This is because, when the underlying *AAL5/ATMSUBINTERFACE* interfaces are managed, IP has an additional symptom set to deduce the *AAL5/ATMSUBINTERFACE* as DOWN instead of arriving to Router DOWN.

So in such cases, when the user expects a router DOWN event, the custom management policy for the AAL5/ATMSUBINTERFACE must be removed.

Creation of NetworkConnection Redundancy Groups (NCRG)

CQ 580365 / SR40971838

NCRGs are simply reference classes and they are not created during discovery.

There is only one relationshipset: ComposedOf/PartOf between members of the NetworkConnection class. If that is populated, the typical events (ReducedRedundancy, AtRisk, AllComponentsDown) work as expected.

Technical notes for IP Manager 9.2.1

There are no technical notes for IP Manager 9.2.1.

IP Manager fixed problems from 9.2.x that apply to 9.3.x

"Fixed problems in IP Manager 9.2.x" on page 61

Fixed problems in IP Manager 9.2.x

Features introduced in the Fixed Problems tables from 9.2.x.release notes are included here.

CQ/SR/Issue number	Problem summary	Description of fix	Found in version
Fixed problems iin IP 9.2.x			
CQ 605787	Smarts IP Domain Manager did not discover Adtran TA5000 devices (1.3.6.1.4.1.664.1.747).	If the device is not discovered with default configuration, then proceed as follows: Adtran TA5000 devices are re-certified in Smarts 9.2.1 Patch 2 and can be discovered with the below settings in place because of lack of compatibility of Adtran TA5000 devices with advance SNMP settings of Smarts IP: Set following parameters in tpmgr-param.conf: #cd ‹BASEDIR›/IP/smarts/bin #./sm_edit ../conf/discovery/tpmgr-param.conf GetBulkPattern-.1.3.6.1.4.1.664.1.747 FALSE DoNotUseASNMPPattern .1.3.6.1.4.1.664.1.747 SingleOIDPerPacketForASNMPPattern .1.3.6.1.4.1.664.1.747 Load the changes into the running domain: #./sm_tpmgr –s ‹DM_Name› --load-conf=tpmgr-param.conf This device will be discovered with PsuedoWire Interfaces, Advanced PsuedoWire Interfaces, and IPHost interfaces. For Advanced PsuedoWire Interfaces and IPHost interfaces, AdminStatus will always be "UP". For PsuedoWire Interfaces, OperStatus is mapped to "DOWN" when the OID adGenAdvancedPseudowireCEMDestinationMac {1.3.6.1.4.1.664.5.70.30.2.1.1.6} returns zero MAC(00 00 00 00 00 00). Windows: \lib\sm_devstat_adtran.dll \lib\sm_health_adtran.dll Unix: /lib/libsm_devstat_adtran.so /lib/libsm_health_adtran.so /conf/discovery/tpmgr-param.conf /conf/discovery/DISCOVERY_ADTRAN.import /rules/discovery/ic-adtran-card.asl /rules/discovery ic-adtran-interface.asl /rules/discovery ic-adtran-ta5000-health.asl /rules/discovery ic-adtran-ta5000-vlan.asl	

CQ 643103 CQ 602815 SR 54481736	ic-serviced daemon error. On Solaris 10, whenever ./ic-serviced start and ./ic-serviced stop commands were run, the following message appeared: root@infruz402 # ./ic-serviced start Starting sm_serviced ... ./ic-serviced: test: unknown operator ==	The error message is because of the following line in the ic-serviced shell script: 109: ["$OS" == "Linux"] && touch $LOCKFILE 129: ["$OS" == "Linux"] && rm -f $LOCKFILE The shell was not accepting the == operator. Hence the above lines are now changed by using a single = operator: 109: ["$OS" = "Linux"] && touch $LOCKFILE 129: ["$OS" = "Linux"] && rm -f $LOCKFILE	9.2.0.1
CQ 643963 CQ 604621 SR 56130730	When the IPV6LinkLocalAddressEnabled attribute in the discovery.conf file is set to TRUE, a network connection in the topology is removed.	When link local IPv6 addresses are discovered, the interface will have more than one IP and this results in the NetworkConnection not getting created. Now the link local IPv6 addresses are ignored, Aggregate link gets created and the NetworkConnection is present under the ComposedOf relation.	9.2
CQ 605086 SR 56335942 CQ 644841	SNMP V3 traps were processed when the credentials did not match the seedfile.	Security is now available for processing SNMP V3 traps both in the Trap Adapter or in the Trap Exploder only if SM_SNMP_V3_STRICT_TRAP_CHECK =TRUE is set in the runcmd_env.sh file. By default, this flag is not available in the runcmd_env.sh file. Hence security check is not in place and this is the default legacy behavior. For the SNMP V3 trap processing, fixed the following (to display warnings): • Seedfile has auth and the SNMP trap has no auth for this user • Seedfile has auth/priv and the SNMP trap has auth/noPriv for this user • Seedfile has auth/priv and the SNMP trap has wrong auth password The env-var SM_SNMP_V3_STRICT_TRAP_CHECK should be set to TRUE to enable the above.	9.1

CQ 644698 CQ 606473 CQ 644621 CQ 644655 CQ 644408 SR 56784012	If an SNMP agent has a community string with commas, call to SNMP_AccessorInterface::reset_ability_to_poll for this agent failed. Error ADSP-INVALID_ADDRESS-Invalid agent address appears in the log.	The problem was that commas were used as delimiters in the agent ID string, so having commas in the parameter values confused the parser. When constructing an agent ID string, a caller will need to escape commas in parameter values. The escaping is done by doubling the comma. Fixed the agent ID parsing routine to handle escaped (doubled) commas.	9.2.1
CQ 644666	When using IP and ESM topology in Watch4n, two instances of a single Nexus 1000 vSwitch get created. The preference is for only one consolidated vSwitch object.	The issue has been resolved by introducing new parameter *DisableNexus1KVConsolidation* in tpmgr-param.conf file for IP and in ESM.import file for ESM. The consolidation can be enabled through this parameter. By default, the consolidation is disabled. • If in IP, *DisableNexus1KVConsolidation* = FALSE in tpmgr-param.conf file, and in ESM ESM_Manager:: ESM-Manager DisableNexus1KVConsolidation = FALSE, then consolidation of Nexus 1000 vSwitch will take place in ESM. This means that the vSwitch name will be same in both in ESM and IP server. • If in IP, *DisableNexus1KVConsolidation* = TRUE in tpmgr-param.conf file, and in ESM ESM_Manager:: ESM-Manager DisableNexus1KVConsolidation = TRUE, then consolidation of Nexus 1000 vSwitch will not take place in ESM. This means that the vSwitch name will be different in ESM and in IP. **Note:** If you are not using EMC M&R (Watch4net), then there is no need to use this parameter, because in Service Assurance Manager, Nexus1KV consolidation is present by default.	9.2.2

CQ 577960 CQ 577931 SR 55817374	VLAN Interfaces were not showing traffic statistics. Performance instrumentation on VLAN interfaces was not created.	A new parameter *ExcludePerfInstDescPattern* is added in the tpmgr-param.conf file. This is used to enable performance instrumentation for all VLAN interfaces and for all SYSOIDs. By default, performance instrumenting for interfaces having VLAN as its description is excluded. *ExcludePerfInstDescPattern* is used to disable creating performance instrumentation for those interfaces which have *VLAN* in the description. Current utilization will be populated only if the Performance instrumentation is created. The description added below is not case sensitive. If in case you want to instrument these interfaces for performance, prefix a "~" (tilde symbol) to the below description pattern. For example, ExcludePerfInstDescPattern-.1.3.6.1.4.1.9.1.1327 ~*VLAN*|~*TUNNEL* The below attribute is SysOID specific. If you want to apply this for the entire topology, then comment the SysOID specific flag and then uncomment the global, that is, the ExcludePerfInstDescPattern flag. This will apply the settings for all the Interfaces/Ports whose description contains VLAN. The following example will enable performance instrumentation for all devices with interfaces/ports having VLAN in its Description: ExcludePerfInstDescPattern ~*VLAN* Note that a system restart is a must for this to take effect. If both the parameters are used, that is, SysOID specific and global, then the priority is given to SysOID specific, and the settings will be applied accordingly. ExcludePerfInstDescPattern-.1.3.6.1.4.1.9.1.1327 *VLAN* ExcludePerfInstDescPattern *VLAN*	9.2

CHAPTER 4
Server Manager 9.2.x Features and Technical Notes

This chapter includes the following sections:

- Server Manager features from 9.2.x that apply to 9.3.x 66
- Server Manager technical notes from 9.2.x that apply to 9.3.x 67
- Server Manager fixed problems from 9.2.x that apply to 9.3.x 72

Server Manager features from 9.2.x that apply to 9.3.x

- "Features in Server Manager 9.2.2" on page 66
- "Features in Server Manager 9.2.1" on page 66l

Features in Server Manager 9.2.2

The following Server Manager 9.2.2 features are incorporated in the *9.3 EMC Smarts Server Manager User and Configuration Guide*:

- Disk rebuild
- Polling and Threshold enhancement for F5 Load Balancer
- Inter-Switch Link (ISL) monitoring

The following Server Manager 9.2.2 features are not supported for 9.3.x:

- SMI-S Provider discovery
- The SpareStoragePool class

Features in Server Manager 9.2.1

There are no enhancements in Server Manager 9.2.1.

Server Manager technical notes from 9.2.x that apply to 9.3.x

- "Technical notes for Server Manager 9.2.2" on page 67
- "Technical notes for Server Manager 9.2.1" on page 71

Technical notes for Server Manager 9.2.2

This section contains special instructions for post-installation. Follow the procedures provided below:

- "Disabling tcp_fusion" on page 67
- "Users with custom code that uses EMC Smarts API" on page 67
- "Enabling Cisco Discovery Protocol (CDP)" on page 67
- "Configuring the runcmd_env.sh file to avoid ASL error messages" on page 70 should be performed as a best practice.
- "Configuring the IP Availability and Performance Manager to discover Microsoft cluster hosts using the cluster node name" on page 70
- "Creating field certification for vSphere (ESX 9.0) discovery" on page 70 should be performed if you need to discover VMware vSphere topology.

Disabling tcp_fusion

To disable tcp_fusion, add the following lines at the bottom of the /etc/system file. Restart the system for the new setting to take effect.

```
Disable tcp fusion
set ip:do_tcp_fusion = 0×0
```

Users with custom code that uses EMC Smarts API

If the code is based on Java, Perl, or C API, modify and retest the code according to the changes in this release. If a dynamic model is used, recompile the code.

The *EMC Smarts Dynamic Modeling Tutorial* explains the concepts and methods of dynamic modeling.

Enabling Cisco Discovery Protocol (CDP)

To create connections between a physical switch and an ESX server, Cisco Discovery Protocol (CDP) should be enabled on both the physical switch and the vSwitch (Standard vSwitch, Distributed vSwitch and Nexus 1000 vSwitch) installed on the ESX server. CDP should be running in both Listen and Advertize mode. If CDP is not enabled in both Listen and Advertize mode, then the connections between the physical switch and the ESX server will not be created.

NOTICE

The ESX hostname should have an entry in DNS server mapping to ESX IP address.

The procedures to enable CDP are:

- "Steps to enable CDP on Standard Virtual Switch installed on an ESX server" on page 68
- "Steps to enable CDP on Distributed Virtual Switch installed on an ESX server" on page 68
- "Steps to configure CDP on Nexus 1000v Switch installed on an ESX server" on page 69

Steps to enable CDP on Standard Virtual Switch installed on an ESX server

To enable CDP on a standard virtual switch installed on an ESX Server:

1. Log in to the Service Console using SSH, remote console, or physical console as root. Verify the current CDP setting for the desired virtual switch. In this example, vSwitch1:

   ```
   [root@server root]# esxcfg-vswitch -b vSwitch1
   down
   ```

   NOTICE

   The output *down* indicates that CDP is currently not implemented.

2. Set the CDP status for a given virtual switch. The possible values are:
   - down
   - listen
   - advertise
   - both

   ```
   [root@server root]#  esxcfg-vswitch -B both vSwitch1
   ```

3. Verify the new setting:

   ```
   [root@server root]# esxcfg-vswitch -b vSwitch1
   both
   ```

Steps to enable CDP on Distributed Virtual Switch installed on an ESX server

To enable CDP on a vNetwork Distributed Switches (vDS) installed on an ESX Server:

1. Connect to vCenter Server using the vSphere Client.
2. In the vCenter Server home page, click **Networking**.
3. Right-click **vDS** and click **Edit Settings**.
4. Go to Properties and select **Advanced**. Select the options to set the CDP settings.

Example

An example of configuring CDP for a Cisco 6500 physical switch:

```
switch#
switch# config terminal
switch(config)# cdp run
switch(config)# interface g1/1
switch(config-if)# cdp enable
switch# show cdp neighbors
```

```
Capability Codes: R - Router, T - Trans Bridge, B - Source Route
Bridge
            S - Switch, H - Host, I - IGMP, r - Repeater, P - Phone
Device ID                    Local Intrfce   Holdtme   Capability
Platform         Port ID
pa-tse-h24.pasl.vmware.com Gig   1/1           121      S
VMware ESXvmnic2
pa-tse-h24.pasl.vmware.com Gig   1/2           121      S
VMware ESXvmnic3
```

Steps to configure CDP on Nexus 1000v Switch installed on an ESX server

To configure CDP on a Nexus 1000v switch installed on an ESX server:

1. Log in to Nexus 1000v VSM using SSH.
2. Execute the following commands:

```
Nexus1000v(config)# cdp enable
Nexus1000v(config)#
Nexus1000v# show cdp neighbors
Capability Codes: R - Router, T - Trans-Bridge, B -
Source-Route-Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater,
                  V - VoIP-Phone, D - Remotely-Managed-Device,
                  s - Supports-STP-Dispute


Device-ID           Local Intrfce Hldtme Capability  Platform
Port ID

bdsc-cs7.lss.emc.com  Eth3/2       124    S I      WS-C2960-24TT
Fas0/23
bdsc-cs7.lss.emc.com  Eth4/1       120    S I      WS-C2960-24TT
Fas0/20
bdsc-cs7.lss.emc.com  Eth4/2       120    S I      WS-C2960-24TT
Fas0/22
bdsc-cs7.lss.emc.com  Eth6/1       179    S I      WS-C2960-24TT
Fas0/7
bdsc-cs7.lss.emc.com  Eth6/2       169    S I      WS-C2960-24TT
Fas0/8
Nexus1000v#
```

Enabling SNMP probe for VMware discovery

To enable SNMP probe for VMware discovery:

1. Use the **sm_edit** utility to edit the DISCOVERY_VMWARE.import file:

```
sm_edit local\conf\esm\DISCOVERY_VMWARE.import
```

2. Uncomment the ICF_SNMPDiscoveryDriver:

```
#
#ICF_SNMPDiscoveryDriver::VMWARE-Probe {
#   ElementTypeGlob = "ESM"
#   Priority = 2000
#   waitForCompletion = TRUE
#   ExecutesDriver = GA_SelectiveDriver::VMWARE-SELECTIVE-Driver
#   #trace = TRUE
#}
```

3. Restart the Domain Manager.

Configuring the runcmd_env.sh file to avoid ASL error messages

After discovering Server Manager and F5 devices through an IP 7.0.3 server, ASL error messages similar to the following may appear in the Server Manager log file:

```
[30-Mar-2010 15:06:54+272ms EDT] t@58 Discovery #10
ASL-W-ERROR_RULE_SOURCE-While executing rule set
'/opt/xxx/EISMHF2/ESM/smarts/rules/dxa/dxa-topology-sync.asl'
 ASL-ERROR_ACTION-While executing action at:
 ASL-CALL_STACK_RULE-   RuleName: GET_DISPLAY_VALUE, Line: 3190
 ASL-CALL_STACK_RULE-   RuleName: IMPORT_ATTRS_AND_RELS, Line: 2875
 ASL-CALL_STACK_RULE-   RuleName: EOF, Line: 1270
 ASL-ERROR_UNDEFINED_VARIABLE_ACCESS-Variable 'remoteValue' has not
been
    assigned a value

[30-Mar-2010 15:06:54+315ms EDT] t@58 Discovery #10
ASL-E-ERROR_LOSTSERVER-Lost contact with remote server --
terminating parser
 ASL-ERROR_GET-While attempting to get property 'PartOf' of object
    'MR_Object::frame.lss.emc.com'
 SM-EIO-Fatal I/O error
```

These ASL error messages are generated as a result of connectivity issues between Server Manager and IP Availability Manager. The best practice is to ensure that the firewall running between IP Availability Manager and Server Manager allows the flow of data that is used by keep alive mechanism. The keep alive mechanism is used between Server Manager and IP Availability Manager. Disabling keepalives should be used as the last resort.

Configuring the IP Availability and Performance Manager to discover Microsoft cluster hosts using the cluster node name

To ensure that Microsoft cluster hosts are properly discovered by the IP Availability and Performance Manager, you must force the IP Availability and Performance Manager to name each cluster host by the cluster node name instead of the DNS name. To do this, you must add the cluster node name to the hosts file on the IP Availability Manager server.

Creating field certification for vSphere (ESX 9.0) discovery

To discover VMware vSphere topology, create a field certification of the vSphere OID in the EMC Smarts IP Availability Manager, as follows:

NOTICE

If you are at IP 7.0.2 Patch 2 and IP 8.1.1 Patch 1, then there is no need of adding the certification manually.

1. At IP *BASEDIR*/smarts/bin, enter the following command:

   ```
   sm_edit conf/discovery/oid2type_Field.conf
   ```

   NOTICE

   After saving the oid2type_Field.conf file, the certification needs to be loaded in ‹*BASEDIR*›/IP/smarts/bin/sm_tpmgr –s ‹IP_Domain› --reloadoid

2. Type the following lines:

```
# VMware 4 ESX server (vSphere )
.1.3.6.1.4.1.6876.4.1 {
    TYPE = Host
    VENDOR = VMWare
    MODEL = ESX4.0
    CERTIFICATION = CERTIFIED
    CONT = Generic-MIB2
    HOSTRESRCS = MIB2

INSTRUMENTATION:
    Disk-Fault = HostResources:DeviceID
    FileSystem-Performance = HostResources:DeviceID
    CPU/Memory = HostResources:DeviceID
    Interface-Fault = MIB2
    Interface-Performance = MIB2
}
```

NOTICE

The System ObjectID .1.3.6.1.4.1.6876.4.1 is for VMware ESX 9.0 and above. If the System ObjectID is .1.3.6.1.4.1.6876.4.1, then certification is required. Else certification is not required.

3. Save and close the file.

NOTICE

Start the agents on the VMware ESX server. In VMware ESX 3.5, two agents — VMware agent and NetSNMP agent should be started. In VMware ESX 9.0, only one agent, that is, the VMware agent should be started.

Technical notes for Server Manager 9.2.1

There are no technical notes for Server Manager 9.2.1.

Server Manager fixed problems from 9.2.x that apply to 9.3.x

"Fixed problems in Server Manager 9.2.x" on page 72

Fixed problems in Server Manager 9.2.x

Features introduced in the Fixed Problems tables from 9.2.x.release notes are included here.

CQ/SR/Issue number	Problem summary	Description of fix	Found in version
Fixed problems in ESM 9.2.x			
CQ 642604 SR 53835358 CQ 601020	F5 - Missing Threshold for Pool or PoolMembers.	Created a threshold group for LoadBalancerPool to set threshold values of AtRiskThreshold. Also created Polling groups for LoadBalancerPool, LoadBPoolMembers, VirtualServer and also a new setting for each which helps in unmanaging analysis of a particular object. Modified files: • conf/init/settings.conf • esm/rules/f5/f5_esm-setting.asl • esm/rules/f5/f5_pool-setting.asl • esm/rules/f5/f5_pool_member-setting.asl • esm/rules/f5/f5_virtual_server-setting.asl • esm/rules/init/init-settings.asl	9.1
CQ 609104 SR 57718300	When the wildcard (*) is used in the Process Monitoring table argument to monitor processes, ESM generates Missing Process events for the process even though the process is running (false positive).	When the process argument in the Process Monitoring table is changed to a wildcard (*), all processes with the same name are grouped and monitored under one application process. The count of process instances in the table is set to the number of processes with the same name. The Missing Process event is generated only when the count of process instances reaches 0.	9.2.1
CQ 608922 CQ 605648 SR 56306992	Processor, FileSystem, and Memory objects were also discovered in ESM. These objects should not be discovered in ESM because these were already discovered and monitored by the IP Manager.	Added a configurable flag EnableCPUFilesystemMemoryDiscovery in ‹Smarts›/conf/esm/ESM.import file to enable or disable the discovery of Processor, FileSystem, and Memory objects. By default this flag is set to TRUE. Set it to FALSE if you want to disable the discovery of Processor, FileSystem, and Memory objects.	9.2.1

CHAPTER 5
MPLS Manager 9.2.x Features and Technical Notes

This chapter includes the following sections:

- MPLS Manager features from 9.2.x that apply to 9.3.x 74
- MPLS Manager technical notes from 9.2.x that apply to 9.3.x 74
- MPLS Manager fixed problems from 9.2.x that apply to 9.3.x..... 75

MPLS Manager features from 9.2.x that apply to 9.3.x

- "Features in MPLS Manager 9.2.2" on page 74
- "Features in MPLS Manager 9.2.1" on page 74

Features in MPLS Manager 9.2.2

There are no new features implemented in MPLS Manager 9.2.2.

Features in MPLS Manager 9.2.1

There are no enhancements in MPLS Manager 9.2.1.

MPLS Manager technical notes from 9.2.x that apply to 9.3.x

- "Technical notes for MPLS Manager 9.2.2" on page 74
- "Technical notes for MPLS Manager 9.2.1" on page 74

Technical notes for MPLS Manager 9.2.2

Users with custom code that uses EMC Smarts API

If the code is based on Java, Perl, or C API, modify and retest the code according to the changes in this release. If a dynamic model is used, recompile the code.

The *EMC Smarts Dynamic Modeling Tutorial* explains the concepts and methods of dynamic modeling.

Technical notes for MPLS Manager 9.2.1

There are no technical notes for MPLS Manager 9.2.1.

MPLS Manager fixed problems from 9.2.x that apply to 9.3.x

"Fixed problems in MPLS Manager 9.2.x" on page 75

Fixed problems in MPLS Manager 9.2.x

Features introduced in the Fixed Problems tables from 9.2.x.release notes are included here.

CQ/SR/Issue number	Problem summary	Description of fix	Found in version
Fixed problems in MPLS 9.2.x			
CQ 601008 SR 53723510	SMARTS is not reporting a notification when LdpAdjacency between two devices is disconnected.	LDP adjacency down event is now getting generated when the OperStatus is set to NOTPRESENT. For both Cisco and Juniper instrumentation classes, the LDPSessionStatusEnum::NOTPRESENT is mapped to IsDownStatus being TRUE.	9.2.1

CHAPTER 6

Network Protocol Manager 9.2.x Features and Technical Notes

This chapter includes the following sections:

- Network Protocol Manager features from 9.2.x that apply to 9.3.x 78
- Network Protocol Manager technical notes from 9.2.x that apply to 9.3.x..... 79
- Network Protocol Manager fixed problems from 9.2.x that apply to 9.3.x 81
- Fixed problems in Network Protocol Manager 9.2.x patch documentation..... 81

Network Protocol Manager features from 9.2.x that apply to 9.3.x

- "Features in Network Protocol Manager 9.2.2" on page 78
- "Features in Network Protocol Manager 9.2.1" on page 78

Features in Network Protocol Manager 9.2.2

There are no new features implemented in Network Protocol Manager 9.2.2.

Features in Network Protocol Manager 9.2.1

RFC 4444 device support for IPv6 IS-IS

CQ 643327

The RFC - 4444 support is only for IPv6 ISIS.

NOTICE

ISIS RFC 4444 MIB support is verified and validated for devices only with IPv6 configured in them. For pure IPv4 devices with RFC 4444 ISIS-MIB support, the devices get discovered by the IP Manager but will not be listed as part of ISIS_TopologyCollection and hence the NPM ISIS server will not discover the devices.

In the NPM 9.2 release, issues were observed on RFC - 4444 ISIS IPv6 discovery and monitoring. Now RFC 4444 followed devices (ISIS IPv6 MIB) are fully supported regardless of device vendors (Cisco and Juniper). Discovery and monitoring support is complete and the ISIS NPM Domain Manager is able to stitch topologies discovered using RFC -4444 MIBs.

Following is the complete OID information for RFC-4444:

```
#####################################
#      ISIS-MIB - RFC 4444
#####################################

#isisAreaAddr
                    {40, ".1.3.6.1.2.1.138.1.1.3.1.1"},
#isisSysLevelType
                    {45, ".1.3.6.1.2.1.138.1.1.1.2"},
#isisISAdjIPAddrAddress
                    {50, ".1.3.6.1.2.1.138.1.6.3.1.3"},
#isisCircIfIndex
                    {51, ".1.3.6.1.2.1.138.1.3.2.1.2"},
#isisRouterSysID
    {55, ".1.3.6.1.2.1.138.1.1.1.3"},
#isisISAdjNeighPriority
                    {60, ".1.3.6.1.2.1.138.1.6.1.1.10"},
#isisISAdjState
                    {65, ".1.3.6.1.2.1.138.1.6.1.1.2"}
```

Network Protocol Manager technical notes from 9.2.x that apply to 9.3.x

- "Technical notes for Network Protocol Manager 9.2.2" on page 79
- "Technical notes for Network Protocol Manager 9.2.1" on page 80

Technical notes for Network Protocol Manager 9.2.2

- "Users with custom code that uses EMC Smarts API" on page 79
- "NPM dump instrumentation utility" on page 79

Users with custom code that uses EMC Smarts API

If the code is based on Java, Perl, or C API, modify and retest the code according to the changes in this release. If a dynamic model is used, recompile the code.

The *EMC Smarts Dynamic Modeling Tutorial* explains the concepts and methods of dynamic modeling.

NPM dump instrumentation utility

The NPM dump instrumentation utility is a tool that can provide useful event information for troubleshooting by EMC Support. Here are sample commands for running the utility:

Note: You must send both the output file plus the repository file (for example, bgp.rps) in order for EMC Support to troubleshoot the problem.

Table 25 Scripts to write the instrumented values to an output file

Commands	Description
ic_dumpinst_bgp.asl	Write the instrumented attribute values for BGP objects to an output file.
ic_dumpinst_eigrp.asl	Write the instrumented attribute values for EIGRP objects to an output file.
ic_dumpinst_isis.asl	Write the instrumented attribute values for IS-IS objects to an output file.
ic_dumpinst_ospf.asl	Write the instrumented attribute values for OSPF objects to an output file.

BGP instrumentation utility command

The command to write the instrumented values is:

```
sm_adapter --output=<LOG_FILE> --server=<BGP-DOMAIN-MANAGER>
   utils/ic_dumpinst_bgp.asl

Where
 * <BGP-DOMAIN-MANAGER> - Name of the BGP Domain Manager to connect to.
 * <LOG_FILE> - Name of log file to write instrumentation objects to.
 * Log file will be located under the <BASEDIR>/smarts/local/log
```

EIGRP instrumentation utility command

The command to write the instrumented values is:

```
sm_adapter --output=<LOG_FILE> --server=<EIGRP-DOMAIN-MANAGER>
   utils/ic_dumpinst_eigrp.asl

Where
 * <EIGRP-DOMAIN-MANAGER> - Name of the EIGRP Domain Manager to connect
   to.
 * <LOG_FILE> - Name of log file to write instrumentation objects to.
 * Log file will be located under the <BASEDIR>/smarts/local/log
```

IS-IS instrumentation utility command

The command to write the instrumented values is:

```
sm_adapter --output=<LOG_FILE> --server=<IS-IS-DOMAIN-MANAGER>
   utils/ic_dumpinst_isis.asl

Where
 * <IS-IS-DOMAIN-MANAGER> - Name of the IS-IS Domain Manager to connect
   to.
 * <LOG_FILE> - Name of log file to write instrumentation objects to.
 * Log file will be located under the <BASEDIR>/smarts/local/log
```

OSPF instrumentation utility command

The command to write the instrumented values is:

```
sm_adapter --output=<LOG_FILE> --server=<OSPF-DOMAIN-MANAGER>
   utils/ic_dumpinst_ospf.asl

Where
 * <OSPF-DOMAIN-MANAGER> - Name of the OSPF Domain Manager to connect
   to.
 * <LOG_FILE> - Name of log file to write instrumentation objects to.
 * Log file will be located under the <BASEDIR>/smarts/local/log
```

Technical notes for Network Protocol Manager 9.2.1

There are no technical notes for Network Protocol Manager 9.2.1.

Network Protocol Manager fixed problems from 9.2.x that apply to 9.3.x

"Fixed problems in Network Protocol Manager 9.2.x" on page 81

Fixed problems in Network Protocol Manager 9.2.x

Features introduced in the Fixed Problems tables from 9.2.x.release notes are included here.

CQ/SR/Issue number	Problem summary	Description of fix	Found in version
Fixed problems in NPM 9.2.2			
CQ 605924 CQ 604957 SR 56310962 CQ 604789 SR 55949384 CQ 604333	OSPFNeighborRelationship was not built for IPv6 and Neighbor down alerts was not notified for both IPv4 and IPv6.	A configurable flag EnableAlertsOnDynamicOSPFLinks is now added to ‹BASEDIR›/NPM/smarts/conf/ospf/ospf.conf. Use sm_edit to modify the value. If this flag is set to TRUE and an SNMP error occurs, Neighbor state will be set as DOWN, otherwise UNKNOWN. Default value is TRUE. libsm_msi.so libsm_rpsm_ospf.so rules/settings.asl	9.2

Fixed problems in Network Protocol Manager 9.2.x patch documentation

Features introduced in the Enhancements section of the 9.2.x NPM patch documentation are included here.

Regular full discovery in BGP is long when new devices are added at fixed intervals in AM Domain

CQ 615584

Issue

Regular full discovery in BGP Domain Manager takes long time when user adds few devices at fixed intervals in AM Domain.

Resolution

Implemented a Smarts driver-based mechanism to detect newly added devices automatically. This mechanism imports only the newly added devices into NPM BGP.

There are three attributes introduced with this fix in conf/bgp/bgp.conf which enable or disable the ability to detect new devices.

- LiteSynchronizationInterval = 3600
- EnableReconfigureLaunchedLiteSynchronization = FALSE
- EnablePeriodicLiteSynchronization = FALSE

By default, lite-synchronization between BGP Domain Manager and underlying AM domain is disabled. Set the above attributes to TRUE and restart the BGP Domain to turn on the lite- synchronization.

CHAPTER 7
VoIP Availability Manager 9.2.x Features and Technical Notes

This chapter includes the following sections:

- VoIP Availability Manager features from 9.2.x that apply to 9.3.x 84
- VoIP Availability Manager technical notes from 9.2.x that apply to 9.3.x 88
- VoIP Availability Manager fixed problems from 9.2.x that apply to 9.3.x..... 89

VoIP Availability Manager features from 9.2.x that apply to 9.3.x

- "Features in VoIP Availability Manager 9.2.2" on page 84
- "Features in VoIP Availability Manager 9.2.1" on page 86

Features in VoIP Availability Manager 9.2.2

- "Discovery and monitoring of VoIP components for AudioCodes MediaGateway devices - Mediant 1000 and Mediant 2000" on page 84
- "Discovery and Monitoring of VoIP Components for ACME SBC 3800 and 4000" on page 85
- "Realms" on page 86

Discovery and monitoring of VoIP components for AudioCodes MediaGateway devices - Mediant 1000 and Mediant 2000

CQ 644770

Discovery and monitoring of VoIP components for AudioCodes MediaGateway devices have been added in VoIP 9.2.2. The devices that are supported are Mediant 1000 and Mediant 2000 AudioCode devices.

Mediant 1000 and Mediant 2000 AudioCodes devices are discovered as "MediaGateways" in the IP Manager. Mediant 1000 and Mediant 2000 MediaGateways devices are a part of VOIP_TopologyCollection.

The IP containment discovery for the MediaGateways devices is based on MIB2.

All the VoIP components for the AudioCodes device are discovered from the AcGateway MIB (1.3.6.1.4.1.5003.9.10.3).

The following components and attributes are discovered as part of VoIP AudioCodes discovery:

- Channels — Attribute discovered is StartingPhoneNumber
- TrunkGroupSettings — Attribute discovered is TrunkGroupID
- IPProfile — Attribute discovered is ProfileName
- TelProfile — Attribute discovered is ProfileName
- ProxyIP — Attribute discovered is ProxyIPAddress
- Tel2IPRouting — Attribute discovered is ProfileID
- IP2TelRouting — Attribute discovered is ProfileID

NOTICE

The above components are discovered as part of "VoipNetworkService" for the AudioCodes devices.

Table 26 on page 85 lists the events generated for AudioCodes MediaGateway VoipNetworkService.

Table 26 Events generated for AudioCodes MediaGateway VoipNetworkService

Managed Object	Event	Condition	Instrumentation source
Channels	Unresponsive	A channel for AudioCodes MediaGateway is unresponsive.	AcGateway MIB
TrunkGroupSettings	Unresponsive	A channel for AudioCodes MediaGateway is unresponsive.	AcGateway MIB
IPProfile	Unresponsive	A channel for AudioCodes MediaGateway is unresponsive.	AcGateway MIB
TelProfile	Unresponsive	A channel for AudioCodes MediaGateway is unresponsive.	AcGateway MIB
ProxyIP	Unresponsive	A channel for AudioCodes MediaGateway is unresponsive.	AcGateway MIB
Tel2IPRouting	Unresponsive	A channel for AudioCodes MediaGateway is unresponsive.	AcGateway MIB
IP2TelRouting	Unresponsive	A channel for AudioCodes MediaGateway is unresponsive.	AcGateway MIB

The above VoIP components are monitored for their status. All components that are added as part of VoIP discovery are imported into SAM. The map console for the device reflects these components for the VoipNetworkService “ComposedOf” relationship.

Discovery and Monitoring of VoIP Components for ACME SBC 3800 and 4000

ACME SBC 3800 and 4000 are the SBCs from ACME Packet. SBCs are the devices in the VoIP domain which are used to exert control over signaling and media streams involved in setting up, conducting, and tearing down telephone calls or other interactive media communications.

ACME SBCs are certified in IP as Switches and are discovered for components like Interface, Cards, and so on based on MIB2. With the new IP implementation, these devices are now part of VOIP_TopologyCollection in IP. VoIP domain further picks up these devices for discovery and monitoring of VoIP components.

Prerequisites:

- ACME packet SBC SNMP agent need to be configured and open 161 port SNMP queries. It needs to respond to .1.3.6.1.4.1.9148 tree OIDs.
- Trap destination needs to be set or configured in SBC to send to the VoIP trap receiver.
- VoIP-AM license with VOIP_ACME feature.

VoIP AM uses SNMP queries and creates/discovers the following components in the VoIP AM topology:

SBC application -> VoIPNetworkService

New traps from APSYSMGMT-MIB are modeled in VoIP 9.2.2 and available in the acme-trapd.conf file.

NOTICE

Because information is not available in SBC MIB regarding associated Audio Codecs with which SBC is associated, relationship with codec devices is not discovered or displayed in the topology. So Audio Codecs association discovery is not supported.
Though SBC is available in HA mode, information is not available in the MIB which describes the redundancy configuration. So redundancy groups will not be shown. Only the VoipCluster class with associated SBCs are displayed.

Realms

Realm is a specific domain in VoIP for Session Initiation Protocol (SIP) authentication, which has its own usernames and password configured. Realms are discovered based on APSYSMGMT-MIB (1.3.6.1.4.1.9148.3.2.1.1). These components are classified as part of VoipNetworkService for the SBCs.

Realms are discovered and monitored as part of VoIP. In this release, SBC certification is based only on SNMP and based on APSYSMGMT-MIB.

The following attributes are instrumented and monitored for the Realms:

- ActiveSessionInbound
- ActiveSessionOutbound
- TotalSessionInbound
- TotalSessionOutbound

"Realms" is the new class that is added in the VoIP domain. SAM has been modified to import the Realms class as part of topology import. The map console for the device reflects these components for the VoipNetworkService -> ComposedOf -> Realms relationship.

apSigRealmStatsTable (1.3.6.1.4.1.9148.3.2.1.2.4) is probed to discover Realms instances. VoIP instruments Realms elements and monitors them for availability.

Table 27 on page 86 lists the events generated related to SBC VoipNetworkService.

Table 27 Events generated related to SBC VoipNetworkService

Managed Object	Event	Condition	Instrumentation source
Realms	Unresponsive	A Realm for ACME SBC is unresponsive.	APSYSMGMT-MIB

Features in VoIP Availability Manager 9.2.1

VoIP release 9.2.1 includes important new features and changes to the existing features:

- "Support Avaya Communication Manager 6.3" on page 87
- "VoIP 9.2.1 now handles Avaya stale phones issue by enabling the "EnableAvayaStalePhonesCleanup" flag in voip.conf . When enabled, in the next discovery, Avaya stale phones get deleted." on page 87

Support Avaya Communication Manager 6.3

In the VoIP 9.2 release, Avaya Communication Manager 6.2 was certified with CLI/SNMP discovery. In this release, Avaya Communication Manager 6.3 is certified with CLI/SNMP discovery.

Support clean up of Cisco and Avaya stale phones

CQ 602317 / SR 54403618 / CQ 643208

VoIP-AM did not clean up Cisco IP Phones from its topology as they are either removed and no longer registered to any Cisco Unified Call Manager clusters. A once instantiated Phone instance stays within VoIP AM repository and this gave inaccurate information to the end user managing Cisco Phones.

To clean Cisco stale phones:

1. sm_edit conf/voip/voip.conf
2. Set the EnableCiscoStalePhonesCleanup = TRUE.
3. Restart the server to have this flag setting take effect. Subsequent CUCM rediscovery should cleanup Cisco stale phones.

```
# If TRUE, stale Cisco phones will be removed from the topology
# Default value is FALSE(By default stale phones are not cleaned).
# EnableCiscoStalePhonesCleanup = TRUE
```

CQ 576438

VoIP-AM did not clean up Avaya IP Phones from its topology because they were removed and no longer registered to any Avaya Communication Manager. An instantiated Phone instance stays within VoIP AM repository and this gave inaccurate information to the end user managing Avaya phones.

VoIP 9.2.1 now handles Avaya stale phones issue by enabling the “EnableAvayaStalePhonesCleanup” flag in voip.conf . When enabled, in the next discovery, Avaya stale phones get deleted.

To clean Avaya stale phones:

1. sm_edit conf/voip/voip.conf
2. Set the EnableAvayaStalePhonesCleanup = TRUE.
3. Restart the server to have the flag setting to take effect. Subsequent Avaya CMs (re)discovery should cleanup Avaya stale phones.

```
# If TRUE, Avaya Stale phones will be removed from the topology
# Default value is FALSE (By default stale phones are not removed).
# EnableAvayaStalePhonesCleanup = TRUE
```

VoIP Availability Manager technical notes from 9.2.x that apply to 9.3.x

- "Technical notes for VoIP Availability Manager 9.2.2" on page 88
- "Technical notes for VoIP Availability Manager 9.2.1" on page 88

Technical notes for VoIP Availability Manager 9.2.2

VoIP ACME SBC Polling Group was not found under the VoIP Network Services Group

CQ 576242

IThe ACME SBC Polling Groups were not found under the VoIP Network Services Group.

If the ACME feature is enabled and if the VoIP server is started, the ACME SBC polling group is created under the VoIP Network Services Group.

If the ACME feature is not enabled, the ACME SBC polling group will not be created. If the feature is enabled later and if an old RPS is used, the ACME SBC polling group will not appear. In this case, run the following commands to load the settings and to create the ACME SBC polling group:

```
dmctl -s VoIP-AM invoke ICF_PolicyManager::ICF-PolicyManager loadSettings
   <VOIP-INSTALL>/smarts/[local]/conf/voip/acme-settings.conf
dmctl -s VoIP-AM invoke ICF_PolicyManager::ICF-PolicyManager reconfigure TRUE
```

How to get cluster information on Cisco and Avaya devices

This is regarding the VoIP topology split feature. To configure topology split, you can get information on the cluster name using the commands provided and update the same in the voip.conf file. Then add the cluster information for Cisco and Avaya devices by uncommenting the lines in voip.conf.

To get information on cluster name, run the following command from the VoIP install InCharge/VoIP/smarts/bin:

Avaya

./sm_snmp -d ‹Avaya Call Server IP address› get
.1.3.6.1.4.1.6889.2.8.2.2.9.0.1.3.6.1.4.1.6889.2.8.2.2.9.0 = ‹Cluster Name›

Cisco

./sm_snmp -d ‹Cisco Call Manager IP address› get
.1.3.6.1.4.1.9.9.156.1.1.2.1.8.1.1.3.6.1.4.1.9.9.156.1.1.2.1.8.1 = ‹Cluster Name›

The output you get is the cluster name that can be used to configure TSM in voip.conf.

Technical notes for VoIP Availability Manager 9.2.1

There are no technical notes for VoIP Availability Manager 9.2.1.

VoIP Availability Manager fixed problems from 9.2.x that apply to 9.3.x

"Fixed problems in VoIP Availability Manager 9.2.x" on page 89

Fixed problems in VoIP Availability Manager 9.2.x

Features introduced in the Fixed Problems tables from 9.2.x.release notes are included here.

CQ/SR/Issue number	Problem summary	Description of fix	Found in version
Fixed problems in VoIP 9.2.x			

CQ 567402	In voip.conf, in ManagedClusters configuration, ClusterID values should be case-insensitive.
	`#ManagedClusters = {`
	`#`
	`{<ClusterID>, <Vendor>},`
	`#`
	`{<ClusterID>, <Vendor>}`
	`#  }`
		Prior 9.2.1 release values were case sensitive, so can miss some clusters if case is sensitive.
		In 9.2.1, ClusterID values can be any case.
			9.2.1