

Release 2.6 Caveats

Caveats describe unexpected behavior in Cisco IOS XE Release 2. Severity 1 caveats are the most serious caveats; severity 2 caveats are less serious. Severity 3 caveats are moderate caveats, and only select severity 3 caveats are included in the caveats document.

This section contains open and resolved caveats for the current Cisco IOS XE maintenance release.

The *Dictionary of Internetworking Terms and Acronyms* contains definitions of acronyms that are not defined in this document:

http://www.cisco.com/en/US/docs/internetworking/terms_acronyms/ita.html

This section consists of the following subsections:

- [Open Caveats—Cisco IOS XE Release 2.6.2, page 167](#)
- [Resolved Caveats—Cisco IOS XE Release 2.6.2, page 176](#)
- [Open Caveats—Cisco IOS XE Release 2.6.1, page 196](#)
- [Resolved Caveats—Cisco IOS XE Release 2.6.1, page 206](#)
- [Open Caveats—Cisco IOS XE Release 2.6.0, page 233](#)

Open Caveats—Cisco IOS XE Release 2.6.2

This section documents possible unexpected behavior by Cisco IOS XE Release 2.6.2

- CSCsu59515
Telnet inside host from outside the host fails when port 23 is statically allocated on a Cisco ASR 1000 Router.
Workaround: None
- CSCtc05275
On a Cisco ASR 1000 Router a false memory leak has been seen within the AAA Memory Stats Tool.
When there are multiple separately allocated attribute lists attached to a request or and event in AAA, the memory stats tool gives wrong information, as it does not account for all the separate lists after freeing the event with the request.
Workaround: None
- CSCtc18663
When running PPPoEoA PTA, one may see some ATM VCs stay in inactive state on a Cisco ASR 1000 Router.
This condition may occur when loading this configuration while bring up the ATM SPA.
Workaround: Is to do a **shut** and then a **no shut** on the interface this should allow for all the ATM VCs to come up.
- CSCtc45832
When tracking stops the data-plane logs out of the PKT-MEM trace log this problem will occur on an ASR 1000 Router Series the sessions will be dropped and the QoS hierarchy will shut down. There also will be pending queue objects waiting to be flushed out in the list.
The following command will show the BQS RM status:

show plat hard qfp act inf bqs stat

In rare conditions, an error may occur for extreme over-subscribed environments. When sending 10G (For example: 5G as priority, and 5G as non-priority) traffic to a 1G interface.

All priority and control packets are dropped by the hardware this occur when the packet buffers are depleted; and when the schedule stops forwarding output packets

Workaround: There is no known workaround to this problem.

- CSCtc69297

Tracebacks has been seen with cli **sh platform hardware qfp active** feature acl tree on the Cisco ASR 1000 Router.

This condition has been seen, when there are a huge number of acls configured on the router.

Workaround: None

- CSCtc76606

The following error message has been observed on the console, when the SPA is out of sync:

SPA_OIR-3-OUT_OF_SYNC

Workaround: None

- CSCtc86844

Idmgr invalid error messages has been seen on a Cisco ASR 1000 Router console.

This instance has been observed after the router has scaled to 32k PPPoL2TP sessions and running traffic with events. However, there are no impact on the router.

Workaround: None

- CSCtc90106

Memory leaks are seen when changing the “fvrf” during traffic flow in the IPSEC_RMAL process on a Cisco ASR 1000 Router. This condition has been observed when changing “fvrf” to show as “fvrf2”, while sending traffic and checking incremental memory leaks the memory leaks are the seen on the router.

Workaround: Do not change fvrf's frequently.

- CSCtc91018

On a Cisco ASR 1000 Router the subinterface counters with Frame Relay Encapsulation can show higher values than the counters on the main interface, when self-pinging the subinterface.

Workaround: None

- CSCtd08709

When one LTS is restricted with CAC calls are not terminating through another LTS.

After configuring call admission control on LTS2 to 1 and making 20 calls through LAC all are go to LTS2 as per the priority, and as call admission configured on LTS2 call should be diverted back to LAC and should terminated on LTS1 which is not happening.

Workaround: Do not restrict call admission control on LTS.

- CSCtd21252

Unified SBC crash has been seen on the ASR 1000 Router Series.

This condition may occur, when configuring a large IPv6 media-address on the router.

Workaround: None

- CSCtd36301

At every session churning of IPv6 PPPoE uses more prefixes for same tunnel and session value. No used IPv6 Prefixes in local IPv6 pool are incremented at every session flap iteration in IPv6 LNS for same tunnel and session value.

This instance may happen, when Local IPv6 prefix pool is used to assign ipv6 address and the sessions are churning at a flap rate of 70 sessions per seconds for 8000 sessions.

Workaround: None

- CSCtd37057

On a heavily loaded Cisco ASR 1000 Router Series, rapid QoS queuing configuration changes involving the removal of existing configuration and addition of new configuration could cause the system to experience temporary resource outage.

The conditions under which this has been observed involve 32000 flapping PPPoE sessions combined with configuration changes on the system.

Workaround: Avoiding rapid and large QoS configuration changes on a heavily loaded system will avoid the problem reported in this caveat.

- CSCtd80542

Loop observed, when configuring SNMP bulk mib walk. The loop has been observed at tunnelInetConfigIfIndex.

This condition has occurred, when scaled configuration includes tunnel interface 2147483647.

Workaround: None

- CSCtd83379

DHCP discover packets are not reaching the server via a Bridge from the client.

This condition have been seen when the Pagent Client is initiating DHCP discover message to the server via the Relay.

The Relay (ASR 1002) is using the unnumber interface to forward the DHCP discover packets to the server. The Bridge to Bridge between Ethernet and serial interfaces are using the same bridge group.

It has been seen that the DHCP discover packets are reaching up to the Bridge interface and the Relay unnumber interface is not receiving the DHCP discover packet.

Workaround: No workaround.

- CSCtd87072

IOSD will restart, when changing tunneling mode in scaled IPSec Sessions on an ASR 1000 Router Series.

This condition has been observed, after IOSD restarts the tunneling mode has changed in a scaled IPSec Session enviroment.

Workaround: None

- CSCtd91950

A Cisco ASR 1000 Router Series with the Lawful Intercept feature configured may reset unexpectedly under certain conditions when streams are modified/**disabled/re-enabled** during traffic flow.

The conditions necessary for this situation to be encountered are multiple MDs, configuration of circuit-id based pre-provisioned stream entries and active PPPoE sessions.

Workaround: There are no known workarounds.

- CSCtd98510

Some of the L2TPv3 Xconnects are not coming up after repeated (5-6) switchovers and OIR.

This instance may occur when an AC is down while sessions are in local state and are not ready.

Workaround: Is to clear L2TP to recover from this problem.

- CSCte43453

QoS accounting Interim record for the parent policy-map class-default class has incorrect packets and bytes stats while under traffic load.

This condition has been seen when PTA session with Model D2.2 QoS has been enabled. QoS accounting has been enabled at the parent policy-map class-default class. While under traffic load, the accounting Interim record has incorrect stats as compared to the QoS stats in the output of show policy-map session.

Workaround: None

- CSCte46896

Following traceback appears on the a Cisco ASR 1000 Router console:

```
%EVENTLIB-3-TIMEHOG: F0: cpp_sp: undefined: 30160ms,
Traceback=1#ad497e64d353fac0e9ed1351f534cf6f evlib:F3B0000+D120 evlib:F3B0000+A838
cpp_common_os:F8E8000+10E2C cpp_common_os:F8E8000+10EDC evlib:F3B0000+DB60
```

When 1K Prefixes with 5 traffic class each prefix is configured. The traceback could appear in the below mentioned scenarios:

MC is already configured for 5K TCs with mode monitor both with traffic turned on and BR is reloaded with BR configs

With MC is already configured for 5K TCs with mode monitor both with traffic turned on and issuing “**clear oer master ***” on the MC.

With MC is already configured for 5K TCs with mode monitor both with traffic turned on and the best utilization value is moved from one link to another.

Workaround: None

- CSCte50863

An fman_fp core is generated when the Template ACL feature is disabled or enabled several times with 4k PPP sessions with per-user ACLs.

This condition has been observed, when bringing up 4000 PPP Sessions terminated on PTA with per-user ACLs. With the template ACL feature enabled, only a few templates are created. Disable the template ACL feature and since there are only 4000 PPP Sessions, TCAM exhaustion by this action is not expected. Enable the template ACL feature again. Repeat until an fman_fp core is generated (usually seen within 10 iterations).

Workaround: Is to tear down PPP Sessions before disabling and enabling the Template ACL feature.

- CSCtf06872

Kernel crash may occur with GETVPN configuration (with 1 GDOI Group and 3 VRF's).

This condition are seen with overnight traffic and the kernel crash may occur within a GETVPN Topology.

Workaround: None

- CSCtf08810

Multicast traffic loss observed in broadband environment.

This condition happens after RP switchover, multicast traffic takes longer to converge.

Workaround: None

- CSCtf16429

Stale object has been seen on RP2 switchover with Route and MPLS flaps.

Workaround: None

- CSCtf23385

When PTA is configured for 32k PPPoEoA or 16k PPPoEoA AutoVC with the following kind of configuration:

```
interface atm 2/0/0.65000 multipoint
range pvc 1/32 1/4033
pvc-in-range 1/32
!
pvc-in-range 1/33
:
: so on 4000 pvc-in-ranges
```

Then when the PTA is unconfigured in the following sequence:

1. First unconfigure all pvc-in-range
2. then unconfigure range pvc
3. unconfigure interface

And reconfigured, it is found that all autovcs on the standby RP do not get created.

The condition is caused due to specific order of unconfig as mentioned above.

Workaround: Do not unconfigure the the above mentioned sequence. Unconfigure the interface only, then this issue is not seen.

- CSCtf43664

Ucode crash on loading EoMPLS configuration on a Cisco ASR 1000 Router.

This condition happens after starting up L2TPv3 while trying to copy EoMPLS configurations on the same interfaces.

Workaround: None

- CSCtf57963

On a Cisco ASR 1000 Router VRF with VRFx inspect vrf-default are added.

This condition may occur when the above option gets **enabled** as soon as the CLI “parameter-map type inspect global” is added with ZB Firewall.

Workaround: None

- CSCtf81635

Trace back warnings are observed in the log.

This has been observed when changing the ACL configuration used by a large number of PPPoE Sessions while some of them are connecting or disconnecting.

Workaround: None

- CSCtf97660

In an ASR 1000 Router configured for CUBE(SP Edition), a SIP session will sometimes terminate with a 503 error message. The PDLOG will indicate a “Socket Write Error”.

This only happens when TCP is being used as the transport for SIP signaling and a media IP address was reconfigured to be used for a Signaling address without reloading.

Workaround: Is to avoid re-using a media IP address for signaling or if required, save the configuration and reload after making the configuration change.

- CSCtf98979

The following error message appears when stale object are seen after RP switchover:

```
6RU_BR2#sh platform software object-manager fp active stale-object
Object identifier: 2085
Description: Route-map name OER_INTERNAL_RMAP
Status: Done
Object identifier: 2090
Description: Feat 6, CG 1, rtmap name OER_INTERNAL_RMAP
Status: Done
```

This condition are seen with dynamic route maps enabled and PFR setup. The switch over is done on the active BR and the stale object is seen.

Workaround: None

- CSCtg01020

IPSec tunnel fails (Phase 2) to establish between two ASR 1000 Routers when site-to-site VPN is configured due to invalid SPI.

Workaround: The IPSec tunnel may come up after issuing a **reload**.

- CSCtg32407

The RP may crash while unconfiguring ATM multipoint interfaces when configured with a different bba groups that has different session limits after bringing up the pppoeoa sessions.

This instance may occur when an ASR 1000 Router has ATM Multipoint interfaces configured with different BBA Groups and the session limits are different.

Workaround: Is to disable the pppoe config under interface level and then unconfigure the bba group.

- CSCtg33275

A Cisco ASR 1000 plogd crashes when reloading Cisco IOS XE 2.6.0 Release when using 6RU with dual RP2.

The following has been observed on the console:

```
core file "MCP-6RU-2_RP_0_plogd_25197.core.gz" is generated during RP0 reload.
```

Workaround: There is no known workaround.

- CSCtg53307

The QoS police functionality might fail if user configures both “police” and “priority <kbps>” in the same traffic class.

This condition may occur when the user configures this unsupported configuration with “police” and “priority <kbps>” in the same traffic class, actually only one police feature is supported per traffic class, and later remove one of the commands, the traffic sent through this class might fail to be policed to the configured rate.

Workaround: Is to only, enable one police feature in the same traffic class.

- CSCtg53599

The %COMMON_FIB-3-FIBIDBINCONS2 error has been logged on the standby RP after sessions are established.

The condition has been seen when ASR1006 with dual RP1 and FP10 installed.

The following error has been logged on the Standby RP multiple times after bringing up 10 PPPoE sessions with Model F QoS:

```
*Apr 29 17:08:42.792: %COMMON_FIB-3-FIBIDBINCONS2: An internal software error
occurred. Virtual-Access2.1 linked to wrong idb Virtual-Access2.1
```

Workaround: None

- CSCtg59328

When IPCP renegotiates for an existing PPPoE session, the new IPv4 address does not get synced up with the standby.

This instance may occur when IPCP renegotiates for an existing PPPoE session, the new IPv4 address does not get synced up with the standby.

Workaround: None

- CSCtg68228

MQC on ATM VC fails to sync up with RP Standby.

This condition may occur when MQC has downloaded from RSIM to an ATM and the AutoVC fails to sync up with RP Standby.

Workaround: None

- CSCtg88218

When establishing 128 ISG IP session with L4R feature that has duration with frequency provisioned on FP40, ucode core is generated with the following traceback:

```
6RU8_L4R_UT#show running-config | b ge
*May 17 21:33:23.092: %CPPHA-3-FAULT: F1: cpp_ha: CPP:0
desc:INFP_INF_SWASSIST_LEAF_INT_INT_EVENT0 det:DRVr(interrupt) class:OTHER sev:FATAL
id:2121 cppstate:RUNNING res:UNKNOWN flags:0x7 cdmflags:0x0
*May 17 21:33:23.092: %CPPHA-3-FAULTCRASH: F1: cpp_ha: CPP 0 unresolved fault
detected, initiating crash dump.
*May 17 21:33:23.093: %CPPDRV-6-INTR: F1:
/tmp/sw/fp/1/0/fpx86/mount/usr/cpp/bin/cpp_driver[6366]: CPP10(0) Interrupt : May 17
16:33:23.086438: :INFP_INF_SWASSIST_LEAF_INT_INT_EVENT0 ner
```

This condition has been observed when sending traffic to establishing 128 ISP IP sessions that have L4Redirect with non-zero duration and non-zero frequency.

Workaround: None

- CSCth03545

A memory leak has been seen on a Cisco ASR 1000 Router when the traffic is sent for a long period of time (6 to 7 hours).

This condition has been observed when 1500 bytes of traffic are sent for a long period of time (6 to 7 hours) then it results in a memory leak followed by router crash.

Workaround: None

- CSCth09005

Active RP crashes under heavy call load and 220 CPS is configured with billing enabled.

This instance may occur under heavy call load and 220 CPS is configured on the Active RP with billing enabled.

Workaround: None

- CSCth15799

When issuing a ping to multicast the process fails from one of the hosts while Multicast Group is configured.

This may occur when GDOI CM is applied to 2 interfaces and there is no local-addr configured.

Workaround: Is to clear crypto gdoi on the GM.

- CSCth27728

After SBC has been configured on an ASR 1000 and a SIP call is made the router crashes.

The conditions has been observed when the “del-prefix 0” instructs SBC to remove the first zero digits from a dialed number, which means not doing anything. SBC does not handle being instructed to remove zero digits from the number and this is the cause of the crash. Removing this from the config should result in the same behavior and avoid the crash.

Workaround: The customer has changed “edit del-prefix 1 add-prefix 64” to “edit del-prefix 0 add-prefix 64”. Instead of this they should just use “edit add-prefix 64”.

- CSCth29934

When Primary SIP OIR on the insertion side was executed, the Protocol-up delay of primary core side IF has been observed.

Workaround: remove a physical line in core-side. After that, insert SIP and then no shut the IF.

- CSCth30370

Traffic drops running Cisco IOS XE 2.6 on the ASR 1000 Router with AAA QoS Policy Accounting feature configured.

Traffic drops are observed while doing ISSU upgrade when running Cisco IOS XE 2.6 with AAA QoS Policy Accounting feature configured.

Workaround: None

- CSCth34753

CPP Crash when shutting down WCCP interface with the following error message displayed on the console:

```
un 8 16:32:35.218 pst: %CPPHA-3-FAULT: F0: cpp_ha: CPP:0
desc:INFP_INF_SWASSIST_LEAF_INT_INT_EVENT0 det:DRVR(interrupt) class:OTHER sev:FATAL
id:2121 cppstate:RUNNING res:UNKNOWN flags:0x7 cdmflags:0x0
```

This conditions has been observed when WCCP is enabled on ASR1002 and data traffic is being redirected to WAE.

Workaround: Is to gracefully shutdown WCCP in the WAE first before shutting down the router WCCP interface.

For example:

```
WAE-91#show wccp status
```


WCCP version 2 is enabled and currently active

```
WAE-91#conf t
WAE-91(config)#no wccp version 2
WCCP clean shutdown initiated
Waiting for shutdown ok (1 seconds) . Press ^C to skip waiting
WCCP clean shutdown wait time expired
WAE-91(config)#end

WAE-91#show wccp status
WCCP is not enabled
WAE-91#
```

The router will stop redirecting traffic to WAE once wccp service is disconnected from WAE.

For example:

```
In WAN1-1002-R46#
Jun 18 13:49:15.310 pst: %WCCP-1-SERVICELOST: Service 61 lost on WCCP Client
10.111.46.10
Jun 18 13:49:15.310 pst: %WCCP-1-SERVICELOST: Service 62 lost on WCCP Client
10.111.46.10
WAN1-1002-R46#
```

- CSCth38187

Traffic is loss with IPv6 Static Route function, the ASR 1000 Router failed after clearing FIB.

This condition occurred while checking the IPv6 Static Route function, some traffic was loss within IPv6 Static Route entries after clearing FIB on the router.

Workaround: None

- CSCth41121

An ASR 1000 Router crashes while processing a renegotiation rejection (reINVITE 491) on a call which is being transcoded.

This condition occurs when a reINVITE is rejected (a renegotiation failure) on a call which is already established and not using a transcoder. The reINVITE was attempting to use a transcoder (the new stream needed transcoding). The trigger for this crash is that the renegotiation adds an extra stream to the call (a new m= line in the SDP) and the reINVITE is rejected.

Workaround: None

- CSCth42453

SIP endpoints with shared line appearance fail to receive incoming call properly after an ASR 1000 Router failover.

This is a SBC (CUBE(SP Edition) problem running on an ASR 1000 platform.

There is no impact to normal SIP endpoint services.

Workaround: None

- CSCth48008

An ASR 1000 ESP may crash due to traffic which is being encrypted.

The exact conditions for this are not yet known. Fragmented GRE traffic which needs to be encrypted may be the trigger.

Workaround: There is no known workaround at this time.

- CSCth49752

EoMPLS remote link status is not shown in the show sub-interface output.

This condition has been when EoMPLS remote link status is not shown in the show interface output when xconnect has binded into the dot1q encapsulated nterface.

Workaround: None

- CSCto03123

Symptoms:

1. A slow memory leak is observed on the cman_fp process on an FP and the cmcc process on a SIP. This issue is seen on all the flavors for FPs and CCs. The leak is of the order of less than 100-122K bytes per day.
2. Additional memory leak can occur when frequent sensor value changes take place.

Conditions: This symptom of the first leak does not occur under any specific condition. The second leak occurs when sensor-related changes take place.

Workaround: There is no workaround.

PSIRT Evaluation: The Cisco PSIRT has evaluated this issue and does not meet the criteria for PSIRT ownership or involvement. This issue will be addressed via normal resolution channels.

If you believe that there is new information that would cause a change in the severity of this issue, please contact psirt@cisco.com for another evaluation.

Additional information on Cisco's security vulnerability policy can be found at the following URL:

http://www.cisco.com/en/US/products/products_security_vulnerability_policy.html

Resolved Caveats—Cisco IOS XE Release 2.6.2

All the caveats listed in this section are resolved in Cisco IOS XE Release 2.6.2

- CSCsc49958

AAA Authentication fallback method to enable password does not work properly if RADIUS server is unavailable.

When the RADIUS server is unavailable, enter any username but then the enable password as the user password.

Workaround: None

- CSCsx45326

This is an enhancement to remove the performance optimization achieved by the ddts# CSCef70161.

This condition happens when the **neighbor <> as-override** command was giving problems. This happened when it was used in an ipv4 VRF without SoO configuration on the PE and higher weight configuration on a particular CE.

Workaround: To get the best performance optimization achieved in ddts# CSCef70161, when **neighbor <> as-override** is configured in ipv4 VRF. Use the SoO feature to isolate specific peer out of update-group.

- CSCsx56362

BGP selects paths which are not the oldest paths for multipath on a Cisco ASR 1000 Router. This causes BGP to unnecessarily flap from multipath to non-multipath as a result of route flaps.

This condition has been observed when:

BGP is configured

More than one equally-good route is available

BGP is configured to use less than the maximum available number of multipaths

Workaround: There is no workaround.

- CSCsy23839

On Cisco ASR 1000 Router Series, CPU utilization of SIP (SPA Interface Processor) may be 100%.

This symptom is observed with the following procedure:

1. Open a terminal window for telnet to ASR 1000.
2. Telnet to ASR 1000.
3. Run the request platform software console attach x/x (login SIP IOS) command.
4. Close the terminal window without exiting from SIP IOS.
5. You can see that the ioscon process is not terminated and its CPU utilization is around 100% by the monitor platform software process command.

Workaround: Resetting the SIP resolves the issue.

- CSCsz45263

ISIS fails to come up due to the redundant 'clns mtu' is being added under the tunnel interface configuration.

This symptom is seen, after reload of the router.

Workaround: None

- CSCsz53438

When ip header compression is configured on the ASR 1000 Router, but not on the corresponding router, an unexpected reload of the embedded systems processor may occur.

This has been seen, when IPHC is configured on the ASR 1000 Router, but not on the router to which it is directly connected.

Workaround: Is to **enable IPHC** on both routers.

- CSCsz60746

A static route configured through an unnumbered interface which is in shutdown state will not show up in RIB even after the interface state is UP.

This condition has been seen when configuring a static route on a shut down interface having ip unnumbered configuration.

Workaround: Is to remove and re-add the static route.

- CSCta11120

When two servers are configured under a group; the first server in the group is inaccessible. At the time of the first request, the request failed over to the second server in the group. The first server in the group comes up; however, the router uses the second server to process the request instead of the first server.

In addition, when a single-connection is configured, switch uses only primary tacacs server for http authorization even though the primary server is down.

This condition has been seen when two servers are configured under a group; the first server in the group is inaccessible. At the time of first request, the request fail-over to second server in the group.

- HTTP Authorization is configured for switch's GUI access
- Multiple TACACS servers are configured with single-connection option
- Primary Tacacs server is down

Workaround: Disable single-connection

- CSCta40318

IOS or the IOSd may crash on a Cisco ASR 1000 Router.

This condition has been observed when ISAKMP CAC (call admission control) is configured, the CAC limit is reached and **debug crypto isakmp** is **enabled**.

Workaround: Is to **undebug crypto isakmp**.

- CSCtb58282

When an ASR 1000 is running Cisco IOS, the device may reload when **show tcp brief** is issued.

This condition has been seen when the following has occurred:

1. The “ip domain lookup” needs to be configured. It is on by default.
2. The ip address of the foreign host in the tcp session needs to have a very long domain name associated with it (in the order of 70 characters, only).
3. The port number of the foreign host needs to be 5 digits long.

When the ip domain lookup is disabled, the problem could still happen if the host has a static entry configured via the “ip host” command.

Workaround: Is to configure “no ip domain lookup”. Or, avoid using **show tcp brief** on the device.

- CSCtc62440

On a Cisco ASR 1000 Router Series, the removal of sub-interfaces may under certain conditions result in MFIB_MRIB-3-FAILED_WIRE_FIND error messages being generated on the Route Processor (RP).

There is no functional impact due to this issue.

Workaround: There are no known workarounds.

- CSCtd09817

ISG L2 Connected DHCP session is terminated on renewal after vrf transfer.

This condition has been seen when an ASR 1000 Router is configured as ISG for L2 connected subscriber session and vrf transfer is done without any change in dhcp class.

Workaround: None

- CSCtd34056

This enhancement request is to allow for **crypto pki crl ca size** to be saved in the ASR 1000 Router config and to not disappear after reload.

Workaround: None

- CSCtd89923

Webex SPA hard disk sectors are corrupted. This condition has been observed when SIP10 is configured with a Webex SPA running release 2.6.0 image that is Soft-OIR'ed. This configuration can potentially corrupt the sectors on the hard disk of the Webex SPA.

Workaround: Is to shutdown the SPA before reloading the SIP10.

- CSCte08821

When **sh l2tp session packets tunnel id** is issued on an ASR 1000 Router with wrong session id, the session id shows as a junk value.

Workaround: None

- CSCte09945

When an Cisco ASR 1000 Router operates in the Unified SBC mode, after a hardware switch over using CLI **redundancy force-switchover**, during the old active RP is booting, issue CLI **no sbc**. Check failure error is observed in the RP console log.

Workaround: No workaround until now.

- CSCte14955

An unexpected reload may happen on the ASR 1000 Router Series. This has seen, when BGP VPNv4 is configured and a neighbor is flapping on the router.

Workaround: None

- CSCte21062

Session churn shows a slow memory leak which manifests during individual session teardown when the one sec accounting accuracy feature is **enabled**.

This condition has been observed when, **subscriber accounting accuracy <VALUE>** is configured, background variables are allocated to support feature messaging. These variables are allocated a small amount of memory which is unfortunately not freed when the session is disconnected. This leads to a small memory leak averaging between 50-60 bytes per session disconnection.

Workaround: Removal of configuration related to subscriber accounting accuracy.

Example: **no subscriber accounting accuracy 1000**

- CSCte37344

The following IOS console message is printed during an attempt to add a non-queueing class to the 2nd level of a 3-level hierarchy, within a QoS policy that is attached to one or more interfaces:

At least one queueing feature needed for every class in the 2nd level policy with 3-level of hierarchy

Subsequent operations, even policy-map removal, will cause failures.

This condition happens when 3-level QoS policy-map is applied to one or more interfaces, when additional 2nd level classes are added (depending on the timing of events), even if the classes have queueing features, the user may see the following console message:

At least one queueing feature needed for every class in the 2nd level policy with 3-level of hierarchy

Once this message appears, subsequent attempts to modify or detach the policy will encounter errors and/or classification will not work correctly.

Can be seen when running 12.2(33)XNF.

Workaround: If the problem occurs, the FP/ESP must be rebooted.

To avoid the problem, remove the QoS policy from all interfaces first, make the policy-map modifications, then re-attach the policy.

Further Problem Description:

The problem occurs because the class-add event **leaks-through** even though the class-add operation is not allowed. From this point forward, IOS and PD layers are out-of-sync, so there are even errors on policy-map detachment and removal.

- CSCte39643

When PfR receives an EIGRP route change, the router may unexpectedly reload.

The symptom is observed with PfR and EIGRP configurations. It is observed some time after PfR receives an EIGRP route change, but before the previous EIGRP route is removed in the routing table, when PfR tries to recycle a previous EIGRP route.

Workaround: There is no workaround.

- CSCte49283

Sometimes the LNS router sends an incorrect NAS-Port value.

The symptom is observed when the LNS router sends a stop accounting-request to the RADIUS server.

Workaround: There is no workaround.

- CSCte64750

Slower PPPoE sessions bring up rate on the Cisco ASR 1000 Router.

This condition was observed when L2TP HA and congestion control has been enabled.

Workaround: None

- CSCte82240

SBC accepts “.” when key_addr_type is “**DIALED_DIGITS**”. This condition can occur, when set exact matching means has been set as:

rpsRtgActionKeyAddrWildcardType to AMB_MW_EXPLICIT_WILDCARD.

This is possible to have a “.” when rpsRtgActionKeyAddrType is set to AMB_MW_ADDR_TYPE_DIALED_DIGITS. However, it is no longer allowed when rpsRtgActionKeyAddrWildcardType is AMB_MW_EXPLICIT_WCARD (which means SBC should perform an explicit match).

Workaround: None

- CSCte82351

The BGP aspath encoding in snmp (bgp4PathAttrASPathSegment) is encoding all aspath information as 32bit. This is not compatible with RFC4273 which defines this oid.

This condition may occur on all IOS versions.

Workaround: None

- CSCte84710

When IPv6 Unicast is enabled on an ASR 1000 Router the following error message is displayed on the console:

error message flag_icmp_error_gen type 1 and code 0 popup

Workaround: None

- CSCte87294

The following L2TP related error with traceback might show up on the Standby RP:

```
L2TP-3-ILLEGAL: ____:_____: ERROR: Unable to reserve session ID 2047
Traceback summary example:
0x11305a7 is in errmsg
0x1f1c056 is in l2tp_errmsg_internal
0x1f1c1e8 is in l2tp_errmsg
0x1f1c33f is in l2tp_error_traceback
0x35ab131 is in l2tp_ha_create_session
0x35b1791 is in l2tp_ha_process_ICRQ_chkpt
0x35b0e90 is in l2tp_ha_process_proto_session_chkpt
0x35b0749 is in l2tp_cpf_process_message
```

```

0x35b0648 is in l2tp_chkpt_q_handler
0x35ae86f is in l2tp_ha_l2tp_msg_handler
0x1f5215e is in l2tp_ha_l2tp_msg_handler_os
0x1f25e73 is in l2tp_mgr_process

```

The session on the active will not be present on the standby.

This only happens when the tunnel goes down and a session id belonging to the tunnel gets reused on the active.

Workaround: Is to Reboot the standby, after bulk sync the Standby RP will match the Active RP.

- CSCte95275

When an ASR 1000 Router is out of free memory, some FMI codes are trying to free up memory out of the NULL chunk that has never been created.

This condition is observed when the ASR 1000 Router is out of free memory.

Workaround: None

- CSCte97814

On an ASR 1000 Router with BGP enabled, a small fixed size chunk memory leak is observed during boot-up. To be exact, it is observed just after config bulk-sync in redundant RP setup.

This symptom is observed on Cisco ASR 1000 Series Routers with a redundant RP setup and BGP enabled.

Workaround: There is no workaround.

- CSCtf01618

A Cisco ASR 1000 Router may unexpectedly reload due to SegV error.

This condition has been observed, when the ASR 1000 Router must be running 12.2(33)XND1 or later XND or 12.2(33)XNE or even later 12.2(33)XN releases and DMVPN is configured with Tunnel Protection.

Workaround: Remove Tunnel Protection.

- CSCtf04257

On an Cisco ASR 1000 running IOS XE 12.2(33)XND1 below message may be seen, when trying to configure a EoMPLSoGRE VC: %SW_MGR-3-CM_ERROR:

Connection Manager Error - provision segment failed [SSS:Eth:<number>] - no resources available.

This condition has been seen on Cisco ASR 1000 Router, running IOS XE 12.2(33)XND1. When destination of VC is changed from original to something else and then changed back to original.

Workaround: None.

- CSCtf13343

Command authorization for commands involving a 4-byte ASN fails. Command accounting for these commands will record an incorrect ASN or ip address.

The following commands are impacted:

Global configuration mode:

router bgp x.y

BGP configuration submode:

neighbor <address> remote-as x.y

ip vrf <vrf name> submode:

route-target <ip address>:X

route-target x.y:z

route-target y:z

If you turn on the relevant AAA debugs, you will see some arguments appear multiple times in a given authorization or accounting request, and others not appear at all.

This problem is seen whenever command authorization and/or command accounting is configured for any of the following commands:

Global configuration mode:

router bgp x.y

BGP configuration submode:

neighbor <address> remote-as x.y

ip vrf <vrf name> submode:

route-target <ip address>:X

route-target x.y:z

route-target y:z

A typical affected configuration in 12.0 and earlier would say:

aaa new-model

aaa authorization commands 15 default tacacs+

A typical affected configuration in 12.0T and later would say:

aaa new-model

aaa authorization commands 15 default group tacacs+

Workaround: You may be able to permit authorization of affected commands by allowing changing you tacacs+ server configuration to permit commands which include repeated arguments.

There is no workaround for the incorrect accounting records.

Further Problem Description: IOS releases not including 4-byte ASN support see a more limited form of this problem where only the first and last byte of the ipv4 address are sent to the AAA server. On such releases, the ASNs are sent as normal.

In Cisco IOS Release 15.0 and later, only the route-target command is tracked by this CSCtf13343. You need to have CSCtg42163 and CSCtg42088 integrated as well in order to get the fix for router bgp x.y and neighbor <address> remote-as x.y respectively.

- CSCtf13704

Memory leaks are seen when Graceful Restart is configured on an ASR 1000 with BGP sessions processing.

The following error message can appear during Graceful Restart:

%BGP-3-NEGCOUNTER

The symptom is observed with non-NSF Graceful Restart on releases with the fix for CSCtd99802.

Workaround: There is no workaround.

- CSCtf15982

While large number of DHCP sessions are coming up, the router may crash due to corrupted chunk header.

This issue happened while large number of unauthenticated sessions were coming up but it may also happen for authenticated sessions. There's no clear condition as to why this has happened.

Workaround: There is no known workaround at this time.

- CSCtf23727

On a dual-RP PE router where a BGP CE peer is connected via a PPP link and the CE peer is also configured for NSR (**neighbor ha-mode sso** command), forwarding for prefixes learned from the CE router may fail after an RP switchover. After the switchover, the affected routes appear in the BGP table without a bestpath and the reason **nexthop inaccessible** listed.

The problem can be seen when all of the following conditions are true:

- The PE router is dual-RP
- The PE router is configured for SSO redundancy mode and is operating in hot standby mode
- The PE is connected to a CE router over an PPP link
- On the PE router, the CE neighbor is configured for BGP NSR
- BGP learns prefixes from the CE and the nexthop addresss for those prefixes are via PPP
- An RP switchover is performed

Workaround: **Shut** and **unshut** the PPP link to the affected CE or **disable** nexthop address tracking (**no bgp nexthop trigger enable**) for VPNv4 and re-enable 1-2 minutes later.

- CSCtf27187

Traffic stops after initiating SPA OIR.

This symptom is observed only when initiating SPA OIR.

Workaround: Is to do a SIP OIR, the traffic should resume.

- CSCtf28793

When an ASR 1000 has the following configuration in BGP:

aggregate-address ip addresss summary-only advertise-map
route-map name suppress-map **route-map name**

Additional configuration are observed:

- Routes matching the suppress-map are not suppressed.
- The aggregate address is advertised.

The above condition are seen with a **reload** or after a hard **clearing** of BGP Peers. Both the advertise-map and suppress-map must be configured.

Workaround: Is to reconfigure the aggregate command, or use an aggregate command without the advertise-map/suppress-map combination.

- CSCtf29685

LNS Router crashed when sending accounting stop request.

This condition is observed when PPPoE setup is configured with LAC and LNS. In addition, when LNS downloads the account the process is stopped with failures when configurations from AAA server with template type are initiated.

Workaround: None

- CSCtf33539

When an ASR 1000 supporting L2TP High Availability and managing a large number of L2TP tunnels as a LAC or an LNS, may spontaneously reload in very rare circumstances, shortly after a stateful switchover or SSO, with a stack trace similar to the following example:

```
0x12f5ce4c is in l2tp_ha_sfo_resync_done
0x12f5ce24 is in l2tp_ha_sfo_resync_done
0x1215bfb0 is in
l2tp_ha_resync_receive_control_packet_os
0x12135920 is in l2tp_manage_ctrl_conn_for_pak
0x121369f8 is in l2tp_process_control_packets
0x1212fc88 is in l2tp_mgr_process_control_packets
0x1212fe0c is in l2tp_mgr_process
```

This condition has been observed when ASR 1000 supports L2TP High Availability and managing a large number of L2TP tunnels as a LAC or an LNS,

Workaround: There is no workaround.

- CSCtf36152

When ASR(LAC) receives StopCCN from LNS due to the lack of resources (L2TP session limit), the ASR (LAC) returns ZLB with bad sequence number.

In this case, the correct Ns/Nr of ZLB should be Ns=1/Nr=1.

Workaround: None

- CSCtf47795

An ASR 1000 Router may crash when **show ip bgp neighbor** command is executed.

Workaround: None

- CSCtf50075

A traffic blackhole can occur on the Cisco ASR 1000 Router Series.

The symptom is observed following **shut/unshut/shut** during redundant forwarding on an interface.

Workaround: There is no workaround.

- CSCtf51834

After a stateful switchover (SSO) on an IOS router supporting L2TP HA, the counter showing the number of L2TP sessions which were destroyed because they were not completely established at the time of the SSO, may be incorrect.

This counter is visible with the command **show l2tp redundancy detail** in the section Sessions destroyed during resync phase.

For example, the sessions destroyed during resync phase:

```
Poisoned:          0
Unestablished:    10    -- This value may be incorrect
Tunnel in resync:  0
```

After a stateful switchover (SSO) on an IOS router supporting L2TP HA.

Workaround: No workaround.

- CSCtf65536

ESP can crash while performing SIP calls using Cube-SP function.

This symptom is observed when hairpinned SIP calls are present, but it is timing related, so it doesn't occur in all cases.

Workaround: There is no workaround.

- CSCtf66633

A Floating static route with the **permanent** keyword may not get installed into the routing table when the primary route goes down.

Workaround: None

- CSCtf69391

Output drops on an interface incrementing apparently due to ISG with session drops.

This condition may happen when Low traffic has been seen on the interface. This appears to be actual packet drops from traffic on an interface.

Workaround: None

- CSCtf70312

When POS PA OIR and HA (using the Active RP crashed) simultaneously resulted in SIP crash.

SPA is OIRed and active-RP is forced crash(using "test crash" cli)simultaneously.

This resulted in the following message seen on new active RP due to SIP reset:

```
%PMAN-3-PROCHOLDDOWN: SIP0: pman.sh: The process mcpcc-lc-ms has been helddown (rc 134)
```

Workaround: None

Further Problem Description: The following trace for the SIP is displayed on the console:

```
<3478133836,4037269920>: %ASR1000_SIP_SPA-3-IPCPORT: SIP0/2: Failed to open  
IPC port 'IPC Master:  
Slot 0/2 ICP', error session in use
```

Reproducibility:

Is very low, unless there is no delay between SPA OIR and RP is forced to crash. This defect cannot be reproducible.

Impact: The SIP is crashed and reloaded with “process mcpcc-lc-ms has been helddown” message.

- CSCtf70851

Input/Output Rate freezes and doesn't get updated. This symptom is observed if the interface is **shut** with the traffic running, the input/output rate gets stuck and doesn't go back to 0.

Workaround: Giving **no shut** on the interface restarts the input/output rate.

- CSCtf71575

CE to CE ping failed over when EoMPLS is configured on a Vlan interface.

This condition has been observed when CE to CE ping failed with EoMPLS configured on a native Vlan interface.

Workaround: None

- CSCtf71998

The follow tracebacks are seen while PPTP sessions are being processed:

```
Mar  9 06:32:40 coltel-gw 231109: Mar  9 03:32:49.413: %SW_MGR-3-CM_ERROR: Connection  
Manager Error - provision segment failed [SSS:PPTP:17175151] - add to database fail.  
Mar  9 06:32:40 coltel-gw 231110: -Traceback= 4B5E50 4B6DDC 4B7504 12D3D34 12D3E84  
2555B20 2555BF8 12D135C 27A8DA4 27A8E54 12D1E80 12D2014 12C2C14 27A8DA4 27A8E54  
12C38A8
```

This condition are observed on LNS when VPDN group with PPTP has been configured.

Workaround: None

Further Problem Description: Tracebacks are thrown every few seconds on LNS with PPTP configurations, while sessions are coming up.

- CSCtf75446

The Cisco ASR 1000 Router console may freeze up in unconfiguring atm subinterface.

Workaround: None

- CSCtf78196

Although tunnel interface has alternative path to an OSPF neighbor, when the primary interface goes down, the tunnel interface goes down for a moment.

The symptom occurs when a tunnel tracks an MTU from higher value to a lower value on the outgoing interface. (It is seen on many images)

Workaround: Statically configure “**ipv6 mtu <mtu>**” on tunnel interfaces.

- CSCtf79163

Asymmetric carrier delay does not work on an ASR 1000 Router.

This condition has been observed when asymmetric carrier delay is configured on the router.

Workaround: Is to use symmetric carrier delay.

- CSCtf80105

When basic SIP-SIP calls are placed using automation scripts, calls start failing due to UDP socket connection error.

The symptom is observed when the router is configured with a dial peer and with SNMP. A dial peer is most likely required to reproduce the issue, but it is possible that a different UDP protocol other than SNMP could also cause the symptom. Once a call failure occurs, all the calls placed later will fail with a UDP socket connection error.

Workaround: Use the following steps:

1. Under sip-ua, configure **connection-reuse** (which is a hidden command).
2. Configure the use of TCP.

- CSCtf80843

On an ASR 1000 Router tracebacks are seen when PBHK do not have a port mapping for an active connection.

This instance would only occur after clearing IP sessions on the router where active PBHK port mappings exist.

Workaround: None

- CSCtf82883

When clearing a VRF route, there is a traffic drop on other VRF routes.

The symptom is observed with an L3 VPN configuration.

Workaround: There is no workaround.

Further Problem Description: Some LTE broker distribution is leaked to other VRFs.

- CSCtf83092

Standby resets continuously while ISSU upgrade from a non-componenterized IOS image to a componenterized IOS image.

The issue is seen with an MPLS VC configuration.

Workaround: There is no workaround.

- CSCtf84237

An Cisco ASR 1000 Router may reload with the following crash decoded tracebacks:.

In this example, the summary traceback has been observed:

```
0x123d7e24 is in vpdn_apply_vpdn_template_pptp
0x1239c100 is in l2x_vpdn_template_find
0x123d81dc is in vpdn_apply_l2x_group_config
0x123cfedc is in vpdn_mgr_call_initiate_connection
0x123cce68 is in vpdn_mgr_event
0x123ce974 is in vpdn_mgr_process_client_connect
0x123cf248 is in vpdn_mgr_process_message
0x123cf368 is in vpdn_call_manager
```

This condition may happen when an invalid tunnel-type VSA is configured as shown in this example:

```
vsa cisco generic 1 string vpdn:tunnel-type=l2tp_bad
```

Workaround: Is to configure a correct tunnel-type VSA in Radius.

- CSCtf86998

In a GETVPN ASR 1000 Router Series deployment, packets on one of the ASR GM router interfaces are not encrypted.

This symptom is observed when GM1 is in passive mode.

Workaround: There is no workaround.

- CSCtf90157

When an ASR 1000 selects link local address instead of global unicast address of unnumbered loopback interface to send ICMPv6, the time exceeds packet over Virtual Access over the interface.

Workaround: There is no workaround.

- CSCtf92423

After switchover on an ASR 1000 Router the Peer routes are learned from PPP and are not in RIB/CEF tables.

This symptom is observed when Switchover with PPP are learned on each of routes.

Workaround: None

- CSCtf93465

In a CUBE(SP Edition) ASR 1000 Router, the following message is seen when trying to enter SBC config mode:

```
SBC: Internal error - SBC configuration cannot be processed.
```

This condition sometimes happens after unconfiguring SBC.

Workaround: The workaround is to do a reload.

- CSCtf95905

An ASR 1000 Router may crash in the BGP HA SSO process. The following error message is shown when the standby RP is booted:

```
%SYS-2-CHUNKBADMAGIC: Bad magic number in chunk header, chunk <hex-addr> data
<hex-addr> chunkmagic <hex-addr> chunk_freemagic <hex-addr> -Process= "BGP HA
SSO"
```

The symptom is observed with the following conditions:

- The router is configured for SSO redundancy mode.
- BGP is configured.
- Some BGP peers have NSR configured (using the **neighbor ha-mode sso** command) and NSR is active for those peers.
- The standby RP is loaded and progresses to hot standby state after NSR sessions are already established on the active RP.

Workaround: Is to configure peers intended to be **enabled** for NSR for passive open only (using the **neighbor transport connection-mode passive** command) and then **enable** NSR on the BGP peers after the router has already reached hot standby state.

- CSCtf98758

Standby RP crashes after replacing the basic configuration of the router with an au3-e3 configuration.

This symptom is observed after initiating the following steps:

1. Configure the router with back-to-back SDH link for full AU3-E3 configurations with SPA-1XCHOC12/DS0.
2. Save the running configuration using **copy run bootflash:au3-e3.conf**
3. Reload the router with config register set to 0x2142. This will get the router running configuration to the basic default configuration.
4. After the router is up with redundancy setup and basic default configuration, execute the config replace command with the target config that was saved in step 1. {Config replace bootflash:au3-e3.con}

Workaround: There is no workaround.

- CSCtf98802

Config replace command when executed in a particular way causes the router to malfunction.

This symptom is observed after the following steps:

1. When we try to remove channelized configuration using config replace command, it will ask for the confirmation of the same as below:

Unprovision clear channel interface ?[confirm]

2. If we put any character other than 'y' or 'n' it will not remove the channel configuration for that particular path.
3. Now, if I try to remove these channels that were not cleared before manually, the system is behaves improperly:

```
Router(config-controller)#au-3 1
%ERROR: Standby doesn't support this command
% Invalid input detected at '^' marker.
```

```
Router(config-controller)#
As you see above system is not allowing to enter into the controller configuration
mode and resulting into "%ERROR: Standby doesn't support this command" message.
```

Workaround: By this point of time only after reload of the router, the situation comes under control and then only we can alter the controller configurations.

- CSCtg04289

After all the detached adjacencies are changed the **congestion sip buffer-tuning** and **congestion sip pool-size InbPoolSize 8000 CLI** command an ASR 1000 Router may reload.

This symptom is observed after viewing the IPS trace which indicates that there is a bug in the CLI code. The CLI deactivates the SIP TM entity before changing the buffer pool size, but then afterwards does not reactivate the SIP TM entity.

Workaround: None

- CSCtg06730

When following parameters of BASEROOT package set by MGC, that overwrites T-MAX configuration on DBE:

root/normalMGCExecutionTime

root/MGCProvisionalResponseTimerValue

This condition is observed when bringing up H.248 session and MGC set the parameters by Modify message.

Workaround: None

- CSCtg11491

System may encounter CPUHOG and an error with the following traceback:

```
%SYS-3-CPUHOG: Task is running for (2302)msecs, more than (2000)msecs (1/1),process =
Exec.
after clearing 4k+ ISG Radius Proxy sessions thru CLI : clear radius-proxy client <ip
addr>
```

This symptom is observed on a Cisco ASR 1000 Series Router when functions as an Intelligent Service Gateway (ISG) Radius Proxy, when thousands of sessions were established.

Workaround: There is no workaround.

- CSCtg12139

On an ASR1006 running IOS 12.2(33)XNF with SPA-2XCT3/DS0 card in slot 0/3SPA-2XCT3/DS0 is configured sends an alarm that the **DS3 Port Admin Down**.

This symptom is observed when an ASR1006 running IOS 12.2(33)XNF with SPA-2XCT3/DS0 card in slot 0/3SPA-2XCT3/DS0 is configured after one of the T1s on DS3 0/3/0 changes state, the **DS3 Port Admin Down** alarm for DS3 0/3/1 (the other DS3 on that card) is clearing and being re-inserted.

Workaround: Is to ignore the alarms as it not affecting any functionality.

- CSCtg12975

Memory leaks are seen due to the Allocation PC (**L2TP mgmt daemo**) and Name (**L2X GRP CLASS NAME**) in **show memory debug leaks** output.

The Memory Leaks occurs when vpdn group config is removed while l2tp tunnel is still up.

Workaround: Is to take down tunnel before removing vpdn group from config.

- CSCtg13217

ICMP Fragmentation required (type 3, code 4) and Host Unreachable Administratively (type 3, code 13) is not sent back if packets are hitting MTU checking, or ACL deny on Egress interface.

This condition is observed when an ASR 1000 Platform is running IOS 12.2(33)XNE.

Workaround: There is no workaround.

- CSCtg13790

An ASR 1000 Router may crash while placing a call with **no call-route p-called-party-id**.

This instance may occur as shown in the following example:

1. sipp--->CUBE (SP)--->-sipp
2. SIP to SIP call
3. Placed a call with INVITE sip:uJVvp1GE4YDaWiEVqCLE7Ql9Y1bph7xF@9.45.39.1:5060 SIP/2.0
4. and uJVvp1GE4YDaWiEVqCLE7Ql9Y1bph7xF is in the Contact header of REGISTER request
5. Validation fails since **call-route p-called-party-id** is disabled
6. CUBE(SP Edition) crashes due to validation failure since **call-route p-called-party-id** is disabled.

Workaround: None

- CSCtg16498

LNS VPDN message is incorrect when receiving CDN from LAC as follows:

```
%VPDN-4-SESSIONERROR: L2TP LNS R102 unable to terminate user cisco@cisco.com; Result 1, Error 0, No disconnect reason given
```

It should start with “%VPDN-4-SESSIONERROR: L2TP LAC”.

This occurs when receiving CDN's result code is “1” and the following is configured on the router:
vpdn logging is enabled

Workaround: There is no workaround.

- CSCtg16516

In a CUBE(SP Edition) ASR 1000 system, the bandwidth limits are reported in statistics even though tman/pol is OFF. This condition occurs when tman/sdr or tman/pdr is set.

Workaround: Don not set tman/sdr or tman/pdr when tman/pol is OFF.

- CSCtg16544

In a CUBE(SP Edition) ASR1000 system, the bandwidth limit for Side B is being reported for Side A, even though no bandwidth limit is set for sideA. This condition occurs when a bandwidth limit is configured for Side B, but not for Side A.

Workaround: Is to configure bandwidth limits for both sides.

- CSCtg18261

The BR fails to learn an application which is configured to learn traffic for a flow with dscp ef set.

This condition is observed when the BR fails to learn an application which is configured to learn traffic for a flow with dscp ef set. **show ip cache v flow** shows that the flows are seen and traffic is going through but the application does not get learned. There are 2 applications configured at same time, one for tcp and 2nd for dscp ef. The 1st one gets learnt but not the second one.

Workaround: None

- CSCtg18726

When using an ASR 1000 Router the route may fail to originate network (type-2) LSA and therefore not to install routes to the routing table.

This condition are seen in design with backup interface, when the following has occurred:

- backup interface has same IP address like primary
- OSPF network type is broadcast

- both, primary and backup interfaces are configured to act as DR

Workaround: Use p2p network type. Do not configure pair of primary/backup interfaces to act as DR.

Further Problem Description: This is day-1 issue, all IOS releases without fix are affected (if configuration matches conditions, note, it's rare configuration).

- CSCtg21602

In this example the following message is displayed on the Active RP Console:

```
%CPPOSLIB-3-ERROR_NOTIFY: F1: cpp_cp: cpp_cp encountered an error -Traceback=
1#ed4b69bc77a25ff35c522388cbb72a96 errmsg:D94E000+2160 cpp_common_os:E0A4000+B920
cpp_common_os:E0A4000+19148 cpp_exmem_mgr:ED1E000+895C cpp_exmem_mgr:ED1E000+9080
cpp_common_os:E0A4000+163D0 cpp_rrm_local_api_lib:EFDF000+3150
cpp_rrm_svr_lib:F004000+53F0 cpp_rrm_svr_lib:F004000+76CC cpp_rrm_svr_lib:F004000+80A4
cpp_rrm_svr_lib:F004000+9810 cpp_dmap:E954000+15264 cpp_dmap:E954000+21BF4
cpp_common_os:E0A4000+
```

This message is displayed right after un-configuring sbc by entering command:

no sbc global db

Workaround: None

- CSCtg23952

Gratuitous ARPs are sent on all interfaces when performing a copy tftp operation.

This condition happens when performing a copy tftp operation.

Workaround: None

- CSCtg25056

OSPF IEFT GR recovery aborts on restarting an ASR 1000 Router when ip address of OSPF enabled interface is also assigned on another interface of the same router which is admin down.

Workaround: Remove duplicate ip addresses from admin down interfaces.

- CSCtg26760

On an ASR 1000 Platform, WCCP redirection does not work after a reload, or following a change in the redirect ACL.

The syslog show a message similar to the following:

```
%FMFP-3-OBJ_DWNLD_TO_CPP_FAILED: F0: fman_fp_image: Batch type 6 ID 318 download
to CPP failed
```

Workaround: The only known workaround is to reboot the router.

- CSCtg30995

Delay of RP switchover associated with NV_BLOCK_INITFAIL message appearing on standby-turned-active RP console. When the manual switchover command, **redundancy force-switchover** and a filesystem command like **copy running-config startup-config** is issued from different consoles of active RP almost at the same time, such a problem may occur.

Workaround: Avoid issuing any filesystem access command simultaneously with the manual RP switchover command. Should the above problem occur, please execute the same filesystem command, say, **copy running-config startup-config** from the standby-turned-active console.

- CSCtg32004

After applying QoS Model C / D2 configuration on PTA and on issuing the cli **show platform hardware cpp active infrastructure bqs status** the output of “**# of Active Schedules**” value is less than the expected one.

This condition is observed ASR 1000 Router with RP2 Processor configured with Model C / Model D2.

Workaround: None

- CSCtg32647

Crypto tunnel fail to come up and the following message is displayed on the console:

```
%CRYPTO-3-IKMP_QUERY_KEY: Querying key pair failed.
```

This condition is observed when running IOS XE 2.6.0. Other releases may be affected.

Workaround: Downgrade to IOS XE 2.4.1. Other versions might work as well.

Further Problem Description: The issue happens while validating the peer certificate.

- CSCtg35230

VPDN sessions are created when SCCRQ and SCCRP have different ip addresses on an ASR 1000 Router.

This condition has been observed once the ip address is downloaded from the AAA server, the change to the IP address on LNS2 while creating two VPDN sessions.

Workaround: None

- CSCtg37082

Following warning messages are seen on active RP upgrade (ISSU) and switchover after upgrading to Release IOS XE 2.6.1 from 2.3.0e:

```
%IMRP-3-IMRP_MSG_CANNOT_RELAY: R1/0: imand: IMRP Peer ios_rp_iosd_slot_1:
cannot relay message to SPA 15/15
%IOSD_IMCC_CAPI-3-MSGDISPATCH: SIP2/1: Unable to dispatch received TDL
```

After ISSU upgrading the active RP to IOS XE Release 2.6.1 from 2.3.0e, when RP switchover is initiated this warning will be seen on the newly active RP.

The warning messages do not affect the general working of the SPA during ISSU when the VLAN Tunnel mode feature is not used.

Workaround: None

- CSCtg38018

An ASR 1000 Router could hang when static session is provisioned on the box upon switchover.

The following conditions has been observed upon switchover:

1. IP static session is provisioned
2. IP session is not HA aware
3. Switch-over is performed

Workaround: Deprovisioned static session before switch-over, after standby takes over, and become active unit. Reprovision static session.

- CSCtg40901

An ASR 1000 Router crash is seen while authenticating with TACACS.

The symptom is observed if the TACACS server does not respond.

Workaround: Is to use multiple connections.

Alternate Workaround: Is to configure a dummy TACACS server.

- CSCtg42998

An IOS XE Router supporting L2TP HA and acting as a LAC, may fail to effectively clear VPDN sessions which were cleared by the Client or LNS device just at the time when a stateful switchover is occurring.

This condition may happen when L2TP HA Route Processor switchover has occurred. In addition this problem may surface around 2 seconds or less, after a stateful switchover.

Workaround: An idle timer may be configured on the LAC.

The following commands should be configure on the relevant Virtual-Template:

interface

ppp timeout idle

ip idle-group access-list inout

or

The sessions can be cleared manually by an operator.

- CSCtg44097

Connect-Info 77 attribute is sent twice in a Pre-Auth Access-Request.

Workaround: None

- CSCtg46605

Due to caching issues the memory is cached and freed as required, this gives deceptive Free memory values when monitoring the ASR 1000 Router. The router had been changed to use the committed memory for monitoring the memory utilization, but users are not able to monitor the committed memory via SNMP.

This condition has been observed when monitoring the memory utilization on an ASR 1000 Router Platform.

Workaround: Is to use the CLI **show platform software status control-processor brief** to retrieve the committed memory.

- CSCtg50288

Standby RP crashes when configuring NAT Pool.

This condition are observed when the Standby RP crashes, while processing the following steps, below:

```
ASR-2RU(config)#ip nat pool pool-60 66.0.0.0 66.0.255.255 prefix-length 16
ASR-2RU(config)#ip nat pool pool-60 66.0.0.0 66.255.255.255 prefix-length 8
%Error, pool size should be maximum 19 bits long
ASR-2RU(config)#ip nat pool pool-60 66.0.0.0 66.0.255.255 prefix-length 16
%Unable to synchronize pool with standby RP
ASR-2RU(config)#
*Apr 28 16:23:14.318 IST: %REDUNDANCY-3-STANDBY_LOST: Standby processor fault
(PEER_NOT_PRESENT)
*Apr 28 16:23:14.318 IST: %REDUNDANCY-3-STANDBY_LOST: Standby processor fault
(PEER_DOWN)
*Apr 28 16:23:14.318 IST: %REDUNDANCY-3-STANDBY_LOST: Standby processor fault
(PEER_REDUNDANCY_STATE_CHANGE)
*Apr 28 16:23:14.386 IST: %IPNAT_HA-3-TRANSMIT: Unable to send via IPC
IPNAT_ADDRPOOL_MSG pool pool-60 id 7; retry queue flush
-Traceback= 1#cc462b2a175cd9784a73925c8c37beb5 :10000000+C49434 :10000000+C497B8
:10000000+2C77E38 :10000000+157B8E8 :10000000+157BEEC :10000000+1551070
:10000000+BB4DEC :10000000+BBB688 :10000000+2B07D90 :10000000+BC989C
```

```
*Apr 28 16:23:21.604 IST: %RF-5-RF_RELOAD: Peer reload. Reason: EHSA standby down
*Apr 28 16:24:06.115 IST: %REDUNDANCY-5-PEER_MONITOR_EVENT: Active detected a standby
insertion (raw-event=PEER_FOUND(4))
```

Workaround: None

- CSCtg50351

Two RADIUS access-request messages are sent upon receiving DHCPv6 SOLICIT message.

This can occur when the ASR1000 has been configured as a BRAS Router with “**ipv6 nd prefix framed-ipv6-prefix**” configured.

Workaround: There is no workaround, unless “**no ipv6 nd prefix framed-ipv6-prefix**” can be configured.

- CSCtg50359

Asymmetric carrier delay values are not updated or present in running-config with in these 2 examples:

1. - set carrier-delay msec 0
 - set carrier-delay up/down msec <any>
 - > uses the configured msec timer but running-config still shows **carrier-delay msec 0**.
 - > when symmetric carrier delay already configured, the configuring of asymmetric delay need be blocked, i.e first the symmetric carrier delay need to unconfigured before configuring any asymmetric values.
2. - set **no carrier-delay msec 0**
 - set **carrier-delay up msec 1**
 - set **carrier-delay down msec 0**
 - > the down timer is not updated and not visible in the running-config

This conditon will be visible from IOS XE 2.5 (12.2(33) XNE releases) and onwards.

Workaround: None

- CSCtg52972

Configuring **ip flow-export template options sampler** on an ASR 1000 Router may stop Netflow from working and this may cause errors. It is not supported, so the command should be rejected at the CLI.

This symptom is observed when configuring the unsupported feature **ip flow-export template options sampler** on an ASR 1000 Router.

Workaround: Is to reload.

- CSCtg57720

Junk characters are seen in the **remote address** part of “**show vpdn tunnel summary o/p**” when sessions are coming up in a new tunnel. Usually seen with 4 or more tunnels.

Workaround: None

- CSCtg62555

An ASR 1000 Router may be out of service after removing ip address from ip portbundle source loopback interface.

This symptom is observed on a Cisco ASR1000 series router when functions as an Intelligent Service Gateway (ISG), when Port Bundle Host Key (PBHK) is enabled on sessions, when thousands of sessions were established and running high rate traffic on both upstream and downstream direction.

Workaround: There is no workaround.

- CSCtg71904

Fragmented UDP packets with ip length ≤ 25 bytes are not getting encrypted.

This condition is observed ASR 1000 forwards fragmented UDP packets with ip length ≤ 25 bytes out in clear text.

Workaround: None

- CSCtg93623

On ASR1000 Router, SBC man/sdr does not do policing properly with these steps:

When performing the following:

1. SBC is deactivated using **no activate**
2. SBC configuration is removed using **no sbc global db**
3. SBC is reconfigured without **control-dscp af11 marker-dscp af12 pdr-coefficient 300**

Workaround: Is to enter these commands between steps 1) and 2) above:

no control-dscp af11 marker-dscp af12 pdr-coefficient 300

- CSCtg94290

Execution time of “**copy run start**” delays to 6-10 minutes, with show tech-support simultaneously in different vty.

Workaround: None

- CSCth04143

When running Traditional Netflow (configured with ip flow ingress/egress), with tcp packets in the Netflow cache, if the "show ip cache ver flow" is issued, the tcp flags information will show up at the bottom of each cache entry next to the FFlags token, but the numeric value of tcp flags will be printed (incorrectly) as 0, in the tabular representation of the cache entry.

This defect is observed when running traditional Netflow, when Netflow cache is populated with TCP packets and when the verbose form of the show command is issued i.e.

show ip cache ver flow

Workaround: The workaround is to read the data from the bottom of each cache entry next to the FFlags token.

Note that the integrity of the data exported to the Netflow collector is not affected.

- CSCth08505

Sometimes PPPoE sessions do not sync up with the Standby RP.

This condition may occur on the first attempt when the PPPoE sessions are established and they fail to sync up with the Standby RP.

Workaround: Reloading Standby RP may resolve this problem.

- CSCth09196

On an ASR 1000 Router core file is generated.

This condition is observed when accessing beyond packet data memory.

Workaround: None

- CSCth10088
On ASR 1000, when an invalid request of AuditValue/Modify missing termination ID is received, Context ID is missing in the ER from DBE as the reply.
Workaround: None
- CSCth11039
The following cli did not count the SBC flow pair statistics correctly:
show sbc global dbf flow-pair statistics
It used to count rejected pinholes if a command is rejected with ERR=421.
Workaround : None
- CSCth15353
There are few incorrect result codes seen in the VPDN system logging.
Workaround: None
- CSCth15629
When an ASR 1000 Router receives a IPv6 Neighbor solicitation, the resulting Neighbor Advertisement may not be seen leaving the router.
Workaround: None
- CSCth30815
The Result Code description for STOPCCN were incorrect.
Workaround: None
- CSCth39877
When L2TP tunnel on a ASR 1000 Router goes down, no syslog message was being logged for L2TP tunnel going down.
Workaround: None
- CSCth47836
An Cisco ASR 1000 Router may crash while processing a RTSP packet when a particular type of packet is received.
This problem might appear while processing a RTSP packet when a particular type of packet is received, it may lead to a crash on the router.
Workaround: None

Open Caveats—Cisco IOS XE Release 2.6.1

This section documents possible unexpected behavior by Cisco IOS XE Release 2.6.1

- CSCta46670
When disabling and enabling **control plane host** a few times this may generate an error message on the Cisco ASR 1000 Router. This has been observed, when configuring **control plane host** followed by **no control plane host** a few times on the router.
Workaround: None is required since there appears to be no functional impact.
- CSCtb84718
Output of show cli "sh crypto gdoi gm acl" does not correctly display as a COOP Key Server.

This has been observed, when COOP Key Servers has been configured on the GM.

Workaround: None

- CSCtb98877

On the ASR 1000 Router Series subsequent call fails after a SIP Session Refresh timeout occurs after an HA switchover in CUBE environment.

This occurs in a back to back CUBE environment:

CUCM1 - SIP - CUBE1 - SIP - CUBE2 - SIP - CUCM2

The CUCM SIP Refresh is set to 90 seconds, and a call is made. HA switchover occurs on CUBE1, and the call is disconnected as expected. The same call is made again, but the originating endpoint on CUCM1 gets a Busy tone, while the terminating endpoint on CUCM2 gets Ringing tone.

CUBE2 sends a 503 Internal error with the following cause code:

Reason: Q.850;cause=38 - [Network out of order]

Workaround: None

- CSCtc54288

When changing the group number associated with a virtual IP address causes hosts to lose contact with the virtual router.

This instance occurs when a virtual IP address is associated with one group, and then that group is unconfigured and the same virtual address used by another group. Since each group is uniquely associated with a virtual MAC address the ARP tables of all hosts that were using the previous group will contain invalid entries. When the interface is shut, while configuring new groups then gratuitous ARPs will be sent to refresh any hosts' ARP tables before the interface is ready to forward traffic. The hosts will not realize that the vIP/vMAC association in their ARP tables are invalid and will be unable to forward traffic via the known (virtual IP) gateway.

Workaround: A delay can be used to stall the VRRP initialization process after unshutting an interface:

vrrp delay minimum

reload

The values used are the number of seconds to delay, which is platform dependent. 30 seconds for interface delay and 300 seconds for reload delay are a good first values to test.

- CSCtc55215

On the ASR 1000 Router, when shape rate is configured as % of an ATM PVC for GRE QoS it is not updated after the PVC rate has changed.

This may occur when changes to the PVC rate and its ATM class has been configured on the router at the same time.

Workaround: None

- CSCtc62440

On a Cisco ASR 1000 Router Series, the removal of sub-interfaces may under certain conditions result in MFIB_MRIB-3-FAILED_WIRE_FIND error messages being generated on the Route Processor (RP). There is no functional impact due to this issue.

Workaround: There are no known workarounds.

- CSCtc69297

Tracebacks has been seen with cli **sh platform hardware qfp active feature acl tree** on the Cisco ASR 1000 Router.

This condition has been seen, when there are a huge number of acls configured on the router.

Workaround: None

- CSCtc78745

When deleting a few tunnels from PE side, when CE and PE are having different number of tunnels the Cisco ASR 1000 Router starts throwing msgs.

The following message has been seen:

```
Oct 27 14:33:40.170 IST: %ACE-3-TRANSERR: ASR1000-ESP(14): IKE trans 0x1D8C; opcode 0x60; param 0x1FD2; error 0xA; retry cnt 0
```

This condition has been observed, when tunnel mismatch between CE and PE are kept for a long time on the router.

Workaround: There are no known workarounds as of now.

- CSCtc96467

STANDBY RP reloads twice with **issu runversion**, while downgrading from Release 2.6 to 2.5.

This instance may occur, when Super Package has been configured with ISSU, which causes the STANDBY RP to reload with **issu runversion**.

Workaround: None

- CSCtd48042

When defining vrfa adjacency the vrfa as singal-address is used, this can start an attack and the EP will show in vrfa blacklist on the Cisco ASR 1000 Router.

This instance may occur, when vrfa adjacency has been defined, but vrfa as singal-address is used on the router.

Workaround: None

- CSCtd84323

Under Unified SBC SIP IPv6 to IPv4 scenario with DTMF digits via INFO method, the following Traceback has been seen on the console following RP failover scenario:

```
"*Dec 14 20:59:14.494: %ASR1000_INFRA-5-IOS_INTR_HISTORY: [5|0] [0:0] [0->0] ra[ 1* 0x0 1* 0x0 ] -Process= "SBC main process", ipl= 0, pid= 314". The traceback causes a temporary outage in service, but SBC does recover without any manual intervention.
```

This traceback has been observed in the following conditions:

1. SIP IPv6 to IPv4 calls
2. DTMF digits transferred via INFO method
3. RP failover has been executed at some point in past.

Workaround: None

- CSCtd87114

When rate-limit is configured after msg-body, it is not shown in show logging on the ASR 1000 Router Series.

Workaround: Is to configure rate-limit before msg-body for logging discriminator.

- CSCtd89923

Webex SPA hard disk sectors are corrupted.

This condition has been observed when SIP10 is configured with a Webex SPA running release 2.6.0 image that is Soft-OIR'ed. This configuration can potentially corrupt the sectors on the hard disk of the Webex SPA.

Workaround: Is to shutdown the SPA before reloading the SIP10.

- CSCtd91015

The Cisco ASR 1000 Router does not roll back to the base image even though the rollback timer has expired for ISSU Superpackage Downgrade from Release 2.6 to 2.5. ISSU Superpackage Downgrade does not finish within the specified “roll back” time, but router does not rollback to the base image. Tracelogs shows that the timer has been expired and a user prompt has appeared. But the prompt does not appear on the console. SPA's will move to “inserted” mode and at certain times STANDBY RP will reload.

Workaround: ISSU will work fine, when rollback time is increased.

- CSCtd92548

On the Cisco ASR 1000 Router Series when issuing “issu runversion” the FP, ccp_cp_svr cores on the new Active RP.

This conditions may occur while doing super package issu, after issuing "issu runversion", cpp_cp_svr cores on the new active RP. This has been seen on the multidimensional scaled environment where as both PPPoEQinQ and PPPoEoA are configured on the same Cisco ASR 1000 Router.

Workaround: None

- CSCte01388

The FMAN FP process may crash on the ASR 1000 Router Series.

This has been observed, when VPN has been configured on the router.

Workaround: None

- CSCte07777

The ASR 1000 Router Series may face HQF clean up issues within a QoS ATN PVP enviroment.

This condition may happen on a Cisco ASR 1000 Router when running pre-released image.

Workaround: None

- CSCte17127

Calls are failing due to an invalid tls certificate or they may be completing when the certificate is invalid.

This issue ties into how long the SBC keeps the tcp and tls connection up and also when the ASR 1000 Router does not revalidate the certificates for a deleted or newly added trust point tls peer. The same applies to the scenario where a certificate has to be replaced.

Workaround: Set the tls idle timer to a value of 3 minutes to minimize the time that the tls peer.

- CSCte28845

With Cisco ASR 1000 Router operating in uSBC mode, all adjacencies are locked in Detached state after an upgrade or change where the SBC must be deactivated and activated.

When SBC is deactivate or activated or the same for one of the adjacencies, the system prints a routing error log.

The problem occurs when there is an digit routing entry in the routing table that is missing the destination adjacency datafill.

In most cases the SBC will not allow this to be configured in the first place without throwing an error but there are some scenarios where this configuration can get into the database without an error.

Workaround: Remove the entry with “no dest adjacency” or “add a dest adjacency” to the entry datafill.

- CSCte48047

On a ASR 1000 Router Series the output from the **sh platform software status control- processor** may incorrectly indicate that the ESP committed memory is greater than 100%. There is no functional impact due to this.

Workaround: There are no known workarounds.

- CSCte49434

Upon doing RP switchover on the Cisco ASR 1000 Router, the following error messages are seen on the RP console:

```
*Jan 18 19:11:13.860: %IOSXE-3-PLATFORM: R1/0: kernel:
/scratch/mcpre/BLD-BLD_V122_33_XNF_ASR_RLS6_THROTTLE_LATEST_20100118_18001
2/os/linux/drivers/binos/i2c/pca9535/pca9535_main.c:set_reg_output_port_0 (line 185): write
pca9535 register at 02 failed
```

```
*Jan 18 19:11:13.882: %CMRP-3-I2C_WRITE: R1/0: cmand: An I2C write has failed because
Input/output error -Traceback= 1#4800cc02ea45b52bc53dc957092af093 errmsg:EDA4000+2160
:10000000+24608 :10000000+4C354 :10000000+4A1AC :10000000+46FEC :10000000+47598
:10000000+492F0 :10000000+49B2C evlib:F15B000+D854 evlib:F15B000+FF74
:10000000+311E4 c:E53C000+1D078 c:E53C000+1D220
```

Conditions: This problem has been seen on the a Cisco ASR 1006 Router with redundant RP's when the command “**redundancy force-switchover**” is issued. This problem has been seen very infrequently.

Workaround: There is no workaround. This problem is not affecting the function of the router.

- CSCte61735

Memory leak has been seen when MQC is configured on the Cisco ASR 1000 Router.

This can occur, when QoS has been configured on the router, in an ISG environment.

For example the following conditions have been observed:

```
interface ATM4/0.1 point-to-point
no atm enable-ilmi-trap
pvc 0/101
class-vc crosshairs
vbr-nrt 500 400 50
dbs enable
service-policy in DefaultIn
service-policy out DefaultOut
!
vc-class atm crosshairs
protocol ppp Virtual-Template1
encapsulation aal5snap

interface Virtual-Template1
ip unnumbered Loopback0
ppp authentication chap
end
```

The memory leak occurs when a link is flapped up and down.

Workaround: None

- CSCte78406

On the Cisco ASR 1000 Router console the following error message has been logged on the new standby RP, when PTA sessions are established:

```
*Feb 2 10:21:36.635: %COMMON_FIB-3-FIBIDBINCONS2: An internal software error occurred.
Virtual-Access2.1 linked to wrong idb Virtual-Access2.1
```

This condition may occur, once PTA sessions are established when performing a RP switchover. After both RPs are synced up with flapped sessions. The error messages are logged on the new standby RP.

Workaround: None

- CSCte78938

Xconnect configuration is rejected after replacing the MPLS xconnect configuration with manual L2TPv3 configuration on the ASR 1000 Router Series.

This condition has been seen, when EoMPLS xconnect is configured, while trying to modify the configuration to use L2TPv3 Xconnect on the router.

Workaround: Do not configure L2TPv3 on an interface which previously was used for EoMPLS.

- CSCte82240

SBC accepts “.” when key_addr_type is “DIALED_DIGITS”. This condition can occur, when set exact matching means has been set as:

```
rpsRtgActionKeyAddrWildcardType to AMB_MW_EXPLICIT_WILDCARD.
```

This is possible to have a “.” when rpsRtgActionKeyAddrType is set to AMB_MW_ADDR_TYPE_DIALED_DIGITS. However, it is no longer allowed when rpsRtgActionKeyAddrWildcardType is AMB_MW_EXPLICIT_WCARD (which means SBC should perform an explicit match).

Workaround: None

- CSCtf04444

L2TPv3 sessions do not come up on a Cisco ASR 1000 Router.

The condition sessions do not come up when copying l2tpv3 xconnect on the same Gige interface that used an MPLS xconnect. earlier.

Workaround: Is to write the new l2tpv3 config on nvram and reload the router.

- CSCtf05408

IP address on a loopback interface is lost on the Cisco ASR 1000 Router Series.

Workaround: Is to reconfigure the loopback interface.

- CSCtf07876

The following error message may appear on the ASR 1000 Router console:

```
%IDBINDEX_SYNC-4-RESERVE: Failed to lookup existing ifindex for an interface on the
Standby, when allocating a new ifindex from the Active.
```

This error message has been observed after SSO.

This problem may occur after deleting some configured TE tunnels and when using SSO.

Workaround: None

- CSCtf13608

ESP Kernel crashes when using Release 2.6 image on the ASR 1000 Router Series.

Workaround: None

- CSCtf13925
 CPP back pressures n2 occurs on the FP10.
 This can impact performance when using IPSEC/SSL on the FP10.
 This condition occurs on every record n2 returns to CPP.
 Workaround: None
- CSCtf16427
 Pending objects seen while churning ANCP sessions on a Cisco ASR 1000 Router.
 This condition are seen when the ASR 1000 Router has been configured with Model F with PPPoEoA ANCP sessions.
 In addition, pending objects are seen after churning the ANCP sessions. Workaround: There is no known workaround.
- CSCtf19360
 Memory Leak are seen when processing Crypto IKMP.
 This condition can occur when Vrf Aware EZVPN has been configured on the ASR 1000 EzVPN-Server
 Workaround: None
- CSCtf27659
 After 3 to 4 RP switchovers with a 480 sec sleep in between each RP switchover, the flow-control ids on the active and standby RP are out of synch with the following error:

```
Feb 23 17:29:27.418: %ASR1000_RP_SPA-3-VC_FLOWID_ALLOC_FAIL: Failed
to allocate a flow control identifier for VC 16871576under interface ATM0/1/0
```

 As a result of which all 32k PPPoEoA sessions do not get established.
 This condition has been seen when RP switchover causes the flow-control ids on the active and standby RP to be out of sync, and thus not all sessions getting established.
 Workaround: Is to reload the router.
- CSCtf36152
 When ASR(LAC) receives StopCCN from LNS due to the lack of resources (L2TP session limit), the ASR (LAC) returns ZLB with bad sequence number.
 In this case, the correct Ns/Nr of ZLB should be Ns=1/Nr=1.
 Workaround: None
- CSCtf51834
 After a stateful switchover (SSO) on an IOS router supporting L2TP HA, the counter showing the number of L2TP sessions which were destroyed because they were not completely established at the time of the SSO, may be incorrect.
 This counter is visible with the command show l2tp redundancy detail in the section Sessions destroyed during resync phase.
 For example, the sessions destroyed during resync phase:

```
Poisoned:          0
Unestablished:    10    -- This value may be incorrect
Tunnel in resync: 0
```

 After a stateful switchover (SSO) on an IOS router supporting L2TP HA.

Workaround: No workaround.

- CSCtf65681

This is for CUBE(SP Edition), an SBC application on the ASR 1000 Router Series.

The SBC service failed after receiving a SIP REGISTER response.

This happens when the response comes with a strange expire value, 0xFFFFFFFF.

Workaround: There is no work around on CUBE(SP Edition). But the customer should isolate the entity sending 0xFFFFFFFF as the expire value and disable it.

That is not a normal vale.

- CSCtf69311

Phase2 tunnel History Table MIB values for VRF Aware IPSec is not fetched on the ASR 1000 Router.

This condition may happen when VRF Aware IPSec configuration on UUT and Phase2 IPSec tunnel is cleared with <clear crypto session>.

Workaround: None

- CSCtf69391

Output drops on an interface incrementing apparently due to ISG with session drops.

Conditions: Low traffic may be seen on the interface. This is alarming in that these appear to be actual packet drops from traffic.

Workaround: None

- CSCtf70393

This is a minor problem with CUBE(SP Edition), a SBC service running on the ASR 1000 Router Series.

The is no reason counter increased when SBC rejecting SIP incoming call with 503 response.

This problem happens when certain internal resource is running out because of ongoing signaling traffic.

Workaround: This does not bring down the SBC service, just missing counter.

- CSCtf70450

Very low performance when reassembling MPLS traffic.

This condition has been seen when enabling 'mpls mtu max', then runnning EoMPLSoGRE traffic.

Workaround: None

- CSCtf74687

MEGACO errors 421 and 430 observed on the RP.

This condition are seen in a RP switchover scenario, under high load.

Workaround: None

Further Problem Description: When calls are not fully setup at the time of switchover will not have been replicated to the standby, resulting in 430 when the MGC sends subsequent messages after the switchover. There is also a timing window where replication information is 'in-flight', and gets lost at the time of switchover, resulting in 421 when the MGC attempts to MODIFY the gate after switchover. Both of these symptoms are normal and expected in the course of operation. The increased TAT is a function of the increased load in this environment.

- CSCtf75746

The ASR 1000 RP2 core when using “no sbc” and this can happen sometimes, not everytime.

This appears to happen because a call is waiting for a buffer to send a message, and then the SBC is deactivated. When the buffer comes in, this assert is seen because the call is no longer exists.

Workaround: None

- CSCtf77225

PBR counters under **sh route-map dynamic** shows inconsistency, sometimes the counters increment and sometimes it shows “0”.

The traffic class being monitored is “INPOLICY”, and traffic is being forwarded appropriately.

```
6RU_BR2#sh route-map dynamic
route-map OER_INTERNAL_RMAP, permit, sequence 0, identifier 922746881
Match clauses:
ip address (access-lists): oer#1
Set clauses:
ip next-hop 200.1.1.2
interface GigabitEthernet0/1/5
Policy routing matches: 0 packets, 0 bytes --->>> no matches
Current active dynamic routemaps = 1
6RU_BR2#sh mpls forwarding-table | i 90.1.1
22      No Label   90.1.1.0/24      84075             Gi0/1/0      14.1.1.2      ---->>
packets being forwarded
6RU_BR2#
```

Workaround: None

- CSCtf84496

ESP might generate fman-fp-image core when remote peers have burst with aggressive IPSec rekey activity.

This condition has been observed when 2K svti crypto session flaps per 30 seconds.

Workaround: None

- CSCtf93465

In a CUBE(SP Edition) ASR 1000 Router, the following message is seen when trying to enter SBC config mode:

```
SBC: Internal error - SBC configuration cannot be processed.
```

This condition sometimes happens after unconfiguring SBC.

Workaround: The workaround is to do a reload.

- CSCtf95136

SIP notify to RFC2833 DTMF interworking failed.

SBC did not correctly signal DTMF interworking capabilities in a flow where SBC is configured to support DTMF NOTIFY for the caller.

This condition has been seen when n 2000K sends from SBC to caller side, there is no call-info header.

Workaround: None

- CSCtg04257

After SBC hit the system congestion the following messages are seen on the console:

```
*Apr  3 02:06:34.956: %SBC-2-MSG-3802-0432-BA2C8E-1302: SBC/SIP:
SBC is currently congested and is rejecting new call requests.
*Apr  3 02:06:37.966: %SBC-2-MSG-3802-0432-BA2C8E-1302: SBC/SIP:
```

SBC is currently congested and is rejecting new call requests.

Then SBC quits to process all the messages. This problem may occur when congestion occurs on a Cisco ASR 1000 Router .

Workaround: None

- CSCtg07737

More than 70 sec multicast traffic loss occurred after performed CC software OIR.

Workaround: None

- CSCtg09182

FMAN traceback messages are seen on the Cisco ASR 1000 Router console:

```
%FMFP_URPF-3-LIST_DOWNLOAD: F1: fman_fp_image: Unicast RPF list create for list 10594
fail to download because No such file or directory.
```

This condition may occur when the ASR 1000 Router has 6000 subscribers Loaded with stateful traffic from an IXIA, 15000 vanilla PPP (pass policy) in V-T 4, 2500 flapping subscribers with FW, and 2500 flapping subscribers with no firewall. The failure is on a 10G interface carrying MPLS.

Workaround: There is no known workaround.

- CSCtg16498

LNS VPDN message is incorrect when receiving CDN from LAC as follows:

```
%VPDN-4-SESSIONERROR: L2TP LNS R102 unable to terminate user cisco@cisco.com; Result
1, Error 0, No disconnect reason given
```

It should start with “%VPDN-4-SESSIONERROR: L2TP LAC”.

This occurs when receiving CDN's result code is “1” and the following is configured on the router:

vpdn logging is enabled

Workaround: There is no workaround.

- CSCtg21602

In this example the following message is displayed on the Active RP Console:

```
%CPPOSLIB-3-ERROR_NOTIFY: F1: cpp_cp: cpp_cp encountered an error -Traceback=
1#ed4b69bc77a25ff35c522388cbb72a96 errmsg:D94E000+2160 cpp_common_os:E0A4000+B920
cpp_common_os:E0A4000+19148 cpp_exmem_mgr:ED1E000+895C cpp_exmem_mgr:ED1E000+9080
cpp_common_os:E0A4000+163D0 cpp_rrm_local_api_lib:EFDF000+3150
cpp_rrm_svr_lib:F004000+53F0 cpp_rrm_svr_lib:F004000+76CC cpp_rrm_svr_lib:F004000+80A4
cpp_rrm_svr_lib:F004000+9810 cpp_dmap:E954000+15264 cpp_dmap:E954000+21BF4
cpp_common_os:E0A4000+
```

This message is displayed right after un-configuring sbc by entering command:

no sbc global dbc

Workaround: None

- CSCtg30995

Delay of RP switchover associated with NV_BLOCK_INITFAIL message appearing on standby-turned-active RP console.

When the manual switchover command, **redundancy force-switchover** and a filesystem command like **copy running-config startup-config** is issued from different consoles of active RP almost at the same time, such a problem may occur.

Workaround: Avoid issuing any filesystem access command simultaneously with the manual RP switchover command. Should the above problem occur, please execute the same filesystem command, say, **copy running-config startup-config** from the standby-turned-active console.

- CSCtg37082

Following warning messages are seen on active RP upgrade (ISSU) and switchover after upgrading to Release IOS XE 2.6.1 from 2.3.0e:

```
%IMRP-3-IMRP_MSG_CANNOT_RELAY: R1/0: imand: IMRP Peer ios_rp_iosd_slot_1:
cannot relay message to SPA 15/15
```

```
%IOSD_IMCC_CAPI-3-MSGDISPATCH: SIP2/1: Unable to dispatch received TDL
```

After ISSU upgrading the active RP to IOS XE Release 2.6.1 from 2.3.0e, when RP switchover is initiated this warning will be seen on the newly active RP.

The warning messages do not affect the general working of the SPA during ISSU when the VLAN Tunnel mode feature is not used.

Workaround: None

- CSCto03123

Symptoms:

1. A slow memory leak is observed on the cman_fp process on an FP and the cmcc process on a SIP. This issue is seen on all the flavors for FPs and CCs. The leak is of the order of less than 100-122K bytes per day.
2. Additional memory leak can occur when frequent sensor value changes take place.

Conditions: This symptom of the first leak does not occur under any specific condition. The second leak occurs when sensor-related changes take place.

Workaround: There is no workaround.

PSIRT Evaluation: The Cisco PSIRT has evaluated this issue and does not meet the criteria for PSIRT ownership or involvement. This issue will be addressed via normal resolution channels.

If you believe that there is new information that would cause a change in the severity of this issue, please contact psirt@cisco.com for another evaluation.

Additional information on Cisco's security vulnerability policy can be found at the following URL:
http://www.cisco.com/en/US/products/products_security_vulnerability_policy.html

Resolved Caveats—Cisco IOS XE Release 2.6.1

All the caveats listed in this section are resolved in Cisco IOS XE Release 2.6.1

- CSCso88429

CME or CUBE will reject an inbound SIP INVITE if Max-Forwards is greater than 70.

The symptoms are observed when a Max-Forwards header field in SIP INVITE is greater than 70.

Workaround: There is no workaround.

Further Problem Description: From RFC 3261: 20.22 Max-Forwards

- CSCsq24672

A call through CUBE may not establish for a Re-Invite-based call flow. The call may drop.

This symptom is observed if the endpoint to which the CUBE is communicating sends a Re-INVITE for a call before it has received an ACK from the other call leg for the original INVITE. CUBE may not forward this Re-Invite to the other call leg, and the call will disconnect.

Workaround: There is no workaround.

- CSCsq57238

An interface is congested. A QoS policy-map is applied to the interface such that one of the traffic-classes receives only infrequent packets. That traffic class is seen to have higher than expected latency. If steady traffic is sent through the same traffic class, then latency is as expected and bandwidth is seen to be shared between traffic classes as per their relative bandwidth guarantees.

The symptoms are observed on any interface, but is most obvious with low speed interfaces such as ATM PVCs with 256k or less bandwidth.

Workaround: If the traffic class with the infrequent traffic is configured with **“priority”**, then latency will be minimized.

- CSCsv98245

The output of the "show ip bgp neighbor x.x.x.x advertised-routes" displays "Originating default network 0.0.0.0" even when default network is not expected to be originated.

This may be seen in the **“show ip bgp neighbor x.x.x.x advertised-routes”** output, even if there are no routes being advertised, as shown in the example:

```
Router# show ip bgp neighbor 10.0.10.10 advertised-routes
BGP table version is 3, local router ID is 10.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
r RIB-failure, S Stale, m multipath, b backup-path, x best-external
Origin codes: i - IGP, e - EGP, ? - incomplete
Originating default network 0.0.0.0
Network          Next Hop          Metric LocPrf Weight Path
Total number of prefixes 0
```

Workaround: No workaround.

- CSCsx66105

Chunk memory leaks at "SADB SA Header" are seen on the Group Member. The memory leaks are seen when ipsec SAs are cleared using the command "clear crypto gdoi"

Workaround: No Workaround

- CSCsz07615

When **reload** cli is issued it takes some time to gracefully bring down the system. In this process it will unusually take longer time for redundancy protocols to notify its peers and cause few timing issues.

This has been only observed when redundancy protocols like HSRP are seen to misbehave.

Workaround: Is to bring down the system instead of gracefully reloading.

- CSCsz69148

When running an embedded syslog manager (ESM) TCL script on the ASR to filter logs, memory leaks in IOSD ipc task and ESM Logger occur. This memory leak applies to RP1 and RP2. Any feature which uses heavy logging (for example, audit logging for firewall features) will exploit this issue readily.

Workaround: There is not current workaround other than to not use the ESM feature.

- CSCsz82950

On the Cisco ASR 1000 Router a peer RP reloads. If any configurations are done using NMS for DCTM MIB, this symptom occurs when unconfiguring the configuration that is created by DCTM MIB configuration.

There is no workaround.

Further Problem Description: DCTM was not HA supported before. HA is supported

- CSCta21525

High CPU is seen when sending SNMP queries to a router just configured to send TRAPS. The host sending the query must have the correct TRAP community.

When configuring: **'snmp-server host 10.13.37.1'**, the router will start to listen to SNMP-requests sent to the router. The router sets the community to public if nothing else is specified.

Workaround: Is to secure SNMP with the use of ACL's & and add a more secure SNMP community for the above command.

- CSCta26492

OSPFv3 does not advertise prefixes related to virtual-access interface on a Cisco ASR 1000 Router. This has been seen when configuring **"ipv6 ospf"** under virtual-template on the router.

Workaround: configuring "redistribute connected"

- CSCta58068

During BGP convergence, CPU spike may be seen on the local PE in an MVPN configuration after conditions.

This condition happens to cause excessive BGP convergence and high CPU utilization (with and without traffic) in an MVPN setup can be varied as:

- Remote PE neighbor switchover
- On local PE, while doing a "clear ip bgp <bgp nbr>".
- On bringing up the local PE
- Large configuration such as one with 300 MDT default tunnels.

Here is an example of an MVPN configuration where this problem can be exhibited:

1. OSPF routing protocol is enabled on all the networks in the topology.
2. Each PE router, has 100 MVRFs defined (between vpn_0 to vpn_99).
3. MDT default is configured on all the mVRFs on the PE routers.
4. PE routers have an iBGP session, ONLY with the RR (route-reflector).
5. eBGP session exists between the Router and PE1, with Router sending 200,010 VPNv4 routes.
6. OSPF session also exists between Router and PE1, with Router sending 100 OSPF routes.

In effect the following states are present in the network:

On PE and RR routers:

1. IGP states = 100 OSPF routes
2. BGP states = 200,010 VPNv4 route

On PE routers ONLY:

1. VRF sessions = 100 VRFs (vpn0 to vpn_99)
2. MDT sessions = 100 SSM session

Workaround: None.

- CSCtb05810

When applying the no distance command, the summary-prefix disappears from the route table. When you check the OSPFv3 database, the summary route exists.

Workaround: Is to configure summary-prefix command again.

- CSCtb62351

The output of '**show ip vrf detail <vrf_name>**' shows the number of routes greater than the actual number of routes in a VRF.

This problem happens when the route count does not get updated correctly when a route with a better distance replaces an existing route only occurs when the route in question is also a major network.

Workaround: Is to issue the command '**clear ip bgp vpnv4 vrf ***'.

- CSCtb62689

When system is congested and a switchover occurs SBC drops calls instead of answer with the SIP congestion message 503 that indicates that the system is congested.

The problem only occurs after a switchover of a congested system.

Workaround: None

- CSCtb86371

TCP packets from client requiring PBHK are silently drop by ISG router.

After a TCP connection is idle with no activity for more than 60 seconds, the PBHK portmapping will be removed to protect the ISG from losing memory on lost connection. However if a client tries to use an existing TCP connection it has with a server after a 60 seconds of inactivity, those packets are dropped, giving a hang TCP connection. The behavior expected by ISG should be to return a TCP RST.

Workaround: Is to force a reset or close the TCP client connection.

- CSCtb89424

In rare instances, a Cisco ASR 1000 Router may crash while using IP SLA udp probes configured using SNMP and display an error message similar to the following:

```
hh:mm:ss Date: Address Error (load or instruction fetch) exception, CPU signal 10, PC = 0x424ECCE4
```

This symptom is observed while using IP SLA on the router.

Workaround: There is no workaround.

- CSCtb94498

The RP resets at IGMP Input on a Cisco ASR 1000 Router.

This issue has been seen only when repeated config and unconfig of the VRFs is performed.

Workaround: The issue will not be seen if there is more than 5secs time delay between the config/unconfig.

- CSCtb96865

When LAC fails to attempt to establish L2TP session, it will add LNS's ip addresses to dead cache entry, even if the reason of failure due to LAC internal error. The LNS's ip-address will be put on busy list for 60 seconds due to this.

This conditon has been seen when LAC has internal failure and fails to create l2tp tunnel to lns.

Workaround: There is no known workaround.

- CSCtc12904

A crash in IPv6 ND processing a timer wheel on the Cisco ASR 1000 Router.

This is a very rare condition that might happen as a neighbor is removed from the ND cache.

Workaround: None

- CSCtc18656

When the NAT box is configured as the Rendezvous Point (RP). This does not allow for source address translation for the encaps packet received from the First Hop Router. NAT box is configured as Rendezvous Point (RP) decapsulates the packet and forwards it to NAT outside without translation which will create incorrect S,G state for a inside local source address on the downstream routers after NAT router.

Workaround: None

- CSCtc31545

Some EIGRP routes may not be installed in the routing table after a link flaps. The route is seen as “active” in the EIGRP topology table, and the active timer is “never”.

This has been seen when a backup path has an EIGRP composite metric of infinity and the primary path is flapping.

Workaround: Is to use realistic metrics for all paths rather than very high delay values which may result in an infinite metric. Once in this condition, the only way to resolve the issue is to clear the neighbors.

- CSCtc35416

VPDN debugs are enabled on the ASR 1000 Router Series. When username conditional debugging is **enabled**, VPDN debugs should not be printed, but a few debug messages are still seen.

Workaround: None

- CSCtc37349

Under router distribute-list prefix command ACL option is shown. Which is not correct.

As a result the access list is configured, which has prevented configuring prefix name with numbers.

This condition are seen when router ospf 10 distribute-list prefix in used on the ASR 1000 Router.

Workaround: None.

- CSCtc42941

On a Cisco ASR 1000 Router the Standby is not coming up.

When a distribute-list is configured, the ACL is created if it does not exist. Then remove the ACL, but the distribute-list configuration that ties to the ACL is not removed. Configure the IPv6 ACL configuration with the same ACL name. Save the configuration and reload it.

Workarounds:

1. When a access list is removed, remove corresponding distribute-list configuration as well.
2. Do not use the same access list name for IPv4 and IPv6.

Further Problem Description: Is to use router bgp 100 distribute-list sample in exit and no ip access-list standard sample ipv6 access-list sample permit any write to the memory.

- CSCtc51539

A Cisco ASR 1000 Router crashes with a “Watch Dog Timeout NMI” error message.

This symptom is observed only on devices configured with Bidirectional Forwarding Detection (BFD). For further information on BFD, consult the following link:

http://www.cisco.com/en/US/docs/ios/12_0s/feature/guide/fs_bfd.html

Workaround: Is to disable BFD.

- CSCtc57356

IOS SLB natpool will source nat addresses outside the range specified in config.

This has been observed when IOS SLB with client nat and a nat pool with same address as start and end range and 31 bit mask.

Workaround: s to USE 2 (two) ip addresses in range and not not allow for other devices including backup IOS SLB or other slb in same network to use the same range.

- CSCtc68037

A Cisco ASR 1000 Router device may experience an unexpected reload as a result of mtrace packet processing.

Workaround: None other than avoiding the use of mtrace functionality.

- CSCtc87430

The following errors are seen on Active RP during RP switchover with scaled sessions(around 5k):

“asr1000 bssess: RPC header processing failed, error=5001”

This conditions are seen in an environment that is using the following setup:

Agilent--->(atm)--->ASR 1000 Router ---->(10GB)----->LNS(c10k)----->Agilent

and with the following configurations as shown in this example:

1. Configure the Cisco ASR 1000 Router as LAC with Model D2.1 QOS
2. Start session bring-up
3. At around 5k session, issue RP SWO
4. Noticed errors on new Active RP

Workaround: None

- CSCtd24065

The output of the command **show subscriber statistics** shows that number of “SHDBs in use” is greater than the total number of unique subscribers for the deployment. This might contribute to issues such as an “out of IDs” message or sessions not coming up.

The symptom occurs for DHCP-initiated sessions either when:

1. Session idle times out followed by a lease expiry or you release the lease.
2. Session is cleared using the **clear subscriber session** command and there is a lease expiry or you release the lease.

Workaround: There is no workaround.

Further Problem Description: This can also contribute to a small amount of observed memory leak. This problem occurs in code branches where IP session HA is not supported. In these branches, the above steps cause a SHDB handle to not be cleared properly when other datastructures are cleared.

- CSCtd25664

ERSPAN session are not sending traffic to the analyser on the Cisco ASR 1000 Router Series. ERSPAN session are not filtering traffic as expected on the router.

This condition has been observed, when the Cisco ASR 1000 Router is running 12.2(33)XND1 and previous versions.

1. ERSPAN configured with vlan filtering
2. ERSPAN configured with vlan sourcing

Workaround: Is to do the following:

1. Filter traffic on the analyser
2. There is no known workaround.

- CSCtd32560

During Cisco ASR 1002 or Cisco ASR 1004 ISSU upgrade from IOS XE 2.3.2 to IOS XE 2.5.0, a loss of QoS functionality can occur on some and all targets.

Loss of QoS functionality has been observed right after RP upgrade and switchover while following Cisco ASR 1002 or Cisco ASR 1004 ISSU procedure. The QoS functionality does not recover on its own and only occurs on policies that are both hierarchical (at least 2-level) and contain policers. The condition can be identified by the following command:

show platform hardware qfp active interface if-name <if_name> info | include QoS

If there is no output returned from this command then there has likely been a QoS service disruption due to this problem.

Workaround: QoS functionality can be resumed on the interface by removing and re-attaching the QoS policy. Alternately, the problem can be avoided by upgrading to IOS XE 2.4.x first (including the ESP). The upgrade path would be IOS XE2.3.2 -> IOS XE 2.4.x -> IOS XE 2.5.x.

- CSCtd33780

When there are two vrfs A (exporting) and B (importing), admin shutting down the neighbor under vrf A while doing:

'sh ip bgp vpnv4 vrf seven <prefix_learned_from_neighbor_shut_under_vf_one>' with auto-more causes the Cisco ASR 1000 Router to crash.

This condition are seen for this timing issue, the following events MUST happen VERY FAST to be reproduced:

1. Add vrf "seven" which imports paths from existing vrf (vrf "one")
2. Admin **shut** the neighbor under vrf "one" while simultaneously doing

'<sh ip bgp vpnv4 vrf seven <prefix_learned_from_neighbor_shut_under_vf_one>'

with auto-more in another terminal

Workaround: No workaround.

- CSCtd36639

The following traceback are seen %ASR1000_INFRA-5-IOS_INTR_OVER_LIMIT while running overnite with 64 groups and 1500 OIF's each with no traffic:

%ASR1000_INFRA-5-IOS_INTR_OVER_LIMIT: IOS thread disabled interrupt for 17 msec

```
-Traceback= 1#cbbad1ee0e5052dcf6fb7a00f91e7a88 :400000+D1A6A9 :400000+40476FD
:400000+404849B :400000+250CF16 :400000+250CD07 :400000+2510305 :400000+15941A7
:400000+41067B0 :400000+41064D6
```

This conditions keeps a scaled configuration running overnight without any traffic.

Workaround: None

- CSCtd43965

The command **snmp context** *<context>* is not available to be configured for EIGRP multicast address-families. Furthermore, on images that support EIGRP multicast address-families but do not have the EIGRP MTR (Multi-Topology Routing) plug in, this command is not available to be configured for EIGRP unicast address-families either. As a result, it is not possible to associate these address-families with a specific SNMP context.

This conditions are affects of any image that has EIGRP release 2.5 or later code and supports EIGRP multicast address-families.

Workaround: No workaround.

- CSCtd48455

When “ip summary-address eigrp ...” advertises a subset of component routes in addition to the summary.

This symptom happens when the number of redistributed prefixes extends beyond approximately 100 routes.

Workaround: Is to do the following:

1. Assign an IP address which is a component of the summary to a connected interface; ie. a loopback interface.
2. Run **clear ip eigrp neighbor soft** (will not reset adjacency) **clear ip eigrp neighbor** (will reset adjacency)

Further Problem Description: This problem reappears if **clear ip route *** is executed.

- CSCtd48480

Memory leak are seen in the function ppp_aaa_apply_peruser_attributes.

This condition are seen while initiating and clearing ppoe call consecutively resulting the memory leak.

Workaround: No workaround.

- CSCtd54970

Internal routes may only be tagged with values less than or equal to 255. With this defect, the tag was allowed to be set to higher values if a route-map was applied outbound on a specified interface.

Whe the outbound route-map is applied to a specific interface with set tag great than 255 for internal routes.

Workaround: Set tag to a value less than or equal to 255.

- CSCtd56668

The Cisco ASR 1000 Router retains Multicast MAC entry in HSRPDA TCAM even after port 'shut'. This may possibly lead to packet duplication.

When virtual MAC is added to HSRPDA TCAM of active RP upon configuring HSRP between two ASR 1000 Routers. This entry persists after port **shut** on the active.

Workaround: There is no known workaround.

- CSCtd61194

When configuring ERSPAN on FastEthernet, gives an error:

“SPAN is not supported on SPA interface”

ASR1K(config-mon-erspan-src)#source interface ?FastEthernet FastEthernet IEEE 802.3GigabitEthernet GigabitEthernet IEEE 802.3z Port-channel Ethernet Channel of interfaces TenGigabitEthernet Ten Gigabit Ethernet ASR1K(config-mon-erspan-src)#source interface fastEthernet 0/3/0 SPAN is not supported on SPA interface (FastEthernet0/3/0)ASR1K(config-mon-erspan-src)#

Conditions: This condition has been observed when running release Version 2.4.1.

Workaround: Is to Use GigabitEthernet.

- CSCtd63242

A traceback or crash may be seen when deleting a subinterface that has IPv6 EIGRP running on it.

This condition are seen when ipv6 eigrp <as> is configured on a subinterface and the a **no interface <subinterface>** is entered.

Workaround: Is to Remove eigrp from the subinterface with a "no ipv6 eigrp <AS>" before deleting the subinterface.

- CSCtd66013

On Cisco ASR 1000 Service Series Routers, Last reload reason: 4 is displayed in show version output after power-cycle.

This symptom is observed after RP switchover and then power-cycle.

Workaround: None. This issue is cosmetic and does not affect traffic and operation on the router.

- CSCtd66189

Route-map with 'set pv6 next-hop peer-address' in 6PE setup sets wrong NH, such as A00:4D:: (for ipv4 address of the peer 10.0.0.77) instead of ::FFFF:10.0.0.77 - ipv4-mapped address. Route-map with 'set pv6 next-hop peer-address' can be either incoming or outgoing

Workaround: None known at this time

- CSCtd68197

Memory leak might be seen in IPv6 RIB Redistribute process.

This happens when the router has ipv6 eigrp neighbor and has some ipv6 eigrp router. Memory leak might be seen with following step.

1. When changing ipv6 address for loopback interface on ipv6 eigrp neighbor router.
2. Shutdown the connected interface on ipv6 eigrp neighbor router. --> reducing the holding of IPv6 RIB Redistribute.
3. No shutdown thNe connect interface on ipv6 eigrp neighbor router. --> returning the holding value.
4. Repeat to 1.

This is ONLY seen when an ipv6 eigrp process is partially configured. If fully configured, it does not occur and this is the easiest workaround to make sure that the eigrp process is configured and there are interfaces participating (even just a loopback) with that eigrp process (**sh ipv6 eigrp <asnum> interface**). Reproducibility is not 100%, but after 2. or 3. operation, sometimes holding value might be increased from previous value of same condition.

Workaround: Is to - configure the IPv6 eigrp process completely - deconfig IPv6 eigrp process that is not running.

- CSCtd69644

ISSU scripts indicate load version failure due to WMA SPA being offline.

While, attempting a superpkg issue upgrade, when we issue an ISSU load version command, the WMA spa seems to dissappear from the output of the "show hw-module subslot all oir" command for a few seconds, only to re-appear later after approx 15-20+ seconds. This causes the ISSU libraries to indicate that not all SPA's came online after loadversion. This issue is not seen with the GE, POS, CT3 SPA's in the system.

Workaround: Suggested workaround is to wait for 15-20 seconds after load version completes, prior to checking the output of **show hw-module subslot all oir** command to verify WMA SPA coming online.

- CSCtd70582

Traffic Class services will remain in **show subscriber session** output under "Policy Information" after traffic class has disconnected by timer events.

This conditions are only seen when Traffic Class is disconnected through an Idle Timer or Absolute Timer expiring.

Workaround: If traffic class service is disconnected through normal (User Intervention), issue is not seen. For Timer disconnected Traffic Class services, no known workaround at this time.

- CSCtd72441

On a Cisco ASR1000 series router, when the command **show platform software wccp <service-id> counters** is executed, the obj_id field in the output in rare situations maybe a large negative number. It is a cosmetic issue and does not affect functionality.

This condition are seen when WCCPv2 is configured on the router and is redirecting traffic. The object id value is greater that 2147483647. The command **show platform software wccp <service-id> counters** is executed

Workaround: There is no workaround. However there is no functionality impact.

- CSCtd73567

The Cisco ASR 1000 Router Series may reload unexpectedly while reassembling a fragmented ip packet.

Workaround: None

- CSCtd75033

Cisco IOS Software is affected by NTP mode 7 denial-of-service vulnerability.



Note Note: The fix for this vulnerability has a behavior change affect on Cisco IOS Operations for Mode 7 packets. See the section Further Description of this release note enclosure.

Cisco IOS Software with support for Network Time Protocol (NTP) contains a vulnerability processing specific NTP Control Mode 7 packets. This results in increased CPU on the device and increased traffic on the network segments.

ntp master <any following commands>

ntp peer <any following commands>

ntp server <any following commands>

ntp broadcast client

ntp multicast client

The following example identifies a Cisco device that is configured with NTP:

```
router#show running-config | include ntp
```

```
ntp peer 192.168.0.12
```

The following example identifies a Cisco device that is not configured with NTP:

```
router#show running-config | include ntp
router#
```

To determine the Cisco IOS Software release that is running on a Cisco product, administrators can log in to the device and issue the show version command to display the system banner. The system banner confirms that the

device is running Cisco IOS Software by displaying text similar to “Cisco Internetwork Operating System Software” or “Cisco IOS Software.” The image name displays in parentheses, followed by “Version” and the Cisco IOS Software

release name. Other Cisco devices do not have the show version command or may provide different output.

The following example identifies a Cisco product that is running Cisco IOS Software Release 12.3(26) with an installed image name of C2500-IS-L:

```
Router#show version
Cisco Internetwork Operating System Software
IOS (tm) 2500 Software (C2500-IS-L), Version 12.3(26), RELEASE SOFTWARE (fc2)
Technical Support: http://www.cisco.com/techsupport
Copyright ) 1986-2008 by cisco Systems, Inc.
Compiled Mon 17-Mar-08 14:39 by dchih
<output truncated>
```

The following example shows a product that is running Cisco IOS Software release 12.4(20)T with an image name of C1841-ADVENTERPRISEK9-M:

```
Router#show version
Cisco IOS Software, 1841 Software (C1841-ADVENTERPRISEK9-M), Version
12.4(20)T, RELEASE SOFTWARE (fc3)
Technical Support: http://www.cisco.com/techsupport
Copyright ) 1986-2008 by Cisco Systems, Inc.
Compiled Thu 10-Jul-08 20:25 by prod_rel_team
<output truncated>
```

Additional information about Cisco IOS Software release naming conventions is available in "White Paper: Cisco IOS Reference Guide" at the following link:

<http://www.cisco.com/warp/public/620/1.html>

Workaround: There are no workarounds other than disabling NTP on the device. The following mitigations have been identified for this vulnerability; only packets destined for any configured IP address on the device can exploit this vulnerability. Transit traffic will not exploit this vulnerability.



Note Note: NTP peer authentication is not a workaround and is still a vulnerable configuration.

* NTP Access Group

Warning: Because the feature in this vulnerability utilizes UDP as a transport, it is possible to spoof the sender's IP address, which may defeat access control lists (ACLs) that permit communication to these ports from trusted IP addresses. Unicast Reverse Path Forwarding (Unicast RPF) should be considered to be used in conjunction to offer a better mitigation solution.

```
!--- Configure trusted peers for allowed access
access-list 1 permit 171.70.173.55
!--- Apply ACE to the NTP configuration
ntp access-group peer 1
```

For additional information on NTP access control groups, consult the document titled "Performing Basic System Management" at the following link:

http://www.cisco.com/en/US/docs/ios/netmgmt/configuration/guide/nm_basic_sys_manage.html#wp1034942

* Infrastructure Access Control Lists

Warning: Because the feature in this vulnerability utilizes UDP as a transport, it is possible to spoof the sender's IP address, which may defeat ACLs that permit communication to these ports from trusted IP addresses. Unicast RPF should be considered to be used in conjunction to offer a better mitigation solution.

Although it is often difficult to block traffic that transits a network, it is possible to identify traffic that should never be allowed to target infrastructure devices and block that traffic at the border of networks.

Infrastructure ACLs (iACLs) are a network security best practice and should be considered as a long-term addition to good network security as well as a workaround for this specific vulnerability. The iACL example below should be included as part of the deployed infrastructure access-list, which will help protect all devices with IP addresses in the infrastructure IP address range:

```
!---
!--- Feature: Network Time Protocol (NTP)
!---
access-list 150 permit udp TRUSTED_SOURCE_ADDRESSES WILDCARD
    INFRASTRUCTURE_ADDRESSES WILDCARD eq 123
!--- Note: If the router is acting as a NTP broadcast client
!---   via the interface command "ntp broadcast client"
!---   then broadcast and directed broadcasts must be
!---   filtered as well. The following example covers
!---   an infrastructure address space of 192.168.0.X
access-list 150 permit udp TRUSTED_SOURCE_ADDRESSES WILDCARD
    host 192.168.0.255 eq ntp
access-list 150 permit udp TRUSTED_SOURCE_ADDRESSES WILDCARD
    host 255.255.255.255 eq ntp
!--- Note: If the router is acting as a NTP multicast client
!---   via the interface command "ntp multicast client"
!---   then multicast IP packets to the mutlicast group must
!---   be filtered as well. The following example covers
!---   a NTP multicast group of 239.0.0.1 (Default is
!---   224.0.1.1)
access-list 150 permit udp TRUSTED_SOURCE_ADDRESSES WILDCARD
    host 239.0.0.1 eq ntp
!--- Deny NTP traffic from all other sources destined
!---   to infrastructure addresses.
access-list 150 deny udp any
    INFRASTRUCTURE_ADDRESSES WILDCARD eq 123
!--- Permit/deny all other Layer 3 and Layer 4 traffic in
!---   accordance with existing security policies and
!---   configurations. Permit all other traffic to transit the
!---   device.
access-list 150 permit ip any any
!--- Apply access-list to all interfaces (only one example
!---   shown)
interface fastEthernet 2/0
    ip access-group 150 in
```

The white paper entitled "Protecting Your Core: Infrastructure Protection Access Control List" presents guidelines and recommended deployment techniques for infrastructure protection access lists and is available at the following link:

http://www.cisco.com/en/US/tech/tk648/tk361/technologies_white_paper09186a00801a1a55.shtml

* Control Plane Policing

Provided under Control Plane Policing there are two examples. The first aims at preventing the injection of malicious traffic from untrusted sources, whilst the second looks at rate limiting NTP traffic to the box.

- Filtering untrusted sources to the device.

Warning: Because the feature in this vulnerability utilizes UDP as a transport, it is possible to spoof the sender's IP address, which may defeat ACLs that permit communication to these ports from trusted IP addresses.

Unicast RPF should be considered to be used in conjunction to offer a better mitigation solution.

Control Plane Policing (CoPP) can be used to block untrusted UDP traffic to the device. Cisco IOS software releases 12.0S, 12.2SX, 12.2S, 12.3T, 12.4, and 12.4T support the CoPP feature. CoPP can be configured on a device to help protect the management and control planes and minimize the risk and effectiveness of direct infrastructure attacks by explicitly permitting only authorized traffic that is sent to infrastructure devices in accordance with existing security policies and configurations. The CoPP example below should be included as part of the deployed CoPP, which will help protect all devices with IP addresses in the infrastructure IP address range.

```
!--- Feature: Network Time Protocol (NTP)
access-list 150 deny udp TRUSTED_SOURCE_ADDRESSES WILDCARD any eq 123
!--- Deny NTP traffic from all other sources destined
!--- to the device control plane.
access-list 150 permit udp any any eq 123
!--- Permit (Police or Drop)/Deny (Allow) all other Layer3 and
!--- Layer4 traffic in accordance with existing security policies
!--- and configurations for traffic that is authorized to be sent
!--- to infrastructure devices
!--- Create a Class-Map for traffic to be policed by
!--- the CoPP feature
class-map match-all drop-udp-class
  match access-group 150
!--- Create a Policy-Map that will be applied to the
!--- Control-Plane of the device.
policy-map drop-udp-traffic
  class drop-udp-class
    drop
!--- Apply the Policy-Map to the
!--- Control-Plane of the device
control-plane
service-policy input drop-udp-traffic
```

In the above CoPP example, the access control list entries (ACEs) that match the potential exploit packets with the "permit" action result in these packets being discarded by the policy-map "drop" function, while packets that match the "deny" action (not shown) are not affected by the policy-map drop function.

- Rate Limiting the traffic to the device

The CoPP example below could be included as part of the deployed CoPP, which will help protect targeted devices from processing large amounts of NTP traffic.

**Warning**

Warning: If the rate-limits are exceeded valid NTP traffic may also bedropped.

```
!--- Feature: Network Time Protocol (NTP)
access-list 150 permit udp any any eq 123
!--- Create a Class-Map for traffic to be policed by
!--- the CoPP feature
class-map match-all rate-udp-class
match access-group 150
!--- Create a Policy-Map that will be applied to the
!--- Control-Plane of the device.
!--- NOTE: See section "4. Tuning the CoPP Policy" of
!--- http://www.cisco.com/web/about/security/intelligence/coppwp_gs.html#5
!--- for more information on choosing the most
!--- appropriate traffic rates
policy-map rate-udp-traffic
class rate-udp-class
police 10000 1500 1500 conform-action transmit
exceed-action drop violate-action drop
!--- Apply the Policy-Map to the
!--- Control-Plane of the device
control-plane
service-policy input drop-udp-traffic
```

Additional information on the configuration and use of the CoPP feature can be found in the documents, “Control Plane Policing Implementation Best Practices” and “Cisco IOS Software Releases 12.2 S - Control Plane Policing” at the following links:

http://www.cisco.com/web/about/security/intelligence/coppwp_gs.html and

http://www.cisco.com/en/US/docs/ios/12_3t/12_3t4/feature/guide/gtrtlmt.html

Further Description: Cisco IOS Software releases that have the fix for this Cisco bug ID, have a behavior change for mode 7 private mode packets.

Cisco IOS Software release with the fix for this Cisco bug ID, will not process NTP mode 7 packets, and will display a message “NTP: Receive: dropping message:

Received NTP private mode packet. 7” if debugs for NTP are enabled.

To have Cisco IOS Software process mode 7 packets, the CLI command <cmd>ntp allow mode private</cmd> should be configured. This is disabled by default.

This is the same as the vulnerability which is described in:

<http://www.kb.cert.org/vuls/id/568372>

Cisco has release a public facing vulnerability alert at the following link:

<http://tools.cisco.com/security/center/viewAlert.x?alertId=19540>

Cisco IOS Software that has support for NTPv4 is NOT affected. NTPv4 was introduced into Cisco IOS Software: 12.4(15)XZ, 12.4(20)MR, 12.4(20)T, 12.4(20)YA, 12.4(22)GC1, 12.4(22)MD, 12.4(22)YB, 12.4(22)YD, 12.4(22)YE and 15.0(1)M.

All other versions of Cisco IOS and Cisco IOS XE Software are affected.

To see if a device is configured with NTP, log into the device and issue the CLI command **show running-config | include ntp**. If the output returnseither of the following commands listed then the device is vulnerable:

- CSCtd86572

DMVPN EIGRP topology hub does not update spoke with new next hop information for prefix with equal cost paths.

This condition may happen when dual hubs are setup where each spoke only peers with one hub and hubs exchange information. Equal cost path next hops are on two spokes off the same hub. Hubs are configured with “**no ip next-hop-self eigrp.**”

Workaround: Is to clear EIGRP neighbor on spoke that has stale next hop.

- CSCtd89787

High CPU under interrupt with a large number of alignment errors.

This condition are seen when issuing **show ip bgp ipv4 mdt vrf <VRF name> neighbors.**

Workaround: None

- CSCtd89963

Cisco IOS BGP session is flapped on the Cisco ASR 1000 Router Series.

When the Cisco ASR 1000 Router received BGP UPDATE messages from its peer with invalid AS4_PATH attribute, BGP session is reset and NOTIFICATION message is sent out with error ‘update malformed’.

The symptom is observed only when using the testtool (routem) to generate BGP UPDATE message with mismatched AS_PATH and AS4_PATH attributes to send to the router.

Workaround: There is no known workaround.

- CSCtd90979

When configuring hierarchical QoS policy-map with percent based rate configuration, the rate calculation might be wrong when the QoS policy is applied to 10 GigabitEthernet interface.

The translation from percent to absolute value (in Kbps) might be wrong when QoS policy is applied to 10 GigabitEthernet interface.

Workaround: To change from using the percent rate to the absolute rate in BPS (bits per second) in parent shaper would avoid running into this issue.

- CSCte02973

Routing protocols like EIGRP may be dropped in the global table.

The symptom is observed when multicast is configured for a VRF and no multicast is configured for the global table.

Workaround: Is to enable “**ip multicast routing**” and create a loopback interface with “**ip pim sparse-mode**” enabled.

Further Problem Description: The problem should not occur for MVPN since this is not a valid configuration, as multicast in the core is a requirement.

- CSCte03142

Enhancement to allow user to configure MAC address for ATM p2p sub-interface.

To be used for RBE when configuring MAC on that ATM p2p subinterface.

Workaround: None

- CSCte09171

When configuring the hardware SIP OIR, we see more v4 and v6 unicast loss than expected on the ASR 1000 Router Series.

This condition may occur when packet loss has been seen during SIP OIR Test conditions.

Workaround: There is no known workaround.

- CSCte20171

HSRP active router send ICMP redirect message that source address set to physical interface IP address.

This condition are observed when Virtual IP address should be used as source address.

Workaround: There is no known workaround.

- CSCte26324

Hidden command "l2tp tunnel busy timeout 0" cannot be configured on the ASR 1000 Router Series.

Workaround: None

- CSCte29212

1. EIGRP summary-address with leak-map configuration is removed after reloading the router.
2. summary "leak-map" option is only selectable after entering admin dist.
3. summaries with the same subnet but different masks overwrite each other.

The following conditions are seen:

#1 & #2. EIGRP summary-address is configured with the default administrative distance and a leak-map.

#3. The same subnet is specified on multiple EIGRP summary-address commands using different masks.

Workarounds: #1 & #2. Use a non-default administrative distance for the summary route: ip summary-address eigrp 1 192.168.0.0 255.255.0.0 4 leak-map leak-routes #3. There is no workaround.

- CSCte38945

Unable to get ping reply from the multicast group configured on loopback interface.

The symptom can occur when there are multiple routes populated in an interface and the interface goes down. All the routers associated with the interface should be removed, but only one is deleted. This results in the ping failure.

Workaround: Shut down the other interfaces associated with the router and enable it again.

- CSCte39004

Traceback seen on Pe2 when ipv6 route on ce2 is removed after ipv6 network is removed on pe2 in 6pe fp environment.

This condition are only seen with IPV6 and MPLS configuration.

Workaround: None

- CSCte46020

When using a nas-port-format which is different from default encoding 4/1/3, the NAS-Port-ID and NAS-Port radius attributes do not reflect the requested encoding. This is for sessions which originate on ATM interfaces only, i.e. PPPoEoA.

Depending on physical interface location, the NAS-Port-ID and NAS-Port radius attributes may not be represented correctly.

Workaround: Physically move (if possible) the interfaces into ports which can be correctly encoded with 4/1/3 bit distribution.

- CSCte50144

A Cisco ASR1002 Router reports incorrect CPU utilization. It reports a low CPU utilisation and also reports an overall utilization lower than the utilization under interrupts.

As shown in this example:

CPU utilization for five seconds: 5%/25%; one minute: 8%; five minutes: 8%

This symptom has been observed on the ASR1002 Router under high CPU utilisation of the RP CPU, caused by excessive rate of punted traffic.

- CSCte50206

Suppress authentication null-username suppresses system accounting messages

If you configure 'aaa authentication suppress null-username', system accounting records will not be generated.

Workaround: None

- CSCte51436

Symptom: Pressing Hold during a SIP-to-SIP call through CUBE(Ent) on the ASR 1000 Router results in intermittent disconnects. The phone behind the ASR CUBE hears a fast busy tone.

When CUBE dial-peers are configured with dtmf-relay of: **“rtp-nte”**, **“sip-notify rtp-nte”**, or none.

ASR CUBE(Ent) version from CCO: asr1000rp2-adventerprise9.02.05.00.122-33.XNE.bin

Workaround: Is to use **“sip-notify”** as the dtmf-relay method.

- CSCte51715

Logs will appear with the following error message:

```
%IOSXE-4-PLATFORM: R0/0: kernel: EXT2-fs warning: checktime reached, running e2fsck is recommended
```

In this condition there is no service impact on the router, only logs are generated frequently.

When a Cisco ASR 1000 Router is running any image before 12.2(33)XNF1 release.

The images for 12.2(33)XNF1 has the fix for this condition.

Workaround: Is to perform the following steps:

1. Drop into Linux shell
2. Issue the following set of commands:

```
/sbin/tune2fs -c 0 -i 0 /dev/sda1
```

```
/sbin/tune2fs -c 0 -i 0 /dev/sda2
```

```
/sbin/tune2fs -c 0 -i 0 /dev/sdb1
```

```
/sbin/tune2fs -c 0 -i 0 /dev/sdb2
```

- CSCte52369

On a Cisco ASR1000 Router, the RADIUS will send a NACK for the first COA request message and Radius Authentication will fail.

This condition are seen when the RADIUS receives "ACCESS-ACCEPT" with 'Unsupported Vendor' attribute

Workaround: Work around is to send the COA request message again.

- CSCte53365

The connected EIGRP-owned global addresses are put into the EIGRP topology database after the IPv6 router eigrp <as> process is configured to “**no shutdown**”.

This symptom is observed when the router is reloaded with an IPv6 EIGRP instance configured “**shutdown**”, then the configuration is changed to “**no shutdown**”.

Workaround: Configure “**shutdown**” then “**no shutdown**” on the interfaces.

- CSCte58468

OSPF process running in global routing table does not declare it self as ASBR router in the router LSA (E-bit is not set) and therefore his external default route is not installed.

This problem happens during reconfiguration after 'router ospf X' is removed from the configuration and later added back.

This problem are seen only if another OSPF process which runs in VRF is configured, but not always. Both processes must share internal datastructures.

Workaround: None

- CSCte58825

There is a crash upon conducting an snmpwalk from “enterprise mib oid 1.3.6.1.4.1”.

The symptom is observed on a Cisco ASR 1000 Series Aggregation Services router that is running Cisco IOS Release 12.2(33)XNE.

Workaround: Configure SNMP view to exclude ipSecPolMap as follows:

```
snmp-server view <view name> iso included
```

```
snmp-server view <view name> ipSecPolMapTable excluded
```

- CSCte60167

EIGRP IPv6 is no longer active on an interface.

This condition has been observed after removing HSRPv6 configuration from the interface.

Workaround: Is to do **shutdown** and **no shutdown** on the interface.

- CSCte62914

Track process timer is scheduled 1000 per second because there is no default timer value for track stub object. There are no known serious side effects, although the extra schedules could degrade performance by a small amount when system is under load.

If track <number> stub object is configured, then a timer is started using 0 as the timer frequency which results in a 1 millisecond timer.

Workaround: Do not use stub object tracking.

- CSCte64090

After Route Processor (RP) switchover, PPPoE traffic may drop though sessions that will stay up.

This condition has been seen when triggering a Route Process (RP) switchover via command **redundancy force-switchover**.

Workaround: None

- CSCte64156

Under certain circumstances, the ROMMON variables may show:

“PLATFORM_MAX_INTERFACES =128K” while there is no “platform max-interface 128k” configured.

This usually occurs after the router has been reloaded.

Workaround: Is to configure “platform max-interface 128k” and then “no platform max-interface 128k”. A reboot is recommended after.

- CSCte69014

Multiple memory leaks are seen when you try to bring up an unauthenticated session:

```
0x11FF95E0    7812    37  AAA Request Data
0x11FFA020    4692    18  SSS AAA auth req
0x11CBD48C    1692    18  AAA AUTHEN Username
```

The symptom is observed when a TAL authorization failure occurs due to access-reject. When you try to bring the session up again this leak is seen.

The following steps are:

1. Configure L2-connected session.
2. Bring up an L2-connected IP session and verify the access-reject event is triggered.
3. Check for these leaks using ""show memory"" commands.

Workaround: There is no workaround."

- CSCte69621

Missig CLI for configuring deny policy options:

crypto ipsec ipv4-deny {clear|deny|jump}

Workaround: None

- CSCte69761

Intermittently the eigrp learned default route (0.0.0.0/0) is deleted from the routing table.

If a router receives a default-network and 0.0.0.0/0 prefix and both are marked as candidate default prefixes in the routing table, and the default-network prefix is lost/deleted, the 0.0.0.0/0 prefix will also be deleted from the routing table even though an EIGRP topology entry remains for the 0.0.0.0/0 prefix.

Workaround: Is to do the following:

1. A **clear ip route *** can be issued and the EIGRP topology table entry for 0.0.0.0/0 be re-installed again into the routing table.
2. Configure ‘no default-information in’ under the router EIGRP process on the router which intermittently loses the 0.0.0.0/0 prefix.
3. Reconfigure the network, if possible, to discontinue use of the ‘ip default-network’ command and rely on the use of the 0.0.0.0/0 prefix.

- CSCte72075

BGP VRF IPv6 session fails to stay up. As soon as a BGP End-Of-RIB (EOR) message is send the peer responds with an update malformed BGP notification and the session is torn down. The session gets stuck in NoNeg state.

This condition Two BGP routers each configured with a VRF that attempt to establish an IPv6 BGP session within the VRF.

Workaround: There is no workaround.

- CSCte72128

After reload, ""cdp enable"" is missing on tunnel interfaces.

The following conditions have been observed:

- Having CDP activated on tunnel interfaces (for ODR usage in DMVPN for example)
- Running XNE1, this has not been seen on 12.2(33)XNE

Workaround: Is to add it manually, after a reload.

- CSCte73093

EIGRP resync is not triggered when modifying inbound with outbound prefix-list or ACL.

This condition has been observed when there is an interface associate with EIGRP distribute-list .

Workaround: Use distribute-list without any associating interface.

- CSCte75406

A crash can occur if the memory is low during the initialization of the OSPF process.

The symptom is observed if the memory is low during OSPF process initialization.

Workaround: There is no workaround."

- CSCte75784

About 500 bytes of memory is leaked for each configured delegate subscriber that has the supported options header configured for it. The memory is lost during each show with run or wr mem.

This condition are seen when configuring the delegate subscriber with the supported options header tags.

Workaround: Do not configure supported options header tags.

- CSCte77136

CLNS routing over GRE tunnels is not working on the ASR1000 Router.

Conditions: CLNS routing over GRE tunnels is configured, specifically with a GRE tunnel as the egress interface (output from the ASR100). In this scenario, CLNS packets are not forwarded via fast switching.

Workaround: I to use the following configuration change (needed on a per interface basis):

"no clns route-cache" "to disable CLNS fast switching"

- CSCte78165

On the Cisco ASR 1000 Router a device may reload when the 'show ip protocol' is issued.

This condition may occur when the Routing protocol is configured, and is trying to redistribute ISIS routes.

Workaround: Do not use the **show ip protocol** command for now.

- CSCte79759

On a DMVPN hub router, NHRP multicast replication entry is not deleted from its replication list when the corresponding nhrp cache entry is deleted.

This problem occurs when a DMVPN spoke is no longer registered with a DMVPN hub router.

Workaround: The workaround is to remove the tunnel interface on the DMVPN hub router and re-add it.

- CSCte83888

When PoD request contains target Acct-Session-Id prepended with NAS-Port-ID it will not be honored.

This condition are seen when PoD prepended with NAS-Port-Id for target session.

Workaround: Is to use only the Session-Id which is located after the, "_" in the Account-Session-ID to specify the session needing disconnect.

- CSCte89787

A Cisco ASR 1000 Router crashes after the Segment Switch Manager (SSM) reports that an invalid segment has been detected:

%SW_MGR-3-INVALID_SEGMENT: Segment Switch Manager Error - Invalid segment - no segment class.

The crash follows this message.

The symptom is observed on a Cisco ASR 1002 that is running Cisco IOS Release 12.2(33)XND1. The crash is caused by a NULL pointer de-reference following the "no segment class" error. The error itself is not fatal and the crash should have been avoided.

Workaround: There is no workaround.

- CSCte92659

On the ASR 1000 Router Series show memory debug leaks would show SSS holding some memory. This condition are seen during longer hours of session flapping.

Workaround:None

- CSCte92745

After removing a user in resource policy the Cisco ASR 1000 Router fails in the first attempt. This condition are seen only with the first attempt and second time gets removed.

Workaround: None

- CSCte92790

Router has unnecessary periodic (50mins) BGP update without topology change.

This condition are seen when configuring BGP Best External feature on EDGE router as ""bgp advertise-best-external"".

Workaround: None

- CSCte94156

When running Release 2.5.1 the Cisco ASR 1000 Router fails to update the PST value in TBAR., causing other

GM to fail sending traffic on the ASR 1000 Router with anti-replay error messages. This happens whenever the local ACL is changed on the GM or by KS failure and recovery.

Workaround: None

- CSCte94237

A crash happened when type cli **sh sbc <non-exist sbc name> sbe sip statistic** ASR 1004-2#sh sbc afaf sbe sip statistics SBC "afaf" has not been configured.

No SIP statistics found.

ASR1004-2#

```
*Feb 10 10:34:40.228 SGT: %SCHED-2-SEMNOTLOCKED: Virtual Exec attempted to unlock an
unlocked semaphore -Traceback= 1#359fd7114b043d9cc307b84aa384d228 :10000000+B95EA0
:10000000+B96224 :10000000+20BDEA4 :10000000+34E3578 :10000000+359D3B4
:10000000+359EF50 :10000000+B02D30 :10000000+B0954C :10000000+283AF14
:10000000+B1766C
```

This will limit the memory address space engine can use to 3400MB, thus limiting the call in progress.

This condition are seen when typing cli `sh sbc <non-exist sbc name> sbe sip statistic` in console.

Workaround: Is to only type that command with a configured sbc name.

- CSCte97907

On a Cisco ASR 1000 Router with RP2 may get out of sync with NTP master every 18 minutes for approximately 1 minute. This may offset the NTP Master which will cause an increase up to -1052.1 msec and the sync will get lost.

This instance has been observed, when NTP is enabled and running apr. 20 minutes.

Workaround: None

- CSCte98082

PPPoE Relay Session is failing to come up on LAC with some specific configuration.

Workaround: None

- CSCtf00427

A router may experience a severe memory leak issue when the following command is configured:

privilege exec level <level> show ip

ospf neighbor

The symptom is observed when running Cisco IOS Release 12.2(33)XNE

or 12.2(33)XNE1. The issue is not platform dependent.

Workaround: Is to reload the router.

- CSCtf01344

IOSD core@chunk_diagnose while doing ASR 1004 ISSU upgrade.

The problem is limited to the 4RU when attempting an ISSU upgrade with VRF-aware IPsec features and an uninitialized webex SPA in the system.

Workaround: Properly initialize Webex SPA before ISSU upgrade.

- CSCtf05183

Tracebacks are seen when changing tunnel mode from gre or ipip to ipsec ipv4 with cdp enabled.

This problem are seen when changing tunnel mode from gre or ipip to ipsec ipv4 with cdp enabled.

Workaround: None

- CSCtf07776

The below traceback can be seen in two scenarios on a Cisco ASR 1000 Router:

- During UUT reload
- After shutting the FRR enabled interface

```
%FRR_OCE-3-GENERAL: un-matched frr_cutover_cnt.
```

```
-Traceback= 40DCB368 40DCB220 40DCB444 40DEC968 40D15FE4 40D1BACC 40D13BD4 40D14810
```

This condition are seen on the router with TE and FRR enabled on interface during the reboot and issue.

Workaround: There is no know workaround.

Further Problem Description: This problem has no impact on forwarding functionalities, but it does have performance impact on unstable network situations.

- CSCtf07907

RP Crash observed when doing RP switchover after deleting some tunnel configuration.

This instance may occur when switchover is to be done after deleting some tunnel config and traffic is flowing in background.

Workaround: None

- CSCtf11997

For the following sbc-sbe configuration:

```
call-policy-set 1
irst-call-routing-table RT-DSTADDR
rtg-dst-address-table RT-DSTADDR
entry 1
match-address ^bus[0-9][a-z] regex
```

The command “no match-address” would NOT delete (i.e. unconfig) the match-address.

This is also observed for the match-address under rtg-src-address-table configuration.

Workaround: The workaround is to delete “entry 1” or “rtg-dst-address-table RT-DSTADDR” or “call-policy-set 1”.

- CSCtf15848

The following error seen on re-configuring channel-groups after switchover

EFC ERROR: spa_efc_config_ds1_channel - channel in use

The following conditions have been seen:

- active RP booted with 8xcht1/e1 and channel-group is configured on t1 controller
- load the standby RP, and do a switchover
- on new active RP, unconfigure and reconfigure the channel-group

Tracebacks with “EFC ERROR: spa_efc_config_ds1_channel -channel in use” seen.

Workaround: before switchover, configure channel-groups on active RP when standby RP is up.

- CSCtf16359

The ASR 1000 Router when configured as GETVPN GM will not make any local GM acl change of removing extended ACL effective, until a new rekey from Key server has been configured.

This condition has been seen when the ASR 1000 router is configured as GETVPN GM.

Workaround: None

- CSCtf17273

A Cisco ASR 1000 Router crashes during startup when receiving an AS_SET attribute from its peer.

This symptom is observed when the BGP peer sends an AS_PATH or AS4_PATH containing an AS_SET attribute.

Workaround: There is no workaround.

- CSCtf19923

IP SLA: icmp-echo detect 300+ msec delay on a Cisco ASR 1000 Router.

This has been seen under corner conditions an incoming packet might be delayed by about 900 msec resulting in incorrect IP SLA. This is the max delay if there are no other packets received after the icmp response occurrence is very rare about 2 to 3 times in 10,000,000 pkts

Workaround: None

- CSCtf21390

When EIGRP for IPv6 is used and a default route summary is entered on an interface (::/0), the appropriate topology table entry and IPv6 route is not created or sent to peers. This conditions are seen when default route summary is entered on an interface (::/0) in EIGRP for IPv6.

Workaround: Is to use a summary other than ::/0, since other summaries beside the default route work correctly.

- CSCtf25514

Policy routing is enabled on Virtual Template interface when ip policy route-map command is downloaded via Radius, while it is not supported via CLI.

Above symptom is seen on Cisco ASR 1000 Routers running IOS of version 122(33)XND, 122(33)XNE and 122(33)XNF.

Workaround: None

- CSCtf26943

Session is not going down after per-user push on a Cisco ASR 1000 Router.

This condition happens after a Per-User Push from RSIM with incorrect QoS attributes Session is not going down.

Workaround: None"

- CSCtf26946

The ASR 1000 Router crashed when mis-input command **no int sbc1** under ASR1004-3(config-sbc-dbe)#

The conditions are seen when the following has occurred:

1. provision sbc1 interface
2. provision dbc and media-address
3. input **no interface sbc1** under ASR1004-3(config-sbc-dbe)#

Workaround: None

- CSCtf32412

Tunnel drops [Phase I and Phase II's] may occur on the ASR 1000 Router.

If a specific phase II pair of SA's is timing out due to configured idle time, then the

ASR 1000 Router will drop the Phase I, all Phase II to that particular peer and create a tunnel drop.

Workaround: Do not use crypto ipsec security-association idle-time.

- CSCtf32693

On Cisco ASR 1000 Router, configuring xconnect on a VLAN, SNMP 64 bit counters are not getting updated.

This condition are seen when one of the vlan on same port have xconnect configuration.

Workaround: There is no workaround.

- CSCtf33363

Port information is missing in nas-port string sent to Radius on the ASR 1000 Router Series. This instance may occur with PPP sessions on the router.

Workarounds: None

- CSCtf33960

Router alert label 1 is deleted upon SSO, so mpls ping over and RA only PW failed.

This issue has been observed in RA only l2vpn configuration.

Workaround: None

- CSCtf36402

The ASR 1000 Router may crash when the user telnets and Transmission Control Block is cleared for that session before entering password. This instance has been observed when AAA Authentication protocol is set to TACACS.

Workaround: Do not clear the Transmission Control Block for a session before entering password.

- CSCtf40702

A Cisco ASR 1000 Router Series with Route Processor 2 Engine may unexpectedly reload do to a SegV crash. This will happen if there is a monitor session configured that uses a source interface with a range. This can either be a crash while configuring via cli or a crash at bootup if the command is in the startup configuring.

Workaround: Don not use the source inter range

CSCtf41171

With QoS policy accounting enabled, using the **clear subscriber session uid <uid>** command to clear a session can result with incorrect packet/byte counts on the generated accounting Stop record.

The following conditions are seen:

1. QOS accounting enabled
2. The SAME accounting group is applied to a class in BOTH the input AND the output policy-maps.
3. **clear subscriber session uid <uid>** is used to clear the session

Under these conditions the packet/byte counts on the generated accounting Stop record may be incorrect.

Workaround: Is to use an alternate method to clear the session, such as **clear pppoe all** or **clear ppp interface <interface>**.

- CSCtf54092

Wrong if index exported to NFC log when polling for ATM sub-interface.

This condition has been seen when ATM sub-interface is configured.

Workaround: None

- CSCtf59446

On a Cisco ASR 1000 Router the Standby Router processor may experience a s/w reset.

This condition has been seen when the ASR 1000 Standby Router processor experiences a s/w reset, after issuing “**show l2tp session interworking username <username>**” on the standby route processor while L2TP tunnels are establishing.

Workaround: Is to issue the command “**show l2tp session interworking username <username>**” in active route processor to get command output without any s/w reset.

- CSCtf59781

When DHCPv4 client sends dhcp discover packet with broadcast flag = off, the performance of Cisco ASR 1000 Router working as DHCPv4 relay is not good compared to if it receives dhcp discover packets from client with broadcast flag = on.

This problem is expected only if the dhcp client is sending discover packet with broadcast flag = off(unicast).

Workaround: None

- CSCtf66271

An ASR 1000 Router running asr1000rp1-adventerprisek9.02.04.02.122-33.XND2.bin, upgraded to asr1000rp1-adventerprisek9.02.06.00.122-33.XNF.bin displays the complete cert chain like:

```
crypto pki certificate chain JUTnetRoot-Pilot
certificate ca 3C5A00179190F2DF23325330195B9B67
308203AE 30820296 A0030201 0202103C 5A001791 90F2DF23 32533019 5B9B6730
0D06092A 864886F7 0D010105 05003071 310B3009 06035504 06130255 53311930
17060355 040A1410 41542654 20436F72 706F7261 74696F6E 311F301D 06035504
0B131646 6F722054 65737420 50757270 6F736573 204F6E6C
:
:
truncated
```

Whereas before upgrade it displayed the same as:

```
crypto pki certificate chain JUTnetRoot-Pilot
certificate ca 3C5A00179190F2DF23325330195B9B67 nvram:ATTCorporati#9B67CA.cer
```

This condition has been seen when ASR 1006 running asr1000rp1-adventerprisek9.02.06.00.122-33.XNF.bin image.

Workaround: None

- CSCtf85471

Tunnel Client Auth-ID mismatch has been seen between the Active and Standby RP when load-balancing profile is used with Radius.

Issue is seen with specific RADIUS profile for load-balancing.

Workaround: The workaround is to disable load-balancing.

- CSCtf91603

In a CUBE(SP Edition) ASR 1000 system configured for H.323 video calls, in some instances where the endpoints send large H.323 TCS messages a video call may be connected with only audio or fail to setup the call entirely.

This condition may occur when TCS message must be large.

Workaround: Is to configure H.323 Video endpoints to advertise fewer capabilities.

- CSCtf95205

The **show sbc global db signaling-flow-stats** reports incorrect value for “Reserved Bandwidth” on the termination whose tman/pol is set as “OFF”.

When the Cisco ASR 1000 Router creates a pinhole based on the received MEGACO ADD request which sets tman/pol=ON with tman/sdr&tman/mbs parameters for one side termination and tman/pol=OFF without tman/sdr&tman/mbs parameters for the other side termination.

Workaround: None

- CSCtg02140

Upon switchover with the webex call flow, media is not preserved.

This problem has been seen when switchover with webex flow Add(A,B), Delete (B), and Add (A,C) has been used.

Workaround: None

- CSCtg02617

The following error has been observed on the ASR 1000 Router console:

Config Sync: Line-by-Line sync verifying failure

The above error has been seen when entering any parser CLI in parser view on the router.

Workaround: None

- CSCtg06681

SIP method profile allows defining a mapping of status-codes. RP2 CUBE(SP Edition) would crash while removing status-code map.

For RP2 CUBE(SP Edition), consider the following configuration:

```
config t
sbc test
sbe
sip method-profile SIPmessage
method INVITE
action as-profile
map-status-code
range 183 value 180      <==== incorrect mapping
end
```



Note That there is only one entry for mapping status-code.

When we try to unconfigure “range 183 value 180” as follows the RP2 CUBE(SP Edition) would crash:

```
config t
sbc test
sbe
sip method-profile SIPmessage
method INVITE
map-status-code
no range 183 value 180    <==== causes crash
end
```

Workaround: The workaround is to unconfigure “map-status-code” and then re-configure it with correct mapping of status-code as follows:

```
config t
sbc test
sbe
sip method-profile SIPmessage
method INVITE
no map-status-code        <==== unconfigure map-status-code
map-status-code           <==== re-configure
range 180 value 183       <==== correct mapping
end
```

- CSCtg11844

Crypto tunnel will not come up and pass traffic on the ASR 1000 Router Series. This condition are seen when the ASR 1000 Router running IOS 12.2(33)XNF code with nat outside and a crypto map on the same interface, if you remove and readd the crypto map from the nat outside interface the

crypto tunnel will not come up and start passing traffic until such time as you remove the "ip nat outside" statement. Once the crypto map is up and running you can read the "ip nat outside" command.

Workaround: Is to remove the nat outside command, get the crypto tunnel up and passing traffic then read the nat outside command back to the interface.

- CSCtg17977

BBA L2TP LNS subscriber sessions are not in sync between Active and Standby RP's. This condition has been observed during ISSU upgrade.

Workaround: There is no known workaround.

- CSCtg27141

UDP Jitter operation is not working on the ASR 1000 Router Series.

This condition has been when using UDP Jitter the operation is not working after time out failure occurs on the router.

Workaround: None

Open Caveats—Cisco IOS XE Release 2.6.0

This section documents possible unexpected behavior by Cisco IOS XE Release 2.6.0

- CSCsw78270

SIP core during ISSU has been observed intermittently, when upgrading, or downgrading to 2.3.0 Release.

This instance may only occur, while executing ISSU runversion on the 6RU, 4RU and 2RU superpackages after upgrading from 2.2.0 to 2.3.0 releases.

Workaround: None

- CSCta24676

On the ASR 1000 Router when an attempt is made to login to the kerberos client, the RP crashes. This is after the clocks of the UUTs are synchronized and the routers are configured with kerberos credentials.

- CSCta46670

When disabling and enabling **control plane host** a few times this may generate an error message on the Cisco ASR 1000 Router. This has been observed, when configuring **control plane host** followed by **no control plane host** a few times on the router.

Workaround: None is required since there appears to be no functional impact.

- CSCtb84718

Output of show cli "sh crypto gdoi gm acl" does not correctly display as a COOP Key Server.

This has been observed, when COOP Key Servers has been configured on the GM.

Workaround: None

- CSCtb98877

On the ASR 1000 Router Series subsequent call fails after a SIP Session Refresh timeout occurs after an HA switchover in CUBE environment.

This occurs in a back to back CUBE environment:

CUCM1 - SIP - CUBE1 - SIP - CUBE2 - SIP - CUCM2

The CUCM SIP Refresh is set to 90 seconds, and a call is made. HA switchover occurs on CUBE1, and the call is disconnected as expected.

The same call is made again, but the originating endpoint on CUCM1 gets a Busy tone, while the terminating endpoint on CUCM2 gets Ringing tone.

CUBE2 sends a 503 Internal error with the following cause code:

Reason: Q.850;cause=38 - [Network out of order]

Workaround: None

- CSCtc13911

Backup tunnels in TE FRR scenario keeps flapping after RP reload.

This may happen only with MIB walk through in tandem.

Workaround: To Do an RP switchover. Without this workaround it takes 2 to 3 hours for the tunnels to become stable.

- CSCtc35744

Configure a user profile with multiple 'Cisco-Avpair=lcp:interface-config=<cmd>'. Create pppoe session. The per-user access-list attributes get downloaded from the AAA server, and the attributes are applied. The show subscriber session detailed output shows that only the first 'Cisco-Avpair=lcp:interface-config=<cmd>' is getting applied on the session. Typical features configured on a PPP Virtual Access on a per-users basis using this vsa would be IPUnnumbered, VRF, Keepalive, Pool name, PBR, multicastjoins, etc.

This issue is seen when a PPPoE session is brought up with a user profile that has more than one 'Cisco-Avpair=lcp:interface-config=<cmd>' configured.

Workaround: The workaround is to configure multiple Cisco-Avpairs in 1 line.

For example:

Cisco-AVPair = lcp:interface-config=ip vrf forwarding vpngreen

Cisco-AVPair = lcp:interface-config=ip unnumbered loopback2

Should be configured like this:

Cisco-AVPair = lcp:interface-config=ip vrf forwarding vpngreen \nip
unnumbered loopback2

- CSCtc44482

When deleting VRF while STANDBY RP comes up, when SSO switchover happens this may cause an error message after STANDBY RP has been reloaded.

The following error message has been seen:

SYS-3-HARIKARI

Workaround: Do not delete VRF prior to STANBY RP achieving the STANDBY HOT state.

- CSCtc54288

When changing the group number associated with a virtual IP address causes hosts to lose contact with the virtual router.

This instance occurs when a virtual IP address is associated with one group, and then that group is unconfigured and the same virtual address used by another group. Since each group is uniquely associated with a virtual MAC address the ARP tables of all hosts that were using the previous group

will contain invalid entries. When the interface is shut, while configuring new groups then gratuitous ARPs will be sent to refresh any hosts' ARP tables before the interface is ready to forward traffic. The hosts will not realize that the vIP/vMAC association in their ARP tables are invalid and will be unable to forward traffic via the known (virtual IP) gateway.

Workaround: A delay can be used to stall the VRRP initialization process after unshutting an interface:

vrrp delay minimum

reload

The values used are the number of seconds to delay, which is platform dependent. 30 seconds for interface delay and 300 seconds for reload delay are a good first values to test.

- **CSCtc55049**

The ASR 1000 Router may crash and reload following a reboot or initial boot from a power-up.

The embedded syslog manager (ESM) needs to be configured along with an ESM script present during an initial boot or reload. Also, redundant RP/FP appears to be the scenario that has the greatest likelihood of encountering the problem.

Workaround: None. However if problem manifests, the subsequent rebooting is very likely to be successful. When stuck in a situation where crashes are repetitive, momentarily pull redundant RP until system stabilizes, and re-insert redundant RP.

- **CSCtc55215**

On the ASR 1000 Router, when shape rate is configured as % of an ATM PVC for GRE QoS it is not updated after the PVC rate has changed.

This may occur when changes to the PVC rate and its ATM class has been configured on the router at the same time.

Workaround: None

- **CSCtc58124**

When traffic is flowing, (S,G) expiry timer should be updated to 3:30 seconds every 2 minutes. In the VRF context, the expiry timer on (*,G) and OIF is updated but on (S,G) is not on the ASR 1000 Router Series. This will work fine in the Global Context on the router.

This happens, when **vrf** is configured on router.

Workaround: There is no workaround.

- **CSCtc62440**

On a Cisco ASR 1000 Router Series, the removal of sub-interfaces may under certain conditions result in MFIB_MRIB-3-FAILED_WIRE_FIND error messages being generated on the Route Processor (RP).

There is no functional impact due to this issue.

Workaround: There are no known workarounds.

- **CSCtc67457**

On the RP2 a crash has been seen with process IKMP.

This has been observed, when GetVPN Group Member is configured with vrf-lite on the RP2.

Workaround: No known workaround.

- **CSCtc69297**

Tracebacks has been seen with cli **sh platform hardware qfp active feature acl tree** on the Cisco ASR 1000 Router.

This condition has been seen, when there are a huge number of acls configured on the router.

Workaround: None

- CSCtc70742

ASRNAT allows removal (unconfiguration) of static entry even when entries has children translations.

Workaround: There is no known workaround.

- CSCtc78745

When deleting a few tunnels from PE side, when CE and PE are having different number of tunnels the Cisco ASR 1000 Router starts throwing msgs.

The following message has been seen:

```
Oct 27 14:33:40.170 IST: %ACE-3-TRANSERR: ASR1000-ESP(14): IKE trans 0x1D8C; opcode 0x60; param 0x1FD2; error 0xA; retry cnt 0
```

This condition has been observed, when tunnel mismatch between CE and PE are kept for a long time on the router.

Workaround: There are no known workarounds as of now.

- CSCtc86866

While unconfiguring IP NHRP, when mapping has been given a different NBMA Address this clears the original address on the ASR 1000 Router Series.

Workaround: None

- CSCtc91018

On a Cisco ASR 1000 Router the subinterface counters with Frame Relay Encapsulation can show higher values than the counters on the main interface, when self-pinging the subinterface.

Workaround: None

- CSCtc96467

STANDBY RP reloads twice with **issu runversion**, while downgrading from Release 2.6 to 2.5.

This instance may occur, when Super Package has been configured with ISSU, which causes the STANDBY RP to reload with **issu runversion**.

Workaround: None

- CSCtd00489

A traceback indicating that the object was being deleted before the ideal exponent is invalidated has been logged on the ASR 1000 Router Series. An schedule object is freed before the ideal exponent is invalidated. This condition is treated as an error because this points to a missing step in cleaning up prior to destroying an object since this can potentially impact the rate accuracy in the future. This issue occurs while an ATM VC schedule is being deleted on the router.

Workaround: There is no known workaround. There are also no negative side effects identified in the systems where this issue has occurred. The system will continue to operate normally.

- CSCtd03743

Ping fails when VFR is removed on the Cisco ASR 1000 Router.

This instance has been observed after removing the VFR, when trying to ping the destination router loopback with larger packet size, then the ping fails on the router.

Workaround: To ping with small packet size of 3000 and 1500 on the router. Do not ping with packet size of 9300 this will cause the router to fail.

- CSCtd21252

Unified SBC crash has been seen on the ASR 1000 Router Series.

This condition may occur, when configuring a large IPv6 media-address on the router.

Workaround: None

- CSCtd22958

With basic SIP calls RTP traffic running (20 CPS, 1200 sustained calls), a physical OIR of ESP, -25% of active call media traffic is lost for a 20-30 second time period and signaling for new calls coming into Unified SBC during this time period are rejected with 500 - Server Internal Error.

Workaround: Avoid physically removal of active ESP, when Unified SBC calls are in session. To use soft OIR instead.

- CSCtd36301

At every session churning of IPv6 PPPoE uses more prefixes for same tunnel and session value.

No used IPv6 Prefixes in local IPv6 pool are incremented at every session flap iteration in IPv6 LNS for same tunnel and session value.

This instance may happen, when Local IPv6 prefix pool is used to assign ipv6 address and the sessions are churning at a flap rate of 70 sessions per seconds for 8000 sessions.

Workaround: None

- CSCtd37057

On a heavily loaded Cisco ASR 1000 Router Series, rapid QoS queuing configuration changes involving the removal of existing configuration and addition of new configuration could cause the system to experience temporary resource outage. The conditions under which this has been observed involve 32000 flapping PPPoE sessions combined with configuration changes on the system.

Workaround: Avoiding rapid and large QoS configuration changes on a heavily loaded system will avoid the problem reported in this caveat.

- CSCtd48042

When defining vrfa adjacency the vrfb as singal-address is used, this can start an attack and the EP will show in vrfa blacklist on the Cisco ASR 1000 Router.

This instance may occur, when vrfa adjacency has been defined, but vrfb as singal-address is used on the router.

Workaround: None

- CSCtd48500

SNMP 64 bit counters not showing traffic. This has been seen on ASR1002 running 12.2(33)XND1 and XND2 after deploying an ATOM Circuit under it.

Workaround: None

- CSCtd49186

After the Cisco ASR 1000 Router has been reloaded this removes saved VASI, Subinterface, Loopback Interface CLIs under Parser View.

This condition has been observed, when accessing parser view for saved interfaces followed by a reload on the router the interfaces were removed.

Workaround: After reload all of saved interface configurations will need to be re-configure under each parser view on the router.

- CSCtd62358

On a Cisco ASR 1000 Router Series, the rapid deletion and re-creation of VASI interfaces may result in failures in the functioning of the VASI interfaces.

Unexpected resets of the ESP have also been observed under these conditions.

Workaround: Delaying of the order of 2 minutes between the deletion and re-creation of VASI interfaces will avoid this problem.

- CSCtd64206

FP crash may occur, when ISG DHCP sessions flap has churned on the Cisco ASR 1000 Router.

This condition has been observed in ISG DHCP stressed environment.

Workaround: Is to lower the scale of both number of sessions and session churn rate.

- CSCtd70901

During RP switchover, a Cisco ASR 1000 Router running 2.6.0 release may experience IPv6 multicast channel zapping latency more than expected.

This condition has been seen, when RP failover from the active to the standby has occurred.

Workaround: There is no known workaround.

- CSCtd80542

Loop observed, when configuring SNMP bulk mib walk. The loop has been observed at tunnelInetConfigIfIndex.

This condition has occurred, when scaled configuration includes tunnel interface 2147483647.

Workaround: None

- CSCtd84323

Under Unified SBC SIP IPv6 to IPv4 scenario with DTMF digits via INFO method, the following Traceback has been seen on the console following RP failover scenario:

```
"*Dec 14 20:59:14.494: %ASR1000_INFRA-5-IOS_INTR_HISTORY: [5|0] [0:0] [0->0] ra[ 1*
0x0 1* 0x0 ] -Process= "SBC main process", ipl= 0, pid= 314". The traceback causes a
temporary outage in service, but SBC does recover without any manual intervention.
```

This traceback has been observed in the following conditions:

1. SIP IPv6 to IPv4 calls
2. DTMF digits transferred via INFO method
3. RP failover has been executed at some point in past.

Workaround: None

- CSCtd87072

IOSD will restart, when changing tunneling mode in scaled IPSec Sessions on the ASR 1000 Router Series.

This condition has been observed, after IOSD restarts the tunneling mode has changed in a scaled IPSec Session environment.

Workaround: None

- CSCtd87205

The Cisco ASR 1000 Router will reload, when flapping up and down VC's after configuring SSO.

This condition has been observed, when the Cisco ASR 1000 Router reloads after a large amounts of flapping has occurred, and SSO has been forced onto the router. The router may reload.

Workaround: Is to slow down the amount of flapping when doing SSO on the router..

- CSCtd90836

During REKEY there are a lifetime of IPSEC sessions that show junk characters on the ASR 1000 Router Series.

This conditions has been observed, during rekey.

Workaround: None

- CSCtd91015

The Cisco ASR 1000 Router does not roll back to the base image even though the rollback timer has expired for ISSU Superpackage Downgrade from Release 2.6 to 2.5. ISSU Superpackage Downgrade does not finish within the specified “roll back” time, but router does not rollback to the base image. Tracelogs shows that the timer has been expired and a user prompt has appeared. But the prompt does not appear on the console. SPA's will move to “inserted” mode and at certain times STANDBY RP will reload.

Workaround: ISSU will work fine, when rollback time is increased.

- CSCtd91950

A Cisco ASR 1000 Router Series with the Lawful Intercept feature configured may reset unexpectedly under certain conditions when streams are modified/disabled/re-enabled during traffic flow.

The conditions necessary for this situation to be encountered are multiple MDs, configuration of circuit-id based pre-provisioned stream entries and active PPPoE sessions.

Workaround: There are no known workarounds.

- CSCtd91986

Under certain conditions a Cisco ASR 1000 Router Series configured with the Lawful Intercept feature may not intercept traffic sent from a PPPoE client though the packets reach the destination as expected. This may happen when a circuit-id based PPPoE session is up and traffic is sent from the PPPoE client. This may happen for both forwarded and PTA PPPoE session cases.

Workaround: There are no known workarounds.

- CSCtd98510

Some of the L2TPv3 Xconnects are not coming up after repeated (5-6) switchovers and OIR.

This condition has been observed ,when AC is down and session is in local and not ready state.

Workaround: Is to clear l2tp recovers the problem.

- CSCte01388

The FMAN FP process may crash on the ASR 1000 Router Series.

This has been observed, when VPN has been configured on the router.

Workaround: None

- CSCte17127

Calls are failing due to an invalid tls certificate or they may be completing when the certificate is invalid.

This issue ties into how long the SBC keeps the tcp and tls connection up and also when the ASR 1000 Router does not revalidate the certificates for a deleted or newly added trust point

tls peer. The same applies to the scenario where a certificate has to be replaced.

Workaround: Set the tls idle timer to a value of 3 minutes to minimize the time that the tls peer.

- CSCte20171

HSRP ACTIVE Router sends ICMP redirect message that the source address is set to a physical interface IP address. The Virtual IP address should be used as source address.

Workaround: None

- CSCte28845

With Cisco ASR 1000 Router operating in uSBC mode, all adjacencies are locked in Detached state after an upgrade or change where the SBC must be deactivated and activated. When SBC is deactivate or activated or the same for one of the adjacencies, the system prints a routing error log.

The problem occurs when there is an digit routing entry in the routing table that is missing the destination adjacency datafill.

In most cases the SBC will not allow this to be configured in the first place without throwing an error but there are some scenarios where this configuration can get into the database without an error.

Workaround: Remove the entry with “no dest adjacency” or “add a dest adjacency” to the entry datafill.

- CSCte42733

When configuring ip verify unicast reverse-Path and no ip verify unicast reverse-path in a virtual-template and then applying to a ppp session which causes a FP core dump. This condition has been observed, when URPF has been configured on the ASR 1000 Router.

Workaround: Is to **enable** and **disable urpf** in the same virtual-template.

- CSCte42926

Some L2 VPN circuits (PW) are missing or stays down after **clear xconnect all**. This condition has been seen in Scaled L2 VPN environments which includes L2 VPN ATM PWs, EoMPLS and Local Switching.

Workaround: Is to **reload** the router.

- CSCte43453

QoS accounting Interim record for the parent policy-map class-default class has incorrect packets and bytes stats while under traffic load. This condition has been seen when PTA session with Model D2.2 QoS has been enabled. QoS accounting has been enabled at the parent policy-map class-default class. While under traffic load, the accounting Interim record has incorrect stats as compared to the QoS stats in the output of **show policy-map session**.

Workaround: None

- CSCte43891

When QoS policy accounting has been enabled, using the **clear subscriber session uid <uid>** command to clear a session can result in incorrect packet/byte counts on the generated accounting Stop record.

This condition has been seen when following has occurred:

1. **qos accounting enabled**
2. The SAME accounting group is applied to a class in BOTH the input AND the output policy-maps.
3. **clear subscriber session uid <uid>** is used to clear the session, under these conditions the packet/byte counts on the generated accounting Stop record may be incorrect.

Workaround: Is to use an alternate method to clear the session, such as **clear pppoe all** or **clear ppp interface <interface>**

- CSCte46020

When using a nas-port-format which is different from default encoding 4/1/3, the NAS-Port-ID and NAS-Port radius attributes do not reflect the requested encoding. This is for sessions which originate on ATM interfaces only, i.e. PPPoEoA.

Depending on physical interface location, the NAS-Port-ID and NAS-Port radius attributes may not be represented correctly.

Workaround: Physically move (if possible) the interfaces into ports which can be correctly encoded with 4/1/3 bit distribution.

- CSCte48047

On a ASR 1000 Router Series the output from the **sh platform software status control- processor** may incorrectly indicate that the ESP committed memory is greater than 100%. There is no functional impact due to this.

Workaround: There are no known workarounds.

- CSCte50863

An fman_fp core is generated when the Template ACL feature is disabled or enabled several times with 4k PPP sessions with per-user ACLs.

This condition has been observed, when bringing up 4000 PPP Sessions terminated on PTA with per-user ACLs. With the template ACL feature enabled, only a few templates are created. Disable the template ACL feature and since there are only 4000 PPP Sessions, TCAM exhaustion by this action is not expected. Enable the template ACL feature again. Repeat until an fman_fp core is generated (usually seen within 10 iterations).

Workaround: Is to tear down PPP Sessions before disabling and enabling the Template ACL feature.

- CSCte55019

The Cisco ASR 1000 Router crashes when the local-address is configured as '0.0.0.0' under crypto keyring <name>.

This issue has been seen, when the ASR 1000 Router is loaded with asr1000rp1-adventerprise9-mz.122-33.1.5.XNF

Workaround: Is to configure the local-address with a valid ip address.

- CSCte55632

On a Cisco ASR 1000 Router Series configured with the WCCPv2 feature and processing WAAS traffic (HHTTP/FTP), a switchover from the active ESP to the standby ESP may under certain conditions cause the ESP to reset unexpectedly.

Workaround: There are no known workarounds.

- CSCte57932

About 10% of the calls will fail with one way audio on the ASR 1000 Router Series.

This instance may occur, when SIP Endpoints behind a NAT who are called from a H323 trunk and about 10% of the calls will fail with one way audio.

Workaround: There is no known workaround.

- CSCte61735

Memory leak has been seen when MQC is configured on the Cisco ASR 1000 Router. This can occur, when QoS has been configured on the router, in an ISG environment.

For example the following conditions have been observed:

```
interface ATM4/0.1 point-to-point
no atm enable-ilmi-trap
pvc 0/101
class-vc crosshairs
vbr-nrt 500 400 50
dbs enable
service-policy in DefaultIn
service-policy out DefaultOut
!
vc-class atm crosshairs
protocol ppp Virtual-Template1
encapsulation aal5snap
```

```
interface Virtual-Template1
ip unnumbered Loopback0
ppp authentication chap
end
```

The memory leak occurs when a link is flapped up and down.

Workaround: None

- CSCte62029

SBC is disabled (via CLI: **no activate**) service does not completely de-activate even though adjacencies, etc. appear to be in down with detached state. Though SBC will re-activate upon executing the **activate** CLI.

This condition can occur upon de-activation if the following exist:

1. billing is **enabled**
2. a Cisco ASR 1000 Router has redundant RP configuration (software or hardware)
3. SBC incoming SIP call-rate of 20 CPS.

Workaround: The following steps can be executed as a workaround once in failed state:

1. **disable** billing via **billing->no activate** CLI
2. **execute no activate** CLI again for SBC application
3. re-activate SBC service via **activate** CLI
4. re-enable billing via **billing->activate**.

- CSCte64156

Under certain circumstances, the ROMMON variables may show “PLATFORM_MAX_INTERFACES =128K” while there is no “platform max-interface 128k” configured.

This usually occurs after router reload.

Workaround: Is to configure “platform max-interface 128k” and then “no platform max-interface 128k”. A reboot is recommended afterwards.

- CSCte66782

VLAN node may become disabled resulting in possible higher latency for priority packets on the ASR 1000 Router Series.

This condition may occur, when configuring model F broadband configuration using ANCP to change the shape rates on individual VLANs, priority propagation at the VLAN node may become disabled resulting in possible higher latency for priority packets.

Workaround: Do not use ANCP to change the VLAN parent shape rate.

- CSCte71456

Self ping packets are sent by the a Cisco ASR 1000 Router on the serial interfaces and are not applied with Egress feature such as firewall has been used.

This has been observed, when packets are sent to a self IP Address and are not applied with the Egress features at CPP for serial interfaces. All other packets should be fine in this same environment.

Workaround: There is no known workaround.

- CSCte74829

On the Cisco ASR 1000 Router, dsx3LineStatusChange Trap has been seen for index 0. This condition has been observed on SPA OIR, or when creating a ds3 interface on the 1xchOC12-POS SPA.

Workaround: None

- CSCte78938

Xconnect configuration is rejected after replacing the MPLS xconnect configuration with manual L2TPv3 configuration on the ASR 1000 Router Series.

This condition has been seen, when EoMPLS xconnect is configured, while trying to modify the configuration to use L2TPv3 Xconnect on the router.

Workaround: Do not configure L2TPv3 on an interface which previously was used for EoMPLS.

- CSCte81385

When **show network-clock** indicates a “valid” BITS clock state as “valid but not present” on the ASR 1000 Router Series. When a “valid” state BITS clock is removed and re-added, then **show network-clock** indicates BITS state as “valid but not present” even though the Active Source indicates as BITS.

Workaround: There is no workaround. This seems to be a display issue with the **show network-clock** cli output due to the fact that BITS is indicated as the Active Source.

- CSCte83888

When PoD request contains target Acct-Session-Id prepended with NAS-Port-ID

it will not be honored. This condition has been observed, when PoD prepended is configured with NAS-Port-Id for target sessions.

Workaround: Is to use only the Session-Id which is located after the, “_” in the Account-Session-ID to specify the session needing disconnect.

- CSCte82240

SBC accepts “.” when key_addr_type is “DIALED_DIGITS”. This condition can occur, when set exact matching means has been set as:

rpsRtgActionKeyAddrWildcardType to AMB_MW_EXPLICIT_WILDCARD.

This is possible to have a “.” when rpsRtgActionKeyAddrType is set to AMB_MW_ADDR_TYPE_DIALED_DIGITS. However, it is no longer allowed when rpsRtgActionKeyAddrWildcardType is AMB_MW_EXPLICIT_WCARD (which means SBC should perform an explicit match).

Workaround: None

- CSCte97907

On a Cisco ASR 1000 Router with RP2 may get out of sync with NTP master every 18 minutes for approximately 1 minute. This may offset the NTP Master which will cause an increase up to -1052.1 msec and the sync will get lost.

This instance has been observed, when NTP is enabled and running apr. 20 minutes.

Workaround: None