

ribbon™

SBC Core Release Notes

Software Release: 09.02.03R002
Document Revision: 01.01
Published: 19 November 2021

Copyright

© 2020-2021 Ribbon Communications Operating Company, Inc. ("Ribbon"). © 2020-2021 ECI Telecom Ltd.. All rights reserved. The compilation (meaning the collection, arrangement and assembly) of all content on this site is protected by U.S. and international copyright laws and treaty provisions and may not be used, copied, reproduced, modified, published, uploaded, posted, transmitted or distributed in any way, without prior written consent of Ribbon Communications Inc.

Disclaimer and Restrictions

The publication is for information purposes only and is subject to change without notice. This publication does not constitute a commitment on the part of Ribbon. While reasonable efforts have been made in the preparation of this publication to assure its accuracy, Ribbon assumes no liability resulting from technical or editorial errors or omissions, or for any damages whatsoever resulting from the furnishing, performance, or use of the information contained herein. Ribbon reserves the right to make changes to this publication and to Ribbon products without notice in its sole discretion. This publication is not meant to define any interfaces between Ribbon products and any third-party hardware or software products.

Warranties

THIS INFORMATION IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NONINFRINGEMENT. IN NO EVENT SHALL THE RIBBON BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OR PERFORMANCE OF THIS INFORMATION, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Compliance with Applicable Laws and Export Control Laws

The information in this publication is subject to all applicable U.S. federal, state, and local laws. The customer use, distribution, and transfer of any technical information shall be in compliance with all applicable export and import laws and regulations. All Ribbon products and publications are commercial in nature; and the use, duplication, or disclosure by the United States Government is subject to the restrictions set forth in DFARS 252.227-7013 and FAR 52.227-19.

Trademarks

The trademarks, logos, service marks, trade names, and trade dress ("look and feel") on this website, including without limitation the RIBBON and RIBBON logo marks, are protected by applicable US and foreign trademark rights and other proprietary rights and are the property of Ribbon Communications Operating Company, Inc. or its affiliates. Any third-party trademarks, logos, service marks, trade names and trade dress may be the property of their respective owners. Any uses of the trademarks, logos, service marks, trade names, and trade dress without the prior written consent of Ribbon Communications Operating Company, Inc., its affiliates, or the third parties that own the proprietary rights, are expressly prohibited.

UNCONTROLLED COPY: The master of this content is stored in an electronic database and is "write protected"; it may be altered only by authorized persons. While copies may be printed, it is not recommended. Viewing of the master electronically ensures access to the current content. Any hardcopies taken must be regarded as uncontrolled copies.

For access to the technical documentation, log in through the Ribbon Support Services website at <https://ribboncommunications.com/services/support-services/technical-documentation>

1. SBC Core 09.02.03R002 Release Notes 4
1.1 MOP to increase vCPUs Prior to Upgrading SBC SWe on VMware or KVM Hypervisor (9.2.3R2) 176


SBC Core 09.02.03R002 Release Notes

Table of Contents

- [About SBC Release Notes](#)
 - [Release Notes Use and Distribution](#)
 - [Associated Ribbon Announcements](#)
 - [Problems or Questions](#)
 - [About SBC Core](#)
 - [Interoperability](#)
 - [H.323-SIP and SIP-H323 Calls](#)
 - [Compatibility with Ribbon Products](#)
 - [New Features](#)
 - [New Features in Release 09.02.03R002](#)
 - [New Features in Previous Releases](#)
 - [Sample Heat Templates Included in This Release](#)
 - [SBC SWe Cloud Requirements for OpenStack](#)
 - [OpenStack Requirements](#)
 - [SBC SWe Requirements for KVM](#)
 - [SBC SWe Requirements for VMware](#)
 - [Requirements for Using the Infrastructure as Code Environment with SBC SWe](#)
 - [Required Software and Firmware Versions](#)
 - [How to Verify Currently Installed Software/Firmware Versions](#)
 - [Software Bundles](#)
 - [SBC 5000 Series \(51x0/52x0\) Firmware](#)
 - [SBC 5400 Firmware](#)
 - [SBC 7000 Series Firmware](#)
 - [SBC Core Operating System Installation Package](#)
 - [SBC Core Application Package](#)
 - [Cloud Service Archive \(CSAR\) Packages for VNFM Deployment on OpenStack](#)
- [Upgrade Notes](#)
 - [09.02.03R002 Upgrade Information](#)
 - [SBC SWe Pre-Upgrade Requirements](#)
 - [VM CPU resource allocation requirements](#)
 - [Disable Call Trace feature prior to LSWU/upgrade](#)
 - [Manually check for Hostcheck Validation Failed message](#)
 - [Preparing for Upgrade \(All Platforms\)](#)
 - [Supported Live Software Upgrade \(LSWU\) Paths](#)
- [Security Vulnerabilities](#)
- [Resolved Issues](#)
 - [Resolved Issues in 09.02.03R002 Release](#)
 - [Resolved Issues in 09.02.03R001 Release](#)
 - [Resolved Issues in 09.02.03R000 Release](#)
 - [Resolved Issues in 09.02.02R006 and 09.02.02R005 Releases](#)
 - [Resolved Issues in 09.02.02R004 Release](#)
 - [Resolved Issues in 09.02.02R003 Release](#)
 - [Resolved Issues in 09.02.02R002 Release](#)
 - [Resolved Issues in 09.02.02R001 Release](#)
 - [Resolved Issues in 09.02.02R000 Release](#)
 - [Resolved Issues in 09.02.01R003 Release](#)
 - [Resolved Issues in 09.02.01R002 Release](#)
 - [Resolved Issues in 09.02.01R001 Release](#)
 - [Resolved Issues in 09.02.01R000 Release](#)
 - [Resolved Issues in 09.02.00R002 Release](#)
 - [Resolved Issues in 09.02.00R001 Release](#)
 - [Resolved Issues in 09.02.00R000 Release](#)
- [Known Issues](#)
 - [Known Issues in Release 09.02.01R001 to 09.02.03R002](#)
- [Known Limitations](#)
- [Performing a Heat Stack Update when userdata is Updated with SSH Keys](#)

About SBC Release Notes

This release note describes new features, the latest hardware and software requirements, known limitations and other pertinent release information for the latest release of SBC Core.

 Please note that all Ribbon bugs reported by customers on a given software release will be fixed in the latest release on that software release branch.

To view and download the latest End of Product Sale (EoPS) and other End Of Life (EOL) notices, navigate to the Resource Library on the corporate website (<https://ribboncommunications.com/company/get-help/resource-library>).


Release Notes Use and Distribution

Ribbon Release Notes are protected under the copyright laws of the United States of America. This work contains proprietary information of Ribbon Communications, Plano, TX 75023, USA. Use, disclosure, or reproduction in any form is strictly prohibited without prior authorization from Ribbon Communications.

Associated Ribbon Announcements

The following Ribbon announcements (formerly known as WBAs) are referenced in this release note:

- **Warning-14-00020748:** Verify system and databases are fully in sync prior to Live Software Upgrade (LSWU). Applies to all SBC platforms (HW, SWe, Cloud) except the SBCs deployed in a Distributed SBC (D-SBC) architecture
- **Warning-21-00029858:** The AWS SBC Might Fail to Come Up Due to a Metadata Query Failure from the Metadata Server.
- **Warning-21-00029859:** Policy Data syncInProgress after Upgrade Revert
- **Warning-21-00029843:** Unable to Access SBC EMA and PM Post LSWU due to server.key File Corruption

 To view/download Ribbon announcements, do the following:

1. Log on to the Ribbon Support Portal (<https://ribboncommunications.com/services/ribbon-support-portal-login>).
2. From the Quick Access menu, click and download the "Ribbon Support Portal User Guide", and navigate to the "ANNOUNCEMENTS tab" section for instructions to search for and view announcements.

Problems or Questions

For problems or questions, contact the Global Support Assistance Center:

Ribbon Support Portal: <https://ribboncommunications.com/services/ribbon-support-portal>

Voice: +1-833-RIBBON1 (1-833-742-2661)

About SBC Core

The SBC Core platforms address the next-generation needs of SIP communications by delivering media transcoding, robust security and advanced call routing in a high-performance, 2RU, and 5RU form-factor devices enabling service providers and enterprises to quickly and securely enhance their network by implementing services like SIP trunking, secure Unified Communications and Voice over IP (VoIP).

For more product information, refer to the section [About SBC Core](#) in the main documentation space.

Interoperability

The SBC Core software interoperates with the following:

- SIP/H.323 compliant IADs and IP-PBXs
- PSX Policy Server Softswitch via SIP redirects and/or Diameter+ protocol
- SBC 9000 through SIP call signaling and Networks MCS protocol

H.323-SIP and SIP-H323 Calls

When using H.323-SIP and SIP-H.323 call flows, an additional Re-invite/Update may get generated towards the SIP side. To suppress this, enable the IP Signaling Profile (IPSP) flag `Minimize Relaying Of Media Changes From Other Call Leg` at the SIP side.

- For CLI IPSP flag details, refer to [Flags - CLI](#).
- For EMA IPSP flag details, refer to [Common Ip Attributes - Flags](#).

**Note**

H.323 is not supported on SBC SWe cloud deployments.

Compatibility with Ribbon Products

**Tip**

When upgrading your network, ensure to upgrade each product to the most current release to take advantage of the latest features, enhancements, and fixes.

**Info**

For complete interoperability details between various Ribbon products, including backwards compatibility, refer to [Ribbon Product Compatibilities](#).

Refer to [SBC 5000-7000-SWe Interoperability Matrix](#) for the latest and minimum compatible product versions supporting this release.

New Features

New Features in Release 09.02.03R002

There are no new features in this release.

New Features in Previous Releases

To view features in previous releases, refer to the following release notes:

- [SBC Core 09.02.03R000 Release Notes](#)
- [SBC Core 09.02.02R002 Release Notes](#)
- [SBC Core 09.02.02R001 Release Notes](#)
- [SBC Core 09.02.01R000 Release Notes](#)
- [SBC Core 09.02.00R001 Release Notes](#)

Sample Heat Templates Included in This Release

To instantiate the SBC instances, the following templates can be used:

Table 1: SBC Heat Templates

Template Name	Description
heatRgNoDhcp. yaml	Used to instantiate no DHCP, IPv4 or IPv6 deployments. The template supports I-SBC, M-SBC, S-SBC, MRFP and SLB node types. This template includes instructions to enable port redundancy.
heatOamNoDhcp. yaml	Used to instantiate an OAM node.
heatRgNoDhcp- TSBC-template. yaml	Used to instantiate a T-SBC node.



Note

Example template files are packaged together in .tar.gz and .sha256 files separate from the SBC Core application installation and upgrade files:

- cloudTemplates.tar.gz
- cloudTemplates.tar.gz.sha256

SBC SWe Cloud Requirements for OpenStack

The system hosting the SBC SWe Cloud must meet the below requirements for OpenStack:

Table 2: Server Hardware Requirements

Configuration	Requirement
Processor	Intel Xeon processors (Nehalem micro-architecture or above) with 6 cores and above (processors should support hyper-threading). <div data-bbox="337 1354 1344 1516" style="border: 1px solid #ccc; padding: 10px;"> <p> Note Ribbon recommends Westmere (or newer) processors for better SRTP performance. These processors have the AES-NI instruction set for performing cryptographic operations in hardware.</p> </div>
RAM	Minimum 24 GiB
Hard Disk	Minimum 100 GB

Network Interface Cards (NICs)	Minimum 4 NICs.
	<p>Note Make sure NICs have multi-queue support which enhances network performance by allowing RX and TX queues to scale with the number of CPUs on multi-processor systems.</p>
	<p>Note The Intel I350, x540, x550, and 82599 Ethernet adapters are supported for configuring as SR-IOV and DirectPath I/O pass-through devices.</p>
	<p>Note The PKT ports must be 10 Gbps SR-IOV enabled ports.</p>
	<p>Note 6 NICs are required to support PKT port redundancy.</p>

The system hosting the SBC SWe must meet the following requirements to achieve the performance targets listed:

Table 3: S-SBC SWe Requirement

S-SBC SWe Requirements for 1000 CPS/120K Signaling Sessions	Notes
32 vCPUs	Due to the workload characteristics, allocate 20 physical cores with two hyper-threaded CPUs from each core to the SBC.
128 GiB RAM	Must be Huge Page memory. The minimum page size is 2048 KiB, but 1048576 is recommended.
100 GB Disk	None
4 vNICs/6 vNICs	<p>Attach MGT0 port to the Management VirtIO Tenant network.</p> <p>HA port has to be on IPv4 VirtIO Tenant network.</p> <p>Attach PKT0 and PKT1 ports to SR-IOV and Provider network.</p> <p>You must have 6 vNICs to enable PKT port redundancy. For more information, refer to the SBC SWe Features Guide.</p>

Table 4: M-SBC SWe Requirement

M-SBC SWe Requirements for 40K Media Sessions	Notes
16 vCPUs	Due to the workload characteristics, allocate 10 physical cores with two hyper-threaded CPUs from each core and from single NUMA node to the SBC.
32 GiB RAM	Must be Huge Page memory. The minimum page size is 2048 KiB, but 1048576 is recommended.
100 GB Disk	None

4 vNICs/ 6 vNICs	<p>Attach MGT0 port to the Management VirtIO Tenant network.</p> <p>HA port has to be on IPv4 VirtIO Tenant network.</p> <p>Attach PKT0 and PKT1 ports to SR-IOV and Provider network.</p> <p>You must have 6 vNICs to enable PKT port redundancy. For more information, refer to the SBC SWe Features Guide.</p>
------------------	---

**Note**

All NIC ports must come from the same NUMA node from which the M-SBC SWe instance is hosted.

OpenStack Requirements

The SBC SWe supports the following OpenStack environments:

- **Newton** with RHOSP 10 and RHEL 7.4
- **Queens** with RHOSP 13 and RHEL 7.5

**Note**

The SBC SWe was tested on OpenStack Queens with RHOSP 13 and RHEL 7.5.

SBC SWe Requirements for KVM

The following table lists the server hardware requirements.

Table 5: KVM Hypervisor Server Hardware Requirements

Configuration	Requirement
Processor	<p>Intel Xeon processors (Nehalem micro-architecture or above) with 6 cores and above (processors should support hyper threading).</p> <div style="border: 1px solid #ccc; padding: 5px; margin-bottom: 5px;"> <p>Note Ribbon recommends using Westmere (or newer) processors for better SRTP performance. These processors have the AES-NI instruction set for performing cryptographic operations in hardware.</p> </div> <div style="border: 1px solid #ccc; padding: 5px;"> <p>Note The supported CPU Family number is 6 and CPU Model number must be newer than 26. Refer to the Intel Architecture and Processor Identification document for more information.</p> </div>
RAM	Minimum 24 GB
Hard Disk	Minimum 500 GB
Network Interface Cards (NICs)	<p>Minimum 4 NICs</p> <div style="border: 1px solid #ccc; padding: 5px; margin-bottom: 5px;"> <p>Note Make sure NICs have multi-queue support which enhances network performance by allowing RX and TX queues to scale with the number of CPUs on multi-processor systems.</p> </div> <div style="border: 1px solid #ccc; padding: 5px;"> <p>Note The Intel I350, x540, x550, x710, and 82599 Ethernet adapters are supported for configuring as SR-IOV and DirectPath I/O pass-through devices.</p> </div>
Ports	<p>Number of ports allowed:</p> <ul style="list-style-type: none"> • 2 Management ports. For more information on the second Management port, refer to Second Management Port for SWe Deployed on KVM Hypervisor. • 1 HA port • 2 Media ports

SBC SWe Requirements for VMware

The following table lists the server hardware requirements:

SBC SWe for VMware – Server Hardware Requirements

Configuration	Requirement
Processor	<p>Intel Xeon processors (Nehalem micro-architecture or above) with 6 cores and above (processors should support hyper threading).</p> <div style="border: 1px solid #ccc; padding: 5px; margin-bottom: 5px;"> <p>Note Ribbon recommends using Westmere (or newer) processors for better SRTP performance. These processors have the AES-NI instruction set for performing cryptographic operations in hardware.</p> </div> <div style="border: 1px solid #ccc; padding: 5px; margin-bottom: 5px;"> <p>Note The supported CPU Family number is 6 and CPU Model number must be newer than 26. Refer to the Intel Architecture and Processor Identification document for more information.</p> </div> <div style="border: 1px solid #ccc; padding: 5px;"> <p>Note ESXi 6.5 and later releases require approximately 2 physical cores to be set aside for hypervisor functionality. The number of VMs which can be hosted on a server must be planned for accordingly.</p> </div>
RAM	Minimum 24 GB
Hard Disk	Minimum 500 GB
Network Interface Cards (NICs)	<p>Minimum 4 NICs, if physical NIC redundancy is not required.</p> <div style="border: 1px solid #ccc; padding: 5px; margin-bottom: 5px;"> <p>Note Make sure NICs have multi-queue support which enhances network performance by allowing RX and TX queues to scale with the number of CPUs on multi-processor systems.</p> </div> <div style="border: 1px solid #ccc; padding: 5px;"> <p>Note</p> <ul style="list-style-type: none"> The following NICs are supported for configuring as SR-IOV and DirectPath I/O pass-through devices. SR-IOV is supported only with 10 Gbps interfaces (x540/82599/x710): Intel I350, x540, x550, x710 and 82599, Mellanox Connect - 4x, Mellanox Connect - 5x, and CISCO VIC. The VMware Enterprise Plus license is required for SR-IOV. </div>
Ports	<p>Number of ports allowed:</p> <ul style="list-style-type: none"> 2 Management ports. For more information on the second Management port, refer to Second Management Port for SWe Deployed in VMware. 1 HA port 2 Media Ports (no port-redundancy) 4 Media ports (with port-redundancy)

Requirements for Using the Infrastructure as Code Environment with SBC SWe

The following tarball file is required to use the IaC environment to deploy SWe N:1 deployments on VMware:

- raf-20.12-iac1.0_sustaining-240.tar.gz

The environment in which you place and expand the IaC tarball must include:

- A Linux server running RedHat Enterprise Linux (RHEL), CentOS 7 or Debian 9
- Python 2.7 or later
- An internet connection for downloading and installing additional files that may be required for your deployment
- Root access on the instance or ability to become root (for example, using sudo su -)
- Access to the vSphere ESXi host IP from the Linux server where the IaC environment is set up.

For more information on IaC, refer to [Using the Ribbon IaC Environment to Deploy SBC SWe on VMware](#).

Required Software and Firmware Versions

The following SBC 51x0/52x0, SBC 5400 and SBC 7000 software and firmware versions are required for this release. For 5xx0, the BIOS is installed during application installation; whereas, for 5400 and 7000, the BMC/BIOS is included in the firmware package and installed during the firmware upgrade.

Table 6: Required Software and Firmware Versions

Components	Software/Firmware	Version
SBC Platform	SBC 51x0/52x0 BMC	V03.22.00-R000 Kernel: 3.10.108 Busybox: v1.27.2 Openssh: 7.9p1 Openssl: 1.0.2n Lighttpd: 1.4.48-r0 Qualys security issues Password encryption method is SHA512 Lighttpd is secured and supports only TLS1.2 cipher.
	SBC 51x0/52x0 BIOS	V2.7.0
	SBC 5400 Firmware	BMC: V03.22.00-R000 BIOS: V1.18.0
	SBC 7000 Firmware	BMC: V03.22.00-R000 BIOS: V2.14.0
SBC Application	Operating System (OS) Version	V08.02.03-R002
	SonusDB	V09.02.03-R002
	SBC Application	V09.02.03-R002



Note

The firmware package of SBC 5400 and 7000 series includes BMC, BIOS, and other binaries. The firmware is upgraded from the BMC.

How to Verify Currently Installed Software/Firmware Versions

Use the EMA to verify the currently installed software and firmware versions.

Log on to the EMA, and from the main screen navigate to **Monitoring > Dashboard > [System and Software Info](#)**.

Software Bundles

The following software release bundles are available for download from the Customer Portal:

- SBC5x7x_9.2
- SBCSWe_9.2

Download the appropriate software packages for your desired configuration from the Customer Portal (<https://ribboncommunications.com/services/ribbon-support-portal-login>) to your PC:



Note

When upgrading from release 9.0 and above, upload the SHA256 checksum file. Otherwise, use the MD5 file.

SBC 5000 Series (51x0/52x0) Firmware

- firmware-5XX0-V03.22.00-R000.img
- firmware-5XX0-V03.22.00-R000.img.md5
- bmc5X00_v3.22.0-R0.rom.md5sum
- bmc5X00_v3.22.0-R0.rom

SBC 5400 Firmware

- firmware-5400-V03.22.00-R000.img
- firmware-5400-V03.22.00-R000.img.md5

SBC 7000 Series Firmware

- firmware-7X00-V03.22.00-R000.img
- firmware-7X00-V03.22.00-R000.img.md5



Note

Execute the Method Of Procedure (MOP) only for upgrading the FPGA image of an SBC 7000 DSP-LC card when the SBC 7000 DSP-LC FPGA version is 0x14. The MOP can be applied at any version time, with the only restriction being that the BMC firmware version is at least 1.25.0. However, if the SBC application is running version V05.01.00R000 or higher, then the DSPs will be set to disabled and transcoding and transrating calls will fail if the SBC 7000 DSP-LC FPGA version is 0x14. Therefore, it is necessary to upgrade the SBC 7000 DSP-LC FPGA if the version is 0x14, before upgrading the SBC to 5.1.0. However, the MOP can be applied if the application version is higher than 5.1.0. Click [Here](#) to view the 550-06210_DSP-LC_FPGA_Upgrade_MOP.

SBC Core Operating System Installation Package

The ConnexIP Operating System installation package for SBC Core:

- sbc-V09.02.03R002-connexip-os_08.02.03-R002_11_amd64.iso
- sbc-V09.02.03R002-connexip-os_08.02.03-R002_11_amd64.iso.sha256
- sbc-V09.02.03R002-connexip-os_08.02.03-R002_11_amd64.iso.md5



Note

Once the ConnexIP ISO procedure is completed, the SBC application package is automatically uploaded to SBC platforms.



Release 9.2 OS Patches

Release 9.2 includes a new set of OS security patches, and also a new version of confD. Release 9.2.1 includes a new set of OS security patches including the fix for CVE-2021-3156: Heap-Based Buffer Overflow in Sudo (Baron Samedit).

SBC Core Application Package

The SBC Application installation and upgrade package for SBC Core:

- sbc-V09.02.03R002-connexip-os_08.02.03-R002_11_amd64.qcow2
- sbc-V09.02.03R002-connexip-os_08.02.03-R002_11_amd64.qcow2.sha256
- sbc-V09.02.03R002-connexip-os_08.02.03-R002_11_amd64.qcow2.md5
- sbc-V09.02.03R002-connexip-os_08.02.03-R002_11_amd64.ova
- sbc-V09.02.03R002-connexip-os_08.02.03-R002_11_amd64.ova.sha256
- sbc-V09.02.03R002-connexip-os_08.02.03-R002_11_amd64.ova.md5
- sbc-V09.02.03-R002.x86_64.tar.gz
- sbc-V09.02.03-R002.x86_64.sha256
- sbc-V09.02.03-R002.x86_64.md5
- sbc-V09.02.03-R002.x86_64.signature

For detailed information on installation and upgrade procedures, refer to [SBC Core Software Installation and Upgrade Guide](#).

Cloud Service Archive (CSAR) Packages for VNFM Deployment on OpenStack

These files are for SBC SWe deployments in the OpenStack cloud using VNFM.

For VNFM deployment, the VNF Descriptor (VNFD) file is provided in a Cloud Service Archive (CSAR) package for the type of SBC cluster being deploying. VNFs are independent and CSAR definitions are imported into the VNFM via an Onboarding mechanism. There is a procedure for producing the required CSAR variant, for different personalities (S-SBC, M-SBC), different interface types (virtio, sriov).

Files required for CSAR creation:

- createVnfmCsar.py
- vnfmSol001VnfdTemplate.yaml
- sbc-V09.02.03R002-connexip-os_08.02.03-R002_11_amd64.qcow2

For detailed information on installation and upgrade procedures, refer to [SBC Core Software Installation and Upgrade Guide](#).

For details on CSAR creation, refer to [Creating a CSAR Package File](#).

Upgrade Notes

**Warning**

A LSWU on an SBC 7000 should only be performed when the total number of active calls on the system is below 18,000. If the criteria is not met, a double failure during the upgrade may occur, thereby losing all active calls. If such a failure occurs, both active and standby SBC services will go down. Contact Ribbon Support immediately.

**Warning**

Customers upgrading from 922R1 using VMware or KVM need to run the following command as root user on both the active and standby instances:

```
touch /opt/sonus/conf/swe/capacityEstimates/.indexMarker
```

This is not required for upgrades from earlier releases.

**Note**

This note is only applicable for SBC N:1 mode running on virtualized platforms like KVM, VMware. Before beginning the upgrade to 9.2.2R1, issue the following commands on OAM, and then save the configuration to the EMS. This is needed because from 9.2.1R2 release onwards, NTP and timezone settings provided during the deployment will be applied during the system boot-up, and to avoid overwriting those settings by the default NTP and timezone values which are stored in config, you must update the NTP and timezone settings in CDB before the upgrade and save the configuration to the EMS.

```
set system ntp serverAdmin 169.xxx.xxx.x state disabled
commit

set system ntp serverAdmin <ntp_server_ip> version version4 minPoll 4 maxPoll 10
commit

set system ntp serverAdmin <ntp_server_ip> state enabled
commit

set system ntp timeZone vsbcSystem zone <timezone_value>
commit
```

where

<ntp_server_ip> is the IP address for NTP server.

<timezone_value> is the timezone value you must set for the system.

**Note**

Once the installation or upgrade completes on the SBC 51x0 and SBC SWe platforms, the copy of the installation package ([SBC Core Installation and Upgrade Package](#)) is automatically removed from the system.

**Note**

Release 9.2 and later requires additional user account security practices for SBC SWe deployments in Openstack cloud environments. During upgrade of SBC SWe cloud instances deployed using Heat templates, you must use a template that includes SSH keys or passwords for the admin and linuxadmin accounts. The example Heat templates have been updated to include information on how to specify this type of data in the userdata section of a template.

**Note**

As an SBC Core password security enhancement, user passwords automatically expire after upgrading to 8.0.x. As a result, users are required to change their passwords upon initial login immediately following the upgrade.

**Note**

Customers using network licensing mode will be converted to node locked mode (formerly legacy mode) after upgrade to the SBC 8.0.0 Release.

**Note**

The SBC 8.0 5xx0 and 7000 platforms may exhibit a 7% degradation of CPU performance relative to earlier releases. This is attributable to the Spectre/Meltdown security patches.

**Note**

In order to take advantage of performance improvements due to hyper-threading refer to the following [MOP to increase the number of vCPUs prior to SBC SWe \(KVM Hypervisor or VMware\)](#) upgrades from pre-07.01.00R000 release to 07.01.00R000 or higher.

**Note**

In the case of a Live Software Upgrade (LSWU) from 6.0.0R000/6.0.0R001/6.0.0F001/6.0.0F002 to 8.0, The action "Perform Pre-Upgrade Checks" from PM is not supported. Please contact Ribbon Support.

**Note**

The number of rules across SMM profiles in a system is limited to 10000, and the number of actions across profiles in a system is limited to 50000.

Ensure the above conditions are met before LSWU.

**Note**

In NFV environments, the method used for upgrades involves rebuilding the instance, which requires additional disk space on the host. The minimum disk space needed for this operation is listed in the table below.

Table 7: Disk Space Requirements

Flavor	Extra Space Required (GB)
S-SBC	80
M-SBC	80
PSX-M	360
PSX-S	360
PSX-Test	360
EMS_SA	150

**Note**

The SBC 51xx and 52xx systems require 24GB of RAM to run 6.x code or higher.

**Note**

SWe SBC software enforces I-SBC instances to run only with a single vNUMA node in order to achieve deterministic performance. SWe SBC VM having >8 vCPUs hosted on dual-socket physical server with VMware ESXi software needs to follow the steps below to correct vNUMA topology before upgrading to latest SWe SBC software:

- Check 'numa.nodeAffinity' settings on VM. It should be either 0 or 1 based on which NUMA node PKT ports are connected. The following command on ESXi host can be used to check PKT port NUMA affinity:

```
vsish -e get /net/pNics/<PKT port name - vmmnicX>/properties | grep "NUMA"
```

If any of the above settings requires modification, follow the steps below on SWe SBC HA system:

- Shutdown the standby instance.
- If 'numa.nodeAffinity' settings are missing, add the following rows:

numa.autosize.once = FALSE

numa.nodeAffinity' = 0 or 1 (based on PKT port NIC affinity)

On ESXi 6.5 and above releases, vSphere web client can be used to add above rows under Edit settings > VM options > configuration parameters > add parameters;

On ESXi 6.0 and below releases, it can be added under Edit > Advanced > general > configuration parameters > add rows using vSphere client.

- Power-on standby instance.
- Once standby instance is up and running, do manual switchover through CLI and repeat above steps to modify 'number of virtual sockets', 'numa.nodeAffinity' and 'numa.autosize.once' settings.

For more information, refer to:

- [Create a VM](#)
- [Creating Virtual Machines using SR-IOV Interfaces](#)
- [Creating a New SBC SWe VM Instance with Direct IO Passthru](#)

**Note**

Before beginning the upgrade on a SBC running code prior to 8.2R0, the following commands on all the DNS Groups needs to be issued if "ednsSupport" is enabled.

Failure statistics are not being mirrored correctly, and the LSWU state may stay in "syncing" if the "ednsFailures " count is non-zero.

1. Issue the following CLI command to clear the EDNS statistics for all DNS Groups in the system.

```
admin@PLUM> request addressContext default dnsGroup DnsGrp dnsServerReset
reason DNS Server statistics are Reset
```

```
[ok][2020-11-06 04:08:13]
```

```
admin@PLUM> show status addressContext default dnsGroup DnsGrp
dnsServerStatistics 2
```

```
{ ipAddress 10.xx.xx.xx; queries 0; timeouts 0; errors 0; referrals 0; totalTcpConnection 0; tcpConnectionFailed 0;
tcpConnectionSuccess 0; tcpConnectiontorndown 0; tcpFallback 0; ednsStatus supported; ednsFailures 0; }
```

```
[ok][2020-11-06 04:08:22]
```

```
admin@PLUM>
```

2. Disable the ednsSupport to stop mirroring of the statistics if the error count is constantly incrementing or likely to increase during the upgrade.

```
set addressContext default dnsGroup DnsGrp ednsSupport disabled
```

Note: The ednsServer stats will be lost/reset during the upgrade.

**Note**

If the TRF/MRB Features are configured and enabled – some calls are unable to be cleared post upgrade if using the TRF/MRB attributes.

The upgrade is successful and calls continue but some calls may fail to clean up release post upgrade. Session KeepAlive and RTP Inactivity functions will clean any stale calls.

Enable the sessionKeepalive or rtpInactivity monitoring to ensure that mirrored calls are cleaned up post upgrade.

```
set addressContext default zone ZONE_AS sipTrunkGroup TG_AS_SIPP signaling timers sessionKeepalive <value>
```

OR

```
set system media mediaPeerInactivity <value>
```

```
set profiles media packetServiceProfile DEFAULT peerAbsenceAction peerAbsenceTrapAndDisconnect
```

**Note**

Upgrade from a pre 8.0 release with globalization support for registration enabled will see a registration drop during an upgrade.

If the following localNumberSupport is enabled, those registrations will be dropped after first switchover during LSWU.

```
% set addressContext <name> zone <name> sipTrunkGroup <name> signaling localNumberSupport <disabled | enabled>
```

09.02.03R002 Upgrade Information

**Warning**

Prior to performing an upgrade to this release, you must remove usernames that do not conform to the SBC user-naming rules to prevent upgrade failure. Upgrade can proceed successfully after removing all invalid usernames. The following user-naming rules apply:

- Usernames can begin with A-Z a-z _ only.
- Usernames cannot start with a period, dash, or digit.
- Usernames can contain a period(.), dash(-), alphabetic characters, digits, or underscore(_).
- Usernames cannot consist of digits only.
- Usernames can contain a maximum of 23 characters.

The following names are not allowed:

```
tty disk kmem dialout fax voice cdrom floppy tape sudo audio dip src utmp video sasl plugdev staff users nogroup i2c dba operator
```

Note: Any CLI usernames consisting of digits only or not conforming to new user naming rules will be removed after performing a restore config in release 9.2.3R002.

**Warning**

Prior to performing an upgrade to the 9.x release, the dnsGroups with type mgmt must be specified/updated with the "interface" field. The steps are included in announcement "W-17-00022847".

If the above MOP is not run, the LSWU process may fail because of duplicate trunk group or zone names.

**Warning**

Prior to performing an upgrade to 9.x release, the duplicate trunk groups or zones must be removed. The steps are included in announcement "W-17-00022689".

If you are upgrading from any SBC version with ePSX configuration to this release, execute the Method of Procedure, [MOP to Reconfigure SBC \(with ePSX\) to External PSX Prior to an Upgrade to 06.00.00R000 Release](#) prior to performing an upgrade. For a list of supported LSWU paths, refer to [Supported Upgrade Paths](#).

SBC SWe Pre-Upgrade Requirements

VM CPU resource allocation requirements

CPU resource allocation requirements for SBC SWe VM are strictly enforced. You must review and verify these VM settings (including co-hosted VMs) against the documented "VM Configuration Recommendations" on the [For VMware](#) page in the [Hardware and Software Requirements](#) section before upgrading.

If you encounter a problem, correct the CPU reservation settings as specified in step 6 of the "Adjust Resource Allocations" procedure on [Creating a New SBC SWe VM Instance with VMXNET3](#):



Set the CPU reservation for the VM so that it equals the physical processor CPU speed, multiplied by the number of vCPUs divided by two.

For example, a configuration of 4 vCPUs with a processor of 2.99 GHz CPU speed, reserve: $2992 * 4/2 = 5984$ MHz

If the VM uses the same number of vCPUs as the number of physical processors on the server, this reservation may not be possible. In this case, reduce the number of vCPUs assigned to VM by **one** and set the CPU reservation to the appropriate value.

When using the `show table system serverSoftwareUpgradeStatus` command during the upgrade, the Standby server's LSWU status will always display "Upgrading" even though the upgrade may have failed due to **host checker validation**. To check if host validation failed for the Standby, check for **HostCheck Validation Failed** message in the `upgrade.out` log.

Disable Call Trace feature prior to LSWU/upgrade

As a prerequisite for SWe LSWU/upgrade, disable the Call Trace feature prior to performing the LSWU/upgrade and re-enable it once the LSWU/upgrade is completed.

Manually check for Hostcheck Validation Failed message

Perform the following procedure on the Standby to check for the **Hostcheck Validation Failed** message in the `upgrade.out` log.

1. Log on to **ESXi** of the Standby SBC SWe.
2. Check in `/opt/sonus/staging/upgrade.out` (this log shows the **Hostcheck Validation Failed** error).
3. Power off the VM.
4. Reduce the number of vCPUs assigned to VM by **one** and set the CPU reservation to the appropriate value.
5. Power on the VM. The SBC SWe successfully upgrades to the latest version 6.2.0.
6. Run the command `show table system serverSoftwareUpgradeStatus` to confirm the successful upgrade.
7. Perform similar procedure for LSWU on Active.

Preparing for Upgrade (All Platforms)



Note

The SBC 8.0 release skips the SRV query if the flag in a DNS NAPTR response from the DNS server indicates to proceed with "A" record query as per RFC 2915/3403. This is a change in behavior from previous releases, where the SBC performed SRV queries irrespective of the "flag" setting returned by DNS Server. If you use DNS NAPTR/SRV/A record query from SBC to determine peer transport address, ensure the DNS Server is configured to return 'S' flag to invoke an SRV query.



Note

In this release, LSWU infrastructure is added to the Platform Manager (PM), providing the ability to perform LSWU upgrades to later releases using the PM. However, this feature is not currently supported in 4.2.x releases and should not be used at this time.



Warning

Customers who are using the SBC to interop with MS Teams need to review and compare their configuration against the latest configuration guide especially the SMM as it might result in call failures after upgrade if the older SMM is left in place. For more information, refer to [SBC 9.2 - MS Teams Solution Guide](#).

Supported Live Software Upgrade (LSWU) Paths



Attention

This release includes all bug fixes implemented in the releases which are documented in the Supported Upgrade Paths table of this release note.

To view bug fixes in previous releases, refer to the release note(s) of interest from the [SBC 5xx0-7000-SWe Documentation Home page](#).

The SBC Core supports Live Software Upgrade from releases listed in the table below:

Table 8: Supported Upgrade Paths

V06.xx	V07.xx	V08.xx	V09.xx
V06.00.00F009	V07.00.00R000	V08.00.00R000	V09.00.00R000
V06.00.00F014	V07.00.00F001	V08.01.00R000	V09.01.00R000
V06.02.00R000	V07.00.00F002	V08.01.00F001	V09.01.00R001
V06.02.00F000	V07.00.00F003	V08.01.00R001	V09.01.00R002
V06.02.01R000	V07.00.00F004	V08.01.00R002	V09.01.00R003
V06.02.01R001	V07.00.00F005	V08.01.00R003	V09.02.00R000
V06.02.01R002	V07.00.00F006	V08.01.00R004	V09.02.00R001
V06.02.01F001	V07.00.00S400	V08.01.00R005	V09.02.00R002
V06.02.01F002	V07.00.00S401	V08.01.00R006	V09.02.01R000
V06.02.01F003	V07.00.00S402	V08.01.00R007	V09.02.01R001
V06.02.01F004	V07.00.00S404	V08.01.00R008	V09.02.01R002
V06.02.01F005	V07.00.00S405	V08.02.00R000	V09.02.01R003
V06.02.01F006	V07.00.00S406	V08.02.00F001	V09.02.01R004
V06.02.01F007	V07.00.00S407	V08.02.00F002	V09.02.01R005
V06.02.01F008	V07.01.00R000	V08.02.00R001	V09.02.01R006
V06.02.01F009	V07.01.00R001	V08.02.00R002	V09.02.02R000
V06.02.01F010	V07.01.00R002	V08.02.01R000	V09.02.02R001
V06.02.01F011	V07.01.00R003	V08.02.01F001	V09.02.02R002
V06.02.01F012	V07.01.00R004	V08.02.01F002	V09.02.02R003
V06.02.02R000	V07.01.00F001	V08.02.01F003	V09.02.02R004
V06.02.02R001	V07.01.00F002	V08.02.02R000	V09.02.02R005
V06.02.02F001	V07.01.00F003	V08.02.02R001	V09.02.02R006
V06.02.02F002	V07.02.00R000	V08.02.02R002	V09.02.03R000
V06.02.02F003	V07.02.00R001	V08.02.02R003	V09.02.03R001
V06.02.02F004	V07.02.00R002	V08.02.02R004	
V06.02.02F005	V07.02.00S400	V08.02.02R005	
V06.02.02F006	V07.02.00S401	V08.02.03R000	
V06.02.02F007	V07.02.00S809	V08.02.03R001	

V06.02.02F008	V07.02.00S810	V08.02.03F001	
V06.02.02F009	V07.02.01R000	V08.02.04R000	
V06.02.02F010	V07.02.01R001	V08.02.04F001	
V06.02.02F011	V07.02.01R002	V08.02.04R001	
V06.02.02F012	V07.02.01R003	V08.02.04R002	
V06.02.02F013	V07.02.01R004	V08.02.04R003	
V06.02.02F014	V07.02.01F001	V08.02.04R004	
V06.02.03R000	V07.02.01F002	V08.02.05R000	
V06.02.03F001	V07.02.01F004	V08.02.05R001	
V06.02.03F002	V07.02.01F005	V08.02.05R002	
V06.02.03F003	V07.02.01S400	V08.02.05R003	
V06.02.03F004	V07.02.01R005	V08.02.05R004	
V06.02.03F005	V07.02.01R006		
V06.02.03F006	V07.02.01R007		
V06.02.03F007	V07.02.01R008		
V06.02.04R000	V07.02.01R009		
V06.02.04F001	V07.02.01R010		
V06.02.04F002	V07.02.02R000		
V06.02.05R000	V07.02.02R001		
	V07.02.02R002		
	V07.02.02R003		
	V07.02.02R004		
	V07.02.02R005		
	V07.02.02F001		
	V07.02.03R000		
	V07.02.03R001		
	V07.02.03R002		
	V07.02.03R003		
	V07.02.03R004		
	V07.02.03S400		
	V07.02.03S401		
	V07.02.04R000		
	V07.02.04R001		
	V07.02.04R002		
	V07.02.04R003		
	V07.02.04R004		
	V07.02.05R000		
	V07.02.05R001		
	V07.02.05R002		

	V07.02.05R003		
	V07.02.05R004		
	V07.02.05R005		
	V07.02.05R006		
	V07.02.05R007		
	V07.02.05R008		

Security Vulnerabilities

The following table displays the security vulnerabilities resolved in this release.

CVE	Risk	Description
CVE-2021-26691	Critical	In Apache HTTP Server versions 2.4.0 to 2.4.46 a specially crafted SessionHeader sent by an origin server could cause a heap overflow
CVE-2018-20060	Critical	urllib3 before version 1.23 does not remove the Authorization HTTP header when following a cross-origin redirect (i.e., a redirect that differs in host, port, or scheme). This can allow for credentials in the Authorization header to be exposed to unintended hosts or transmitted in cleartext.
CVE-2021-31870	Critical	An issue was discovered in klibc before 2.0.9. Multiplication in the calloc() function may result in an integer overflow and a subsequent heap buffer overflow.
CVE-2019-18218	Critical	cdf_read_property_info in cdf.c in file through 5.37 does not restrict the number of CDF_VECTOR elements, which allows a heap-based buffer overflow (4-byte out-of-bounds write).
CVE-2021-31873	Critical	An issue was discovered in klibc before 2.0.9. Additions in the malloc() function may result in an integer overflow and a subsequent heap buffer overflow.
CVE-2021-31872	Critical	An issue was discovered in klibc before 2.0.9. Multiple possible integer overflows in the cpio command on 32-bit systems may result in a buffer overflow or other security impact.
CVE-2021-3246	High	A heap buffer overflow vulnerability in msadpcm_decode_block of libsndfile 1.0.30 allows attackers to execute arbitrary code via a crafted WAV file.
CVE-2020-35452	High	Apache HTTP Server versions 2.4.0 to 2.4.46 A specially crafted Digest nonce can cause a stack overflow in mod_auth_digest. There is no report of this overflow being exploitable, nor the Apache HTTP Server team could create one, though some particular compiler and/or compilation option might make it possible, with limited consequences anyway due to the size (a single byte) and the value (zero byte) of the overflow
CVE-2021-29154	High	BPF JIT compilers in the Linux kernel through 5.11.12 have incorrect computation of branch displacements, allowing them to execute arbitrary code within the kernel context. This affects arch/x86/net/bpf_jit_comp.c and arch/x86/net/bpf_jit_comp32.c.
CVE-2020-35524	High	A heap-based buffer overflow flaw was found in libtiff in the handling of TIFF images in libtiff's TIFF2PDF tool. A specially crafted TIFF file can lead to arbitrary code execution. The highest threat from this vulnerability is to confidentiality, integrity, as well as system availability.
CVE-2020-25672	High	A memory leak vulnerability was found in Linux kernel in llcp_sock_connect
CVE-2020-25671	High	A vulnerability was found in Linux Kernel, where a refcount leak in llcp_sock_connect() causing use-after-free which might lead to privilege escalations.
CVE-2021-33560	High	Libcrypt before 1.8.8 and 1.9.x before 1.9.3 mishandles ElGamal encryption because it lacks exponent blinding to address a side-channel attack against mpi_powm, and the window size is not chosen appropriately. This, for example, affects use of ElGamal in OpenPGP.
CVE-2021-32399	High	net/bluetooth/hci_request.c in the Linux kernel through 5.12.2 has a race condition for removal of the HCI controller.
CVE-2021-26690	High	Apache HTTP Server versions 2.4.0 to 2.4.46 A specially crafted Cookie header handled by mod_session can cause a NULL pointer dereference and crash, leading to a possible Denial Of Service
CVE-2021-3483	High	A flaw was found in the Nosy driver in the Linux kernel. This issue allows a device to be inserted twice into a doubly-linked list, leading to a use-after-free when one of these devices is removed. The highest threat from this vulnerability is to confidentiality, integrity, as well as system availability. Versions before kernel 5.12-rc6 are affected

CVE-2021-28660	High	rtw_wx_set_scan in drivers/staging/rtl8188eu/os_dep/ioctl_linux.c in the Linux kernel through 5.11.6 allows writing beyond the end of the ->ssid[] array. NOTE: from the perspective of kernel.org releases, CVE IDs are not normally used for drivers/staging/* (unfinished work); however, system integrators may have situations in which a drivers/staging issue is relevant to their own customer base.
CVE-2021-2388	High	Vulnerability in the Java SE, Oracle GraalVM Enterprise Edition product of Oracle Java SE (component: Hotspot). Supported versions that are affected are Java SE: 8u291, 11.0.11, 16.0.1; Oracle GraalVM Enterprise Edition: 20.3.2 and 21.1.0. Difficult to exploit vulnerability allows unauthenticated attacker with network access via multiple protocols to compromise Java SE, Oracle GraalVM Enterprise Edition. Successful attacks require human interaction from a person other than the attacker. Successful attacks of this vulnerability can result in takeover of Java SE, Oracle GraalVM Enterprise Edition. Note: This vulnerability applies to Java deployments, typically in clients running sandboxed Java Web Start applications or sandboxed Java applets, that load and run untrusted code (e.g., code that comes from the internet) and rely on the Java sandbox for security. This vulnerability does not apply to Java deployments, typically in servers, that load and run only trusted code (e.g., code installed by an administrator). CVSS 3.1 Base Score 7.5 (Confidentiality, Integrity and Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:H/PR:N/UI:R/S:U/C:H/I:H/A:H).
CVE-2021-21702	High	In PHP versions 7.3.x below 7.3.27, 7.4.x below 7.4.15 and 8.0.x below 8.0.2, when using SOAP extension to connect to a SOAP server, a malicious SOAP server could return malformed XML data as a response that would cause PHP to access a null pointer and thus cause a crash.
CVE-2021-0512	High	In __hidinput_change_resolution_multipliers of hid-input.c, there is a possible out of bounds write due to a heap buffer overflow. This could lead to local escalation of privilege with no additional execution privileges needed. User interaction is not needed for exploitation.Product: AndroidVersions: Android kernelAndroid ID: A-173843328References: Upstream kernel
CVE-2021-23133	High	A race condition in Linux kernel SCTP sockets (net/sctp/socket.c) before 5.12-rc8 can lead to kernel privilege escalation from the context of a network service or an unprivileged process. If sctp_destroy_sock is called without sock_net(sk)->sctp.addr_wq_lock then an element is removed from the auto_asconf_splist list without any proper locking. This can be exploited by an attacker with network service privileges to escalate to root or from the context of an unprivileged user directly if a BPF_CGROUP_INET_SOCKET_CREATE is attached which denies creation of some SCTP socket.
CVE-2021-22555	High	A heap out-of-bounds write affecting Linux since v2.6.19-rc1 was discovered in net/netfilter/x_tables.c. This allows an attacker to gain privileges or cause a DoS (via heap memory corruption) through user name space
CVE-2021-31618	High	Apache HTTP Server protocol handler for the HTTP/2 protocol checks received request headers against the size limitations as configured for the server and used for the HTTP/1 protocol as well. On violation of these restrictions and HTTP response is sent to the client with a status code indicating why the request was rejected. This rejection response was not fully initialised in the HTTP/2 protocol handler if the offending header was the very first one received or appeared in a footer. This led to a NULL pointer dereference on initialised memory, crashing reliably the child process. Since such a triggering HTTP/2 request is easy to craft and submit, this can be exploited to DoS the server. This issue affected mod_http2 1.15.17 and Apache HTTP Server version 2.4.47 only. Apache HTTP Server 2.4.47 was never released.
CVE-2020-35523	High	An integer overflow flaw was found in libtiff that exists in the tif_getimage.c file. This flaw allows an attacker to inject and execute arbitrary code when a user opens a crafted TIFF file. The highest threat from this vulnerability is to confidentiality, integrity, as well as system availability.
CVE-2021-33909	High	fs/seq_file.c in the Linux kernel 3.16 through 5.13.x before 5.13.4 does not properly restrict seq buffer allocations, leading to an integer overflow, an Out-of-bounds Write, and escalation to root by an unprivileged user, aka CID-8cae8cd89f05.
CVE-2021-23134	High	Use After Free vulnerability in nfc sockets in the Linux Kernel before 5.12.4 allows local attackers to elevate their privileges. In typical configurations, the issue can only be triggered by a privileged local user with the CAP_NET_RAW capability.
CVE-2021-31871	High	An issue was discovered in klibc before 2.0.9. An integer overflow in the cpio command may result in a NULL pointer dereference on 64-bit systems.
CVE-2021-33034	High	In the Linux kernel before 5.12.4, net/bluetooth/hci_event.c has a use-after-free when destroying an hci_chan, aka CID-5c4c8c954409. This leads to writing an arbitrary value.
CVE-2019-11324	High	The urllib3 library before 1.24.2 for Python mishandles certain cases where the desired set of CA certificates is different from the OS store of CA certificates, which results in SSL connections succeeding in situations where a verification failure is the correct outcome. This is related to use of the ssl_context, ca_certs, or ca_certs_dir argument.

CVE-2020-25670	High	A vulnerability was found in Linux Kernel where refcount leak in llcp_sock_bind() causing use-after-free which might lead to privilege escalations.
CVE-2021-30641	Medium	Apache HTTP Server versions 2.4.39 to 2.4.46 Unexpected matching behavior with 'MergeSlashes OFF'
CVE-2020-1934	Medium	In Apache HTTP Server 2.4.0 to 2.4.41, mod_proxy_ftp may use uninitialized memory when proxying to a malicious FTP server.
CVE-2021-0129	Medium	Improper access control in BlueZ may allow an authenticated user to potentially enable information disclosure via adjacent access.
CVE-2020-25673	Medium	A vulnerability was found in Linux kernel where non-blocking socket in llcp_sock_connect() leads to leak and eventually hanging-up the system.
CVE-2021-28964	Medium	A race condition was discovered in get_old_root in fs/btrfs/ctree.c in the Linux kernel through 5.11.8. It allows attackers to cause a denial of service (BUG) because of a lack of locking on an extent buffer before a cloning operation, aka CID-dbcc7d57bffc.
CVE-2021-3564	Medium	A flaw double-free memory corruption in the Linux kernel HCI device initialization subsystem was found in the way user attach malicious HCI TTY Bluetooth device. A local user could use this flaw to crash the system. This flaw affects all the Linux kernel versions starting from 3.13.
CVE-2021-30002	Medium	An issue was discovered in the Linux kernel before 5.11.3 when a webcam device exists. video_usercopy in drivers/media/v4l2-core/v4l2-ioctl.c has a memory leak for large arguments, aka CID-fb18802a338b.
CVE-2019-11236	Medium	In the urllib3 library through 1.24.1 for Python, CRLF injection is possible if the attacker controls the request parameter.
CVE-2021-28971	Medium	In intel_pmu_drain_pebs_nhm in arch/x86/events/intel/ds.c in the Linux kernel through 5.11.8 on some Haswell CPUs, userspace applications (such as perf-fuzzer) can cause a system crash because the PEBS status in a PEBS record is mishandled, aka CID-d88d05a9e0b6.
CVE-2020-26558	Medium	Bluetooth LE and BR/EDR secure pairing in Bluetooth Core Specification 2.1 through 5.2 may permit a nearby man-in-the-middle attacker to identify the Passkey used during pairing (in the Passkey authentication procedure) by reflection of the public key and the authentication evidence of the initiating device, potentially permitting this attacker to complete authenticated pairing with the responding device using the correct Passkey for the pairing session. The attack methodology determines the Passkey value one bit at a time.
CVE-2021-29265	Medium	An issue was discovered in the Linux kernel before 5.11.7. usbip_sockfd_store in drivers/usb/usbip/stub_dev.c allows attackers to cause a denial of service (GPF) because the stub-up sequence has race conditions during an update of the local and shared status, aka CID-9380afd6df70.
CVE-2020-26147	Medium	An issue was discovered in the Linux kernel 5.8.9. The WEP, WPA, WPA2, and WPA3 implementations reassemble fragments even though some of them were sent in plaintext. This vulnerability can be abused to inject packets and/or exfiltrate selected fragments when another device sends fragmented frames and the WEP, CCMP, or GCMP data-confidentiality protocol is used.
CVE-2020-26139	Medium	An issue was discovered in the kernel in NetBSD 7.1. An Access Point (AP) forwards EAPOL frames to other clients even though the sender has not yet successfully authenticated to the AP. This might be abused in projected Wi-Fi networks to launch denial-of-service attacks against connected clients and makes it easier to exploit other vulnerabilities in connected clients.
CVE-2020-1927	Medium	In Apache HTTP Server 2.4.0 to 2.4.41, redirects configured with mod_rewrite that were intended to be self-referential might be fooled by encoded newlines and redirect instead to an unexpected URL within the request URL.
CVE-2021-20292	Medium	There is a flaw reported in the Linux kernel in versions before 5.9 in drivers/gpu/drm/nouveau/nouveau_sgdma.c in nouveau_sgdma_create_ttm in Nouveau DRM subsystem. The issue results from the lack of validating the existence of an object prior to performing operations on the object. An attacker with a local account with a root privilege, can leverage this vulnerability to escalate privileges and execute code in the context of the kernel.

CVE-2021-2369	Medium	Vulnerability in the Java SE, Oracle GraalVM Enterprise Edition product of Oracle Java SE (component: Library). Supported versions that are affected are Java SE: 7u301, 8u291, 11.0.11, 16.0.1; Oracle GraalVM Enterprise Edition: 20.3.2 and 21.1.0. Easily exploitable vulnerability allows unauthenticated attacker with network access via multiple protocols to compromise Java SE, Oracle GraalVM Enterprise Edition. Successful attacks require human interaction from a person other than the attacker. Successful attacks of this vulnerability can result in unauthorized update, insert or delete access to some of Java SE, Oracle GraalVM Enterprise Edition accessible data. Note: This vulnerability applies to Java deployments, typically in clients running sandboxed Java Web Start applications or sandboxed Java applets, that load and run untrusted code (e.g., code that comes from the internet) and rely on the Java sandbox for security. This vulnerability does not apply to Java deployments, typically in servers, that load and run only trusted code (e.g., code installed by an administrator). CVSS 3.1 Base Score 4.3 (Integrity impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:N/UI:R/S:U/C:N/I:L/A:N).
CVE-2021-31829	Medium	kernel/bpf/verifier.c in the Linux kernel through 5.12.1 performs undesirable speculative loads, leading to disclosure of stack content via side-channel attacks, aka CID-801c6058d14a. The specific concern is not protecting the BPF stack area against speculative loads. Also, the BPF stack can contain uninitialized data that might represent sensitive information previously operated on by the kernel.
CVE-2021-34693	Medium	net/can/bcm.c in the Linux kernel through 5.12.10 allows local users to obtain sensitive information from kernel stack memory because parts of a data structure are uninitialized.
CVE-2020-26137	Medium	urllib3 before 1.25.9 allows CRLF injection if the attacker controls the HTTP request method, as demonstrated by inserting CR and LF control characters in the first argument of putrequest(). NOTE: this is similar to CVE-2020-26116.
CVE-2020-36322	Medium	An issue was discovered in the FUSE filesystem implementation in the Linux kernel before 5.10.6, aka CID-5d069dbe8aaf. fuse_do_getattr() calls make_bad_inode() in inappropriate situations, causing a system crash. NOTE: the original fix for this vulnerability was incomplete, and its incompleteness is tracked as CVE-2021-28950.
CVE-2021-29155	Medium	An issue was discovered in the Linux kernel through 5.11.x. kernel/bpf/verifier.c performs undesirable out-of-bounds speculation on pointer arithmetic, leading to side-channel attacks that defeat Spectre mitigations and obtain sensitive information from kernel memory. Specifically, for sequences of pointer arithmetic operations, the pointer modification performed by the first operation is not correctly accounted for when restricting subsequent operations.
CVE-2021-38208	Medium	net/nfc/lcp_sock.c in the Linux kernel before 5.12.10 allows local unprivileged users to cause a denial of service (NULL pointer dereference and BUG) by making a getsockname call after a certain type of failure of a bind call.
CVE-2021-28688	Medium	The fix for XSA-365 includes initialization of pointers such that subsequent cleanup code wouldn't use uninitialized or stale values. This initialization went too far and may under certain conditions also overwrite pointers which are in need of cleaning up. The lack of cleanup would result in leaking persistent grants. The leak in turn would prevent fully cleaning up after a respective guest has died, leaving around zombie domains. All Linux versions having the fix for XSA-365 applied are vulnerable. XSA-365 was classified to affect versions back to at least 3.11.
CVE-2021-28950	Medium	An issue was discovered in fs/fuse/fuse_i.h in the Linux kernel before 5.11.8. A "stall on CPU" can occur because a retry loop continually finds the same bad inode, aka CID-775c5033a0d1.
CVE-2021-29650	Medium	An issue was discovered in the Linux kernel before 5.11.11. The netfilter subsystem allows attackers to cause a denial of service (panic) because net/netfilter/x_tables.c and include/linux/netfilter/x_tables.h lack a full memory barrier upon the assignment of a new table value, aka CID-175e476b8cdf.
CVE-2021-29647	Medium	An issue was discovered in the Linux kernel before 5.11.11. qrtr_recvmsg in net/qrtr/qrtr.c allows attackers to obtain sensitive information from kernel memory because of a partially uninitialized data structure, aka CID-50535249f624.
CVE-2020-7071	Medium	In PHP versions 7.3.x below 7.3.26, 7.4.x below 7.4.14 and 8.0.0, when validating URL with functions like filter_var(\$url, FILTER_VALIDATE_URL), PHP will accept an URL with invalid password as valid URL. This may lead to functions that rely on URL being valid to mis-parse the URL and produce wrong data as components of the URL.
CVE-2020-36311	Medium	An issue was discovered in the Linux kernel before 5.9. arch/x86/kvm/svm/sev.c allows attackers to cause a denial of service (soft lockup) by triggering destruction of a large SEV VM (which requires unregistering many encrypted regions), aka CID-7be74942f184.
CVE-2021-3573	Medium	A use-after-free in function hci_sock_bound_ioctl() of the Linux kernel HCI subsystem was found in the way user calls ioctl HCIUNBLOCKADDR or other way triggers race condition of the call hci_unregister_dev() together with one of the calls hci_sock_blacklist_add(), hci_sock_blacklist_del(), hci_get_conn_info(), hci_get_auth_info(). A privileged local user could use this flaw to crash the system or escalate their privileges on the system. This flaw affects the Linux kernel versions prior to 5.13-rc5.

CVE-2021-31916	Medium	An out-of-bounds (OOB) memory write flaw was found in list_devices in drivers/md/dm-ioctl.c in the Multi-device driver module in the Linux kernel before 5.12. A bound check failure allows an attacker with special user (CAP_SYS_ADMIN) privilege to gain access to out-of-bounds memory leading to a system crash or a leak of internal kernel information. The highest threat from this vulnerability is to system availability.
CVE-2021-33910	Medium	basic/unit-name.c in systemd prior to 246.15, 247.8, 248.5, and 249.1 has a Memory Allocation with an Excessive Size Value (involving strdupa and alloca for a pathname controlled by a local attacker) that results in an operating system crash.
CVE-2020-24586	Low	The 802.11 standard that underpins Wi-Fi Protected Access (WPA, WPA2, and WPA3) and Wired Equivalent Privacy (WEP) doesn't require that received fragments be cleared from memory after (re)connecting to a network. Under the right circumstances, when another device sends fragmented frames encrypted using WEP, CCMP, or GCMP, this can be abused to inject arbitrary network packets and/or exfiltrate user data.
CVE-2021-38209	Low	net/netfilter/nf_conntrack_standalone.c in the Linux kernel before 5.12.2 allows observation of changes in any net namespace because these changes are leaked into all other net namespaces. This is related to the NF_SYSCTL_CT_MAX, NF_SYSCTL_CT_EXPECT_MAX, and NF_SYSCTL_CT_BUCKETS sysctls.
CVE-2020-24588	Low	The 802.11 standard that underpins Wi-Fi Protected Access (WPA, WPA2, and WPA3) and Wired Equivalent Privacy (WEP) doesn't require that the A-MSDU flag in the plaintext QoS header field is authenticated. Against devices that support receiving non-SSP A-MSDU frames (which is mandatory as part of 802.11n), an adversary can abuse this to inject arbitrary network packets.
CVE-2020-29374	Low	An issue was discovered in the Linux kernel before 5.7.3, related to mm/gup.c and mm/huge_memory.c. The get_user_pages (aka gup) implementation, when used for a copy-on-write page, does not properly consider the semantics of read operations and therefore can grant unintended write access, aka CID-17839856fd58.
CVE-2021-22898	Low	curl 7.7 through 7.76.1 suffers from an information disclosure when the '-t' command line option, known as 'CURLOPT_TELNETOPTIONS' in libcurl, is used to send variable=content pairs to TELNET servers. Due to a flaw in the option parser for sending NEW_ENV variables, libcurl could be made to pass on uninitialized data from a stack based buffer to the server, resulting in potentially revealing sensitive internal information to the server using a clear-text network protocol.
CVE-2021-21781	Low	An information disclosure vulnerability exists in the ARM SIGPAGE functionality of Linux Kernel v5.4.66 and v5.4.54. The latest version (5.11-rc4) seems to still be vulnerable. A userland application can read the contents of the sigpage, which can leak kernel memory contents. An attacker can read a process's memory at a specific offset to trigger this vulnerability. This was fixed in kernel releases: 4.14.222 4.19.177 5.4.99 5.10.17 5.11
CVE-2021-2341	Low	Vulnerability in the Java SE, Oracle GraalVM Enterprise Edition product of Oracle Java SE (component: Networking). Supported versions that are affected are Java SE: 7u301, 8u291, 11.0.11, 16.0.1; Oracle GraalVM Enterprise Edition: 20.3.2 and 21.1.0. Difficult to exploit vulnerability allows unauthenticated attacker with network access via multiple protocols to compromise Java SE, Oracle GraalVM Enterprise Edition. Successful attacks require human interaction from a person other than the attacker. Successful attacks of this vulnerability can result in unauthorized read access to a subset of Java SE, Oracle GraalVM Enterprise Edition accessible data. Note: This vulnerability applies to Java deployments, typically in clients running sandboxed Java Web Start applications or sandboxed Java applets, that load and run untrusted code (e.g., code that comes from the internet) and rely on the Java sandbox for security. This vulnerability does not apply to Java deployments, typically in servers, that load and run only trusted code (e.g., code installed by an administrator). CVSS 3.1 Base Score 3.1 (Confidentiality impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:H/PR:N/UI:R/S:U/C:L/I:N/A:N).
CVE-2021-22924	Low	libcurl keeps previously used connections in a connection pool for subsequent transfers to reuse, if one of them matches the setup. Due to errors in the logic, the config matching function did not take 'issuercert' into account and it compared the involved paths *case insensitively*, which could lead to libcurl reusing wrong connections. File paths are, or can be, case sensitive on many systems but not all, and can even vary depending on used file systems. The comparison also didn't include the 'issuer cert' which a transfer can set to qualify how to verify the server certificate.
CVE-2020-24587	Low	The 802.11 standard that underpins Wi-Fi Protected Access (WPA, WPA2, and WPA3) and Wired Equivalent Privacy (WEP) doesn't require that all fragments of a frame are encrypted under the same key. An adversary can abuse this to decrypt selected fragments when another device sends fragmented frames and the WEP, CCMP, or GCMP encryption key is periodically renewed.

Resolved Issues

Resolved Issues in 09.02.03R002 Release

The following Severity 1 issues are resolved in this release:

Table 9: Severity 1 Resolved Issues

Issue ID	Sev	Problem Description	Resolution
SBX-113305 SBX-114425	1	<p>PortFix SBX-113305: DNS recovery packet count of DSCP as "0".</p> <p>Impact: The DNS probe packets are sent from the SBC with a DSCP value of "0".</p> <p>Root Cause: The configured DSCP values were not set for DNS probe keepalive packets.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Configure DNS server with a DSCP value. 2. Ensure the DNS server blacklisting is enabled. 3. Ensure the DNS query is configured DNS server. 4. Shut down the DNS server during this process. <p>Result: Observe that the configured DNS server is blacklisted and sending DNS probe packet towards a DNS server with configured DSCP value.</p>	<p>The code is modified to set a configured DSCP value for packets, including DNS probe packets.</p> <p>Workaround: None.</p>
SBX-114012 SBX-114320	1	<p>PortFix SBX-114012: The SMM is not working when the From contains an escape character.</p> <p>Impact: The SMM may fail to parse a display name in double quote string when the string has more than one escape characters next to each other of a header.</p> <p>Root Cause: Logical error in parsing logic that causes a parsing failure.</p> <p>Steps to Replicate: Configure a rule to access double quote string display name of uri header, or any parameter.</p> <p>Ensure there is an input display name, such as From: "PhoneLite \\\"test1\\\"test2 " < sip:[field0]@dns.com>;tag=testing</p>	<p>The code is modified to address the issue.</p> <p>Workaround: None.</p>
SBX-109056 SBX-114568	1	<p>PortFix SBX-109056: The SBC 5400 MS Teams CQ has transfers issues.</p> <p>Impact: When a call with DLRBT enabled undergoes more than one transfer with REFER signaling and the final transferee does not accept the call, the SBC does not correctly reconnect the call back to the transferor.</p> <p>Root Cause: For an A-to-B call, if B, for example, successfully REFERS the call to C, and then C attempts to REFER the call to D (or back to B), the SBC plays a ringtone towards A while the REFER is taking place.</p> <p>However, when D rejects the call with a 480 "Temporarily not available" message, the SBC does not correctly disable the call context that is playing the tone, nor does it correctly re-enable the A-to-C call association.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. With DLRBT configured on ingress and egress of SBC. Establish an A-to-B call through the SBC. 2. Send a REFER from B to refer the call to C and complete the refer related signaling. This results in an A-to-C call. 3. Send a REFER from C to refer the call to D. Do not answer the call at D but reject it with a 480 "Temporarily not available" message. 4. Verify that the SBC stops playing tone towards A and re-establishes the A-to-C call again and media can correctly flow between A and C. 	<p>The code is modified to disable the tone playing towards A when a 480 is received and then correctly re-enable the call association back to the transferor.</p> <p>Workaround: None.</p>
SBX-114315 SBX-114427	1	<p>PortFix SBX-114315: An SBC failover occurred because of an SCM Process core.</p> <p>Impact: An SCM core can occur if the hostName in the From Header is longer than 64 bytes.</p> <p>Root Cause: The code is copying the host name into an array that is not long enough to hold it. This results in memory corruption.</p> <p>Steps to Replicate: Make a call where the hostName in the From Header is longer than 64 bytes.</p>	<p>The code is modified to use the correct size when copying the host name.</p> <p>Workaround: None.</p>

The following Severity 2 issue is resolved in this release:

Table 10: Severity 2 Resolved Issues

Issue ID	Sev	Problem Description	Resolution
SBX-112547 SBX-114405	2	<p>PortFix SBX-112547: The SBC routes call to 2nd DNS record if 503 is received after 18x.</p> <p>Impact: The SBC routes an INVITE to next DNS record if 503 is received after 18x.</p> <p>Also, even when the dnsCrankback flag is disabled, on getting 503/INVITE timeout case, the SBC reroutes an INVITE to next DNS record.</p> <p>Root Cause: Due to a design defect, the SBC tries the next DNS record even if an error response is received after an 18x.</p> <p>Also, retrying the next DNS record during a 503/timeout when dnsCrankback flag is disabled is legacy behavior. This behaviour is changed to retry the next DNS record only when the dnsCrankback flag is enabled.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Configure the DNS server with two routes for the FQDN route. 2. Make an A-to-B call. 3. The SBC sends an INVITE to first DNS record of and B sends 18x and the 503. 	<p>The code is modified to not apply the DNS crankback procedure if an error response is received after a 18x. Additionally, the default behavior of handling timeout/503 when the dnsCrankback is disabled is updated as part of this fix.</p> <p>With this fix, the SBC retries for the next DNS record only when dnsCrankback is enabled to ensure the error responses match the reasons configured in the crankback profile.</p> <p>Workaround: None.</p>
SBX-113053 SBX-114355	2	<p>PortFix SBX-113053: Race Condition when 183 and 200 OK (INVITE) are received simultaneously. 200 OK is not relayed to other side.</p> <p>Impact: Race condition between internal processing of the 200 OK for an UPDATE and a received 183 Session Progress.</p> <p>Root Cause: After a 200 OK (UPDATE) is received, followed immediately by receipt of 183 Session progress, a subsequent 200 OK (INVITE) containing P-Early-Media header is queued indefinitely at egress when downstreamForking is enabled and earlyMedia forkingBehaviour is pemPriority.</p> <p>Steps to Replicate: Egress SIP trunk group has downstreamForking is enabled and earlyMedia forkingBehaviour is pemPriority.</p> <p>Make a call so that egress signaling is as follows: --> INVITE (with SDP and P-Early-Media: supported) <-- 100 Trying <-- 183 Session Progress (with SDP and no P-Early-Media) --> PRACK <-- 200 OK (for PRACK) --> UPDATE (with SDP) <-- 200 OK (for UPDATE with SDP) <-- 183 Session Progress (no SDP but P-Early-Media: sendrecv) --> PRACK <-- 200 OK (for PRACK) <-- 200 OK (for INVITE) (This message gets queued forever).</p>	<p>The code is modified to eliminate the race condition.</p> <p>Workaround: None.</p>
SBX-113614 SBX-114546	3	<p>PortFix SBX-113614: The ActiveRegCount is greater on the INGRESS side.</p> <p>Impact: When registration(s) take longer than 90 seconds to complete, the ingress zone's activeRegs count can increment, but not decrement, when the registration terminates. This may cause the ingress zone's activeRegs count to grow incorrectly, and never reach ZERO (even after all registrations are terminated).</p> <p>Root Cause: The SIPFE and SIPSG RCB allocation/deallocation become out of sync, indirectly causing the ingress zone's activeRegs count to increment and not decrement.</p> <p>Steps to Replicate: Perform the endless REGISTER/403, or REGISTER/503, or REGISTER/603 sequence.</p>	<p>The code is modified to handle registration(s) that take longer than 90 seconds to complete. In this case, the SIPFE restarts the registration bind timer and avoids "prematurely" deallocating the RCB so that the FE and SG RCB(s) stay "in-sync".</p> <p>Workaround: None.</p>
GSX-54479 SBX-114528	4	<p>PortFix GSX-54479: The isup-oli parameter is sending a single digit OLI value in isup-oli parameter for values 0-9.</p> <p>Impact: The SBC sends all OLI digits as a single digit.</p> <p>Root Cause: This functionality is not within spec.</p> <p>Steps to Replicate: Make a call that uses OLI parameter.</p>	<p>The code is modified to add a leading 0 for OLI digits 0-9.</p> <p>Workaround: None.</p>

Resolved Issues in 09.02.03R001 Release

The following Severity 1 issues are resolved in this release:

Table 11: Severity 1 Resolved Issues

Issue ID	Sev	Problem Description	Resolution
SBX-113954 SBX-104348	1	<p>PortFix SBX-104348: STUN was received on the SBC and the SBC considers these packets are arriving on another IP. And as a result, the STUN wasn't processed.</p> <p>Impact: The SBC passes the STUN packets to wrong IP address when it received them on the alternate IP.</p> <p>Root Cause: The SBC always uses the primary IP address of LIF to send out the packet, even when it receives the STUN packets on the alternate LIF IP.</p> <p>Steps to Replicate: STUN received on alternate IP of LIF and the SBC does not respond correctly, which appears as the SBC not processing it.</p>	<p>The code is modified so the XRM traverses the IP address list to match the address where it received the packet, instead of directly using the primary IP to send the packets.</p> <p>Workaround: None.</p>
SBX-114073 SBX-113853	1	<p>PortFix SBX-113853: The SBC experienced an intermittent crash.</p> <p>Impact: When multiple duplicate Diversion headers are received and stiProfile is enabled and configured on the ingress trunk group, the SBC dumps core.</p> <p>The issue only occurs if six or more Diversion headers are received but fewer than five of them are unique.</p> <p>Root Cause: Removing duplicated Diversion headers from information sent to the PSX causes crash.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Configure the stiProfile, enable, and assign to ingress trunk group. 2. Send an SIP INVITE containing six Diversion headers but only two of them are unique. In other words, send in two Diversion headers repeated three times. 	<p>The code is modified to remove the duplicate Diversion headers from information sent to the PSX.</p> <p>Workaround: None.</p>

The following Severity 2 issue is resolved in this release:

Table 12: Severity 2 Resolved Issues

Issue ID	Sev	Problem Description	Resolution
SBX-114019	2	<p>The DBG logs was rolling rapidly with UacSendUpdate messages.</p> <p>Impact: The SBC DBG logs were filling up quickly with the following log message.</p> <p>114 10192021 114119.823662:1.01.03.16703.MAJOR .SIPSG: UacSendUpdate - local answer for UPDATE: LegSide=Ingress</p> <p>Root Cause: The code was mistakenly printing logs at the MAJOR level.</p> <p>Steps to Replicate: The log is generated during a call where the SBC tries to send an UPDATE out to ingress leg while the PRACK is pending for previous 18x sent.</p> <p>The following control also needs to be disabled on the trunk group the messages are being sent on. set addressContext default zone <zone name> sipTrunkGroup <TG name> signaling doNotAutoAnswer <enable/disable>.</p>	<p>The code is modified to only print the logs at an INFO level.</p> <p>Workaround: None.</p>

Resolved Issues in 09.02.03R000 Release

The following Severity 1 issues are resolved in this release:

Table 13: Severity 1 Resolved Issues

Issue ID	Sev	Problem Description	Resolution
----------	-----	---------------------	------------

SBX-112996 SBX-113115	1	<p>PortFix SBX-112996: There was a SCM core dump.</p> <p>Impact: A rare race condition triggers a bug that caused SCM to core.</p> <p>Root Cause: A bug in the code causes an attempt to access freed memory. This Gateway-to-Gateway bug has existed for a very long time, but was not encountered until now.</p> <p>The bug is triggered by a rare race condition.</p> <p>Steps to Replicate: This bug was found through code inspection.</p> <p>The bug is triggered by a rare race condition that is not reproducible.</p>	<p>The code is modified to avoid accessing freed memory.</p> <p>Workaround: There is no workaround.</p>
SBX-112366 SBX-112738	1	<p>PortFix SBX-112366: A customer's redundancy had an unexpected switchover.</p> <p>Impact: The SBC is coring in the IPsec/IKE code when a call is using an IKE Protection Profile that was configured without any algorithms.</p> <p>Root Cause: The SBC coredumped in the IPsec/IKE code due to an attempt to de-reference a NULL pointer. The pointer is NULL because the IKE Protection Profile being used was configured without any algorithms.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Configure an IKE Protection Profile without any algorithms. 2. Run a call that will use this profile. 	<p>The code is modified to check for a NULL pointer and take the correct error path when the pointer is NULL.</p> <p>Workaround: When configuring the IKE Protection Profiles, ensure that an algorithm is configured for this profile.</p>
SBX-112349 SBX-112847	1	<p>PortFix SBX-112349: The SBC releases a 132 Release code (MODULE FAILURE) message when running an audit call.</p> <p>Impact: An Async CMD Error reported by XRM that triggered call being torn down with release code 132 when running a call audit.</p> <p>Root Cause: Ribbon lacks the necessary level of detail to conclusively determine the root cause.</p> <p>Steps to Replicate: Since we do not know the exact condition and what types of resources were involved, we cannot provide detailed test steps for this case.</p> <p>Suggest running regular regression tests to attempt to reproduce the issue.</p>	<p>The code is modified for LE2MCAST, LE2MCAST_RTCP, LE2SPLITTER. The changes provided are based on source code inspection.</p> <p>Workaround: N/A</p>
SBX-112677 SBX-112715	1	<p>PortFix SBX-112677: The SCMP cored in Late media passthrough calls over a GW-GW.</p> <p>Impact: There was a coredump for a late media passthrough case.</p> <p>Root Cause: Dereferences a NULL Pointer.</p> <p>Steps to Replicate: For the SBC GW-GW setup.</p> <p>Enable the LM passthrough flag.</p> <p>Run the following procedure:</p> <ol style="list-style-type: none"> 1. Ingress sends LM Invite. 2. Egress sends 18x with SDP. 3. Ingress sends SDP in Prack. 4. Once call is established, send LM re-Invite request from Egress. 5. Ingress responds with 200 Ok and receives ACK with SDP. 6. Ingress sends LM re-Invite request towards egress. 	<p>The code is modified to stop the coredump.</p> <p>Workaround: None.</p>
SBX-107747 SBX-112649	1	<p>PortFix SBX-107747: Cyclic switchover tests - Observed a PRS Process core dump on the M-SBC1.</p> <p>Impact: If an M-SBC instance transitions to active shortly after the instance comes up as standby, when certain performance stats are collected, the PRS process will crash.</p> <p>Root Cause: When the instance comes up as a standby, the active instances send their metavariable information to the new standby instance. This information is sent twice for each active instances. The first metavariable information for each active instance is not received by the PRS process because the PRS process is not ready to receive it. If the instance switches to active before the second set of metavariable information is received, then the new active will not have metavariable information from the former active, which causes the PRS crash.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Bring up an MSBC 4-1 cluster. 2. Restart the standby instance. 3. Immediately after the standby instance is fully up, restart one of the active instances. 4. After the standby instance is fully up as the active instance, in the CLI run the following command: <pre>unhide debug show table global icmpGeneralGroupCurrentStatistics</pre>	<p>The code is modified so the CHM process queues the first set of metavariable information until the PRS process is ready to receive it, and then sends the information to the PRS process.</p> <p>Workaround: None.</p>

<p>SBX-112445 SBX-112880</p>	<p>1</p>	<p>PortFix SBX-112445: Conference calls fail when taking a recorded (SIPREC) path, but they work when SIPREC settings are not enabled.</p> <p>Impact: The SBC cleanup call during hold retrieval for the transferred call. This issue is observed only when the SIP Rec is enabled.</p> <p>Root Cause: Invalid operation on media resource during call hold/unhold for the transferred call leads call to cleanup. This issue is observed only when the SIP Rec is enabled.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Enable the SIP Rec. 2. Make a blind/attended call transfer and after successful transfer, perform a call hold and unhold. 	<p>The code is modified to not perform any invalid operation on media resources during call hold/unhold for the transferred call.</p> <p>Workaround: NA.</p>
<p>SBX-112561 SBX-112836</p>	<p>1</p>	<p>PortFix SBX-112561: Regexp string was not exported by "user-config-export"</p> <p>Impact: In the Regexp, the SMM is lost while exporting a configuration using the user-config-export CLI command.</p> <p>Root Cause: The XML formatting is trimming empty node values.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Create the SMM Profile using the following command strings: <pre> set profiles signaling sipAdaptorProfile ESMM state disabled set profiles signaling sipAdaptorProfile ESMM advancedSMM disabled set profiles signaling sipAdaptorProfile ESMM profileType messageManipulation set profiles signaling sipAdaptorProfile ESMM rule 1 applyMatchHeader one set profiles signaling sipAdaptorProfile ESMM rule 1 criterion 1 type message set profiles signaling sipAdaptorProfile ESMM rule 1 criterion 1 message set profiles signaling sipAdaptorProfile ESMM rule 1 criterion 1 message messageTypes requestAll set profiles signaling sipAdaptorProfile ESMM rule 1 criterion 2 type header set profiles signaling sipAdaptorProfile ESMM rule 1 criterion 2 header set profiles signaling sipAdaptorProfile ESMM rule 1 criterion 2 header name WWW-Authenticate set profiles signaling sipAdaptorProfile ESMM rule 1 criterion 2 header condition exist set profiles signaling sipAdaptorProfile ESMM rule 1 criterion 2 header numberOfInstances number 1 set profiles signaling sipAdaptorProfile ESMM rule 1 criterion 2 header numberOfInstances qualifier equal set profiles signaling sipAdaptorProfile ESMM rule 1 criterion 3 type parameter set profiles signaling sipAdaptorProfile ESMM rule 1 criterion 3 parameter set profiles signaling sipAdaptorProfile ESMM rule 1 criterion 3 parameter condition exist set profiles signaling sipAdaptorProfile ESMM rule 1 criterion 3 parameter paramType generic set profiles signaling sipAdaptorProfile ESMM rule 1 criterion 3 parameter name realm set profiles signaling sipAdaptorProfile ESMM rule 1 action 1 type header set profiles signaling sipAdaptorProfile ESMM rule 1 action 1 operation regsub set profiles signaling sipAdaptorProfile ESMM rule 1 action 1 headerInfo headerValue set profiles signaling sipAdaptorProfile ESMM rule 1 action 1 from set profiles signaling sipAdaptorProfile ESMM rule 1 action 1 from type value set profiles signaling sipAdaptorProfile ESMM rule 1 action 1 from value 172.xx.xx.xxx set profiles signaling sipAdaptorProfile ESMM rule 1 action 1 to set profiles signaling sipAdaptorProfile ESMM rule 1 action 1 to type header set profiles signaling sipAdaptorProfile ESMM rule 1 action 1 to value WWW-Authenticate set profiles signaling sipAdaptorProfile ESMM rule 1 action 1 regexp set profiles signaling sipAdaptorProfile ESMM rule 1 action 1 regexp string " " set profiles signaling sipAdaptorProfile ESMM rule 1 action 1 regexp matchInstance all commit </pre> 2. Run the following CLI command to export the configuration: <pre> user-config-export test.xml /profiles/signaling/sipAdaptorProfile </pre> 3. Delete the exported profile from the CLI with the following command: <pre> delete profiles signaling sipAdaptorProfile ESMM commit user-config-import test.xml </pre> 4. Run the following command: <pre> show configuration profiles signaling sipAdaptorProfile display set match string </pre> 5. Check if the Regexp string field, and restore the original value. 	<p>The code is modified to use the external tool "xmllint" to format the XML.</p> <p>Workaround: Manually edit the field in the XML file after an export and before an import to correct the issue.</p>

SBX-111126 SBX-112531	1	<p>PortFix SBX-111126: In the case of delayed offer, the SBC changes the attribute "a=sendonly" into "a=recvonly" before relaying the 200 OK INVITE.</p> <p>Impact: The direction attribute in 200OK sent in response to late media reinvoke is set to an incorrect value when the direction attribute received in 200OK from other side is a value other than a=sendrecv.</p> <p>Root Cause: The datapathmode on the side receiving the late media invite is incorrectly copied from the SBC local datapathmode from the answer leg PSP receiving 200OK when sendSbcSupportedCodecsInLateMediaReinvite flag is enabled.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Enable the sendSbcSupportedCodecsInLateMediaReinvite on the TG towards UAC. 2. Send a direction attribute a=sendonly/recvonly/inactive from the UAS. 3. Send the 200OK to UAC should have direction attribute as expected (a=sendonly/recvonly/inactive respectively) and media is as per expectation. <pre> UAC-----SBC-----UAS INV(no sdp)-> INV(a=sendrecv)-> <-200OK(a=sendonly/recvonly/inactive) <-200OK(a=sendonly/recvonly/inactive) ACK(sdp)-> ACK(no sdp)-> </pre>	<p>The code is modified so the datapathmode in the active PSP on late media offer side is recomputed when the sendSbcSupportedCodecsInLateMediaReinvite flag is enabled.</p> <p>Workaround: Disable the sendSbcSupportedCodecsInLateMediaReinvite IPSP flag on TG receiving late media reinvoke.</p>
SBX-112624	1	<p>The SBC is unable to recognize a PRACK message, when call hunts to a secondary route.</p> <p>Impact: The PRACK was rejected with a 481 when the end-to-end PRACK is active if a call is cranked-back following a successful end-to-end PRACK on the earlier route.</p> <p>Root Cause: If end-to-end PRACK is performed on the first route and then a late crank-back occurs, subsequent PRACK is rejected because stale information is present from the first route.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Ensure the end-to-end PRACK is supported. 2. Make a SIP-SIP call that routes and 18x is received on the first route. As a result, the PRACK procedure is successful. 3. Have the call fail on the first route, e.g. with 404, and the SBC configured to crank-back to a second route. <p>The 18x is received on the second route. When the calling party responds with PRACK, the call has a 481 in response.</p>	<p>The code is modified so that end-to-end PRACK information is started fresh for each route.</p> <p>Workaround: None.</p>
SBX-112307	1	<p>A customer's SBC 7000 SIPREC metadata did not contain an 'AOR' parameter value.</p> <p>Impact: In the SIPREC, the metadata sender and receiver's AOR fields for the tel URI were not populated properly.</p> <p>Root Cause: Due to defect in the design, these fields were populated incorrectly.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Configure the SBC with SIP Rec server. 2. Run a basic call with "tel" URI in From and To header. 	<p>The code is modified to populate as per requirements/standard.</p> <p>Workaround: Not Applicable.</p>
SBX-106316	1	<p>The call established prior to a switchover fails when put on hold after a switchover.</p> <p>Impact: The call established prior to a switchover with a dynamic payload type fails when put on hold after switchover.</p> <p>Root Cause: After a switchover, while processing the hold request, it was unable to locate the codec due to an issue in reconstruction of some flags in PSP data.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Have a HA setup. 2. In IPSP, enable "Minimize relay of media changes from other call leg" flag and "lockdown preferred codec" is enabled. 3. Establish a call with dynamic codec like AMR. 4. Perform a switchover, once standby becomes active send a call hold request. 5. After a short period of time, send an call-un-hold request. <p>result: Call hold invite should be successful with 200 ok for that and call un-hold should also be successful.</p>	<p>The code is modified so that necessary flag for dynamic payloads in PSP data are reconstructed properly.</p> <p>Workaround: Disable "Lock Down Preferred Codec" and enable "Relay Data Path Mode Changes".</p>

SBX-111498	1	<p>The SBC Ladder diagram (.TRC file) shows a response from the wrong IP for a 400 Bad Request.</p> <p>Impact: When an INVITE (with malformed syntax) is sent to a leg1 (pkt1) SIP signaling IP for the SBC, the SBC responds with a 400 Bad request. In the TRC log, for a 400 Bad request PDU, the Local IP/port is printed with a leg 0 (pkt0) SIP Signaling IP Address instead of leg 1 SIP Signaling IP Address. This occurs only when the SBC sends error response message.</p> <p>Root Cause: During formatting of Error message, the code is incorrectly using the pkt0 SIP Signaling IP Address instead of a pkt1 SIP Signaling IP Address resulting in the incorrect value being printed in the TRC log.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Set up a SIP-SBC-SIP. 2. Enable trace logging with level4 and key "peerIpAddress" 3. Run a SIP-SIP call, send INVITE with malformed syntax to pkt1 SIP signaling IP address such that the SBC sends a 400 Bad Request. 4. Verify that in the TRC log, for 400 Bad request PDU, the local IP value is pkt 1 SIP signaling IP address, same source IP where the INVITE message lands. 	<p>The code is modified to correctly read the SIP Signaling IP Address of the leg where the call lands.</p> <p>Workaround: None.</p>
SBX-106037	1	<p>The SBC 7000 is sending a re-INVITE before receiving an ACK to the 200 OK - multi-stream call.</p> <p>Impact: The SBC removes transcodable codecs when the application media m-line is present with port as 0. This results in the SBC sending a re-INVITE without SDP.</p> <p>Root Cause: By default, the SBC does not support transcoding on multi-stream calls. The problem here was that the code incorrectly identified that there was an application stream present even when the port was set to 0. This made it appear as a multi-stream call and resulted in the transcodable codecs being removed.</p> <p>Steps to Replicate: Test Case 1: GW-GW Scenario</p> <ol style="list-style-type: none"> 1. Make a G711-G711 GW-GW call. The SBC will send BYE to peers. The call will fail. Test Case 2: Non GW-GW Scenario 2. Configure below PSPs on ingress and Egress legs Ingress PSP: G729, G711U, G711A This Leg – G711U, G711A, EFR, EVRC, Other Leg – G711U, G711A, Enable Transcode if different pktsize, dtmf, silence suppression and honor preference. Egress PSP: G711SS This Leg – G711U, G711A, Other Leg – G711U, G711A, Enable Transcode if different pktsize, silence suppression 3. Make a G711-G711 call. <p>The user observes logs in TRC file and the SBC will send BYE to peers. The call will fail.</p>	<p>The code is modified so that the stream with a port 0 is effectively ignored when processing transcodable codecs.</p> <p>Workaround: Enabling "AllowAudioTranscodeForMultiStreamCalls" flag on both ingress and egress PSPs will solve the issue.</p>
SBX-112557	1	<p>The password policy is not recognized by the SBC.</p> <p>Impact: The password policy not recognized by the SBC. By default, the admin rules were taken.</p> <p>Root Cause: Password policy was assigned based on roles of the user.</p> <p>Steps to Replicate: Go to Administration/Application Management. In Configure Password Rules tab, check by enable or disable "Use Separate Password Rules for Administrators" where you are able to change the password policy or not in change password model.</p>	<p>The code is modified to use the state variable in PasswordRule class</p> <p>Workaround: Add new variable named "state" to PasswordRule Class as a workaround.</p>
SBX-111743	1	<p>The SBC block list IP with the wrong port.</p> <p>Impact: The ARS block lists port 0, when an INVITE contains a Route header without a port number, and 503 response contains a Retry-After header.</p> <p>Root Cause: The code did not support a case where no port number is provided.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Configure ARS Retry-After profile. 2. Enable callRouting useRouteSet received in ingress and egress trunk groups. 3. Use the UAC to send an INVITE containing Route without port number: Route: <sip:172.xx.xxx.xx;r;dtg=SBX_EXAMPLE> <p>The UAS responds with 503 Service Unavailable with Retry-After header use CLI to check egress zone sipArsStatus.</p>	<p>The code is modified so when no port number is set, we now retrieve the port number from the transaction control block.</p> <p>Workaround: Define an SMM rule that adds the port number (5060) if the Route header in the initial INVITE does not contain a port number.</p>
SBX-113471	1	<p>Multiple SBC core dumps were detected.</p> <p>Impact: There was forking hashtable corrupted memory.</p> <p>Root Cause: Possibility of forking control fails to remove from hashtable when the resource is cleaned up.</p> <p>Steps to Replicate: Unable to reproduce. Run high load with a NAPT, upstream forking and registration and taking a SIP signaling port down/up.</p>	<p>The code is modified to ensure the forking call is removed from hashtable when free.</p> <p>Workaround: None.</p>

SBX-112322	1	<p>The SBC cored when the subscribe crankback failure.</p> <p>Impact: When a relay Subscribe tried to crank back for the second route and if the SBC does not find trunk group, the SBC may core.</p> <p>Root Cause: The SBC cores due to duplicated memory freed.</p> <p>Steps to Replicate: Configure two routes for the relay Subscribe. The second route is invalid.</p> <p>After the first route exhausted, the SBC try to cranked back for the second route. The SBC fails to find trunk group for the second route.</p>	<p>The code is modified to prevent duplicated memory to be freed.</p> <p>Workaround: Correct the second route.</p>
SBX-113183 SBX-107753	1	<p>PortFix SBX-107753: After the LSWU, the apache2 not coming up, resulting in an EMA or PM access issue.</p> <p>Impact: After an upgrade, the Apache did not start.</p> <p>Root Cause: Apache did not start due to a timeout while requesting a password. The Apache would need to be restarted manually especially in case of a failed upgrade.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Perform an upgrade (LSWU/standalone). 2. If upgrade has failed, verify that the apache has come up properly with default keys and cert. 3. If upgrade has succeeded, verify that first Apache came up properly with default keys, then after app start, the installed keys and certs are being used by the Apache. 	<p>The code is modified to find any occurrence of the one-time-issue of server key getting corrupted in the future.</p> <p>Workaround: In case of a failed upgrade or apache startup failure restart apache or apply the workaround mentioned in an issue relating to the apache2 failing after an LSWU.</p>
SBX-113426 SBX-112447	1	<p>PortFix SBX-112447: The SIPRec recording failed for an EGRESS call with GCID 291515755.</p> <p>Impact: The SIPREC sessions fail with the following "MAJOR" logs when a CS call is going for call a modification with re-INVITE and a corresponding re-INVITE is triggered for this towards the SRS while a previous SRS SIP transaction is pending.</p> <p>159 08192021 154751.374314:1.02.00.26912.MAJOR .SIPSG: sipsgRec.c (-5062) 529606794. [SipSgSendSRSMetadataUpdateCmd] SipSgSendSdpCmd Failed with status 491 185 08192021 154751.375973:1.02.00.26913.Minor .SIPSG: SIPRec Recording failed for EGRESS call with GCID 291518377 for Recorder 10.xx.xxx.xx:xxxx. Trunkgroup of Recorder: NICE_PRI_TG</p> <p>Root Cause: The SBC always attempted to send another offer towards SRS even when a offer-answer was pending and this resulted in SRS session failure.</p> <p>Steps to Replicate: Use the following set up to reproduce the issue.</p> <ol style="list-style-type: none"> 1. Main Call (CS) session established. 2. SIPREC session (RS) is established. 3. CS call goes for modification with re-INVITE (Hold/Codec change). A SIPREC re-INVITE is triggered (based on CS re-INVITE). 4. The SRS does not respond for SIPREC re-INVITE. 5. Before the SRS answers the re-INVITE, a CS call triggers yet another modification with re-INVITE (unhold/codec change). 6. The SBC attempts to trigger SIPREC RE-INVITE again even when the previous transaction is pending and results in SRS session failure. 	<p>When a SIP Offer-Answer towards SRS is in progress, SBC shall not attempt to send another offer and instead shall queue the latest event until the current offer-answer is complete.</p> <p>Workaround: None.</p>
SBX-113196	1	<p>A SCM Process coreump occurred on the SIGABRT.</p> <p>Impact: When the INVITE message contains trunk group/trunk context parameter information and the context is an FQDN, then the SBC can use this information to formulate the IP Peer information if there is none supplied from the PSX. This can result in an SCM process coredump, if multiple routes are returned from the PSX.</p> <p>Root Cause: The internal processing code did not account for multiple routes and it was always freeing memory that was possibly assigned to the first route. However, when processing the second or subsequent routes in a crankback scenario, the memory on the first route freed once when routing the call on the second route and then again when the call was actually released. This second freed memory was invalid and results in a core dump.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Pass in the trunk group and trunk context parameters that use FQDN information. 2. Ensure the PSX returns 3 routes and no IP Peer IP or FQDN information. 3. Return a crankback status code in response to the INVITEs sent on route 1 and route 2. 4. Release the call. 	<p>The code is modified to correctly manage the memory on the route being used for the call. The code change frees and allocates memory consistently for a route to avoid freeing the memory twice.</p> <p>Workaround: Remove the trunk group/context using the SMM from the received INVITE.</p> <p>Ensure the PSX only returns one route when processing trunk group /context information.</p>

The following Severity 2-4 issues are resolved in this release:

Table 14: Severity 2-4 Resolved Issues

Issue ID	Sev	Problem Description	Resolution
----------	-----	---------------------	------------

SBX-112857 SBX-113197	3	<p>PortFix SBX-112857: Increase the VNFM response timeout.</p> <p>Impact: By default, the VNFR logic had a 5 second timeout to get responses from VNFM for the REST requests that it sent, but it has been observed that in certain networks this timeout value is not enough. This leads to the VNFR timing out while the VNFM was still trying to send back a response and the VNFR never reported ready state in the VNFM.</p> <p>Root Cause: A delay of 5 seconds was originally thought to be long enough but across multiple networks, this was insufficient.</p> <p>Steps to Replicate: Try a relocation operation with VNFM across multiple sites and ensure the VNFR status returns to ready.</p>	<p>The code is modified so the timeout is now increased to 20 seconds.</p> <p>Workaround: None.</p>
SBX-112487 SBX-113200	2	<p>PortFix SBX-112487: The LI connection to media server was not connected on a switchover.</p> <p>Impact: When using a PCSILI (P-com-session-info flavour of LI) in an N:1 M-SBC setup and running the SBC release 8.2 or later, if there is a switchover, the connection to the LI server might not be re-established.</p> <p>Root Cause: The standby instance was trying to maintain information about the operational status of the interface used for LI processing. But this was only happening correctly for the first instance. The status of the other instances was being misinterpreted and lost. So during the transition to active the LI code thought the interface was not available and did not try to establish the connection to the LI server.</p> <p>Steps to Replicate: In an N:1 setup perform switchovers from each active instance to the standby and check that the connection to the LI server is restored.</p>	<p>The code is modified to collect the status of the interface after the instance is transitioned from standby to active, so that accurate interface information is available for LI processing.</p> <p>Workaround: Bouncing the interface (e.g. ifup/ifdown that is used for LI connection), will help to recover the connection.</p>
SBX-109006 SBX-110819	2	<p>PortFix SBX-109006: The DSP DHC had a Failure and failed to coredump.</p> <p>Impact: The FPGA based DSP Health Check (DHC) fails and as a result, no DSP coredump is collected.</p> <p>Root Cause: The root cause for DHC failure is not known. However, subsequent coredumps failed due to a lack of reception of the DSP BOOTP packets, which is a hard failure. It is speculated that both issues are related.</p> <p>Steps to Replicate: Instrument the code to replicate the issue.</p>	<p>Reboot the node instead of running an application restart to address the issue.</p> <p>Note: This is not a fix for the original issue, but just a proposed workaround to prevent or reduce further failures.</p> <p>Workaround: None.</p>
SBX-112971 SBX-113352	2	<p>PortFix SBX-112971: The diameter process cored continuously when creating the diamNode entry.</p> <p>Impact: The diameter process acts as a client and server. When it acts as server, on accepting TCP connection from the remote diameter client, the diameter process crashes due to an invalid operation.</p> <p>Note: There is no use case where the diam process accept connection from remote diameter client, so this scenario not covered in our earlier testing. This issue observed during misconfiguration.</p> <p>Root Cause: On accepting TCP connection from the remote diameter client, the diameter process tries to perform an invalid operation (null pointer access) and due to this, the diameter process crashes.</p> <p>Steps to Replicate: Configure the diameter node in the SBC and simulate TCP client application that connects the diameter server (on port 3868).</p>	<p>The code is modified so the diameter process does not perform any invalid operation.</p> <p>Workaround: Do not configure same IP address for the peer and for diameter node.</p>
SBX-113413 SBX-113535	2	<p>PortFix SBX-113413: The call failure with an extra INVITE is sending.</p> <p>Impact: If A calls B over the SBC, the SBC sends INV with sendrecv to B. B sends 180/200 with rcvonly. After call is established, the SBC sends a re-INVITE to B with data path mode as sendonly.</p> <p>Root Cause: As part of a previous issue, a behavior change was introduced so that the SBC sends a hold re-INVITE (received from the network) to the other end. This was only for the cases where call is already on hold but data path mode is changed.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Send an INVITE from A to B. The SBC sends sendrecv to egress. 2. From B, send a 180 with rcvonly. 3. From B, send a 200 Ok with rcvonly. <p>A to B is connected.</p> <p>When call is connected, there should not be an re-INV towards egress leg (with rcvonly).</p>	<p>The code is modified so in the case when the SBC receives a re-INVITE from network and data path mode is changed, it triggers a new re-INVITE to the other leg.</p> <p>In other cases, it would be same behavior (prior to SBX-111611).</p> <p>Workaround: None.</p>

SBX-107396 SBX-112247	2	<p>PortFix SBX-107396: Error "SYS ERR - Error 0x3b Line 666" was being printed frequently.</p> <p>Impact: SYS ERR repeatedly logged in SYS logs.</p> <p>Root Cause: The SYS ERR seems to only occur for audio encoding type 0x3b, which is SPEEX_8. But we don't have enough information to determine the call flow that triggered it.</p> <p>Steps to Replicate: The steps cannot be reproduced.</p>	<p>The code is modified to replace the SYS ERR with a major level debug message with the GCID and audio encoding type to help future debugging.</p> <p>Workaround: None.</p>
SBX-110592 SBX-112526	3	<p>PortFix SBX-110592: There was a misleading TRC message when issuing the callTrace action command start command.</p> <p>Impact: The misleading TRC message when issuing a "callTrace action command start command".</p> <p>Root Cause: There was no alarm/log message indicating a global call trace start action command.</p> <p>Steps to Replicate: Issue the callTrace action command and observe the Trace log.</p> <p>admin@vsbxsus2> request global callTrace action command start</p> <p>Sonus Networks, Inc.00000000016000000000000000000000128V08.02.06A002 00000000000000000000000000000000TRC2021082014593700000000000000 072 08202021 145941.298910:1.01.00.00002.Info .NRM: Call Trace Reset</p>	<p>The code is modified to display correct message indicating call trace being reset.</p> <p>Workaround: None.</p>
SBX-110755 SBX-112791	2	<p>Portfix SBX-110755: The ACL rules configured with ipInterface are disabled on every SBC start/restart in SBC5400/10GB pkt configuration.</p> <p>Impact: The IPACL rules created with ipInterface defined are not getting installed on the SBC 5400 systems when there is no license present, or after a license is added.</p> <p>Root Cause: The NP will not install an IPACL rule with an ipInterface defined if that interface is not UP and active. The SBC will retry to install failed rules when it sees the port come up. In this scenario the timing is off, however, it tries too soon as the physical port is up but the associated ipInterface is not yet up.</p> <p>Steps to Replicate: The test involves a SBC 5400 without a license for its 10G packet port. The IPACL rules with ipInterface defined on that packet port will then fail to install. Once the license is successfully added, the IPACL rules should be installed.</p>	<p>The code is modified to add an additional delay before attempting to reinstall the failed IPACL rules when the port comes up and to look for an additional failure that starts a timer and try again.</p> <p>Workaround: Do not use IPACL rules with ipInterface defined.</p>
SBX-110490 SBX-113016	2	<p>PortFix SBX-110490: The IMS preconditions are failing due to UPDATE message race condition.</p> <p>Impact: An UPDATE from the calling party is queued by the SBC because the SBC is waiting for 200 OK for the previous update sent to called party, and when the 200 OK is received, the SBC did not forward the queued Update.</p> <p>Root Cause: Precondition parameters received as part of UPDATE from ingress endpoint is not updated in the egress CCB.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Configure preconditions to transparent on both legs. 2. Configure transmitpreconditions to supported on both legs. 3. Run the script so that there is difference in precondition parameters in an INVITE and UPDATE messages from ingress. <pre>invite: a=des:qos mandatory local sendrecv a=curr:qos local none a=des:qos optional remote sendrecv a=curr:qos remote none (AS script)183 in Progress: a=curr:qos local sendrecv a=curr:qos remote none a=des:qos optional local sendrecv a=des:qos mandatory remote sendrecv a=conf:qos local sendrecv (IAD script)update : a=curr:qos local sendrecv a=curr:qos remote sendrecv a=des:qos mandatory local sendrecv a=des:qos optional remote sendrecv</pre> <ol style="list-style-type: none"> 4. Send the update from the Ingress when the SBC is waiting for 200 of the previous update from the ingress. <p>Verify that the SBC sends the queued update, once it receives the 200 ok for the 1st update.</p>	<p>The code is modified so the precondition level of support was set to transparent on both legs in configuration.</p> <p>When the SBC receives precondition parameters as part of an UPDATE and preconditions are transparent, then the SBC saves the precondition parameters into the egress CCB. When the queued update message is being processed there is check for preconditions flag and comparison of precondition variables is done to send the update to called party.</p> <p>Workaround: Enable the "Disable media lockdown" flag in IPSP for egress leg so that the SBC does not send the first update.</p>

SBX-109364 SBX-111831	2	<p>PortFix SBX-109364: The SBC ERE was unable to handle "sonusdomain.IP" as a keepalive response from DNS.</p> <p>Impact: The SBC ERE was unable to handle "sonusdomain.IP" as a keepalive response from DNS.</p> <p>Root Cause: A format Error (dns_rcode_formerr) return code of a DNS server was unable to whitelist a server that was block listed.</p> <p>The DNS server is not getting whitelisted from a block list since the Format Error (dns_rcode_formerr) return code is not handled in the code.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Configure the ARS profile. 2. Make a DNS server query with invalid format. 3. Verify the alarm to check that the server is on a block list. 4. Check after sending an ICMP message, the server is on a permit list. 	<p>The code is modified to fix the issue.</p> <p>Workaround: None.</p>
SBX-112429 SBX-112818	2	<p>Portfix SBX-112429: The NrmGetCallCount was returning max global callcount, instead of the current stable calls.</p> <p>Impact: The I-SBC sends the total (cumulative) number of stable calls instead of current stable calls to the SLB in its utilization message. This could result in the SLB dropping messages if it thinks the I-SBC has reached 100% utilization.</p> <p>Root Cause: The I-SBC was providing the wrong number of stable calls to the SLB.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Bring up the SLB-ISBC setup. 2. Run a call load. 3. Check in the SLB about current utilization of the I-SBC. 	<p>The code is modified to provide the current number of stable calls in the utilization message to the SLB.</p> <p>Workaround: None.</p>
SBX-111740 SBX-112570	2	<p>Portfix SBX-11174: The SBC takes 72 seconds to send BYE to transferor after call termination.</p> <p>Impact: The SBC is taking 72 seconds to send BYE to the transferor when the transferee disconnects before a transfer target accepts the call.</p> <p>Root Cause: One of the disconnect events towards transfer target is getting ignored in the call control state machine. As a result, a release timer of 70 seconds is getting triggered later and the SBC is sending BYE towards transferor.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. A and B call is connected. 2. B sends a REFER to the SBC. 3. The SBC sends a new INV to transfer target C. 4. C sends 180 Ringing to the SBC. 5. A sends BYE to the SBC before C answers the call. 	<p>The code is modified so that the SBC sends a disconnect immediately towards the transferor.</p> <p>Workaround: None.</p>
SBX-98627 SBX-113514	2	<p>PortFix SBX-98627: The Min-SE header is getting added to the ingress metaDataProfile even though it is not in the initial INVITE.</p> <p>Impact: For the SIPREC, the Min-SE and Session-Expires headers are added to the metaDataProfile even though it is not received in the initial INVITE.</p> <p>Root Cause: When receiving an INVITE, the SBC was incorrectly inserted Min-SE and Session-Expires headers as part of an incoming INVITE.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Configure the SIPREC set up and Min-SE and Session-Expires in sipRecMetadataProfile. 2. Receive an incoming INVITE without the support timer and no Min-SE, no Session-Expires. 	<p>The code is modified so if an incoming INVITE is missing support for a Timer, the SBC resets the value of the Min-SE and Session-Expires only if it is received.</p> <p>Workaround: Use the SMM to delete the Metadata in the SIPREC INVITE.</p>

<p>SBX-105688 SBX-110152</p>	<p>2</p>	<p>Portfix SBX-105688: The TAP ID of an ingress target was not getting embedded in the CCID for IMSLI (both leg interception).</p> <p>Impact: When a call has matched the ingress and egress LI criteria/target for a SIP Out of Dialog message and if each of the criteria was configure with different TAP IDs over the X1 interface, only one of the TAP IDs was processed by the SBC.</p> <p>This resulted in an incorrect TAP ID being sent out for one of the intercepted leg in the Correlation-ID (CCID) and in the TAP ID AVP field (202) of the X2 messages.</p> <p>Root Cause: The SBC always stored only one TAP ID, the TAP ID present in the 0th index of the LI criteria table that is returned by the PSX and it did not store the TAP IDs corresponding to ingress and egress criteria separately.</p> <p>As a result, whenever both legs are intercepted for OOD messages, one of leg would have the incorrect TAP ID in the CCID for X2 messages.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Configure the SBC with the PSX and EMS for IMSLI Lawful interception. 2. Configure the IMSLI for both Leg and set targets using EMS with different TAP ID. 3. Make a subscribe request with the same targets configured. 4. Verify that TAP ID received for both legs are correct. 	<p>The code is modified to store and process the TAP ID for both the ingress and egress target criteria.</p> <p>Workaround: None.</p>
<p>SBX-109103 SBX-111454</p>	<p>2</p>	<p>PortFix SBX-109103: There was an incorrect UDP port used in the Record-Route and Contact on the core side.</p> <p>Impact: The SBC populates an invalid port in the Contact/Record-Route header towards the IMS Core side in call progress (180/200) responses.</p> <p>Root Cause: The signaling engine in the SBC, during a call, progress command modifies the ports without validating the direction of message flow and causes the SBC to populate an invalid port in the SIP responses.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Send an IMS-AKA registration request from a UE device. 2. Ensure the IMS-AKA registration is successful. 3. Let the registered user receive call from IMS network. 4. When UE accepts the call, we can observe in the call progress response sent from the SBC towards the IMS Core will have the invalid port in contact /Record-Route headers. <p>Flow: IMS Core --(SIP UDP)--> SBC ---(SIP IPsec TCP)--> Peer (UE)</p>	<p>The code is modified to identify the direction before changing ports, change only when the flow is towards the UE.</p> <p>Workaround: Use the SMM to change the ports in Contact and Record-Route headers.</p>
<p>SBX-101255 SBX-112998</p>	<p>2</p>	<p>PortFix SBX-101255: The OAM should not configure the same IP for two different pkt interfaces.</p> <p>Impact: The OAM is allowing the same IP and alternate IP to be configured under the same address context.</p> <p>Root Cause: During an address context, the interface creation OAM was allowing the configure of the same IP address (IpV4/IpV6) and alternate IP address for different interfaces.</p> <p>Steps to Replicate: The meta table is added for ipVarV4/ipVarV6/altIpVars value. When creating the ipInterface table, we should also observe ipInterfaceIpVarVMeta(hidden)</p> <p>The meta table is created for ipVarV4/ipVarV6/altIpVars. For all test cases, the following table can be observed:</p> <p>Case 1: Create ipInterfaceGroup,ipInterface and assign ipVarV4 Create 2nd ipInterfaceGroup,ipInterface and assign same ipVarV4 value, confd should throw an error for unique value in a address context.</p> <p>Case 2: Create ipInterfaceGroup,ipInterface assign ipVarV4 and altIpVars same as the one in ipVarV4 should throw an error for unique value.</p> <p>Case 3: Create ipInterfaceGroup,ipInterface assign ipVarV4 and altIpVars Create 2nd ipInterfaceGroup,ipInterface assign ipVarV4 and same altIpVars value as previous confd should throw error for unique value.</p> <p>Case 4: Create ipInterfaceGroup,ipInterface assign ipVarV4 and altIpVars Delete ipInterfaceGroup,ipInterface ipVarV4 Delete ipInterfaceGroup,ipInterface altIpVars Delete ipInterfaceGroup,ipInterface All combination of delete should work</p> <p>Case 5: The same test to be done for ipVarV6.</p> <p>Any other tests are related to ipVarV4/ipVarV6/altIpVars creation/deletion</p>	<p>The code is modified to have a unique IP address (IpV4/IpV6) and alternate IP address for an address context. New meta data is created for validating uniqueness efficiently.</p> <p>Workaround: Do not configure the ipVarV4/ipVarV6/altIpVars</p>

<p>SBX-111019 SBX-113030</p>	<p>2</p>	<p>PortFix SBX-111019: The SBC should not allow to configure the same IP for two different pkt interfaces.</p> <p>Impact: The SBC should not allow to configure same IP for two different pkt interfaces.</p> <p>Root Cause: After a playback file pushed to the SC nodes, the dummy validate is invoked before creation of the addresscontext table, ipInterfaceIpMetaVar table to be moved out of addresscontext yang.</p> <p>Steps to Replicate: The meta table is added for ipVarV4/ipVarV6/altIpVars value. When creating the ipInterface table, we should also observe ipInterfaceIpVarVMeta(hidden)</p> <p>The meta table is created for ipVarV4/ipVarV6/altIpVars. For all test cases, the following table can be observed:</p> <p>Case 1: Create ipInterfaceGroup,ipInterface and assign ipVarV4 Create 2nd ipInterfaceGroup,ipInterface and assign same ipVarV4 value, confd should throw an error for unique value in a address context.</p> <p>Case 2: Create ipInterfaceGroup,ipInterface assign ipVarV4 and altIpVars same as the one in ipVarV4 should throw an error for unique value.</p> <p>Case 3: Create ipInterfaceGroup,ipInterface assign ipVarV4 and altIpVars Create 2nd ipInterfaceGroup,ipInterface assign ipVarV4 and same altIpVars value as previous confd should through error for unique value.</p> <p>Case 4: Create ipInterfaceGroup,ipInterface assign ipVarV4 and altIpVars Delete ipInterfaceGroup,ipInterface ipVarV4 Delete ipInterfaceGroup,ipInterface altIpVars Delete ipInterfaceGroup,ipInterface All combination of delete should work</p> <p>Case 5: The same test to be done for ipVarV6. Once saveAndActivate is called data should reflect in SBC as well.</p> <p>Case 6: Upgrade testing to be performed like from 9.x to 10.x. case 7: error case for upgrade testing, configure ipVarV4/ipVarV6/altIpVars and use same values for different ipInterface within an address context and post upgrade should get an error log for contacting sonus customer service to make the values unique.</p> <p>Run any tests related to ipVarV4/ipVarV6/altIpVars creation/deletion</p>	<p>The code is modified to move ipInterfaceIpMetaVar out of addresscontext yang.</p> <p>Workaround: None.</p>
<p>SBX-111956 SBX-113152</p>	<p>2</p>	<p>PortFix SBX-111956: An early media case in SIPREC re-INVITE is not triggered with latest values for metaDataSource=fromLatest.</p> <p>Impact: Call modification events like codec change/hold/un-hold on the CS that occur before the RS Session is established is not propagated to the RS Call (SIPREC Session).</p> <p>Root Cause: The SBC does not attempts to trigger SIPREC re-INVITE with updated session characteristics when initial SIPREC INVITE transaction itself is pending with SRS and Main Call received a re-INVITE.</p> <p>Steps to Replicate: Run the following procedure to recreate the scenario:</p> <ul style="list-style-type: none"> • Main Call (CS) session established. • SIPREC INVITE (RS) is sent towards SRS. • SRS does not respond for SIPREC INVITE. (SBC - SRS INVITE transaction is pending). • CS call goes for modification with RE-INVITE. (Hold/Codec change). • SRS answers the initial SIPREC INVITE after Main CS Call RE-INVITE. 	<p>The code is modified to queue the SIPREC RE-INVITE event when an initial SIPREC INVITE transaction is pending.</p> <p>When the initial SIPREC INVITE transaction is completed, the queued SIPREC RE-INVITE is sent towards the SRS (This contains the latest session characteristics).</p> <p>Workaround: None.</p>
<p>SBX-110985 SBX-111275</p>	<p>2</p>	<p>PortFix SBX-110985: If the VM name in the upper right corner of the "Configuration Manager" EMA/EMS SBC Manager is clicked, the browser freezes.</p> <p>Impact: If the VM name in the upper right corner of the "Configuration Manager" EMA/EMS SBC Manager is clicked, the browser freezes.</p> <p>Root Cause: The issue occurred because the peer setup was unreachable and to get the peer system info, the curl command that is executed was taking default timeout to get a connection timeout.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Log into EMS. 2. Launch the SBC Node. 3. Click on Monitoring -> System and Software info tab. 	<p>The code is modified to address the issue.</p> <p>Workaround: Refresh the UI page.</p>

SBX-112082 SBX-112551	2	<p>PortFix SBX-112082: There was a runtime error: member access within null pointer of type 'struct CC_SG_ALERTING_UIIND_MSG_STR in CcSgAlertHndl.</p> <p>Impact: Run a basic G711U pass-thru call. After the call got completed, ASAN runtime error is observed in the system logs indicating that the code is taking the address of a field within a null pointer. This could potentially lead to coredumps if using the SIP recording capabilities other than SIPREC.</p> <p>Root Cause: When a call Progress/Alerting message is received from the network, the code was taking the address of a field within the pointer.</p> <p>Steps to Replicate: Run a basic call.</p>	<p>The code is modified to validate that the pointer is not null taking the address of a field within the pointer.</p> <p>Workaround: None.</p>
SBX-113278 SBX-113572	2	<p>PortFix SBX-113278: Need a configurable item added to indicate a legacy meshed network on GW link to lower MAX_ICM value</p> <p>Impact: A call sent from an SBC running 9.0 or above to an older GSX may cause the GSX to core.</p> <p>Root Cause: In 9.0, the PDUs sent from the originating GW to the destination GW in a GW-GW call has increased in size. When the destination GW is an older GSX, the PDU may be larger than the buffer allocated on the GSX for receiving this PDU.</p> <p>In this case, the GSX does not expect a PDU this large and this will result in the GSX overwriting the allocated buffer causing memory corruption that eventually leads to a core.</p> <p>Steps to Replicate:</p> <p>Send an SBC-GW-GW-GSX call involving text (T140) streams.</p> <p>As a result, the GSX may crash.</p> <p>To test the fix:</p> <ol style="list-style-type: none"> 1. Load the new code. 2. Enable oldGsxFixSupport "set global signaling oldGsxFixSupport" 3. Send an SBC-GW-GW-GSX call involving text (T140) streams (this is just 1 way to cause the large PDU to be sent). <p>The GSX should not crash.</p>	<p>The code is modified so the new configuration parameter is enabled if there are any older GSXs in the network:</p> <pre>set global signaling oldGsxFixSupport</pre> <p>When this parameter is enabled, the SBC reduces the max size of the PDUs that can be sent over a GW-GW connection to older GSXs and older SBCs.</p> <p>Workaround: There is no workaround.</p>
SBX-112708 SBX-113492	3	<p>PortFix SBX-112708: The NOTIFY XML body for 200 OK for INVITE does not contain new values that were updated between 180 and 200.</p> <p>Impact: The SBC was not sending latest received P-Asserted-Identity in the call NOTIFY XML body.</p> <p>Root Cause: The SBC was not considering the latest received P-Asserted-Identity and was always sending the first received message in call NOTIFY XML body.</p> <p>Steps to Replicate: Steps:</p> <ol style="list-style-type: none"> 1. Enable the call Notification feature both legs. 2. Send different PAI headers in the 18x and 200 OK. <p>Expected Result:</p> <p>The SBC sends ringing NOTIFY with the XML body contains PAI element that received in 18x and while sending a connected NOTIFY should consider PAI received in 200 OK.</p>	<p>The code is modified to consider the latest received P-Asserted-Identity identity for populating a call NOTIFY XML body.</p> <p>Workaround: None.</p>
SBX-110997 SBX-111367	2	<p>PortFix SBX-110997: There are media problems due to same SSRC after hold /MOH/resume and a previous fix's flag was not helping.</p> <p>Impact: The SBC is not changing local SSRC during a HOLD/RESUME of a Transcoded call.</p> <p>Root Cause: During a HOLD/RESUME of a Transcoded call, the NP was not updated with the new SSRC generated by the SBC, and this resulting in the old SSRC being used in media packets.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Run a basic Basic transcoded HOLD/RESUME SRTP call with the flag enabled. 2. Ensure that during a HOLD, only data-path-mode is changed to sendonly and recvonly respectively to trigger modify scenario and rest of the media attributes remains same. 3. Validate the media packets for change in the SSRC. 	<p>The code is modified to update NP with the latest SSRC generated during mid-call modification due to HOLD/RESUME of Transcoded calls.</p> <p>Workaround: If possible, switch to pass-through calls since this issue is specific to transcoded calls only.</p>

<p>SBX-113048 SBX-113389</p>	<p>2</p>	<p>PortFix SBX-113048: Edited the External IP Interface group name disappearing in Visual First Call View as a diagramSetup. page</p> <p>Impact: Edited the External IP Interface group name disappearing in Visual First Call View as a diagramSetup page.</p> <p>Root Cause: The existing system did not handle save functionality when an external "IP Interface Group" was not empty.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Login to EMA as an admin user. 2. Go to Configuration->Configuration Wizards->Visual First Call Setup. 3. Configure the fields in Carrier and save. 4. Make changes to IP Interface Group and Save. 5. Navigate to All Admin and then navigate back to Configuration->Configuration Wizards->Visual First Call Setup <p>After navigating back to Visual First Call Setup, the changes were saved i.e. Carrier type and IP Interface Group should be visible.</p>	<p>The code is modified for the External "IP Interface Group" for both empty and non-empty cases.</p> <p>Workaround: None.</p>
<p>SBX-111721 SBX-111722</p>	<p>2</p>	<p>PortFix SBX-111721: The EVS encoder picks up the initialCodecMode as the bitrate after switch to Compact Format.</p> <p>Impact: The EVS encoder picks up the initialCodecMode as the current bitrate if a change in packet format from Header Full to Compact is triggered.</p> <p>Root Cause: The root cause is the re-initialization of EVS encoder with bitrate as the initialCodecMode on change in packet format.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Run a EVS<=>AMRWB call. br=5,9-24.4; bw=nb-wb 2. Stream a WB pcap with 13.2Kbps packets with CMR byte for 13.2 CA mode for the first 10s. After, the pcap has 24.4 Kbps packets without the CMR. <p>Results: The mode of operation changes from 24.4Kbps to 13.2 Kbps with CA enabled on receiving the CMR. Once the Endpoint starts sending 24.4 Kbps packets without CMR, the SBC switches to Compact Mode while maintaining its bitrate as 13.2 Kbps with CA enabled.</p> <p>Prior to the fix, when switching to the Compact mode, the SBC would use 24.4 Kbps (initialCodecMode) as the bitrate.</p>	<p>The code is modified to the re-initialize the EVS encoder with bitrate as the localCodecMode rather than the initialCodecMode.</p> <p>Workaround: None.</p>
<p>SBX-111896 SBX-112373</p>	<p>2</p>	<p>PortFix SBX-111896: Automatic daily updates in HFE script for Azure/GCE were causing the network to reset.</p> <p>Impact: Networking on the HFE can be reset by automatic updates.</p> <p>Root Cause: Ubuntu, by default, has an automatic package updaters. Some package upgrades can cause the networking to be reset on the HFE node.</p> <p>Steps to Replicate: Check if the following timers are enabled:</p> <ul style="list-style-type: none"> • sudo systemctl status apt-daily.timer • sudo systemctl status apt-daily-upgrade.timer 	<p>The code is modified to disable the timer that triggers the automatic updates. Package upgrades should only occur out of hours.</p> <p>Workaround: Run the following commands to disable the updates:</p> <ul style="list-style-type: none"> • sudo systemctl stop apt-daily.service • sudo systemctl stop apt-daily.timer • sudo systemctl stop apt-daily-upgrade.timer • sudo systemctl stop apt-daily-upgrade.service • sudo systemctl disable apt-daily.service • sudo systemctl disable apt-daily.timer • sudo systemctl disable apt-daily-upgrade.timer • sudo systemctl disable apt-daily-upgrade.service
<p>SBX-110291 SBX-112568</p>	<p>2</p>	<p>Portfix SBX-110291: AddressSanitizer: The ASAN detected a heap-buffer-overflow-SipSgIncomingCallNfy (unsigned int, sip_addr_str*, sip_msgbody_str*, sip_options_str*, unsigned int, unsigned int, bool) /sonus/p4/ws/jenkinsbuild/sbxAsan100/marlin/.</p> <p>Impact: The ASAN detected "AddressSanitizer: heap-buffer-overflow" while copying the contents of the SIP-I base version string internally.</p> <p>Root Cause: The SIP code was always copying a fixed amount of memory when reading the SIP-I base and version strings to internal memory.</p> <p>Steps to Replicate: Run a SIP-I call flow with INVITE and CANCEL messages.</p>	<p>The code is modified to copy exactly string length size of the SIP-I base and version content to avoid reading past the end of memory buffers as this can cause core dumps.</p> <p>Workaround: None</p>

SBX-105367	2	<p>During a sRTP and fax call, the sRTP context is removed for G711 fax pass-thru.</p> <p>Impact: The sRTP context is removed on the leg that has a G711 fax when there is a t38 to g711 fax call.</p> <p>Root Cause: The sRTP context is not updated for G711-fax.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Run an A(SRTP)-B(RTP) CALL. 2. Run a leg A is G711 fax and B is T38. B sends a re-INVITE with t38 fax. 3. Check the re-INVITE that goes out with SRTP to endpoint A. 4. A party sends 200 ok with SRTP. 5. End the call. 	<p>The code is modified so the the SRTP context is updated only if the calleg has G711 fax.</p> <p>Workaround: None.</p>
SBX-112905	2	<p>The SBC was answering on-hold call with the sendrcv then re-inviting.</p> <p>Impact: On an incoming INVITE onhold, the SBC responds to the 18x offhold and later the re-INVITE onhold.</p> <p>Root Cause: There was a logical error that converts from an inactive to sendrcv when the first 18x is sent out. This issue was introduced a previous fix listed in the 9.2.1 release.</p> <p>Steps to Replicate:</p> <p>On the ingress, configure the Minimize flag, and uncheck relay data path mode.</p> <p>Incoming Invite onhold, egress response 18x.</p> <p>The SBC send 18x to ingress with sendrcv.</p>	<p>The code is modified to apply for subsequent 18x only.</p> <p>Workaround: Enable the relay data path mode.</p>
SBX-110780	2	<p>The SBC was sending Empty Packets when playing an Announcement.</p> <p>Impact: Announcement packets are empty for two stage calls if the NAPT enabled on the ingress.</p> <p>Root Cause: The ARM was not getting indication to start the announcement once NAPT learning was complete.</p> <p>Steps to Replicate: Set up a two stage call and enable the NAT on ingress.</p>	<p>The code is modified to send an indication to start the announcement once NAPT learning was complete</p> <p>Workaround: Disable the NAT.</p>
SBX-111349	2	<p>An SBC 7000 dual crash within 1 minute (old Active side - Fm, new Active - Scm).</p> <p>Impact: The application on both active and standby switched over and the application reset in a short order of each other. Core files are created on both sides of the HA pair.</p> <p>Root Cause: The issue occurred during multiparty call processing where the SBC tries to determine a whether message was sent for the ingress or egress call segment. The code was accessing the pointer to the multi party call resource after it was freed up and set to NULL.</p> <p>Steps to Replicate: This issue is not reproducible. Potentially due to a race condition in the code.</p>	<p>The code is modified so the multi party call pointer is not NULL before reading from it to avoid the core dump.</p> <p>Workaround: There is no known workaround for this issue.</p>
SBX-112953	2	<p>The SRTP license counter did not incremented when the called side omits an SIP 18x response.</p> <p>Impact: The license count for the SRTP license was not updated on a call from SRTP to RTP, if no 18x response message is received.</p> <p>Root Cause: The SRTP license usage is not updated at the ingress when response to an INVITE is only 200 OK message. Therefore, if the egress is not using a SRTP, the call does not consume a license.</p> <p>Steps to Replicate: Make an SIP-SIP call where the ingress leg uses SRTP and egress leg uses RTP. The called side does not sent any 18x response.</p>	<p>The code is modified to correctly update the license count at ingress, even when no 18x is sent.</p> <p>Workaround: None.</p>

SBX-110359	2	<p>No INFO Logging Warning is provided when enabling the info debug/sys.</p> <p>Impact: The SBC CLI/EMA does not provide any warning when enabling INFO logging, indicating impact to system performance and call processing.</p> <p>Root Cause: There was no code to provide the warning to user.</p> <p>Steps to Replicate: Run the following script to reproduce the issue:</p> <pre>admin@sbxsus12% set system admin sbxsus12 cliSetWarningOnEnablingInfoLevelLogging [disabled,enabled] (disabled): enabled [ok][2021-08-19 17:12:13] [edit] admin@sbxsus12% com Commit complete. admin@sbxsus12% set oam eventLog typeAdmin debug filterLevel info [ok][2021-08-19 17:14:51] [edit] admin@sbxsus12% com The following warnings were generated: 'oam eventLog typeAdmin': Enabling INFO level logging can be service impacting. Do you want to continue? Proceed? [yes,no] no Aborted: by user [ok][2021-08-19 17:15:41] [edit] admin@sbxsus12% com The following warnings were generated: 'oam eventLog typeAdmin': Enabling INFO level logging can be service impacting. Do you want to continue? Proceed? [yes,no] yes Commit complete. [ok][2021-08-19 17:15:46] admin@sbxsus12% set oam eventLog typeAdmin system filterLevel info [ok][2021-08-19 17:16:37] [edit] admin@sbxsus12% com The following warnings were generated: 'oam eventLog typeAdmin': Enabling INFO level logging can be service impacting. Do you want to continue? Proceed? [yes,no] yes Commit complete. [ok][2021-08-19 17:16:40]</pre>	<p>The code is modified to provide a warning to the user when enabling INFO level logging for system or debug logs. To maintain backward compatibility and not break any customer scripts, a user needs to enable this extra prompt.</p> <p>Workaround: None.</p>
SBX-112060	2	<p>STIR/SHAKEN services were failing in the second dip in 2-stage call (SIP-SIP).</p> <p>Impact: The SBC was not sending STIR/SHAKEN parameters, such as identity headers, to the PSX in the trigger request portion of a two-stage call or processing the STIR/SHAKEN parameters in the response.</p> <p>Root Cause: The SBC was missing code to handle the interaction between two-stage calls and STIR/SHAKEN logic.</p> <p>Steps to Replicate: Make a two-stage call where the INVITE contains identity headers and check they are sent to the PSX in both the policy request and trigger request.</p>	<p>The code is modified to send STIR/SHAKEN parameters such as identity headers to the PSX in the trigger request portion of a two-stage call and process the STIR/SHAKEN parameters in the response.</p> <p>Workaround: None.</p>
SBX-110746	2	<p>The SIP ACK was missing in TRC, although the SBC sent one.</p> <p>Impact: The 200 OK for the INVITE received on the egress call leg contains a Record-Route header with an FQDN. The SBC resolves the Record-Route FQDN through the DNS and routes the ACK to the resolved target. The SIP ACK PDU that the SBC sent to the egress call leg is missing in the TRC log.</p> <p>Root Cause: The root cause is in SipsDnsLookupRspForTCB() when the DNS response comes back. SIP_CALL_STR *pstCall is set as NULL when SipsTXNDownstreamMsgToFsm is called. Due to the NULL check of pstCall in later functions, the SipSgTraceDumpSipPdu() is not called to log ACK PDU in TRC log.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Set up: SIP-SBX-SIP. 2. Create a dnsgroup and configure the external DNS IP Address. Attach dnsgroup with egress zone. 3. Create a global calltrace filter with "level1" and key "calling". 4. Start the global calltrace and roll trace log event. 5. Run SIP-SIP call. 200 OK for the INVITE received on the egress call leg contains Record-Route with an FQDN. 6. Verify that egress ACK PDU is logged in TRC log. 	<p>The code is modified so in the function SipsDnsLookupRspForTCB, SIP_CALL_STR* pstCall is fetched and use in SipsTXNDownstreamMsgToFsm.</p> <p>Workaround: None.</p>

SBX-109614	2	<p>There were Error logs in DBG file after upgrade on 9.2.1</p> <p>Impact: The SIPFE was reporting two kinds of major level error messages in the DBG logs:</p> <ol style="list-style-type: none"> port range id not found from SipFePortRangeCnxSendDataReq() e.g. SipFePortRangeCnxSendDataReq 366] PORT_RANGE: id not found 412765 port range id not found from SipFeRedundMirrorHashSync() e.g. SipFeRedundMirrorHashSync 1372] PORT_RANGE: failed to find Connection id in HAsh[4194719] <p>Root Cause: Port Range Support was added in a previous feature that involves four subsystems, SIPSG, SIPFE, SIPCM and XRM, among 3 processes. There were several message exchanges between the four subsystems before a port range connection is fully established.</p> <p>Under certain circumstances, it is possible to hit some race conditions and find major DBG messages in DBG logs.</p> <p>Steps to Replicate: The steps cannot be replicated.</p>	<p>The code is modified to print extra call related data in the current major level debug messages to help triage more if the problem occurs in the future.</p> <p>When the issue occurs again, the call ID, signaling port Id, etc., is collected to help identify the call. The SipFeRedundMirrorHashSync() log is reduced from MAJOR level as well.</p> <p>Workaround: N/A</p>
SBX-111541	3	<p>Reloading previous configuration fails when the LDAP delayedSync is enabled.</p> <p>Impact: The delayed sync leaf used to accept only a future time if a configuration export and import is performed and if the system time is greater than the delayed sync time during import than import fails.</p> <p>Root Cause: The delayed sync leaf used to accept only a future time if a configuration export and import is performed and if the system time is greater than the delayed sync time during import then import fails. It fails due to the check on the delayed sync leaf value that should be greater than the current system time.</p> <p>Steps to Replicate: Use the following configuration:</p> <ol style="list-style-type: none"> Add a profile with a future time value on the delayed sync leaf. Export the configuration. Clear the db. Import the configuration when the system time is greater than the configured delayed sync time. Import would fail without this fix. 	<p>Remove the restriction on the delayed sync leaf to be greater than the current system time to address the issue.</p> <p>Workaround: Modify the delayed sync time in the export CLI.</p>
SBX-112267	2	<p>Correct the 92x serialization code for a registration control block.</p> <p>Impact: When the registration control block is made redundant and the active /standby instances are running different software releases e.g. during upgrade, the redundancy logic might not work correctly if the registration control block contains a P-Charging-Vector (PCV) and/or P-Charging-Function-Addresses (PCFA) information. The redundancy logic will work correctly post upgrade when serialization of the redundancy data is not required.</p> <p>Root Cause: The parameter lengths of the PCV and PCFA redundant parameters was not calculated correctly. This can lead to the redundancy code being unable to process all the redundancy information correctly.</p> <p>Steps to Replicate: Run registration related tests with PCV and PCFA header parameters in an older release and then upgrade to 92R2 or later.</p>	<p>The code is modified to correctly set the length of these parameters to avoid problems with redundancy.</p> <p>Workaround: None.</p>
SBX-112493	2	<p>The SBC is not decrypting the IPsec packet when a large PDU was sent over IPv6 from Strongswan.</p> <p>Impact: The SBC is not decrypting the IPsec packet when a large PDU was sent from Strongswan IPsec endpoint.</p> <p>Root Cause: The SBC SWe NP has an issue in last fragment data length processing w.r.t. This StrongSwan fragmented first and encrypted the next combination IPsec packets handling, which caused an ESP trailer offset being incorrect and resulted in call failures.</p> <p>Steps to Replicate: Make large SIP PDU calls with the StrongSwan IPsec endpoint.</p>	<p>The code is modified to work for all combinations.</p> <p>Workaround: Racoon/Navtel/SBC IPsec endpoints can be used if possible as a workaround.</p>
SBX-113546	2	<p>Policy server transactions were failing while running a load (3xx redirection with light dip) with 75K Number Translation criteria.</p> <p>Impact: The postgres db is utilizing more CPU when a dmpm translation is executed at the service layer and as a result, impacting the performance.</p> <p>Root Cause: Using a dmpm rule, either the calling or the called number can be modified at the prerouter layer. The PES previously analyzed the called and calling number completely even though it was exclusive to the called number or the calling. During an analysis of the call, the PES used to perform direct DB fetch operations and this caused the postgres process to use more CPU.</p> <p>Steps to Replicate: Create a dmpm rule service to modify the calling number. And then perform performance testing. The postgres process will utilize more CPU.</p>	<p>The code is modified so that, the PES analyzes either only the called number or only the calling number based on the rule type.</p> <p>Workaround: none.</p>

SBX-113610	2	<p>A PES Process core dumps when deleting an adProfile entry.</p> <p>Impact: The PES Process dumps core when the AD profile is deleted.</p> <p>Root Cause: While synchronizing the data from remote domain controller, the PES fetches the AD profile once from the cache and uses the object for various decision making. So, if the AD profile is deleted, the object reference becomes NULL, and it dumps core.</p> <p>Steps to Replicate: To create an AD profile, create its dependent profiles.</p> <ol style="list-style-type: none"> 1. Create An AD Attribute Profile. set profiles adAttributeMapProfile DEFAULT_AD_ATTRIBUTE_PROFILE adAttributeList adAttribute1 adAttributeName cn set profiles adAttributeMapProfile DEFAULT_AD_ATTRIBUTE_PROFILE adAttributeList adAttribute2 adAttributeName displayName commit 2. Create a domain controller profile set global servers domainController dc0 userName DcUsername password DcPassword primaryAddress DcIpAddress searchScope base=engineering, dc=some,dc=com ldapQueryCriteria cn= commit 3. Create an AD profile: set profiles adProfile DEFAULT_AD_PROFILE delayedSync 2021-09-29T14:00:00 syncInterval 1440 sync enable adServerList 1 dcServer dc0 commit <p>After 1-2 minutes of creating the AD profile, delete it.</p>	<p>The code is modified so the AD profile is not deleted. If, for some reason, the synchronization needs to be stopped, the sync option can be disabled on the AD profile.</p> <p>Workaround: If the AD profile is no more required, do not delete the AD Profile but disable the sync option on it. It will stop all the future auto sync operations.</p>
SBX-112603	2	<p>SBX-107111: Configured the DM/PM Rule not getting executed when the system is loaded with 75,000 number Translation criteria.</p> <p>Impact: When there are more than 3,000 number translation criteria (NTC) objects are created, the PES cache goes for rebuild of the entire cache. The cache rebuild was not caching the NTC objects correctly.</p> <p>Root Cause: During a rebuild of the NTC cache, a static variable was not reinitialized to zero. Since the static variable was not initialized to zero during a cache rebuild, it was using the last modified value, and this was causing the cache to not rebuild with the all the NTC objects.</p> <p>Steps to Replicate: Create more than 3,000 NTC objects and make a call to match the NTC key. It will fail to find a match in the cache.</p>	<p>The code is modified to initialize the static variable to zero, when a fresh cache rebuild is triggered.</p> <p>Workaround: None.</p>
SBX-112670	2	<p>The SBC is sending an unexpected INVITE towards the UAS.</p> <p>Impact: In an SBX-GW-GW-SBX call flow (as described in the test setup), the SBC was sending unexpected re-INVITE to the network. The re-INVITE was internally generated in order to handle the change in SDP parameters triggered by the SMM rules.</p> <p>Root Cause: This was a side effect of for a previous fix in the 9.2.1 release.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Create a GW-GW set up. 2. From UAS send 183 with a=sendonly along with image and text lines with sendrecv. 3. SBC sends 183 with audio line as a=sendonly (That is removed by the SMM). 4. Send a 180 from the UAS that the SBC sends to UAC. 5. Send a 200 from UAS without the sdp. 6. The SBC sends a re-INVITE to ingress with audio m line sendonly along with image and text port as 0 (SMM changes the audio line to sendrecv). 7. The SBC gets 200 for re-INVITE from UAC with audio line sendrecv. <p>After step 7, we should not see a re-INVITE towards the UAS.</p>	<p>The code is modified to take care of the case when holding a re-INVITE is actually received from a network or if it was locally generated. Only in a case where it was from network, we propagate the re-INVITE to the other end.</p> <p>Workaround: None.</p>

SBX-113385	2	<p>An unexpected 504 response was received, after a crankback to "500".</p> <p>Impact: When the SIP trunk group control UseNonDefaultCauseCodeForARSBlacklist is enabled and the outgoing message cannot be sent due to address unreachable status, the crankback to the next end point was not occurring.</p> <p>Root Cause: The code was incorrectly updated during 9.2.2 to try and map from SIP to CPC using internal 7xx status codes which are generated by the stack for events such as timeouts. This led to the disconnect reason code being mapped to 0 that did not exist in the crankback profile instead of 168 which was expected. So the crankback did not occur and the call was released with 504.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Run a test case with multiple trunk groups. 2. Send in a subscribe message that is routed through the SBC and respond back with 500. <p>The message should be attempted to the next end point, but if that end point is blocklisted due to pathcheck configuration the message should be sent to the next again end point.</p>	<p>The code is modified to only consider the standard SIP status codes when trying to map from SIP to the CPC so that the correct disconnect reason code of 168 is used to trigger a crankback.</p> <p>Workaround: None.</p>
SBX-113766	2	<p>The SCM process crash was observed when trying to show the ARS blocklisted entries.</p> <p>Impact: When too many peers are in ARS blocklisted table(more than 90 entries) and the "show status addressContext default zone <ZONE_NAME> sipArsStatus" command was issued, the SCM process crashed.</p> <p>Root Cause: The code was incorrectly indexing to a negative location in an array that led to an invalid memory access and a core dump.</p> <p>Steps to Replicate: Block list multiple peers and then run the status command. This was only seen once in lab testing.</p>	<p>The code is modified to avoid accessing the negative array location to avoid the core dump.</p> <p>Workaround: None.</p>
SBX-90989	2	<p>EMA CDR Viewer: SIP Ladder diagram tool presenting incomplete packets</p> <p>The SBC CDR Viewer is enhanced to overcome SIP Ladder diagram size limitations to ensure large messages display.</p> <p>Impact: Incomplete packets display in the SIP Ladder diagram tool because the EMA did not know to append the next block to the previous PDU in the tool, and dropped the secondary packet block.</p> <p>Root Cause: Due to size limitations, large packets are getting split into two blocks, causing the second block to get prepended with information about the filter when it is written to the TRC log.</p> <p>Steps to Replicate: Use call trace filter to capture some large messages i.e. more than 2K in size and check that they are displayed correctly in the SIP ladder diagram.</p>	<p>The code is modified to ensure large messages display in the SIP Ladder diagram.</p> <p>Workaround: None.</p>

Resolved Issues in 09.02.02R006 and 09.02.02R005 Releases

The following issue is resolved in these releases:

Table 15: Resolved Issues

Issue ID	Sev	Problem Description	Resolution
----------	-----	---------------------	------------

SBX-113278	2	<p>The SBC is enhanced with a configurable parameter to indicate a legacy meshed network exists on the GW-GW links. This will ensure the MAX_LCM value used to pack a message to send to older GSXs is the same size.</p> <p>Impact: A call sent from an SBC running software between 9.1.0 and 9.2.2R4 to any GSX (without the fix GSX-61814) may cause the GSX to core. The max buffer to send to the GSX was changed from 15K bytes to 10K bytes.</p> <p>Root Cause: PDUs sent from the originating SBC to the destination GSX in a GW-GW network are bigger than the buffer on the remote GSX.</p> <p>When the destination GW is a GSX, the PDU may be larger than the buffer allocated on the GSX. When the GSX encounters oversized PDUs, it overwrites the allocated buffer, causing memory corruption and ultimately a coredump may occur.</p> <p>Steps to Replicate:</p> <p>Send an SBC-GW-GW-GSX call involving text (T140) streams. The GSX may crash.</p> <p>To test the fix:</p> <ol style="list-style-type: none"> 1. Load the new code. 2. Enable the oldGsxSupport flag: set global signaling oldGsxSupport enable 3. Send an SBC-GW-GW-GSX call involving text (T140) streams (this is just one way to send a large PDU). <p>The GSX should not crash.</p>	<p>The code is modified to add a new global signaling parameter that you must enable if there are any GSXs in the network without the fix for GSX-61814.</p> <pre>set global signaling oldGsxSupport <disable enable></pre> <p>When this parameter is enabled, the SBC reduces the maximum size of the PDUs that can be sent over a GW-GW connection to GSXs to prevent sending a PDU larger than the buffer size allocated on the GSX, which can cause memory corruption. Once the GSX has fix for GSX-61814, it will send its maximum buffer size supported in the GW OPEN ACK (at which time you can disable this flag).</p> <p>Any SBC running 9.1 or higher software supports this new AVP, and subsequently knows what to send to another SBC over the GW link. Calls will fail before getting cleared by the GSX with a 503, and will now fail on the ingress side with a 500 and advance to the next route.</p> <table border="1" data-bbox="699 491 1284 747"> <thead> <tr> <th data-bbox="699 491 777 562">New Flag</th> <th data-bbox="777 491 1284 562">Description</th> </tr> </thead> <tbody> <tr> <td data-bbox="699 562 777 747">oldGsxSupport</td> <td data-bbox="777 562 1284 747"> <p>When the flag is enabled, the SBC reduces the maximum size of the PDUs sent over a GW-GW connection to a GSX to prevent the PDU sizes from becoming larger than the buffer size allocated on the GSX, which can cause memory corruption.</p> <ul style="list-style-type: none"> • disabled (default) • enabled </td> </tr> </tbody> </table> <p>Workaround: There is no known workaround for this issue.</p>	New Flag	Description	oldGsxSupport	<p>When the flag is enabled, the SBC reduces the maximum size of the PDUs sent over a GW-GW connection to a GSX to prevent the PDU sizes from becoming larger than the buffer size allocated on the GSX, which can cause memory corruption.</p> <ul style="list-style-type: none"> • disabled (default) • enabled
New Flag	Description						
oldGsxSupport	<p>When the flag is enabled, the SBC reduces the maximum size of the PDUs sent over a GW-GW connection to a GSX to prevent the PDU sizes from becoming larger than the buffer size allocated on the GSX, which can cause memory corruption.</p> <ul style="list-style-type: none"> • disabled (default) • enabled 						

Resolved Issues in 09.02.02R004 Release

The following Severity 1 issue is resolved in this release:

Table 16: Severity 1 Resolved Issues

Issue ID	Sev	Problem Description	Resolution
SBX-107747 SBX-112648	1	<p>PortFix SBX-107747: During Cyclic switchover tests, a PRS Process coredump was observed on the MSBC1.</p> <p>Impact: The application on both active and standby switched over and application reset in short order of each other. Core files were created on both sides of the HA pair.</p> <p>Root Cause: The issue occurs during multiparty call processing, when the SBC tries to determine whether a message was destined for an ingress or egress call segment. As a result of running multiparty call processing, the code was accessing the pointer to the multi party call resource after it was freed up and set to NULL.</p> <p>Steps to Replicate: This issue is not reproducible. Potentially due to a race condition in the code.</p>	<p>Check the multi party call pointer is not NULL before reading from it to avoid the coredump.</p> <p>Workaround: There is no known workaround for this issue.</p>

SBX-112561 SBX-112836	1	<p>PortFix SBX-112561: The regexp string "\r\n" was not exported by "user-config-export".</p> <p>Impact: The \r\n in regex is lost while exporting a configuration using the user-config-export CLI command.</p> <p>Root Cause: XML formatting trims empty node values.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Create an SMM Profile using the following commands: set profiles signaling sipAdaptorProfile ESMM state disabled set profiles signaling sipAdaptorProfile ESMM advancedSMM disabled set profiles signaling sipAdaptorProfile ESMM profileType messageManipulation set profiles signaling sipAdaptorProfile ESMM rule 1 applyMatchHeader one set profiles signaling sipAdaptorProfile ESMM rule 1 criterion 1 type message set profiles signaling sipAdaptorProfile ESMM rule 1 criterion 1 message set profiles signaling sipAdaptorProfile ESMM rule 1 criterion 1 message messageTypes requestAll set profiles signaling sipAdaptorProfile ESMM rule 1 criterion 2 type header set profiles signaling sipAdaptorProfile ESMM rule 1 criterion 2 header set profiles signaling sipAdaptorProfile ESMM rule 1 criterion 2 header name WWW-Authenticate set profiles signaling sipAdaptorProfile ESMM rule 1 criterion 2 header condition exist set profiles signaling sipAdaptorProfile ESMM rule 1 criterion 2 header numberOfInstances number 1 set profiles signaling sipAdaptorProfile ESMM rule 1 criterion 2 header numberOfInstances qualifier equal set profiles signaling sipAdaptorProfile ESMM rule 1 criterion 3 type parameter set profiles signaling sipAdaptorProfile ESMM rule 1 criterion 3 parameter set profiles signaling sipAdaptorProfile ESMM rule 1 criterion 3 parameter condition exist set profiles signaling sipAdaptorProfile ESMM rule 1 criterion 3 parameter paramType generic set profiles signaling sipAdaptorProfile ESMM rule 1 criterion 3 parameter name realm set profiles signaling sipAdaptorProfile ESMM rule 1 action 1 type header set profiles signaling sipAdaptorProfile ESMM rule 1 action 1 operation regsub set profiles signaling sipAdaptorProfile ESMM rule 1 action 1 headerInfo headerValue set profiles signaling sipAdaptorProfile ESMM rule 1 action 1 from set profiles signaling sipAdaptorProfile ESMM rule 1 action 1 from type value set profiles signaling sipAdaptorProfile ESMM rule 1 action 1 from value 172.12.34.567 set profiles signaling sipAdaptorProfile ESMM rule 1 action 1 to set profiles signaling sipAdaptorProfile ESMM rule 1 action 1 to type header set profiles signaling sipAdaptorProfile ESMM rule 1 action 1 to value WWW-Authenticate set profiles signaling sipAdaptorProfile ESMM rule 1 action 1 regexp set profiles signaling sipAdaptorProfile ESMM rule 1 action 1 regexp string "\r\n" set profiles signaling sipAdaptorProfile ESMM rule 1 action 1 regexp matchInstance all commit 2. Run the CLI command to the export configuration: user-config-export test.xml /profiles/signaling/sipAdaptorProfile 3. Delete the exported profile from CLI with the command: delete profiles signaling sipAdaptorProfile ESMM commit user-config-import test.xml 4. Run the following command: show configuration profiles signaling sipAdaptorProfile display set match string 5. Check if the regex string field, restores the original value. 	<p>Use the external XML tool called xmllint.</p> <p>Workaround: Manually edit the field in xml file after export.</p>
SBX-112766	1	<p>Details on "LAST RESTART REASON" under command "show table system serverStatus" are incorrect.</p> <p>Impact: The LAST RESTART REASON is not updated with a correct reason of last restart.</p> <p>Root Cause: Reading LAST TESTART REASON function is called twice and in second call, it is set to default.</p> <p>Steps to Replicate: Check the LAST RESTART REASON from CLI show table system serverStatus and call a different type of restart and verify LAST RESTART REASON set properly or not.</p>	<p>The code is modified to set the reading LAST TESTART REASON function to call only once.</p> <p>Workaround: None.</p>

The following Severity 2-4 issues are resolved in this release:

Table 17: Severity 2-4 Resolved Issues

Issue ID	Sev	Problem Description	Resolution
SBX-111349 SBX-112883	2	<p>PortFix SBX-111349: The SBC 7000 has a dual crash within 1 minute (old Active side - Fm, new Active - Scm).</p> <p>Impact: The application on both active and standby are switched over and application reset in short order of each other. Core files are created on both sides of the HA pair.</p> <p>Root Cause: Accessing a NULL pointer that leads to a coredump.</p> <p>Steps to Replicate: Test switchover scenarios to replicate the issue.</p>	<p>The code is modified to address the issue. The ccbPtr->ccMultiCallMsgStrPtr is checked for not being NULL.</p> <p>Workaround: None.</p>

SBX-112857	3	<p>The VNFM response is timing out due to the timeout value length.</p> <p>Impact: By default the VNFR logic had a 5-second timeout to get responses from VNFM for the REST requests that it sent, but it has been observed that in certain networks this timeout value is not enough. This lead to the VNFR timing out while the VNFM was still trying to send back a response and the VNFR never reported ready state in the VNFM.</p> <p>Root Cause: A 5-second delay was originally considered adequate, but after taking multiple networks into account, the delay proved to be insufficient.</p> <p>Steps to Replicate: Try a relocation operation with VNFM across multiple sites and ensure the VNFR status returns to ready.</p>	<p>The code is modified so the timeout is increased to 20 seconds.</p> <p>Workaround: None.</p>
SBX-112493 SBX-112785	2	<p>PortFix SBX-112493: The SBC is not decrypting IPsec packet when large PDU sent over IPv6 from Strongswan.</p> <p>Impact: The SBC is not decrypting IPsec packet when large PDU sent from Strongswan IPsec endpoint.</p> <p>Root Cause: The SBC SWe NP has an issue in the last fragment data length processing. This StrongSwan was fragmented first and encrypted the next combination IPsec packets handling, which caused ESP trailer offset being wrong and call failures.</p> <p>Steps to Replicate: Make large SIP PDU calls with the StrongSwan IPsec endpoint.</p>	<p>The code is modified on the last segment length to work for all combinations.</p> <p>Workaround: The customer and the SBC IPsec endpoints can be used if possible.</p>
SBX-112153 SBX-112627	4	<p>PortFix SBX-112153: An investigation found commits are not saved/activated on many SBCs.</p> <p>Impact: Warnings displayed about inactive configurations on the managed VM.</p> <p>Root Cause: The managed VM was checking the status on the OAM shared drive and reporting the inactive configuration warnings.</p> <p>Steps to Replicate: Test on OAM and manage the VM.</p>	<p>The code is modified to only display the warning on OAM nodes.</p> <p>Workaround: None.</p>
SBX-112117	2	<p>There is a Wall LRA ISBC switchover after the applying ACLs.</p> <p>Impact: Configuring certain combination of ACLs may exceed the hugepage requirement than available in the system. This will cause the ACL configuration to fail and SWe_NP to exit.</p> <p>Root Cause: The code did not account for available hugepages during an ACL configuration.</p> <p>Steps to Replicate: Bring up setup in ISBC profile. Apply the ACL configuration used by customer.</p>	<p>The code is modified to limit the hugepage memory use during an ACL build and ensure it is within the allocated limit.</p> <p>Workaround: Use more than 64G RAM.</p>

Resolved Issues in 09.02.02R003 Release

The following issue is resolved in this release:

Table 18: Resolved Issues

Issue ID	Sev	Problem Description	Resolution
SBX-105370	1	<p>The Active Register per TG shows more than the total stable registration.</p> <p>Impact: Under some condition(s), the ingress zone's activeRegs count can be incremented and not decremented when the registration terminates.</p> <p>The causes the ingress zone's activeRegs count to grow incorrectly, and never reach ZERO (even after all registrations are terminated).</p> <p>Root Cause: The SIPFE and SIPSG RCB allocation/deallocation become out of sync, indirectly causing the ingress zone's activeRegs count to increment and not decrement.</p> <p>Steps to Replicate: Perform REGISTER/401 - REGISTER/403 until the ingress zone activeRegs count issue is detected.</p>	<p>The code is modified to correctly handle the SIPFE's registration bind timer expiration.</p> <p>Workaround: None.</p>

Resolved Issues in 09.02.02R002 Release

The following Severity 1 issues are resolved in this release:

Table 19: Severity 1 Resolved Issues

Issue ID	Sev	Problem Description	Resolution
----------	-----	---------------------	------------

<p>SBX-111688 SBX-112126</p>	<p>1</p>	<p>PortFix SBX-111688: The Request-URI and TO fields in a SIP INVITE are incorrect and overwritten.</p> <p>Impact: When the Diversion header presented in an Ingress INVITE and had no RN in Request-URI, the egress RURI would use the RN in username field, after a PSX LNP dip.</p> <p>Root Cause: The fix for a previous issue broke the previous RURI username setting.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Set up the SBC with an external PSX to make LNP calls. 2. Pay attention to the disableRn flag in egress IPSP in PSX and flag UseGAPwhenRnisDisabled in the SBC's egress TG. Test different combination of < disableRn, UseGAPwhenRnisDisabled > 3. Ingress INVITE needs to contain Diversion header and no RN in the Request-URI. <p>Sample message could be found in the reproduced DBG logs and finalFix DBG logs.</p>	<p>The code is modified to ensure the RURI always contains the correct username.</p> <p>Note: When the egress IPSP has "SIP TO Header Mapping" set to "Called Number", the TO header's username is the ingress TO URI. The number will not be globalized. If the customer would like it to be globalized, the globalization profile of the egress TG could be modified to globalize TO URI. Another way is using SMM to modify it.</p> <p>Workaround: Use the SMM and use the Warning-21-00029918</p>
<p>SBX-111223 SBX-112194</p>	<p>1</p>	<p>Portfix SBX-111223: Memory leak since upgrading to 8.2.5 R0.</p> <p>Impact: High memory utilization.</p> <p>Root Cause: Two leaks of the same structure exist:</p> <ol style="list-style-type: none"> 1. A memory leak is seen when a relayed SUSCRIBE message is cranked back and an Alternate Server cannot be found. 2. The code that handles relayed messages may cause a leak in certain error scenarios. <p>Warning-21-00029922 pertains to this issue.</p> <p>Steps to Replicate: This is a memory leak that may be triggered if a relayed SUBSCRIBE is cranked back and then an alternate route is not found.</p> <ol style="list-style-type: none"> 1. This memory leak that may get triggered if a relayed SUBSCRIBE is cranked back and then an Alternate Server is not found. 2. We were unable to reproduce this error scenario which causes a leak of the Relay Control Block. 	<ol style="list-style-type: none"> 1. The code is modified to take the correct path when an Alternate Server cannot be found while handling a relayed message that has been cranked back. 2. Code has been added to start a timer when we begin processing a relayed message. In error cases, the Relay Control Block will be free when the timer expires - preventing the structure from leaking. <p>Workaround: None.</p> <p>Once the memory utilization reaches 91%, an automatic switchover occurs. To avoid an automatic switchover, Ribbon recommends performing a manual switchover during a low traffic period.</p>
<p>SBX-112309</p>	<p>1</p>	<p>An LSWU on SWe (VMware/KVM) through the Platform Manager fails with the additional reboot of the standby.</p> <p>Impact: An LSWU from 09.02.02R001 to any higher release 1:1 HA SWe (KVM/VMware) through the Platform manager fails because the standby instance undergoes an additional reboot during the upgrade procedure.</p> <p>Root Cause: The Index Marker file was missing on the standby instance prior to the LSWU procedure.</p> <p>The Index Marker file was missing on the standby instance due to when the application comes up with the standby role, the Index Marker is created only when there is a mismatch in the calculated and DB populated indices.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Create 1:1 HA SWe on KVM/VMware. 2. After the HA application comes in sync, run clearDB on the standby. 3. Perform an LSWU on the 1:1 HA SWe. 	<p>The code is modified so the Index Marker file is created on the standby instance irrespective of processor index mismatch in the calculated and DB populated indices.</p> <p>Workaround: Before initiating the upgrade procedure (from PM or CLI), the user needs to run the following command as a root user from the linux shell of the active and standby instance.</p> <pre>touch /opt/sonus/conf/swe/capacityEstimates/.indexMarker</pre>

SBX-112229	1	<p>The term-ioi is not set in the STOP record for a customer call flow.</p> <p>Impact: When running a JJ90.30 to JJ90.30 call flow, the term-ioi value from the PCV header is not being stored in the ingress protocol variant string when the transitPCV control in the JJ90.30 interworking profile is enabled.</p> <p>Root Cause: The original development work was only storing the term-ioi value in the egress protocol variant string.</p> <p>Steps to Replicate:</p> <p>Make JJ90.30 to JJ90.30 call. transitPcv is set to enabled in the customer's Interworking Profile.</p> <p>Call 1:</p> <pre>INVITE --> (PCV with icid-value and orig-ioi) <--- 180 (PCV with icid-value, orig-ioi, term-ioi) <-- 200 OK (PCV with icid-value, orig-ioi, term-ioi) --> BYE</pre> <p>Call 2:</p> <pre>INVITE --> (PCV with icid-value and orig-ioi) <--- 180 (PCV with icid-value, orig-ioi, term-ioi) <-- 200 OK (PCV with icid-value, orig-ioi, term-ioi) <-- BYE</pre> <p>After a call, verify the CDR records. In both START and STOP records, the term-ioi is filled in ingress PVSD.</p>	<p>The code is modified to ensure the term-ioi value is stored in the ingress protocol variant string.</p> <p>Workaround: None.</p>
SBX-112381 SBX-1124488	1	<p>PortFix SBX-112381: The SBC had a Crash-Application in the rebooting loop.</p> <p>Impact: A PRS Process core is occurring when the code is processing an ICE STUN packet and there are more than 20 Teams clients ringing. This could potentially happen in a simultaneous call group pickup scenario.</p> <p>Root Cause: The root cause of the core is a bug in the code that handles the incoming STUN packet. This code is overwriting the end of array by writing too many entries into the array. This results in memory corruption and eventually a core.</p> <p>Steps to Replicate: Run a call group pickup scenario with more than 20 users all having their Teams clients ringing at the same time.</p>	<p>The code is modified to prevent it from writing too many entries into the array.</p> <p>Workaround: Disable media-bypass or reduce the number of users in the call group.</p>

The following Severity 2-3 issues are resolved in this release:

Table 20: Severity 2-3 Resolved Issues

Issue ID	Sev	Problem Description	Resolution
SBX-111751	2	<p>The SCM Process is coredumping.</p> <p>Impact: The SBC is relaying multiple reason header instances in the amount of the quantity squared.</p> <p>Root Cause: Logical error when multiple instances of a reason header on the same line. The SBC is relaying in the amount of quantity squared.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Configure the transparency Reason header. 2. Make a SIP-SIP call, Ingress sends BYE with multiples Reason header instances on the same line. The SBC sends instances in an amount of the quantity squared to the Reason header on Egress. <p>Note: If the ingress sends a large number of instances on the same line.</p> <ul style="list-style-type: none"> • For old (pre 8xx), the SBC may core. • For 8xx and above, the SBC may be unable to send BYE on Egress. 	<p>The code is modified to address the issue.</p> <p>Workaround: Disable the transparency of a reason header.</p>

SBX-111932	2	<p>A restore on revision 1 failed in the standby OAM node.</p> <p>Impact: A restore on revision 1 failed on the standby intermittently.</p> <p>Root Cause: When the restore revision is requested, both OAM nodes restart. In a corner case, there is possibility of both starting as active and later one of the nodes will go down and not come up as standby.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Create a scenario with 1:1 OAM for the SSBC/MSBC/MRFP. 2. In automation, keep the cleanStart as 1 so that the automation will trigger restore revision 1. <p>Both OAM should come up.</p>	<p>The code is modified to handle the corner case.</p> <p>Workaround: Manually reboot the standby.</p>
SBX-112039	2	<p>Uploaded onfig backup file is not available at gconfig.</p> <p>Impact: The uploaded config backup file is not available on the gconfig.</p> <p>Root Cause: The uploaded file was not moved in the config directory because the system name was not correct.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Log in to the EMA. 2. Go to Administration -> System Administration -> File Upload. 3. Upload the file. 	<p>The code is modified to read the correct system name.</p> <p>Workaround: None.</p>
SBX-112085	3	<p>The SFTP transfer is not working to CDR server that only supports AES128_CTR.</p> <p>Impact: One customer has a CDR server that offers AES128_CTR for the server-to-client encryption. When the SBC offers AES256_CTR as the preferred option, the SBC encodes the packets with AES256_CTR, and receives the packets with AES128_CTR encryption. This scenario causes a decoding issue on the SBC that results in the termination of the connection.</p> <p>Root Cause: An invalid encrypted data length is received from the CDR server.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Rollover the accounting file. 2. Verify that the CDR transfer is successful. 	<p>The code is modified to offer up the AES128_CTR as the preferred encryption option, so both sides are using AES128_CTR.</p> <p>Workaround: Disable the AES256 on the CDR server.</p>
SBX-112267 SBX-112332	2	<p>PortFix SBX-112267: The correct 92x serialization results in test failures.</p> <p>Impact: When the registration control block is made redundant and the active/standby instances are running different software releases, the redundancy logic might not work correctly if the registration control block contains P-Charging-Vector (PCV) and/or P-Charging-Function-Addresses (PCFA) information. The redundancy logic will work correctly post upgrade when the serialization of the redundancy data is not required.</p> <p>Root Cause: The parameter lengths of the PCV and PCFA redundant parameters was not calculated correctly. This can lead to the redundancy code being unable to process all the redundancy information correctly.</p> <p>Steps to Replicate: Run registration related tests with PCV and PCFA header parameters in an older release and then upgrade to 922R2 or later.</p>	<p>The code is modified to correctly set the length of these parameters to avoid problems with redundancy.</p> <p>Workaround: None.</p>
SBX-111659 SBX-112378	2	<p>PortFix SBX-111659: The intercept is not working when P-Com.Session-Info is received in INVITE and tones are enabled.</p> <p>Impact: When the lawful intercept target is specified in the P-com-session-info header of the INVITE message, the intercept does not occur if an announcement resource is applied to the call and then freed up. There is no issue if the P-com-session-info header is received in the backward direction.</p> <p>Root Cause: This occurs when using P-early-media with tone profile and media monitoring profile applied to the call flow and the B-party sends P-early-media with a=inactive to start the media monitoring and does not provide RTP or does not send 200 OK prior to the monitoring timer expiring which results in the SBC playing RBT via an announcement resource. When the media path is later cut through and the announcement resource is released the intercept information was mistakenly freed up and the intercept did not occur.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Run a lawful intercept call flow using P-com-session-info header in the INVITE to indicate the target intercept point. The call should be configured with P-Early-Media, tone profile and monitoring profile to check for RTP packets received. 2. In response to the egress INVITE, send back an P-early-media header with a=inactive and no RTP packets prior to the monitoring timer expiring so the SBC plays RBT. 3. Answer the call and check that media is being sent to the intercept server. 	<p>The code is modified to ensure the intercept information is not freed up when freeing the announcement resource.</p> <p>Workaround: None.</p>

SBX-112487	2	<p>The LI connection to media server was not connected on a switchover.</p> <p>Impact: When using the PCSILI (P-com-session-info flavour of LI) in an N:1 M-SBC setup and running the SBC release 8.2 or later if there is a switchover, the connection to the LI server might not be re-established.</p> <p>Root Cause: The standby instance was trying to maintain information about the operational status of the interface used for LI processing. But this was only happening correctly for the first instance. The status of the other instances was being misinterpreted and lost. So during the transition to active, the LI code thought the interface was not available and did not try to establish the connection to the LI server.</p> <p>Steps to Replicate: In an N:1 setup perform switchovers from each active instance to the standby and check that the connection to the LI server is restored.</p>	<p>The code is modified to collect the status of the interface after the instance transitions from standby to active so that the accurate interface information is available for LI processing.</p> <p>Workaround: Bounce the interface e.g. ifup/ifdown that is used for LI connection would help to recover the connection.</p>
------------	---	--	--

Resolved Issues in 09.02.02R001 Release

The following Severity 1 issues are resolved in this release:

Table 21: Severity 1 Resolved Issues

	Issue ID	Sev	Problem Description	Resolution
1	SBX-108317 SBX-111780	1	<p>PortFix SBX-108317: The SBC switchover and isolation of node/SVWSBCD.</p> <p>Impact: The PRS core was truncated because a second ABRT was sent to the process.</p> <p>Root Cause: The PRS core was truncated because a second ABRT was sent to the process while it was in the process of writing the core.</p> <p>Steps to Replicate: This issue is not reproducible.</p>	<p>The code is modified to prevent two ABRT signals from being sent to the same process.</p> <p>Workaround: There is no workaround.</p>
2	SBX-111050 SBX-111595	1	<p>PortFix SBX-111050: There was a coredump during a reboot on a customer SBC.</p> <p>Impact: A misconfiguration of GW Signaling Ports, in which there is a Secondary GW Signaling Port configured but no Primary GW Sig Port configured, and can lead to a core and switchover.</p> <p>Root Cause: When there is no Primary GW Signaling Port configured but there is a Secondary Signaling Port configured, we enter a code path that attempts to dereference a NULL pointer (pointer to Primary Sig Port) when it is attempting to get the dscp value to use for the socket.</p> <p>Steps to Replicate: An SVT engineer should:</p> <ol style="list-style-type: none"> 1. Configure an SBC with a Secondary GW Signaling Port without configuring a Primary GW Signaling Port. i.e. gwSigPort of SBC1 should be in primary mode and SBC2 should be in secondary mode 2. Send a GW-GW call from a remote GW through this SBC using the address of the Secondary GW Signaling Port on an SBC. <p>Without the fix, this should cause a core. With the fix, there should be no issue.</p>	<p>The code is modified to prevent it from attempting to dereference a NULL pointer (pointer to Primary Sig Port) when it is looking up the dscp value to use for the socket.</p> <p>Workaround: The workaround is to avoid configuring a Secondary GW Signaling Port without configuring a Primary GW Signaling Port.</p>
3	SBX-110199 SBX-111474	1	<p>PortFix SBX-110199: There was dbg log flooding with MAJOR .VNFR: * (VAgent) message</p> <p>Impact: Event Logs are flooded with error message, when one of the OAM node or Managed VM goes down.</p> <p>Log Message: MAJOR .VNFR: *(VAgent): send error. rc=-1, error=Resource temporarily unavailable</p> <p>Root Cause: The health check message send fails, when an instance is down.</p> <p>Steps to Replicate: Run an LSWU from VNFM on an ISBC SWE instance in 1 to 1 HA mode.</p>	<p>The code is modified so the messages are only logged on the first error in health check. It is reset, when a message send succeeds.</p> <p>Workaround: The error message stops automatically once the node is recovered and starts to respond to the health check messages.</p>
4	SBX-110884 SBX-111318	1	<p>PortFix SBX-110884: The logs are getting printed in openclovis.</p> <p>Impact: If condition present, SCM is filling openclovis logs with this error message: nprintf buffer too small. Need 306 Have 300 File: /sonus/p4/ws/release /sbx5000_V08.02.02R001/marlin/SIPSG/sipsgRegAgent.c Line: 2527</p> <p>Root Cause: The error message is being logged because the code is attempting to write a debug log into a local buffer that is not big enough.</p> <p>Steps to Replicate: No testing is required. This is no risk with/without fix.</p>	<p>The code is modified to increase the size of the local buffer.</p> <p>Workaround: There is no workaround.</p>

5	SBX-109243 SBX-111274	1	<p>PortFix SBX-109243: The SBC sending 481 for CANCEL in dialog transparency.</p> <p>Impact: When the dialog transparency is enabled and call loops back to the SBC, the SBC does not handle a CANCEL sent from ingress endpoint.</p> <p>Root Cause: The SBC expects a CANCEL to have callinfo header.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Reproduce the issue. Dialog Transparency is enabled and the Call Loops back In. After adding PRACK, Ingress sends CANCEL without callinfo header. The SBC sends a 481. 2. With a fix, the SBC sends 200OK for CANCEL and relays to egress leg. 	<p>The code is modified to ensure the SBC finds the correct SG to handle the CANCEL even when the CANCEL does not have callinfo</p> <p>Workaround: None.</p>
6	SBX-111100	1	<p>Host check validation failing - Required GB RAM 6 but found 0.03125.</p> <p>Impact: The few host machines in the AWS data center use different units for RAM, HostCheck script assumes that available memory is given in MB. The HostCheck script is fixed to calculate available memory based on unit of memory.</p> <p>Root Cause: The HostCheck script does not have code to consider memory unit.</p> <p>Steps to Replicate: Create VM with >=32 GB RAM, application should come up without complaining about memory in VM.</p>	<p>The code is modified to calculate available resources based on units.</p> <p>Workaround: Create instance/VM that has < 32 GB RAM.</p>
7	SBX-107133 SBX-111046	1	<p>Portfix SBX-107133: Max FD limit is reached in SLB beyond 7,04,000 TLS connections (Access Registrations) tried.</p> <p>Impact: Max FD limit is reached in the SLB beyond 7,04,000 TLS connections (Access Registrations) tried.</p> <p>Root Cause: Observing the FD congestion at high load due to FD limits reached.</p> <p>Steps to Replicate: Check the updated FD limits in /etc/security/limits.conf.</p>	<p>The code is modified for the max FD for the KVM and SLB.</p> <p>Workaround: None.</p>
8	SBX-110350	1	<p>The SBC is forming invalid packet where it is adding 00 in around all headers and parameters.</p> <p>Impact: The SBC puts a NULL termination in every parameter and header of the message. In this case, there was a parser error on a parameter when the DNS translation was required.</p> <p>Root Cause: During the parsing of the message after a DNS query, if a parser error occurs, then incorrect termination is put/left in the message when sent on the wire.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. An incoming INVITE had a known syntax error related to alert_info, configuration was to parse through the filterProfile and not transparently forward alert_info. This resulted in parse error and bad formatting when DNS was executed. 2. Configure the alertInfo to transparently pass on egress, and DNS function doing fqd/IP swap executes correctly and egress message is sent with good formatting. 	<p>The code is modified to:</p> <ol style="list-style-type: none"> 1. Log parse an error line number and pdu message. 2. Apply the filterProfile when parsing the message. The customer needed to add filterProfile to transparently forward the parameter failing parsing on egress to avoid parse error and avoid incorrect parameter termination. <p>Workaround: On the Ingress, apply an SMM to rename alert_info to an unknown (x-alert_info), and rename back on the egress.</p>

The following Severity 2-3 issues are resolved in this release:

Table 22: Severity 2-3 Resolved Issues

Issue ID	Sev	Problem Description	Resolution
1	2	<p>The file upload is not working in the Platform Manager.</p> <p>Impact: The file upload is not working in the PM.</p> <p>Root Cause: The request URL length was too high as per-server configuration because that requested call was failed.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Log in to the Platform Manager. 2. Go to Administration -> System Administration -> File Upload. 3. Upload a file and save it. 	<p>The code is modified to make a request call successful.</p> <p>Workaround: We can apply the same changes in a customer setup also and restart the apache server.</p>

2	SBX-111700	2	<p>Getting badRidCb errors during LSWU from 7.2.4R0 to 9.2.2A003</p> <p>Impact: An upgrade from an older version to 9.2.0R0 and newer, the NP reported badRidCb error continuously: MAJOR .IPM: *NP 0 error counter badRidCb incremented: cnt 2310 last error 0x10002.</p> <p>Root Cause: In version 9.2.0R0, Ribbon Protect streaming support was added. During an LSWU, new fields related to Ribbon Protect streaming were all initialized to zero in internal data structures, including the rbnType. But in NP, rbnType = 0 triggered NP to send RTCP_APP related statistics to Protect server, which caused the errors.</p> <p>Steps to Replicate: Run an LSWU from any older version to 9.2.2 with call loads.</p>	<p>The code is modified to initialize the rbnType to 'BRM_RBBN_PROTECT_STREAM_DISABLE' so that NP does not generate RTCP_APP to Protect server.</p> <p>Workaround: No workaround.</p>
3	SBX-111211 SBX-111652	2	<p>PortFix SBX-111211 to 9.2.x - Get "violates foreign key constraint" error when assigning timeRangeProfile to a Route in EMA</p> <p>Impact: Route creation fails when the name of time range profile contains numerals.</p> <p>Root Cause: During the route creation, validation of time range profile was failing as it was containing numerals.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Create a time range profile with a name containing number ex test_1 2. Create a route with test_1 as time range profile. 3. Route creation should be successful. 	<p>The code is modified to consider numerals as valid a character.</p> <p>Workaround: Create a Route from the CLI.</p>
4	SBX-108616 SBX-111571	2	<p>PortFix SBX-108616: Fora Late Media Call, the SBC is not sending a second UPDATE towards ingress when the DLRBT is enabled.</p> <p>Impact: In a late media convert call scenario, when the DLRBT is enabled, the SBC is not sending UPDATE towards the ingress with the list of codecs received from UAS.</p> <p>Root Cause: The minimizing of media changes from other call leg functionality is suppressing the triggering of the UPDATE towards UAC.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. The UAC sends INV without SDP to the SBC. 2. The SBC sends INV with PCMA PCMU G729 101. 3. The UAS sends 180 with PCMA G729. 4. The SBC sends 180 with PCMA G729. 5. The UAC sends PRACK with G729. 6. The SBC starts playing tone with G729 towards ingress EP. 7. The UAS sends 200 OK (for INV) without SDP. 	<p>The code is modified to identify this particular call scenario, such that, the SBC sends an UPDATE towards the UAC if the tone codec is different from the end to end negotiated codec.</p> <p>Workaround: None.</p>
5	SBX-111502 SBX-111566	2	<p>Portfix SBX-111502: Metavar exchange continues even after the allocation is failed in the SBC.</p> <p>Impact: The CHM processes the metavar request/response messages even when the application is going down, which results in failure messages.</p> <p>Root Cause: The CHM does not check if the application is going down while processing the messages from the other nodes.</p> <p>Steps to Replicate: Bring up the SBC in N:1 mode and then shutdown the application on one of the nodes while another node is starting the metavar exchange in CHM activation.</p>	<p>The code is modified to ensure that application is not going down before processing the metavar messages from other nodes.</p> <p>Workaround: None.</p>
6	SBX-111469 SBX-111565	2	<p>Portfix SBX-111469: There are some problem causing MRFPs cannot recover from a restart</p> <p>Impact: The standby SBC gets the already allocated metavars of an active node while coming up in an 'active' role.</p> <p>Root Cause: The standby SBC is not checking for the condition to see if its the only node running in the redundancy group before allocating the last exchanged metavars of active node.</p> <p>Steps to Replicate: Bring up the N:1, trigger a switchover so that assigned standby becomes active and while assigned active node is coming up in 'standby' role. Restart assigned standby and ensure the proper metavars are allocated to assigned standby node while its coming up in 'active' role.</p>	<p>The code is modified to check for the already allocated metavars before allocating the metavars to self node while coming up in the 'active' role.</p> <p>Workaround: Restart assigned standby node if already allocated metavars are allotted to this node.</p>
7	SBX-111310 SBX-111562	2	<p>Portfix SBX-111310: Standby OAM is rebooting after orchestration.</p> <p>Impact: The Standby OAM is restarting after a launch.</p> <p>Root Cause: The Standby OAM is sending the serf event for 'startingStandby' with an older timestamp (which is fetched even before standby OAM is ready to come up) that results in discarding of the serf event by the active OAM node as its prior to member-join event for standby OAM node.</p> <p>Steps to Replicate: Bring up both the active and standby OAM at the same time and then while standby OAM is coming up, disconnect and connect the HA port on active OAM to trigger member-failed and member-join events.</p>	<p>While sending the 'startingStandby' serf event, get the current timestamp so that its not discarded by the active OAM node.</p> <p>Workaround: Restart the standby OAM application so that the startingStandby event is generated again with a current timestamp.</p>

8	SBX-109811 SBX-111520	2	<p>PortFix SBX-109811: The SBC uses port number RTP+1 for RTCP instead of the learned RTCP port number if the RTCP NAT learning completes before RTP NAT learning.</p> <p>Impact: The SBC sends the RTCP packets to a destination port number RTP+1 for the RTCP instead of the learned RTCP port number if the RTCP NAT learning completes before RTP NAT learning.</p> <p>Root Cause: The SBC overwrites the learned RTCP port number with the RTP+1 port number if the RTCP is learned before RTP.</p> <p>Steps to Replicate: Send RTCP packets first then RTP packets to the SBC. After call is connected, verify the RTCP port number, if RTCP is learned before RTP.</p>	<p>The code is modified to use correct RTCP learned port when RTCP learning occurs before RTP, instead of comparing the RTP and RTCP addresses from callLeg structure.</p> <p>Workaround: No workaround.</p>
9	SBX-89177 SBX-111468	2	<p>PortFix SBX-89177: A call is torn down upon a Hold (OA expiry).</p> <p>Impact: A call flow involving an early media, tone and late Crankback results in a call tear down by the SBC while processing the HOLD INVITE.</p> <p>Root Cause: While processing a 200 OK for INVITE, one of the call processing module ends up setting wrong App-Context-Id as active. Later, while processing a HOLD INVITE, this causes a failure since this module is not working on the latest App-Context-Id.</p> <p>Steps to Replicate: Call Scenario:</p> <ol style="list-style-type: none"> 1. Party A calls Party B through the SBC. 2. Party B sends 183 with SDP resulting in media cut-through at the SBC. 3. Later B sends 480. 4. The SBC is configured for the Crankback and as a result sends INVITE to Party C. 5. Party C sends 183 with SDP. 6. Party C sends 180 without SDP resulting in tone. 7. Party C sends 2xx with SDP for INVITE. 8. Party C sends HOLD INVITE. <p>Expected behavior: The SBC successfully process the HOLD INVITE.</p> <p>Actual Behavior (without a fix): The SBC tears down the call while processing HOLD INVITE.</p>	<p>The code is modified to have the correct app Context Id as active while processing 200 OK for INVITE.</p> <p>Workaround: Disable the Tones at the SBC.</p>
10	SBX-110948 SBX-111397	2	<p>PortFix SBX-110948: Load configuration overwrites the local processor index values.</p> <p>Impact: When we perform a load config operation on a 1:1 HA pair (SWe/Cloud), it ends up overwriting the local processor index estimates with the ones that are present in the config dump.</p> <p>This leads to two issues:</p> <ol style="list-style-type: none"> 1. An incorrect index estimates being populated in the DB. 2. There are discrepancy in estimations between active and standby instances. <p>Root Cause: The load config operation results in overriding the local processor index estimates with the ones that are present in the config dump. This results in standby consuming incorrect indices present in the DB thereby causing the discrepancy in estimates of the active and standby instances.</p> <p>The DB should always reflect the estimated processor indices calculated by the active instance in 1:1 HA pair.</p> <p>Steps to Replicate: On 1:1 HA SWe/Cloud instances, perform the following steps:</p> <ol style="list-style-type: none"> 1. Run the following command as root user on the active and standby instances. The following command should give same output on the active and standby instances. cat /opt/sonus/conf/swe/capacityEstimates/.index.txt 2. Perform the load config operation. 3. Run the following command as root user on the active and standby instances. The following command will give different output on the active and standby instances. cat /opt/sonus/conf/swe/capacityEstimates/.index.txt <p>With the fix, this issue will not be observed.</p>	<p>The code is modified to ensure that the DB reflects the estimated processor indices calculated by the active instance in 1:1 HA pair.</p> <p>If the processor indices values stored in the DB does not match with the indices calculated by the standby, then the standby goes for a reboot. From the next boot on-wards, the standby uses the indices stored in the DB for the session estimations.</p> <p>The previously mentioned procedure ensures that the indices and sessions estimates are same on the active and standby instances.</p> <p>Workaround: Run the following commands as root user on the active and standby instances.</p> <ol style="list-style-type: none"> 1. The rm -f /opt/sonus/conf/swe/capacityEstimates/.indexMarker 2. Run a reboot
11	SBX-111363	2	<p>The SBC VNF cannot get to Ready after a VNF Migration.</p> <p>Impact: The VNF does not get ready after a migration to the new VNF.</p> <p>Root Cause: When the new VNF is trying to send the curl request to the VNF, the request is getting rejected. When the VNF gets deployed, it contains the VNF data in its user data and it makes one VNF IPs allowed list. This allowed list is not getting updated for the new VNFs. So, when VNF migrates to the new VNF, new VNF IP is not present in the allowed list.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Launch VNF with https. 2. Deploy a VNF and scale out the VNF. Migrate the VNF to the new VNF. 3. The VNF state should be read and available. 	<p>The code is modified so whenever the new VNF sends a request to VNF, it is accepted.</p> <p>Workaround:</p> <ol style="list-style-type: none"> 1. Reboot the setup. 2. Add new VNF IP in the /opt/sonus/conf/instanceLca.json before the application comes up. 3. Check the state of the VNF on VNF.

12	SBX-99253 SBX-111358	2	<p>SBX-99253: Customer ECGI to CA Mapping Enhancement.</p> <p>Impact: The SMM Switch operation is not working after an SBC reboot.</p> <p>Root Cause: The switchIndex is not being set properly during a config restore.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Configure SMM Rule consisting of switch operation. 2. Run the Call. 3. Perform an SBC reboot. 4. Run the call (SMM will not be executed). 	<p>The code is modified to set correct the SwitchIndex during a config restore.</p> <p>Workaround: After a reboot, we can delete the SMM profile and configure again.</p>
13	SBX-111339	2	<p>CDR Viewer Download all button not working</p> <p>Impact: The CDR Viewer download all button was not working.</p> <p>Root Cause: An issue was caused when we changed the CSP. Due to that change, the content was made more strict and blob type were not allowed in Firefox and IE, resulting in a download failure.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Log in to the EMA. 2. Navigate to CDR Viewer. 3. Select the Sip Ladder and click Download All. 4. Download is successful in Firefox/IE/Chrome. 	<p>The code is modified to allow the same to further processing.</p> <p>Workaround: No workaround.</p>
14	SBX-110956 SBX-111291	2	<p>PortFix SBX-110956: Introduced the Upsampler for EVS in a EVS(wb)<=> NB codec scenario.</p> <p>Impact: The WB will not be used for EVS encoding and as a result, the channel aware mode cannot be enabled in an IDP NB scenario though negotiated through the SDP.</p> <p>Root Cause: There is an absence of an Upsampler in the Upstream path for EVS when WB is negotiated in an IDP NB scenario.</p> <p>Steps to Replicate: Run a EVS<=>G711 call with bw=wb; ch-aw-recv-5;br=5.9-13.2.</p> <p>In this case, the EVS encoder produces WB packets with channel aware mode enabled.</p>	<p>The code is modified for the EVS when the WB is negotiated in an IDP NB scenario.</p> <p>Workaround: None</p>
15	SBX-109300 SBX-111289	2	<p>PortFix SBX-109300: The SBC should not accept cmr byte from discarded packets.</p> <p>Impact: The SBC was accepting the CMR from discarded packets.</p> <p>Root Cause: The CMR was being processed before Packet Validation.</p> <p>Steps to Replicate: Run and EVS<=>EVS call. Let the peer send 32Kbps packets having CMR for 24.4Kbps.</p> <p>In this case, the 32Kbps packets are discarded as we do not support transcoding beyond 24.4 Kbps for EVS. Also since the packets are discarded CMR for 24,4Kbps is also not honored.</p>	<p>The code is modified to validate the packet first followed by the CMR processing.</p> <p>Workaround: None.</p>
16	SBX-109304 SBX-111286	2	<p>PortFix SBX-109304: The 13.2 wb cmr is not accepted when the SBC is operating in channel aware mode.</p> <p>Impact: The channel aware mode once enabled will always be enabled when the EVS encoder operates at 13.2Kbps and bandwidth WB. The only way to disable channel aware mode is to use a bitrate other than 13.2Kbps.</p> <p>Root Cause: The code to disable the channel aware mode if enabled on receiving on 13.2 WB CMR was absent.</p> <p>Steps to Replicate: Run a EVS to G711 call with br=13.2, ch-aw-recv-5; bw=wb. Stream a pcap having CMR for 13.2 WB from the peer.</p> <p>The call starts with 13.2 Kbps with Channel Aware mode being enabled. On receiving the CMR, the channel aware mode will be disabled while the bitrate is still 13.2.</p>	<p>The code is modified to address the issue.</p> <p>Workaround: None.</p>
17	SBX-110439 SBX-111280	2	<p>PortFix SBX-110439: The DSP rejects CMR requesting channel-aware if localCodecMode is set to a value other than 13.2.</p> <p>Impact: The SBC was rejecting a Channel Aware Mode CMR when the current mode of operation of EVS encoder was not 13.2Kbps WB.</p> <p>Root Cause: A conditional check of the current Bandwidth and current Bitrate to honor a Channel Aware Mode CMR.</p> <p>Steps to Replicate: Run a EVS to AMRWB call with the following SDP:</p> <p>br=13.2-24.4; bw= nb-wb; ch-aw-recv=0</p> <p>End point sends a CMR for Channel Aware Mode.</p> <p>In this case, the call starts with EVS encoder encoding packets as 24.4Kbps and on receiving the CMR, the rate changes to 13.2Kbps with Channel Aware mode enabled.</p>	<p>The code is modified to check if 13.2Kbps is a part of the activeCodecSet and whether WB is present in the bandwidth range negotiated to honor a Channel Aware Mode CMR.</p> <p>Workaround: None.</p>

18	SBX-110299 SBX-111278	2	<p>PortFix SBX-110299 to 9.2 - MRFP does not active EVS partial redundancy mode during the call when the remote offers "ch-aw-recv=0"</p> <p>Impact: When the ch-aw-recv=0 is negotiated through the SDP, a CMR request for the Channel Aware mode with offset 2 and the priority HIGH was not being processed.</p> <p>Root Cause: The root cause of this issue was a wrong initialization of the channel aware mode offset and priority in the code.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Run a EVS to G711 call with br=13.2; bw=wb;ch-aw-recv=0 2. Stream a pcap with CA CMR for priority HI and offset 2. <p>Prior to the fix the CMR is not processed. With the fix, the CMR is processed and also honored.</p>	<p>The code is modified to address the issue.</p> <p>Workaround: None.</p>
19	SBX-111270	2	<p>The CDR record for DSP insertion/rejection reason field in 9.2.2R000 test execution has marked the call as "Transcoding" instead of "Transcoding due to DTMF".</p> <p>Impact: The CDR DSP insertion reason is showing as transcoded instead of the DTMF for transcoded call with difference in DTMF parameters on the ingress and egress leg.</p> <p>Root Cause: The wrong code is present and overwrites the DSP insertion reason as transcoded.</p> <p>Steps to Replicate: Make a successful transcoded call with AMR codec on both sides and the DTMF parameter must be different in both ingress and egress.</p>	<p>The code is modified to set the DSP insertion reason based on the answered call leg.</p> <p>Workaround: None.</p>
20	SBX-111256	2	<p>The system level baseUdpPort is allowed to set higher than the range set at the IPTG level from the CLI.</p> <p>Impact: It is possible to set the system media mediaPortRange baseUdpPort value to a greater value than the value assigned to all the associated trunk groups, which is invalid.</p> <p>Example: set addressContext <name> zone <name> sipTrunkGroup <name> media mediaInterfaceGroupName <PIG name>media mediaPortRange baseUdpPort 1030</p> <p>set system media mediaPortRange baseUdpPort 1050</p> <p>Root Cause: With the introduction of the SBC and realms, the validation code was incorrectly looking for realm data to validate the system level baseUdpPort against rather than the zone /address context/trunk group information that is used in SBC configurations.</p> <p>Steps to Replicate: Try to run the configuration below and check that an error report is now generated.</p> <p>commit Aborted: 'system media mediaPortRange': System base UDP port cannot be greater than mediaPortRange configured at a trunk group</p>	<p>The code is modified to correctly use the zone/address context/trunk group information to validate the baseUdpPort information on the SBC configuration.</p> <p>Workaround: Manually check the values on the trunk group before assigning the system level value.</p>
21	SBX-111218	2	<p>Invalid value for "45.20 Reason Header value Q850" in the ACT record.</p> <p>Impact: The invalid values populated in the CDR sub field 20 in ATTEMPT record.</p> <p>Root Cause: In the Signaling Application, during Call Failure the value passed to CAM module to populate in the CDR for the sub field bit 20 was out of Q850 range. This is due to not mapping the internal CPC cause code to the corresponding Q850 standard values.</p> <p>Steps to Replicate: Disable the Reason Header Q850 flag, and make a call. Reject the call with 607, where SIPTOCPC profile mapped as 607-216 post call release check the CDR Sub field - 20.</p>	<p>The code is modified to take care of mapping to right Q850 values.</p> <p>Workaround: None.</p>
22	SBX-111177	2	<p>The SBC incorrectly interprets RTCP packets as RTP when using DLRBT.</p> <p>Impact: The RBT (ring back tone) can be terminated early when the DLRBT (dynamic local ring back tone) is enabled and RTCP packets arrive with the B-party.</p> <p>Root Cause: When the DLRBT functionality was initiated it was not informed that RTCP could be received. This resulted in RTCP packets being treated as RTP and caused the SBC to think RTP was learned. As soon as SDP was received in 183, the SBC triggered cut through and the RBT was stopped even though the call was not answered. This left the A-party with silence until the call was answered.</p> <p>Steps to Replicate: Make a PSTN to MS Teams call with DLRBT enabled. Leave the call in ringing state with MS Teams for a long time and check that the ring tone is continually generated.</p>	<p>The code is modified to correctly identify RTCP packets and not use this as an indication to media being learned so that the ring tone continues to be played until real RTP packets arrive.</p> <p>Workaround: If RTCP is not required on the egress leg then disable it. If RTCP is required there is no work around.</p>

23	SBX-111163	2	<p>Export/import of the syslog configuration fails.</p> <p>Impact: The user-config-import command fails if the customer has configured the SBC to send the Linux level logs to a remote syslog server through the platformRsyslog configuration.</p> <p>Root Cause: The child configuration objects under the platformRsyslog configuration were not being applied in the correct order during the import, which led to the configuration validation logic thinking there was no syslog server configuration and failing.</p> <p>Steps to Replicate: Apply syslog configuration under the platformRsyslog and check that it can be exported and imported without errors.</p>	<p>The code is modified to correctly define the configuration order for platformRsyslog configuration so there are no errors during the import.</p> <p>Workaround: If you edit the syslogState line in the xml file and change it from enabled to disabled, the import will complete correctly. Then manually apply the CLI command to enable the syslogState once the rest of the configuration is imported.</p> <pre><platformRsyslog> <syslogState>disabled</syslogState> <linuxLogs> <platformAuditLog>enabled< /platformAuditLog> <consoleLog>enabled</consoleLog> <sftpLog>enabled</sftpLog> <kernLog>enabled</kernLog> <userLog>enabled</userLog> <daemonLog>enabled</daemonLog> <authLog>enabled</authLog> <syslogLog>enabled</syslogLog> <nntpLog>enabled</nntpLog> <cronLog>enabled</cronLog> </linuxLogs> <servers> <serverNo>server1</serverNo> <remoteHost>2607:f160:10:4043:ce:ff0:0:35</remoteHost> <protocolType>udp</protocolType> </servers> </platformRsyslog></pre>
24	SBX-110984 SBX-111144	2	<p>PortFix SBX-110984: The EVS CMR not honored when in dtx period.</p> <p>Impact: If the SBC receives a CMR request when it is operating in DTX mode, the CMR was not being processed.</p> <p>Root Cause: The rate or bandwidth change received through a CMR was not put into effect after coming out of the no transmission period.</p> <p>Steps to Replicate: Run an EVS<=>EVS call with dtx enabled. On the Ingress leg, send a valid CMR while there is no media being sent on the Egress Leg. Send a CMR on the ingress leg when it is in the "no transmission period"</p> <p>Once the Ingress Leg comes out of the no tx period, the bitrate or the bandwidth as requested by the CMR should be put into effect</p>	<p>The code is modified to address the issue.</p> <p>Workaround: None.</p>
25	SBX-111120	2	<p>Customer Physical Server had a reboot failure.</p> <p>Impact: In a SBC Redundancy Group if an instance contains an SM core that did not cause service outage, the instance can go down when a switchover occurs between two other instances of same RG.</p> <p>Root Cause: The coreHandler/FacHandler process gets called to dump the core during a coredump on SBC if FAC feature is enabled. This process registered itself in sysIdList that is unintended and causes SM to crash while handling instance down event for other instances of same RG.</p> <p>Steps to Replicate: Create a RG, run it for few days and wait for an SM core due to the NTP, which does not cause the SBC to go down. Perform a switchover in same RG between two other nodes. The current node should not go down with fix build.</p>	<p>The code is modified to ensure the FacHandler does not register in sysIdList.</p> <p>Workaround: None.</p>
26	SBX-110850 SBX-111020	2	<p>PortFix SBX-110850: Alter CPU Affinity of Processes in 2vcpu Scenario</p> <p>Impact: There are intermittent crashes of SWe_NP/SWe_UXPAD processes in 2 vcpu deployments.</p> <p>Root Cause: The SWe_NP and SWe_UXPAD DPDK processes running on the same core caused memory corruption in mempools.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Bring up a 2 vcpu VM. 2. Run a heavy transcode call load for long duration. 	<p>The code is modified to launch the SWe_NP and SWe_UXPAD processes on different cores.</p> <p>Workaround: There is no workaround for this issue.</p>

27	SBX-111004	2	<p>The SCM Process cores on the customer PSBCs PS40, PS41, and PS42.</p> <p>Impact: A SCM Process coredump was observed because of Segmentation fault when a call was intercepted with PCSILI and media monitoring enabled.</p> <p>Root Cause: The SCM Process dumped core while pushing request for splitter resources and media monitoring was enabled.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Configure the SBC and PSX for basic audio call. 2. Configure the customer CLI on Egress Leg for basic call. 3. Enable Media Monitoring on Egress leg for basic call. 4. Enable the Tones on Ingress Leg. 5. Perform an interception based on PCSI_LI. 	<p>The code is modified to avoid the coredump.</p> <p>Workaround: None.</p>
28	SBX-109211 SBX-110945	2	<p>PortFix SBX-109211: There was a design gap for AMR around dynamic codec transcoding.</p> <p>Impact: The SBC is configured with restricted AMR/AMR-WB codec in Egress Route PSP along with Transcode "If Different DTMF PT" enabled then consider Egress endpoint selects the first pass thru restricted AMR/AMRWB payload in Offer in its 200 O.K answer. In such a case, if the DTMF attribute of the payload doesn't match then the SBC will try to match the Answered AMR /AMR-WB codec with the transcode-able AMR/AMR-WB codec. In such a case the DTMF attribute will match but the mode-set match fails that causes a call failure.</p> <p>Root Cause:The SBC is configured with restricted AMR/AMR-WB codec in Egress Route PSP along with Transcode "If Different DTMF PT" enabled then consider Egress endpoint selects the first pass thru restricted AMR/AMRWB payload in Offer in its 200 O.K answer. In such a case, if the DTMF attribute of the payload doesn't match then the SBC will try to match the Answered AMR/AMR-WB codec with the transcode-able AMR/AMR-WB codec. In such a case the DTMF attribute will match but the mode-set match fails that causes a call failure.</p> <p>The partial matching of AMR/AMR-WB codec in a Offer Answer cycle was allowed early, and that was what caused the issue.</p> <p>Steps to Replicate: The SBC receives an INVITE with (AMR PT=96 and mode-set=1 and 2833 DTMF PT=101). Egress route is configured with AMR mode-set=0, 1, 2, 3, AMR PT=126 and DTMF PT=110 Egress route is configured to transcode for Difference in 2833 PT This causes the SBC to send out INVITE to UAS as follows: 96 (AMR mode-set=1) 126 (AMR mode-set=0,1,2,3) 110 (DTMF PT)</p>	<p>The code is modified so if the SBC transcodes the call due to Different DTMF, the transcoding is allowed and first offered payload type is used.</p> <p>Workaround: None.</p>
29	SBX-110009 SBX-110876	2	<p>PortFix SBX-110009: The values in callsEstimate.txt differ between the Active and Standby SWe.</p> <p>Impact: The session capacity estimate values differ in active and standby VMs upon an LSWU upgrade of 1:1 the SWe HA pair.</p> <p>Root Cause: The following sequence of events resulted in the issue:</p> <ol style="list-style-type: none"> 1. The processor indices are re-calculated upon LSWU upgrade, which differs from the processor indices calculated by the SBC VM in the earlier version. 2. The session estimates are calculated based on these new processor indices. 3. Post upgrade, once the application comes up, the processor indices that are stored in DB (from the older version VM), are restored back to the /opt/sonus/conf/swe /capacityEstimates/.index.txt file. <p>As a result, the Active and Standby VMs obtained different session estimates upon upgrade.</p> <p>Steps to Replicate: After subjecting the 1:1 SWe HA to LSWU upgrade, content of the following file in Active and Standby VMs would differ: /opt/sonus/conf/swe/capacityEstimates/.callsEstimate.txt</p>	<p>The code is modified to retain the processor indices during the LSWU procedure.</p> <p>Workaround: After completing the LSWU upgrade, once the active(VM-A) and standby(VM-B) VMs are in sync, following a set of operations would make sure that session capacity estimates are identical in Active(VM-A) and Standby(VM-B) VMs:</p> <ol style="list-style-type: none"> 1. Reboot the Standby (VM-B) VM. 2. The standby (VM-B) VM comes up with correct session estimates. Wait for the completion of application sync. 3. Now, reboot the Active(VM-A) VM. This results in application failover. Application running on the Standby (VM-B) VM takes over the active role. 4. Active(VM-A) VM comes up with the correct session estimates and comes in sync with the application running on the Standby(VM-B) VM.
30	SBX-105890 SBX-110766	2	<p>PortFix SBX-105890: On a disk failure, the correct failover is not triggered.</p> <p>Impact: On an SBC that was running 6.2 code, the disks on the SBC got into a bad state and would not process read or write operations. Automatic attempts to reboot the SBC from the application code failed and manual intervention was required to recover.</p> <p>Root Cause: Later software releases already had better handling for this sort of issue but the code was printing multiple logs as part of the automatic reboot process and it was suspected that these could have gotten hung and the system could not get to the reboot command.</p> <p>Steps to Replicate: The issue was not reproducible.</p>	<p>The code is modified to remove the logs to avoid the potential or not getting to actual reboot command in the code.</p> <p>Workaround: The SBC needs to be manually power cycled to recover.</p>

31	SBX-109591 SBX-110172	2	<p>Portfix SBX-109591: Reject the INVITE with 100Rel when TG flag rel100Support is disabled and E2E Prack is disable on that leg after PSX DIP.</p> <p>Impact: The SBC does not tear down the call if the INVITE contains a Require: 100rel and the rel100Support flag is disabled on the ingress sipTrunkGroup, as per RFC3262.</p> <p>Root Cause: When the rel100Support flag is disabled and INVITE contains Require: 100rel, the SBC was not rejecting the Invite with 420 Bad extension. This scenario was not handled.</p> <p>Steps to Replicate: Set this flag: set addressContext default zone ZONE_IAD sipTrunkGroup TG_IAD signaling rel100Support disabled</p> <p>When the INVITE is received with Require: 100rel and endToEndPrack is disabled, the SBC should reject the call with a 420 Bad extension and the SBC should send header Unsupported: 100rel toward the ingress.</p>	<p>The code is modified so that when rel100Support flag is disabled and endToEndPrack is disabled.</p> <p>If the INVITE contains a Require: 100rel, the SBC will reject the INVITE with 420 Bad extension and the SBC will send header Unsupported: 100rel toward the ingress.</p> <p>Workaround: None.</p>
32	SBX-108410 SBX-109379	2	<p>PortFix SBX-108410: [ASAN]: sanitizer.CE_2N_Comp_ScmProcess_3.8866: ==CE_2N_Comp_ScmProcess_3==8866==ERROR: AddressSanitizer: heap-use-after-free on address 0x6190001c77dd at pc 0x558bcc9ff877 bp 0x7fea305f4e00 sp 0x7fea305f45b0</p> <p>Impact: The ASAN reported a "AddressSanitizer: heap-use-after-free" error when a Subscribe request had a NULL character in a quoted string. ie: From: "00 test"<sip:user1@host></p> <p>Root Cause: Invalid access of the freed memory occurred. Accessing memory after it is freed can cause unexpected behavior that may result in core dumps.</p> <p>Steps to Replicate: Run the codenomicomn subscribe-notify suite.</p>	<p>The code is modified so the SBC now logs a parser error If the SBC receives NULL character in a quoted string.</p> <p>Workaround: None.</p>
33	SBX-107798	2	<p>The one way audio when forked leg is MS Teams without ICE/NAT.</p> <p>Impact: For an incoming call that is forked to multiple destinations and DLRBT is enabled, there is a possibility of audio flowing only in one direction when the call gets established.</p> <p>Root Cause: When a 180 ringing is first received from one of the forked destinations, this triggering the SBC to play ringback tone, but if the call subsequently receives 18x messages and RTP media (for media cut through) from a different destination, the SBC fails to activate the RTP media flow resources correctly at the ingress leg of the call because these resources are already activated for playing the ringback tone for the other fork.</p> <p>Steps to Replicate: Set up ----- The SBC configured with call forking such that call from ingress (A) to be forked to egress (B) and egress (C). The DLRBT should be enabled on ingress and egress trunk toneAndAnnouncementProfiles.</p> <p>Procedure -----</p> <ol style="list-style-type: none"> 1. Initiate the call with an INVITE from A that should be forked to B and C. 2. From the egress endpoint B respond with 180 and then from egress endpoint C respond with 180. 3. From egress endpoint C respond with 183 with SDP. 4. Send the RTP media packets back from C to cause media CUTTHRU. 5. From egress endpoint C respond with 200 OK to connect the call. <p>Expected Results -----</p> <ol style="list-style-type: none"> 1. The INVITE sent to B and C. 2/3. The SBC sends back 180 towards A and starts to play ringback tone towards A. 4/5. The call is established, the SBC sends CANCEL towards B, ringback tone should stop and media should flow in both directions between A and C. <p>Before the fix, media was not flowing from A to C.</p>	<p>The code is modified to re-activate the media resources after the call is answered on one of the forks and the remaining forks have been cleared.</p> <p>Workaround: Not available, but the issue does not occur if the DLRBT is not enabled.</p>
34	SBX-106601	3	<p>PortFix SBX-106609: Add a check for /boot partition free space in pre-checks.</p> <p>Impact: There was an upgrade failure due to insufficient disk space on /boot partition.</p> <p>Root Cause: There was an upgrade failed due to failure to copy the new boot images as part of upgrade due to insufficient disk space in /boot partition.</p> <p>Steps to Replicate: Upgrade to the fix build and ensure upgrade is successful.</p>	<p>The code is modified as part of pre-checks to ensure a minimum of 40MB free disk space is available on /boot partition.</p> <p>Workaround: None.</p>

35	SBX-110292 SBX-110789	3	<p>PortFix SBX-110292: A Registration structure update is needed to prevent hijacking/hacking of the user after being rejected with a 403 Forbidden error.</p> <p>Impact: After a hack, the registers rejected with 403 Calls were not working.</p> <p>Root Cause: Both a normal and hacker user had very similar Register messages, except that the hacker's Auth header did not include the correct username. Due to this username mismatch, the SBC rejected the hacker with a 403 message, set the Register into 'challenged' state and later deleted the message.</p> <p>Steps to Replicate: There are 8 combinations of test cases. The major difference is in the username field Auth header.</p> <p>Test1:</p> <ol style="list-style-type: none"> 1. Proper registration: Reg->401Reg (with Auth)200 2. Hacker registration: Reg->401Reg (with hacker Auth)403 3. Proper user refresh 4. Hacker registration: Reg->401Reg (with hacker Auth)403 <p>Check for all fields and expiration timer.</p> <p>Test2:</p> <ol style="list-style-type: none"> 1. Proper registration: Reg->401Reg (with Auth)200 2. Hacker registration: Reg->401Reg (with hacker Auth)403 3. Proper user refresh 4. Hacker registration: Reg (with hacker Auth)403 <p>Test3:</p> <ol style="list-style-type: none"> 1. Proper registration: Reg->401Reg (with Auth)200 2. Hacker registration: Reg (with hacker Auth)403 3. Proper user refresh 4. Hacker registration: Reg->401Reg (with hacker Auth)403 <p>Test4:</p> <ol style="list-style-type: none"> 1. Proper registration: Reg->401Reg (with Auth)200 2. Hacker registration: Reg (with hacker Auth)403 3. Proper user refresh 4. Hacker registration: Reg (with hacker Auth)403 <p>Test5:</p> <ol style="list-style-type: none"> 1. Proper registration: Reg (with Auth)200 2. Hacker registration: Reg->401Reg (with hacker Auth)403 3. Proper user refresh 4. Hacker registration: Reg->401Reg (with hacker Auth)403 <p>Check for all fields and expiration timer.</p> <p>Test 6:</p> <ol style="list-style-type: none"> 1. Proper registration: Reg (with Auth)200 2. Hacker registration: Reg->401Reg (with hacker Auth)403 3. Proper user refresh 4. Hacker registration: Reg (with hacker Auth)403 <p>Test 7:</p> <ol style="list-style-type: none"> 1. Proper registration: Reg (with Auth)200 2. Hacker registration: Reg (with hacker Auth)403 3. Proper user refresh 4. Hacker registration: Reg->401->Reg(with hacker Auth)403 <p>Test 8:</p> <ol style="list-style-type: none"> 1. Proper registration: Reg (with Auth)200 2. Hacker registration: Reg (with hacker Auth)403 3. Proper user refresh 4. Hacker registration: Reg (with hacker Auth)403 	<p>The code is modified so when the SBC rejects the hack because of a username mismatch in the Auth header, it reverts back to previous details and state (that is complete).</p> <p>Workaround: None.</p>
36	SBX-103963 SBX-107387	3	<p>PortFix SBX-103963: Both SBCs restarted at the same time and both mounted drbd0.</p> <p>Impact: Both SBCs restarted at the same time and both mounted drbd0.</p> <p>Root Cause: The PRS Process was not updating the syncStatus flag value, due to which the the standby was also rebooting thinking the sync is not complete yet.</p> <p>Steps to Replicate: When both the the nodes are up and running, restart the standby. And while the standby is coming up, run a switchover from active CLI. The switchover should be successful.</p>	<p>The code is modified to use SMA API to check syncStatus instead of PRS and CHM local syncStatus flags.</p> <p>Workaround: None.</p>

Resolved Issues in 09.02.02R000 Release

The following Severity 1 issues are resolved in this release:

Table 23: Severity 1 Resolved Issues

	Issue ID	Sev	Problem Description	Resolution
1	SBX-110842 SBX-110869	1	<p>PortFix SBX-110842: A switchover on the SBC 5000.</p> <p>Impact: The SAM process core dumped when the SIP code was attempting to lookup a internal SIP Registration Control Block.</p> <p>Root Cause: The core was the result of SIP code attempting to using a invalid index when accessing an array element for the purposes of finding a internal SIP Registration Control Block.</p> <p>Steps to Replicate: We have been unable to reproduce this issue. The root cause was found by code inspection and has been corrected.</p>	<p>The code is modified to prevent the use of an invalid index when accessing an array element.</p> <p>Workaround: There is no known workaround.</p>
2	SBX-110833	1	<p>There are call trace logging issues while all Call Trace filters are disabled.</p> <p>Impact: Messages traced at level 4 for sageTracing may erroneously have FILTER=<name> in the trace header. For the sageTracing, the filter name should be blank.</p> <p>With the sageTracing enabled, 70% of the received INVITES are traced at the level 4 and 0.5% of calls have all their messages traced at level 4. The sageTracing is tracing collected for the Strategic Analysis Gap Execution.</p> <p>Root Cause: The filter name messages traced for sageTracing should be blank but instead uses the 64th entry in the filter names table.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. The sageTracing is enabled. 2. Create at least 64 call trace filters and enable them. 3. Delete the trace filters. 4. Send INVITE messages. 	<p>The code is modified to filter the name blank.</p> <p>Workaround: None.</p>
3	SBX-106871 SBX-110165	1	<p>Portfix SBX-106871: The SBC application times out while checking a callDetailStatus.</p> <p>Impact: The CLI 'show table global callCountStatus' timed out after a call setup failure due to the codec license not being available.</p> <p>Root Cause: The call setup failure in early stage caused an internal out of sync of the call resource, that triggers a handler of 'show table global callCountStatus' timeout when accessing the call resource.</p> <p>Steps to Replicate: Create a call in the SBC with no available codec license of the call.</p>	<p>The code is modified to clear the call resources in all related internal modules upon call failure in early stage to address the issue.</p> <p>Workaround: None.</p>
4	SBX-100881 SBX-109813	1	<p>Portfix SBX-100881: The SBC sends a release to the H323 side.</p> <p>Impact: The SBC is sending a call release to the H323 side.</p> <p>Root Cause: During the codec selection on H323 side, the SBC ran into some offer-answer collision due to ACK SDP from SIP side. As a result, the SBC is terminating the call.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Make a call from SIP to H323 and the call gets connected. 2. Send a late media Re-INV from SIP EP to the SBC. 3. The SBC sends 200 OK with SDP to SIP EP. 4. SIP EP sends ACK with SDP to the SBC. 5. The SBC sends a FACILITY to H323 EP. 6. The SBC terminates the call. 	<p>The code is modified so that the codec selection on the H323 side occurs independently during a modify offer-answer is in progress.</p> <p>Workaround: None</p>
5	SBX-105657 SBX-109924	1	<p>PortFix SBX-105657: The Customer SBC Upgrade failed.</p> <p>Impact: The LSWU upgrade failed from 7.2.2R0 to 8.2.3R0.</p> <p>Root Cause: The /tmp partition was running out of disk space and the restoration failed during an upgrade.</p> <p>Steps to Replicate: Run a LSWU upgrade from 7.2.2R0 to 9.2.2.</p>	<p>The code is modified to ensure there is enough space available.</p> <p>Workaround: Free up some space in /tmp directory and re-run the upgrade.</p>

6	SBX-110247 SBX-110309	1	<p>PortFix SBX-110247: The "sonusSbxNodePolicerMinorAlarmNotification" alarm is generated by the SBC.</p> <p>Impact: Packets arriving on a unallocated port do not get marked as rogue media.</p> <p>Root Cause: The issue occurs only on ports that was used and disabled. Packets arriving on this port get marked as grace-media packets. This is because a very high value of grace period is being incorrectly programmed for the udp port instead of the default 4 seconds. The high value is because of endianness.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Reboot the VM. 2. Make a single normal call and wait for the call to end. Note down RTP UDP port of the SBC used in the call. 3. After the call has ended, stream UDP packets to the same RTP UDP port previously used in the SBC call. We will not see any unallocated port rogue media entry for the stream. 4. Stream the UDP packets to a UDP port that has not been previously used in the SBC call. We will see rogue media entry for the stream. 	<p>The code is modified to perform the appropriate endianness translation on the affected fields to arrive at the correct value.</p> <p>Workaround: No workaround. This does not affect normal media processing, or associated statistics.</p>
7	SBX-110323 SBX-110823	1	<p>PortFix SBX-110323: A SWe media Issue occurred after an active reboot</p> <p>Impact: Media issue after switchover.</p> <p>Root Cause: Before a switchover, the loopback call was set up with 2:xx:xx:x:x (vbsc1's pkt port mac address) for both src and dst MAC address. After switchover, from packet capture, we found that the SRC MAC address= 2:xx:xx:x:x (vbsc2's pkt port mac address). As a result, using the srcMac and dstMac caused a media issue.</p> <p>Steps to Replicate: Use the following configurations to test:</p> <ul style="list-style-type: none"> • Set up loopback and normal calls for both SBC HW and SWe platforms. • Do port and SBC switchovers. • Use single and alternate media IP addresses on the loopback ports. 	<p>The code is modified to use loopback flag mirrored to standby BRM, and overwrite both srcMac and dstMac if loopback flag is set to true. This is done only on SWe when enabling or modifying the RID.</p> <p>Workaround: No workaround.</p>
8	SBX-109200 SBX-110633	1	<p>PortFix SBX-109200: Missing the CDR for Teams call transfer scenario.</p> <p>Impact: For an MS Teams blind transfer scenario with tonesAsAnnouncements enabled, where the MS Teams user A establishes a call to PSTN user B, then MS Teams user A establishes a call to user C and then invokes a blind transfer to establish the call between B and C. When the call is subsequently cleared one of the CDR STOP records for the call is not generated.</p> <p>Root Cause: When the tonesAsAnnouncements are enabled, after a blind transfer is initiated from user A, the subsequent call clearing logic is not correctly deactivating the internal announcement resources causing the call to not fully clear internally which results in STOP record not being generated.</p> <p>Steps to Replicate: With the SBC configured for MS Teams having DLRBT with announcementBasedTones enabled.</p> <p>Procedure -----</p> <ol style="list-style-type: none"> 1. Establish a call from MS Teams User A to PSTN user B. 2. Establish a call from MS Teams user A to PSTN user C. 3. Initiate a blind transfer from MS Team user A, between user B and user C. 4. Clear the call and check CDR records. <p>Expected Results -----</p> <p>Calls, transfers, and announcements should be successful.</p> <p>The SBC should generate START and STOP CDR records for A-B and A-C call.</p>	<p>The code is modified to correctly deallocate announcement resources for this call scenario so subsequent call clear can complete successfully and generate the correct CDR stop records.</p> <p>Workaround: If tonesAsAnnouncements are disabled the call scenario works as expected.</p>

9	SBX-107976 SBX-110218	1	<p>Portfix SBX-107976: Disable the FAC feature in M-SBC, and SLB.</p> <p>Impact: Non-SIP SBC personalities should not have FAC feature enabled.</p> <p>Root Cause: Earlier, the SBC did not have personality check when setting core pattern.</p> <p>Steps to Replicate: Verify the user defined core pattern is not set in /proc/sys/kernel/core_pattern when instances are spawned with the personalities msbc,slb,mrfp irrespective of facState. %set system faultAvalancheControl facState <enabled/disabled></p> <p>User Defined Pattern:</p> <pre>cat /proc/sys/kernel/core_pattern /opt/sonus/sbx/bin/CoreHandler -f /var/log/sonus/sbx/coredump/core.1.%e_%p.%t -t %i -p %p</pre> <p>Default Core Pattern:</p> <pre>cat /proc/sys/kernel/core_pattern /var/log/sonus/sbx/coredump/core.1.%e_%p.%t</pre>	<p>The code is modified to not enable the user defined core pattern for personality type M-SBC, and SLB.</p> <p>Workaround: Disable the FAC Feature. %set system faultAvalancheControl facState disabled</p>
10	SBX-107583 SBX-108381	1	<p>PortFix SBX-107583: There is one way audio on hairpinned calls.</p> <p>Impact: A one way audio problem is seen in a AMR-AMR transcoded call with DTMF interworking when "Different2833PTType" is enabled. This problem is visible for all HD codecs (AMR, AMR-WB, EVS, SILK) that uses dynamic payload in SDP offer.</p> <p>Root Cause: A one way audio problem is seen in a AMR-AMR transcoded call with DTMF interworking when "Different2833PTType" is enabled. This occurs because SBC incorrectly configures the DSP channel with an incorrect payload type, not the payload type which was negotiated during the Offer Answer exchange on the Egress leg.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Enable "Different2833PTType" and TransCode conditional in PSP and set Preferred DTMF payload to 102 on both the PSP. 2. Send AMR payload 96,AMR-WB with DTMF 110 (8k) from Ingress. 3. Ensure Egress answer sends AMR 96 and DTMF 102 in answer. 4. Ensure that the call is transcoded and there is no way audio problem. 	<p>The code is modified so the SBC configures the DSP with pass thru payload type.</p> <p>Workaround: A code fix has been made.</p>
11	SBX-107954 SBX-110166	1	<p>Portfix SBX-107954: The SBC CPXA coredump after EVS+T140 MRFP call.</p> <p>Impact: Occasionally, when the show global callDetailStatus command is executed in the CLI, the CPX process coredumps.</p> <p>Root Cause: If information is not available for a particular call, the data returned did not have a valid type for the Ip addresses returned. This caused a timeout out in confd because the CPX did not return information for the global callDetailStatus.</p> <p>Steps to Replicate: The steps cannot be reproduced.</p>	<ol style="list-style-type: none"> 1. The code is modified to return valid information for a call when information is not available. 2. The code is also modified to send a response to the CPX process when the application does not return proper data so that the CPX process will not crash. <p>Workaround: Do not run the global callDetailStatus command.</p>
12	SBX-108557 SBX-109023	1	<p>PortFix SBX-108557: The SBC continuously core dumps for SCM process since the upgrade to V09.02.01R000.</p> <p>Impact: ScmProcess may coredump due to memory corruption.</p> <p>Root Cause: There is code that is using an invalid pointer when writing to a buffer. This code was only added recently in 9.2.1R0.</p> <p>Steps to Replicate: This problem is triggered by the receipt of an invalid PDU and/or an SMM rule to reject the incoming Invite and an early ATTEMPT record was attempted to be written. The issue is random and depends on what info is NOT available when trying to write the accounting record, therefore the issue may not reproduce all the time.</p>	<p>The code that uses in invalid buffer pointer (which was added recently) has been removed.</p> <p>Workaround: If there is SMM rule (ignore/reject the Invite), then the SMM rule needs to be disabled until patch is applied.</p>

13	SBX-108126 SBX-108271	1	<p>PortFix SBX-108126: Observed a SAM Process crash in another active node during SIPFE crash testing.</p> <p>Impact: The SAM Process coredumps due to a buffer overflow.</p> <p>Root Cause: The key:value length of incoming Message is max 255 bytes. But in the yang spec, the buffer length was defined as 23 bytes. As a result of this issue, the display was truncated and also crashing due to buffer overflow.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Configure the testbed for the fault avalanche testing. 2. Trigger the crash by sending fac-sipfe-0 in the From header, where From(calling Party) is larger than 23 bytes. <p>Expectation:</p> <ol style="list-style-type: none"> 1. The current active should crash the SAM Process. 2. Standby will take over as active and blocks the SIP message that triggered crash. 	<p>The code is modified to handle 255 bytes string, so the buffer overflow does not happen and display is not truncated.</p> <p>Workaround: None.</p>
14	SBX-108388 SBX-108449	1	<p>PortFix SBX-108388: The SBC set "user=phone" on Dummy History-Info header.</p> <p>Impact: When interworking between diversion header and history info header if the diversion count is more than 2 then the SBC generates dummy history info headers. These dummy history info headers contain user=phone which some end points reject.</p> <p>Example of a dummy history info header.</p> <p>History-Info: sip:unknown@unknown.invalid;cause=302;user=phone;index=1.1;mp=1</p> <p>Root Cause: The code was mistakenly adding the user=phone attribute even when no phone number was present in the dummy history info header.</p> <p>Steps to Replicate: Configure the SBC to support interworking from diversion header with a count of 5 to history info headers and check that the dummy history info headers do not include user=phone.</p>	<p>The code is modified not to include user=phone in the dummy history info headers.</p> <p>Workaround: Use SMM to remove the user=phone attribute from the dummy history info headers.</p>
15	SBX-109464 SBX-109465	1	<p>PortFix SBX-109464: Using a leadership algorithm workaround for an old openclavis issue can cause a core dump.</p> <p>Impact: Thesafplus_gms process crashes when coming out of split brain.</p> <p>Root Cause: There was incorrect/inconsistent data results in the code asserting.</p> <p>Steps to Replicate: This problem is not easily reproducible and is caused by HA link instability/flapping.</p>	<p>The code is modified so that the data is consistent.</p> <p>Workaround: The HA link stability.</p>
16	SBX-109224 SBX-110169	1	<p>Portfix SBX-109224: AWS: The upgrade is failing with an error message " instance is not reachable " even though the upgrade status is "success :true"</p> <p>Impact: The AWS upgrade was failing due to instance non-reachability. In actuality, the VNFC was up and running but the VNFR showed VNFC status non-reachable.</p> <p>Root Cause: The VNFR-VNFC communicates using a ZMQ and VNFR updates the VNFC health check using the same ZMQ thread mechanism. There is a Zmqclient and ZmqServer. After an undetermined point in time, the VNFC(ZmqClient) is waiting for an infinite time to get a response from the VNFR(ZmqServer). This is leading to blocking state.</p> <p>Steps to Replicate: Run an upgrade/revert operations on 9.2.2/10.0.0.</p>	<p>The code is modified so the ZMQ communicates in non-blocking mode.</p> <p>Workaround: For the workaround, the user needs to run the pre check command from vnfr_intf before upgrade/revert operations and needs to take action provided as a part of output table.</p>
17	SBX-109922 SBX-110157	1	<p>Portfix SBX-109922: The buffer used while printing PSP's was running out. As a result, this is an error.</p> <p>Impact: The buffer used while printing PSP's was running out because of large PSP. As a result, this is an error.</p> <p>Root Cause: The buffer used while printing PSP's was running out. As a result, this is an error.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Run the suite with the RTCP Interworking Enhancements. 2. The warning above in the Opensaf logs should not come. 	<p>The code is modified to accommodate bigger strings.</p> <p>Workaround: None.</p>

18	SBX-109336 SBX-109925	1	<p>Portfix SBX-109336: The SBC 5110: BIOS not updated to 2.07 post-upgrade as indicated in Release Notes.</p> <p>Impact: The BIOS is not upgrading part of the SBC application upgrade.</p> <p>Root Cause: The flashroom utility is using /dev and /proc file system to upgrade BIOS. these file systems are un-mounted part of grub2 config upgrade.</p> <p>Steps to Replicate: Upgrade the SBC app from 7.x (BIOS=2.6) to >= 8.1.x.(BIOS = 2.7). The BIOS should upgrade part of an application upgrade.</p>	<p>Mount /dev and /proc file system before calling BIOS upgrade</p> <p>Workaround:</p> <ol style="list-style-type: none"> 1. Stop SBC App <ol style="list-style-type: none"> a. #sbxstop 2. Go to /opt/sonus/bios-update/ <ol style="list-style-type: none"> a. #cd /opt/sonus/bios-update/ 3. Run the script to update BIOS. <ol style="list-style-type: none"> a. # ./updateBIOS.sh bios5X00_v2.7.0-R0.rom <p>This scripts will take few minutes to update BIOS, after its done reboot host. In next boot it will boot up with new BIOS.</p>
19	SBX-108127 SBX-110403	1	<p>Portfix SBX-108127: 9.1r3 to 10.0 LSWU failed with the reason "Failed_to_install_or_configure_database"</p> <p>Impact: The LSWU failed with the reason "Failed_to_install_or_configure_database".</p> <p>Root Cause: Ownership for the files in /var/log/postgresql is getting changed at the time of the upgrade.</p> <p>Steps to Replicate: Perform an upgrade from any version to 9.x</p>	<p>The code is modified to restore ownership of the log files in /var/log/postgresql to postgres.</p> <p>Workaround: None.</p>
20	SBX-108237 SBX-108947	1	<p>PortFix SBX-108237: Performance: Observed SAM and SCM process core dump for FAC enabled execution on SBC 5400</p> <p>Impact: The SCM experiences a core dump when the Invite gets a transaction timeout and the 200 Ok for Invite is received at the same time.</p> <p>Root Cause: In the problem scenario, the SBC is trying to send ACK for the 200 OK, the SCM cores when creating a SIP Transaction for ACK Message.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Run a Basic Invite call flow. 2. Do not answer for Invite on the UAS side. 3. After 32 seconds send 200 OK for Invite from the UAC side. <p>A race condition might occur and result in core dump.</p>	<p>The code is modified not to send ACK in this scenario as the CseqType and CseNumber are unknown and 0 respectively.</p> <p>Workaround: Not Applicable.</p>
21	SBX-107690 SBX-108962	1	<p>PortFix SBX-107690: There are SBC call failures observed on T140 load with various MAJOR logs in DBG.</p> <p>Impact: A call fails due to RID Enable errors. (where RID = receiver ID and is mapped to an allocated resource). ThenDBG log shows many BrmAsynCmdErrHdr logs with cmd 0x30 (RID Enable): MAJOR .BRM: *BrmAsynCmdErrHdr: ERROR NpMediaIntf cmd 0x30 gcid 0x2128915e</p> <p>Root Cause: When the RTCP Generation is disabled, the RTCP RID for the call is expected to be disabled by Network Processor (NP).</p> <p>With the introduction of SBX-86241 "Streaming RTCP packets to Protect Server", we now have a case where RTCP Generation is enabled to stream RTCP packets to Protect Server but RTCP packets are not generated for the call and RTCP RID for the call is not enabled.</p> <p>When the RTCP Generation is disabled, the NP uses rtcpMode=RTCP Terminate to indicate that the RTCP RID needs to be disabled.</p> <p>This is incorrect since rtcpMode=RTCP Relay Monitor also has the RTCP RID enabled and NP is expected to disable the RTCP RID.</p> <p>If a call has the rtcpMode=RTCP Relay Monitor and RTCP Generation is disabled, we leak this particular RTCP RID resource.</p> <p>The next time, we try to allocate this particular RTCP RID, NP returns an error indicating RTCP RID is already allocated.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Enable Packet Service Profile //Resources/profiles/media /packetServiceProfile/rtcpOptions /generateRtcpForT140IfNotReceivedFromOtherLeg. 2. On SBC, set //System/media/mediaRtcpControl /t140RtcpMonitorInterval to 20 seconds. 3. Start T140 call but do not send RTCP packet. Terminate call in 10 seconds. 4. If you keep making this type of call, you will eventually see the Enable RID error. DBG log will have the BrmAsynCmdErrHdr entry. 	<p>The code is modified so the Bus Resource Manager (BRM) knows whether RTCP RID is allocated and needs to be disabled by NP. The code is also modified so that BRM indicates RTCP RID needs to be disabled or not.</p> <p>When NP receives the command, it uses this new parameter to decide if RTCP RID should be disabled or not.</p> <p>Workaround: Disable RTCP and disable RTCP termination.</p>

22	SBX-109384 SBX-110027	1	<p>Portfix SBX-109384: The 8.2.4F001 SBC 5400 Global level SMM stopped to work.</p> <p>Impact: The Global SMM function stopped working post upgrade and similar issue was observed after multiple switchovers.</p> <p>Root Cause: The Global SMM configurations were not being restored.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Attach an SMM at global level (for 200ok of pathcheck options). 2. Perform multiple switchovers or LSWU to another build. 3. Check if the SMM is being applied. 	<p>Restore the Global SMM configurations during multiple switchovers to address the issue.</p> <p>Workaround: Delete the profile set at global level and configure it again.</p>
23	SBX-107430 SBX-108140	1	<p>PortFix SBX-107430: URI parameter setting order of R-URI does not meet RFC 3966 requirements.</p> <p>Impact: In the request URI, the isub parameter was after the NPDI parameter, which does not meet the RFC 3966 requirement. (ie. npdi; isub=1111)</p> <p>Root Cause: In the code (sipsgCallProcEngine.c), the isub parameter was getting populated after NPDI and other user parameter.</p> <p>Steps to Replicate: SipSgMapIsdnSubAddrToSipParam this function is responsible for populating isub parameter to request uri, this moved prior to all other parameter populating functions.</p>	<p>The code is modified so the isub parameter populating function is moved before all other parameter populating functions.</p> <p>Workaround: NA</p>
24	SBX-108270 SBX-110278	1	<p>PortFix SBX-108270: The S-SBC sends a DNS SRV before completing all NAPTR query.</p> <p>Impact: The S-SBC sends DNS SRV before completing all NAPTR query.</p> <p>Root Cause: A DNS agent request timeout occurs after 10 seconds. Because the DNS agent sends the DNS lookup (NAPTR) request to DNS Client process, the DNS client process then queries to an external DNS server for NAPTR records. The DNS agent waits until 10 seconds then timeout's for each DNS lookup. As a result, the DNS request is failed with the first DNS server and query is still in process for a second DNS server.</p> <p>Steps to Replicate: Configure 2 DNS server's</p> <ol style="list-style-type: none"> 1. Configure DNS global configuration to: admin@SBX136% show global servers DNS global retransmissionCount 14; retransmissionTimer 500; set addressContext default dnsGroup DNSGrp1 negativeDnsCacheSupport disabled 2. Configure dnslookupTimeoutTimer to 10 sec. 3. Stop the DNS server. 4. Make a call. <p>Observation:</p> <ol style="list-style-type: none"> 1. DNS should retransmit DNS NAPTR query to DNS server1 and after should move to DNS server2. 2. After 10 sec, NAPTR query shall stop. 3. The SBC should send SRV/A query to DNS Server2. 	<p>The code is modified to:</p> <ol style="list-style-type: none"> 1. Add new configuration timer flag "dnslookupTimeoutTimer" for DNS lookup. 2. Once the lookup timer expires, the DNS client will stop sending same query to DNS server (Ex: if NAPTR query is re-transmitted towards the DNS server, after timer expire. NAPTR query will be stopped). <p>Workaround: None.</p>
25	SBX-107517 SBX-108990	1	<p>PortFix SBX-107517: Import configuration results in two different final status depending where you look at it.</p> <p>Impact: The postgres DB restore fails that leads to provisioning issues such as: The "show configuration" does not display a complete configuration related to policy entities like sipTrunkGroup, IpPeer etc. The "delete" operation fails for sipTrunkGroup and other policy related entities. The exportConfig does not contain policy related entities.</p> <p>Root Cause: The PostgresDB restore is failing due to restricted permission.</p> <p>Steps to Replicate: Set up required: 1to1 S-SBC virtual platform and register it within Direct-Single type of cluster on EMS.</p> <ol style="list-style-type: none"> 1. Configure 2 TG and saveAndActivate revision (e.g. revision 2). 2. Configure couple of more CLIs and save AndActivate (e.g revision 3). 3. Restore revision 2 either from CLI or EMS. 4. After restore is successful. try to delete a TG. Observed Application failure. <p>After implementing the fix, use the 4 steps and as a result, the delete was successful.</p>	<p>The code is modified to permit an appropriate permission to postgres user while restoring the pg_policy dump.</p> <p>Workaround: Use the following steps: #cd /opt/sonus/sbx/scripts</p> <p>Update the oam-config.sh on active and standby the SBC as following:</p> <pre>replace \$CHMOD -R 660 \$extractDir/ with \$CHMOD -R 665 \$extractDir/ accessRights=\$STAT -c %a \$SONUS_TMP_DIR \$CHMOD 775 \$SONUS_TMP_DIR</pre> <p>find \$RM -rf \$extractDir and add following line below this</p> <pre>\$CHMOD \$accessRights \$SONUS_TMP_DIR</pre> <p>Restore a revision that has pg_policy_xx.dump with correct Trunkgroup, ipPeer etc. entries.</p>

26	SBX-110354 SBX-110831	1	<p>Portfix SBX-110354: After upgrading from V08.02 or V09.02 to 10.00.00R000 successfully, a REVERT from 10.00.00R000 to V08.02.04R000 is failing.</p> <p>Impact: Reverting from the higher software version to lower software version fails.</p> <p>Root Cause: The OAM does not upload restored revision against lower software version (ie. restore of config tar ball that was created on lower software version before upgrade). After following the MOP for downgrade when OAM comes up, it downloads the last revision uploaded on EMS (in this case config revision for higher software version). As a result, it fails to start the service because of a version mismatch.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Spawn OAM + S-SBC on V08.02.04R000. 2. Create multiple revision (e.g revision 1, 2, 3, 4). 3. Upgrade to software version V10.00.00R000. The revision 5 gets created from revision 4, Create revision 6, 7, 8. 4. Restore to revision 4 -> revision 9 will get created and uploaded to EMS. 5. Downgrade to revision V08.02.04R000. <p>Expected O/P: The OAM should come up with revision 10 created from revision 9.</p>	<p>The code is modified to upload a configuration when a revision that is created on a lower software version is restored. And after a downgrade, it picks the revision with software version where the downgrade is done.</p> <p>Workaround: None.</p>
27	SBX-106475 SBX-110899	1	<p>Portfix SBX-106475: The Cloud SBC instances are not coming up after a fresh installation or rebuild and the MGMT I Pis unreachable in BLR-PC3.</p> <p>Impact: The SBC instances are not coming up after a fresh installation or rebuild and the mgmt ip is unreachable when using the X710 sriov vfs for pkt interfaces.</p> <p>Root Cause: The X710 driver reset has been known to experience a delay in completing, thus causing a cloud-init service failure. For example, after the OS boot (as part of renaming the interfaces by sonusdev), the X710 driver is restarted. When the X710 does not immediately reset, the next cloud-init service fails, and the system is not in proper state.</p> <p>Steps to Replicate: The SBC instances should come up properly after fresh installation or rebuild and mgmt ip should be reachable with all supported sriov vfs for pkt interfaces.</p>	<p>The code is modified so after the renaming is done and the network driver is restarted, wait for some seconds to check if all the nic ports are up and running in sonusdev itself, before exiting the service. This blocks all other dependent service from starting. Once the nic port are up, proceed with the system bring up. With these code changes, the nic does not come up then, we log it as critical log and expect the user to check on it.</p> <p>Workaround: None.</p>
28	SBX-109597 SBX-110158	1	<p>Portfix SBX-109597: Observed SCM process crash on the newly active S-SBC during SWO testing.</p> <p>Impact: While the SBC is running in sensitive mode, observed the SCM Process core during a switchover testing under the 1000 cps call load.</p> <p>Root Cause: As part of switchover, all the connected calls are rebuilt on a newly active SBC. But, there may be a chance that the SBC can hold few stale entries related to active calls in the system, which will be cleaned up after some time. If any of these stale entries receive an unexpected events, that leads to a core dump.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Run 1000 cps load with 120 CHT. 2. Perform a switchover by killing the PRS Process. 	<p>The code is modified to identify these events for the stale entries and ignore them to avoid a core dump when the SBC running in the sensitive mode.</p> <p>Workaround: None.</p>
29	SBX-108225 SBX-109351	1	<p>PortFix SBX-108225: Observed core files on fresh installed VMware SWe.</p> <p>Impact: The SSREG and SLWRESD processes a core dump on VMware SWe due to the time being set in past.</p> <p>Root Cause: When the SBC is installed using ISO and then app is installed, the ntp sync is occurring while the application is running and in case the time is set in the past, it causes SSREQ and SLWRESD processes to core dump.</p> <p>Steps to Replicate: Launch with the NTP IP other than a default IP, check if a core dump occurs due to time being set in the past.</p>	<p>The code is modified to update the /etc/ntp.conf before the SBC comes up so the runntpd can access the NTP server IP.</p> <p>Workaround: No workaround.</p>

30	SBX-108219 SBX-110162	1	<p>Portfix SBX-108219: Observed a SAM Process coredump in SLB for 1000 cps Peering Call load.</p> <p>Impact: In a load run, there was a core being observed in the SLBDEBUG call when it hits the default states as part of the SBC message processing.</p> <p>Root Cause: In the SLB debug call, there was new line () being passed as part of a string.</p> <p>Steps to Replicate: Not reproducible. The issue seen once when running. Run 1000 cps INVITE call load with PRACK enabled.</p>	<p>The code is modified so the new line () is not part of the debug string.</p> <p>Workaround: There is no workaround identified for this issue.</p>
31	SBX-110642 SBX-110995	1	<p>PortFix SBX-110642: A coredump occurs after a blind transfer call.</p> <p>Impact: The DNS process is coring when TEAMS blind transfer call is made.</p> <p>Root Cause: The NULLError check was missing before accessing DNS Group pointer in DNS module. This occurs when DNS Group is not attach to Zone. So get function for the DNS group pointer returns NULL.</p> <p>Steps to Replicate: PSTN to TEAMS TEAMS blind transfer to PSTN2 Expected Result: There should not be a coredump during the BT.</p>	<p>The code is modified so now the DNS group is fetched after setting the current DNS group to default DNS group, this is case when the DNS group IP not provided.</p> <p>Workaround: No workaround.</p>
32	SBX-108612 SBX-110161	1	<p>Portfix SBX-108612: An SCM Process core dump occurred when INVITE with message size of 31700 bytes is sent</p> <p>Impact: Receipt of SIP messages with Content-Type: multipart/mixed and more than 10 content entries cause coredump, if a message manipulation rule is active with criterion on message body.</p> <p>Root Cause: Missing handling for unexpected message received.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Configure and apply a message manipulation rule with criterion messageBody. 2. Send in INVITE with Content-Type: multipart/mixed and more than 10 content entries. 	<p>The code is modified so that if more than 10 content entries are present, the message is not considered for SMM rule actions with criterion on message body.</p> <p>Workaround: None.</p>
33	SBX-109327 SBX-111048	1	<p>Portfix SBX-109327: Some CDR Events in ACT files have timestamps that are not current.</p> <p>Impact: Duration Field (+24 Duration of SIP Registration/Subscription Context) is not displayed for some register transaction.</p> <p>Root Cause: Duration Field (+24 Duration of SIP Registration /Subscription Context) is only updated in case of de-register for Successful transaction.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Configure SBC for Registration/Subscription. 2. Perform following Registration. A -----REGISTRATION----->SBC-----REGISTRATION---->B A<-----404-----SBC<-----404-----B A -----REGISTRATION----->SBC-----REGISTRATION---->B A<-----200OK-----SBC<-----200OK-----B 3. Check CDR for "Duration Field" (+24 Duration of SIP Registration/Subscription Context) 	<p>The code is modified to display the duration of SIP Registration /Subscription context for every CDR event.</p> <p>Workaround: None.</p>
34	SBX-106760 SBX-107967	1	<p>PortFix SBX-106760: Performance: While running IMS AKA Registrations on SWe KVM, cannot achieve 1000 subscribers with 30 RPS properly.</p> <p>Impact: The SBC was not able to achieve the IMS AKA registrations even at a lower rate (30 registrations per second).</p> <p>Root Cause: For one of the IPsec features, the XRM code was modified to perform route lookup and next hop destination MAC resolution during IPsec SA setup. This was causing a delay in setting up IPsec SAs during load causing timeouts during the IMS AKA registration.</p> <p>Steps to Replicate: Run the IMS AKA Registration load at 30rps.</p>	<p>The code is modified to not invoke the route loopup and destination MAC resolution code that is causing delay in setting up IPsec SA, which addresses the issue.</p> <p>Workaround: None.</p>

35	SBX-110066 SBX-110491	1	<p>Portfix SBX-110066: Observed a SAM Process core dump in SLB while testing Performance with SLB (TLS Peering Calls) - One time Occurrence.</p> <p>Impact: A SAM Process core was observed on a shutdown.</p> <p>Root Cause: A race condition in the process shutdown code occurred allowing access to an invalid pointer, and causing a core dump during the shutdown of the SAM Process.</p> <p>Steps to Replicate: Restart the SBC and look for a SAM Process core.</p>	<p>The code is modified to avoid the race condition that led to the core.</p> <p>Workaround: No workaround, but since this core is during shutdown, there is no impact to normal operations of the SBX.</p>
36	SBX-109953 SBX-109984	1	<p>PortFix SBX-109953: Performance: Observed NP_WRK0 Core on Active Instance on OpenStack while using VIRTIO.</p> <p>Impact: The SWE_NP can crash during a sbxrestart or sbxstop in the extra small memory profile.</p> <p>Root Cause: There is corruption due to incorrect use of global variable in an extra small memory profile.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Bring up the SBC in extra small memory profile. 2. Run a sbxrestart or sbxstop. 	<p>The code is modified to address the issue.</p> <p>Workaround: None.</p>
37	SBX-108310 SBX-109287	1	<p>PortFix SBX-108310: Various link monitor issues on the SBC SWEs with port redundancy enabled.</p> <p>Impact: On the VMware platform with Intel 82599 SR-IOV packet port VFs, the link does not get restored upon toggling the underlying physical link.</p> <p>Root Cause: Because of VMware PF driver shortcoming, DPDK ixgbev PMD code does not get link up indication from link speed status register for SR-IOV VFs from Intel 82599 card.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Bring up port redundancy setup on VMware with SR-IOV 82599 VFs packet port. 2. Toggle the carrier link state on the NIC by cable pull or interface link toggle on host through the IP commands. 3. Observe the link status is not restored. 	<p>The code is modified to detect the physical link loss from the NACK status set on the hardware mailbox register, specifically for the VMware platform.</p> <p>Workaround: No workaround. Only reboot clears the link state.</p>
38	SBX-106858 SBX-110390	1	<p>Portfix SBX-106858: The SBC clears a call on receipt of 200OK answer - CFNA call.</p> <p>Impact: The SBC clears call on receipt of answer for a MRF transcoded call involving egress UPDATE followed by tone play.</p> <p>Root Cause: During tone play, the SBC detaches the MRF resource from the call since tone is played by SBC. However, the media endpoint resource facing MRF stays with the call that results in failure later while assigning MRF resource back to the call during answer processing.</p> <p>Steps to Replicate: The following events lead to failure in this call flow:</p> <ol style="list-style-type: none"> 1. 183 with SDP from egress leads to cut-thru (MRF transcoded call). 2. Then, egress peer sends SIP UPDATE to the SBC. 3. Followed by 180 from egress which arms RTP monitoring at the SBC. 4. Followed by RTP failure notification after 2 seconds which starts Tone. 5. Followed by answer which results in failure. 	<p>Detach the media endpoint resource facing MRF from the call during tone play to address the issue.</p> <p>Workaround: None.</p>
39	SBX-107492 SBX-110168	1	<p>Portfix SBX-107492: Observed a SCM Process core dump after 2 hours of load run with EVRCB to G711 transcode calls.</p> <p>Impact: During call load, the SIPFE module was crashing while writing date to a mapped address.</p> <p>Root Cause: The write to memory mapped address for files was causing a delay and leading to health check time outs causing core dumps</p> <p>Steps to Replicate: Run a 100 cps SIP-SIP call load for about 2-6hrs period of time.</p>	<p>The code is modified from file to Shared memory, which solved the problem</p> <p>Workaround: None.</p>

40	SBX-110184 SBX-110381	1	<p>Portfix SBX-110184: Performance: While pumping Emergency call load on top of normal call load, HPC calls are getting rejected with 488 sip Error due to dsp overload even though resources were reserved on Openstack D-SBC HA.</p> <p>Impact: The HPC calls are rejected with 488 even though the SBC has sufficient DSP resources reserved for such calls through the highPriorityCompressionReservation configuration.</p> <p>Root Cause: In a D-SBC setup, when DSP allocation request is received, the T-SBC applies the call admission policy to know whether the call can be admitted or not before proceeding with DSP allocation. In this process, the T-SBC was not considering if the allocation request was for HPC call or normal call. As a result, the HPC calls were treated like normal call leading to rejections if this T-SBC is running under high load.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Reserve DSPs for HPC calls. 2. Run a normal call transcoding at 200cps. 3. Run HPC transcoded calls at 5cps. 4. Ensure the HPC calls work as long as reserved DSP resources are not exhausted. 	<p>The code is modified to give preferential treatment to HPC calls during call admission policy at the T-SBC.</p> <p>Workaround: No workaround.</p>
41	SBX-107137	1	<p>A DEADLOCK was detected for sysID 85, task CHM.</p> <p>Impact: Health check timeout resulted in deadlock for CHM and process crash to recover.</p> <p>Root Cause: The CDB read call took longer than a health check interval resulting in CHM process being crashed to recover.</p> <p>Steps to Replicate: The problem could not be reproduced.</p>	<p>The code is modified so the health check timeout is increased to 30 now from previous 10. The longer health check interval allows for unexpected cases where processing can take longer than expected especially on SWE systems.</p> <p>Workaround: None.</p>

42	SBX-110303	1	<p>The MGMT Port Status was showing the wrong status.</p> <p>Impact: An Unconfigured Management port is shown as UP when there is no cable connected to the port.</p> <p>Root Cause: The SBC did not keep track of management port status if the management port is not configured/used.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Show the issue. <ol style="list-style-type: none"> a. Leave SBC 5400 mgt2 port unconnected. b. Display mgmtPortStatus - the unconnected mgt2 shown as admnEnabledPortUp. This is the issue. <pre>admin@YF01> show table system ethernetPort mgmtPortStatus CE PORT IF NEGOTIATED DUPLEX RX ACTUAL TX ACTUAL RX TX RX NAME NAME INDEX MAC ADDRESS SPEED LINK STATE MODE BANDWIDTH BANDWIDTH PACKETS PACKETS ER ----- ----- YF01 mgt0 1 0:10:6b:3:c8:d speed1000Mbps admEnabledPortUp full 232 0 45 3 0 YF01 mgt1 2 0:10:6b:3:c8:e speed1000Mbps admEnabledPortUp full 0 0 6 2 0 YF01 mgt2 3 0:10:6b:3:c8:f unknown admEnabledPortUp unknown 0 0 0 0 0 YF01 mgt3 4 0:10:6b:3:c8:10 speed1000Mbps admEnabledPortUp full 0 142 2 42 0 [ok][2021-05-18 10:32:15]</pre> 2. Show the fix. <ol style="list-style-type: none"> a. Leave SBC 5400 mgt2 port unconnected. b. Display mgmtPortStatus - the unconnected mgt2 shown as admnEnabledPortDown (this is the fix). <pre>admin@YF01> show table system ethernetPort mgmtPortStatus CE PORT IF NEGOTIATED DUPLEX RX ACTUAL TX ACTUAL RX TX NAME NAME INDEX MAC ADDRESS SPEED LINK STATE MODE BANDWIDTH BANDWIDTH PACKETS PACKETS ----- ----- YF01 mgt0 1 0:10:6b:3:c8:d speed1000Mbps admEnabledPortUp full 248 0 271 6 YF01 mgt1 2 0:10:6b:3:c8:e speed1000Mbps admEnabledPortUp full 0 286 27 103 YF01 mgt2 3 0:10:6b:3:c8:f unknown admEnabledPortDown unknown 0 0 0 0 YF01 mgt3 4 0:10:6b:3:c8:10 speed1000Mbps admEnabledPortUp full 0 0 0 8 [ok][2021-05-18 13:32:27] admin@YF01></pre> c. Display alarms and make sure there is no alarm on mgt2 (un-configured management port) <pre>admin@YF01> show table alarms currentStatus ALARM CLEAR ID TYPE TIMESTAMP INITIAL TIMESTAMP COUNT DESC ----- ----- 1237 AUTOMATIC 2021-05-18T17:27:53-00:00 2021-05-18T17:27:53-00:00 1 System Policer Alarm Level: Minor, Policer [ok][2021-05-18 13:32:30] admin@YF01></pre> 	<p>Keep track of the management port state changes even if there is no Management IP interfaces on the management ports. Do not generate management port down event if there is no Management IP interfaces on the management port.</p> <p>Workaround: None. Optionally, connect unused management ports to Ethernet switch to reflect the port UP status.</p>
43	SBX-109893	1	<p>The SBC frequently performs a switchover with a core dump after a 9.2.1R0 upgrade.</p> <p>Impact: An SCM experienced a core dump in the code that is executed only when the STI feature is enabled.</p> <p>The core dump was the result of code in SipSgStiCopyDisplayNametoCPC() attempting to de-reference a NULL pointer. The fix is to add a check for NULL before attempting de-reference the pointer.</p> <p>Root Cause: The root cause is that there is code in an STI specific function that attempted to de-reference a NULL pointer. A NULL pointer check is missing in this code.</p> <p>Steps to Replicate: Specific steps to reproduce this issue are not known. The root cause and the fix were found through code inspection and core analysis.</p>	<p>The code is modified to add a check for NULL before attempting de-reference the pointer.</p> <p>Workaround: The only known workaround is to disable STI.</p>

44	SBX-109055	1	<p>Error in EMA when modifying and saving sipAdapterProfile.</p> <p>Impact: Error in EMA when modifying and saving sipAdapterProfile.</p> <p>Root Cause: The scenario comes up when profile name has ellipses in datatable. If profile name is very big, ellipses (...) will be added at the end, due to this the actual profile name is not passed back during save operation.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Login in to EMA. 2. Navigate to SIP Adaptor Profile. 3. Select the profile that has ellipses (...) and modify then save it. 4. Profile gets saved with any error. 	<p>The code is modified to match if ellipses present in profile name, getting the title of selected entry from datatable. The title is not available in profile name without ellipses thus taking text from selected entry.</p> <p>Workaround: No workaround.</p>
45	SBX-110003	1	<p>There was a SBC 5400 core dump after MSRP call with direct media.</p> <p>Impact: A PRS core was encountered when customer was running MSRP with direct media and MSRP Multiplexing was enabled.</p> <p>Root Cause: The root cause is that the MRM code is using an invalid index to get a pointer to an array element.</p> <p>Steps to Replicate: This is not reproducible as it is most likely triggered by a race condition.</p>	<p>The code is modified the ensure that MRM uses the correct index when attempting to get a pointer to an array element.</p> <p>Workaround: The only known workaround is to avoid running MSRP with direct media and MSRP Multiplexing.</p>
46	SBX-107983	1	<p>The LSWU from V07.02.00-R002 to V09.02.00-R001 has failed with the error "CpxIntfErrorExit".</p> <p>Impact: Upgrade fails and the SBC fails to come up due to missing SNMP configuration.</p> <p>Root Cause: The targetAdrsParams for standard V2 trap had been removed from the configuration and this information was expected to be present during the upgrade processing.</p> <p>Steps to Replicate: Perform upgrade with SNMP configuration in place and check for errors in the CE_Node.log file.</p>	<p>The code is modified to identify the SNMP upgrade failures.</p> <p>Workaround: Manually add back in targetAdrsParams for traps.</p>
47	SBX-110059	1	<p>The SBC upgrade failed as the SNMP trapTarget with a name containing white-space.</p> <p>Impact: If any SNMP trapTarget name contains a space, upgrade to 9.2.1 will fail.</p> <p>Root Cause: A trap target is configured with a name containing a space character, which is allowed for this table.</p> <p>Steps to Replicate: On 8.2 software, use CLI to create a trapTarget with space character: set oam snmp trapTarget "test target" ipAddress 1.2.3.4 state enabled Perform an upgrade to 9.2.1.</p>	<p>The code is modified so that it can cope with trapTarget names with space characters.</p> <p>Workaround: Prior to upgrade ensure no trapTarget names have space, recreate with a new name not containing space.</p>
48	SBX-109174	1	<p>The SBC unable to load config after upgrading from 7.2.3S400 to 10.0 A006</p> <p>Impact: After upgrading a cloud SBC from a pre 9.2.0 version, unable to reload the config when the 'SystemName' is not set to 'vsbcSystem'</p> <p>Root Cause: The configuration filename pre 9.2.0 uses 'vsbcSystem' instead the of the defined 'System Name', meaning the SBC is unable to find the file to load it.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Save a configuration on a cloud SBC pre 9.2.0 (e.g. AWS). 2. Upgrade the SBC. 3. Attempt to load the configuration file. 	<p>The code is modified to address the issue.</p> <p>Workaround: Change configuration filename to be that of the set 'SystemName'. e.g. If 'SystemName' is set as 'aws-sbc' change: config_vsbcSystem_20210414_035150.tar.gz -> config_aws-sbc_20210414_035150.tar.gz</p>

49	SBX-110639 SBX-111102	1	<p>PortFix SBX-110639: When the call is dipped the invite sends only the LRN.</p> <p>Impact: When ingress INVITE's DIVERSION header does not present and TO header is different from RURI, i.e., it has redirecting original number, the egress RURI will take translated number, i.e., LRN, as RURI in egress INVITE, after LNP dip.</p> <p>Root Cause: The RURI should take LRN in egress, while DIVERSION header in ingress presents. But the code change for a previous issue has made RURI equal to LRN even DIVERSION header was missing.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Set up SIP to SIP LNP call, on SBC and external PSX system. 2. The TO header username (phone number) is different from RURI. 3. There is no DIVERSION header. <p>This should reproduce the problem and test the fix.</p> <p>Please go through all regression test related to SBX-106067. It's important NOT to break any functionalities in SBX-106067.</p>	<p>The code is modified the condition to only set RURI equal to LRN while DIVERSION header presents.</p> <p>Workaround: None.</p>
----	-------------------------	---	--	--

The following Severity 2-3 issues are resolved in this release:

Table 24: Severity 2-3 Resolved Issues

	Issue ID	Sev	Problem Description	Resolution
1	SBX-104733 SBX-110295	2	<p>Portfix SBX-104733: SCM Process core dump during healthcheck.</p> <p>Impact: The SCM Process core dumped when too many set operations executed in a single commit.</p> <p>Root Cause: The Call Processing tasks takes longer than 10 seconds when subscribing to the changes made to SIP Trunk Group.</p> <p>Steps to Replicate: Configure 12 trunk groups. Modify 11 trunk groups with a single commit command and the CLI will deliver the following error message: Aborted: 'addressContext default zone ZONE_IAD sipTrunkGroup': Too many set operations between commits. Revert the session and retry(max set operations 10).</p> <p>Again modify 10 TG in single commit: O/P: commit complete</p>	<p>The code is modified per commit is changed to 10 from earlier value of 50.</p> <p>Workaround: Ribbon recommends performing a commit after every CLI transaction.</p>
2	SBX-104619 SBX-109911	2	<p>PortFix SBX-104619: The FM process core dumped.</p> <p>Impact: The FM Process crashed and wrote core dump.</p> <p>Root Cause: The FM Process tried to read the /proc/meminfo file, which should always exist, but it got a file not found error.</p> <p>Steps to Replicate: It is not known how to reproduce this as this should never happen, so defensive code added to prevent null read/write.</p>	<p>The code is modified so when we cannot read the /proc/meminfo file, we return the last good value read instead of a NULL to prevent the crash.</p> <p>Workaround: None.</p>
3	SBX-109177 SBX-109282	2	<p>Portfix SBX-109177: The SbcSftp throws an error as "Failed to add ACL".</p> <p>Impact: The sbcsftp application does not remove the ACL created correctly.</p> <p>Root Cause: The permissions to delete the ACL gets lost while lowering permissions to ensure the sbcsftp can only access the current user's files.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Run SbcSftp as linuxadmin. Verify no 'Failed setting permissions' error messages appear. 2. Verify Cleanup Script: <ol style="list-style-type: none"> a. Run SbcSftp but kill the the process before it completes, with 'kill -9 \$PID'. b. As root, get the name of the ACL: '\$CPSI -d acl grep SbcSftp'. c. Wait 15 minutes. d. Verify ACL removed: '\$CPSI -d acl grep SbcSftp'. 	<p>The code is modified so that we can raise the permissions again once SFTP is complete.</p> <p>Workaround: None.</p>

4	SBX-109968 SBX-110233	2	<p>PortFix SBX-109968: The 822R0 SCM Process core dumped.</p> <p>Impact: The SCMPProcess core dumped on a SWe.</p> <p>Root Cause: Active SCM sent a Redundancy message to Standby SCM with a Registration index that was outside of its MAX range. The error handling code caught this and attempted to clean up the Registration Control Block but had a problem because the Control Block had not yet been initialized.</p> <p>Steps to Replicate: This problem is not reproducible.</p> <p>It is most likely triggered by a race condition which results in the Active and Standby to have different numbers for max number of Registrations (on a SWe the max number of Registrations comes from the file callEstimates.txt and in this case, the files on the active and standby do not match).</p> <p>This appears to be a SWe specific issue (the callEstimates.txt file is only used in SWE setups).</p>	<p>The code is modified to initialize the Registration Control Block immediately after allocating it, before any validation checks, so that the error handling code can clean up the RCB without causing a core.</p> <p>Workaround: There is no known workaround.</p>
5	SBX-110215 SBX-110387	2	<p>Portfix SBX-110215: A coverity issue.</p> <p>Impact: Code fix for an SCM core dump (SBC-108237) when the Invite gets a transaction timeout and the 200 Ok for Invite is received at the same time needed some cleanup. Coverity issue: CID:338630</p> <p>Root Cause: In the problem scenario, the SBC is trying to send ACK for the 200 OK, when the SCM cores while creating a SIP Transaction for ACK Message.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Run a Basic Invite call flow. 2. Do not answer for Invite on the UAS side. 3. After 32 seconds, send 200 OK for Invite from the UAC side. <p>A race condition might occur and result in core dump.</p>	<p>The code is modified not to send ACK in this scenario as the CseqType and CseNumber are unknown and 0 respectively.</p> <p>Workaround: None.</p>
6	SBX-101432 SBX-109029	3	<p>PortFix SBX-101432: Question about the DSP insertion.</p> <p>Impact: Customer was seeing the "DSP insertion/rejection reason" CDR fields set to "Rejected codec unlicensed" and wanted to know why.</p> <p>Root Cause: The documentation did not provide any guidance as to when this value could appear.</p> <p>Steps to Replicate: The customer call scenario is unknown.</p>	<p>The code is modified to provide guidance on when it could happen as per the scenario below:</p> <p>These are the cases when DSP Rejection reason can be set to "Rejected codec unlicensed"</p> <ol style="list-style-type: none"> 1. The SBC does not have a DSP. 2. The SBC has DSP but codec licenses absent. 3. The SBC has DSP and codec licenses but transcoding not enabled for the codecs. <p>In any of the above cases if passthrough is possible and if the call is successful then DSP Rejection reason updates after a successful offer answer.</p> <p>This value might appear in the case where the offer /answer exchange did not complete.</p> <p>As the customer was unable to provide details on the specific call flow additional info level logging and call trace logs is modified as well to provide more details for the future.</p> <p>Workaround: None.</p>
7	SBX-108143 SBX-110227	2	<p>Portfix SBX-108143: Set FAC as enabled by default.</p> <p>Impact: The FAC was enabled by default in 9.2 but was disabled in release 9.2.1. It is now enabled by default in further releases of 9.2.x and 10.0.</p> <p>Root Cause: The FAC was temporarily disabled by default until performance issues were fixed.</p> <p>Steps to Replicate: Run the FAC impacts performance on high cps, and perform load test.</p>	<p>The code is modified so the shared memory is used instead of memory mapped files.</p> <p>Workaround: Keep the FAC disabled if using high cps until it is upgraded to fixed version.</p>
8	SBX-101239 SBX-110571	2	<p>PortFix SBX-101239: The congestion seen in a SBC with no traffic.</p> <p>Impact: The SBC 5400 reporting congestion even when no calls active.</p> <p>Root Cause: On the 5400 server, the EMA has allocated 4 CPUs for java and there are times when these can all run hot even when no calls on the system. The minimum number of hot CPUs to trigger congestion on the 5400 was set to 4 so it could report without any calls.</p> <p>Steps to Replicate: Enable SIP ladder diagram and CDR viewer on check for congestion reports on an idle system.</p>	<p>The code is modified to avoid congestion when no calls. This is similar to the other 5xxx models where the hot CPUs is one more than the number of java CPUs in EMA.</p> <p>Workaround: Disable SIP ladder diagram and CDR viewer services to reduce the java CPU usage.</p>

9	SBX-105436 SBX-109736	2	<p>Portfix SBX-105436: The CDR issues for REGISTER.</p> <p>Impact: There were issues while writing CDRs for REGISTER method.</p> <p>Root Cause: The issues below were present w.r.t REGISTER CDRs:</p> <ol style="list-style-type: none"> CDRs were not generated for REGISTER in case of crankback scenarios. TG name was not updated properly for REGISTER CDRs. Disconnect reason was not updated properly for REGISTER CDRs. <p>Steps to Replicate:</p> <ol style="list-style-type: none"> REGISTER an endpoint through SBC with/without crankback. Check CDRs. Entries should be proper. 	<p>The code is modified for:</p> <ol style="list-style-type: none"> Generating EVENT records for REGISTER in case of crankback scenarios. Correcting the egress TG name in the EVENT CDR for REGISTER. Correcting the disconnect reason in the EVENT CDR for REGISTER. <p>Workaround: None.</p>
10	SBX-108198 SBX-110388	2	<p>Portfix SBX-108198: Coverity issues in nrsPktLifCsv.c</p> <p>Impact: There was a coverity fix that caused the issue.</p> <p>Root Cause: Fix added for the Coverity error.</p> <p>Steps to Replicate: The issue cannot be reproduced.</p>	<p>The code is modified to address the issue.</p> <p>Workaround: None.</p>
11	SBX-108058 SBX-110927	2	<p>Portfix SBX-108058: The SBC CpxAppProc memory leak.</p> <p>Impact: During the SBC startup processing there is a small memory leak.</p> <p>Root Cause: During the startup processing, the SBC is allocating memory while reading configuration information and it is not being freed up correctly at the end of the provisioning steps.</p> <p>Steps to Replicate: This issue was only observed while running with ASAN images in the engineering lab as the amount of memory leaked is small and cannot be checked.</p>	<p>The code is modified to correctly free the memory allocated while processing the configuration data.</p> <p>Workaround: None.</p>
12	SBX-105856 SBX-110160	2	<p>Portfix SBX-105856: The Wrong Version in URL after a Restore to an older version.</p> <p>Impact: The wrong version in URL after a Restore to an older version.</p> <p>Root Cause: URL is picked from CDB from path "/system/objectStoreParameters/uri" that is independent of software version of revision.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> Spawn OAM act, standby and one SSBC with V10.00.00A008. Create multiple revisions (e.g., rev1, 2, 3, 4, 5) Upgrade cluster to V10.00.00A009. <pre>[root@VSBCSYSTEM-2-vsbc2 172.12.34.567 evlog]# cat /mnt/gfsvol1/config-versions.txt 5,V10.00.00A008,oam_config_20210510T042937.tar.gz,0,2021-05-10T04:29:37,downloaded_from_ems 6,V10.00.00A009,oam_config_20210510T043100.tar.gz,1,2021-05-10T04:31:00,REBOOT_REQ,auto_save_after_software_change [root@VSBCSYSTEM-2-vsbc2 172.12.34.567 evlog]#</pre> <ol style="list-style-type: none"> Manually reboot and replace the SM Process with a fix. Restore revision 3. <pre>7,V10.00.00A008,oam_config_20210510T045034.tar.gz,6,2021-05-10T04:50:34,RESTORE,restored_from_3 8,V10.00.00A009,oam_config_20210510T045311.tar.gz,6,2021-05-10T04:53:11,REBOOT_REQ,auto_save_after_software_change</pre> <p>Observation: Revision 7 uploaded with correct URL.</p> <ol style="list-style-type: none"> Remove revision 7 from /mnt/gfsvol1 and restore revision 7. Revision successful and newly created revision uploaded on EMS with correct URL <pre>11,V10.00.00A008,oam_config_20210510T055545.tar.gz,20,2021-05-10T05:55:45,RESTORE,restored_from_9 12,V10.00.00A009,oam_config_20210510T055745.tar.gz,20,2021-05-10T05:57:45,REBOOT_REQ,auto_save_after_software_change</pre>	<p>The code is modified to update the URL with software version of revision being restored.</p> <p>Workaround: None.</p>
13	SBX-108356 SBX-110006	2	<p>PortFix SBX-108356: The SBC has various runtime errors found in np.log.</p> <p>Impact: Internal ASAN builds report runtime errors on SWE platform related to left shift of signed integers.</p> <p>Root Cause: Appropriate integer casts were missing in code, which caused ASAN runtime warnings related to bit shift.</p> <p>Steps to Replicate: Bring up ASAN build on SWE platform and observe if np.log contains ASAN runtime warnings related to bit shifts.</p>	<p>The code is modified to fix the runtime warnings in ASAN builds.</p> <p>Workaround: No workaround.</p>

14	SBX-102925 SBX-110143	2	<p>Portfix SBX-102925: The issues in Ova\Qcow2 installation</p> <p>Impact: The SBC SWe installed through the ISO and SBC SWe created through the image launch method have different prompts and root password.</p> <p>Root Cause: While installing through the image launch method, the SBC SWe was getting configured the cloud SBC way w.r.t root password and prompt.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Install the SBC SWe using the ISO. Check the root password and CLI prompt. 2. Instantiate the SBC SWe using Image launch. Check root password and CLI prompt. <p>Results from step 1 and step 2 should be same.</p>	<p>The code is modified to ensure both image launch method and ISO install method of the SBC SWe has no root password and same CLI prompt.</p> <p>Workaround: The user can manually remove root password on ISO installed SWEs using below command: usermod root -p ""</p>
15	SBX-110179 SBX-110217	2	<p>Portfix SBX-110179: The LeakSanitizer: detected memory leaks.</p> <p>Impact: A memory leak occurs when a logical management interface is added or modified.</p> <p>Root Cause: A confd cursor element was not closed when exiting the routine that validates the logical management interface being added or changed and this resulted in a memory leak.</p> <p>Steps to Replicate: Make changes to the logical management interfaces and check for memory increasing in the Cpx process.</p>	<p>The code is modified to ensure the associated memory is freed to avoid the leak.</p> <p>Workaround: None.</p>
16	SBX-105763 SBX-110036	2	<p>PortFix SBX-105763: Move the dmesg monitoring from SM to a platform cron job</p> <p>Impact: Move dmesg monitoring from SM to a platform cronjob</p> <p>Root Cause: Since the dmseg can be large for long-running SBCs (thus, take longer to dump logs), the function can take longer causing an SM healthcheck.</p> <p>Steps to Replicate: Check for "/var/log/sonus/tmp/dmesgErrorMarker.tmp" after manually running the script</p>	<p>The code is modified to run every minute as a cron job to find i/o and filesystem errors in dmesg. If error is found, it'll create a marker file in tmp, which can be later used by other script to check sanity of the system.</p> <p>Workaround: None</p>
17	SBX-94852 SBX-110154	2	<p>Portfix SBX-94852: Security, Audit and other logs modifiable (including deletion).</p> <p>Impact: Administrator users are able to delete or modify evlog files on the SBC.</p> <p>Root Cause: Users belonging to upload group(like admin) had write access on the evlog dir, which allows them to delete/modify log files.</p> <p>Steps to Replicate: Use SFTP to login to the evlog directory as user admin, and attempt to modify/delete log files.</p>	<p>The code is modified that prevents the admin from having writing access on the files owned by other users.</p> <p>Workaround: None.</p>
18	SBX-108435 SBX-109815	2	<p>Portfix SBX-108435: The SBC CpxApp Process dumps core in the standby OAM.</p> <p>Impact: The CPX core dump occasionally writes when the SBC is stopped.</p> <p>Root Cause: Replication Engine thread is not stopped when the CPX receives deactivate request.</p> <p>Steps to Replicate: It is rarely reproducible.</p>	<p>The code is modified to address the issue.</p> <p>Workaround: None.</p>
19	SBX-106543 SBX-110144	2	<p>Portfix SBX-106543: The SBC experiences a core dump while running UAS Notify Request.</p> <p>Impact: The SBC core dumps while processing the UAS Notify Request within the XML body.</p> <p>Root Cause: In API SipsCheckForAnyBody(), the loop where the string pointer pch is incremented each time was being accessed before validating for an out of bounds condition that caused the crash.</p> <p>Steps to Replicate: Send a NOTIFY with XML body from UAC towards the SBC.</p>	<p>The code is modified to add a out of bounds check before accessing the pointer value.</p> <p>Workaround: None.</p>
20	SBX-108370 SBX-110816	2	<p>Portfix SBX-108370: The top2 on the SBC core takes lot of CPU cycles.</p> <p>Impact: The SBC performance monitoring tools like top2 at times take 20% of a CPU core there by reducing total available CPU resources for management activities on SBC.</p> <p>Root Cause: Introduced nice values to sbxPerf commands but still top2 taking around 20-30% at times. As a result, need to check ways to optimize it to take less CPU using /root/.top2rc.</p> <p>Steps to Replicate: Install fix build and by using top command make sure CPU utilization of top2 should not be much CPU utilized.</p>	<p>The code is modified to take less CPU utilization of the SBC performance monitoring tools like top2.</p> <p>Workaround: None.</p>

21	SBX-109681 SBX-110923	<p>2 Portfix SBX-109681: Low MOS score for UNENCRYPTED_SRTP calls</p> <p>Impact: In SRTP for unencrypted, Authenticated case, SRTP packets were being discard at NP due to authentication key mismatch.</p> <p>Root Cause: In SRTP for unencrypted, authenticated combinations key size is required in NP to derive the session authentication keys. Since the cipher key size was not being pass to NP, session authentication key was wrongly calculated.</p> <p>Steps to Replicate: On the receipt of an SDP offer with crypto attributes, if 'enable SRTP' is Configured in packet service profile, a common crypto suite is found in the crypto attribute received in the offer and in the packet service profile, and session parameters if included are allowed, the offer will be accepted. In the answer, the SBC includes the same tag used in the offer.</p> <p>Test Setup on the SBC:</p> <ol style="list-style-type: none"> 1. Endpoint1 Configured with SRTP profile(SRTP/SHA-1-80). Disable all Session Parameters in packet service profile. Allow Fallback enabled. 2. Endpoint2 Configured with SRTP profile(SRTP/SHA-1-80).Enable Session Parameters UNENCRYPTED_SRTP in packet service profile. Allow Fallback enabled. <p>Procedure:</p> <p>Endpoint1 sends an offer SDP with crypto attribute SRTP/SHA-1-80 and with Session Parameters UNENCRYPTED_SRTP Endpoint2 replies with crypto attribute SRTP/SHA-1-80 and with Session Parameters UNENCRYPTED_SRTP</p>	<p>The code is modified so the SBC application is passing cipher key size also to NP to calculate session authentication keys.</p> <p>Workaround: Workaround is to not send unencrypted SRTP. Use cases with Encrypted SRTP and authentication no issue is seen.</p>
22	SBX-105900 SBX-110229	<p>2 Portfix SBX-105900: Resize log volume on every boot.</p> <p>Impact: The log volume is not being resized on every boot.</p> <p>Root Cause: Currently, we build the filesystem on cinder volume only on the first boot.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Create and attached a cinder volume of 50 GB. 2. Bring down the SBC, resized cinder volume to 100 GB. 3. Launch the SBC with the cinder volume attached. <p>Check if the cinder volume is resized to 100 GB.</p>	<p>Check if file system is already built. If the file system is not built, resize the volume to address the issue.</p> <p>Workaround: None.</p>
23	SBX-108190 SBX-110223	<p>2 Portfix SBX-108190: SBC: The callTracing does not work after reverting a switchover.</p> <p>Impact: You cannot enable the callTrace on calls after a switchover under the following circumstances: When maximum calltrace count is reached before the switchover, and all calls with calltrace enabled are terminated after the switchover.</p> <p>Root Cause: The internal calltrace state was not properly synchronized to the standby node that caused no new calls with calltrace ON request can be enabled.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. On a 1+1 (active/standby) SBC set up, set the calltrace count to 10. 2. Create 10 calls with calltrace enabled. 3. Perform a switchover. 4. Terminate the 10 calls after switchover. 5. Perform a switchover. 6. Make new calls with calltrace ON requested in the ADD termination command. <p>These new calls should have calltrace enabled.</p>	<p>The code is modified so that calltrace works after a a switchover.</p> <p>Workaround: Reboot both SBC nodes.</p>
24	SBX-109084 SBX-109214	<p>2 PortFix SBX-109084: The IPv6 SBC traps was not received to the EMS.</p> <p>Impact: For trapTargets, certain IPv6 addresses are sent to the incorrect IPv6 address.</p> <p>Root Cause: If the IPv6 address, converted to the form of 16 decimal octets separated by periods have a length greater than 47 digits, the address will be truncated.</p> <p>Steps to Replicate:</p> <p>Provision an OAM SNMP trapTarget with a customer IP.</p> <p>Observe the tailf snmp.log that the actual address sent to.</p>	<p>The code is modified so the buffer used to store the converted IPv6 string is stored in a buffer of 64 characters, which accommodates any IPV6 address.</p> <p>Workaround: None.</p>

25	SBX-109862 SBX-110925	2	<p>Portfix SBX-109862: The SBC is not rejecting the SRTP call when disallowSrtpStream is enabled.</p> <p>Impact: The SBC is not rejecting the SRTP call when disallowSrtpStream is enabled in ingress PSP and Invite is received with only SRTP stream.</p> <p>Root Cause: This is specific call flow when disallowSrtpStream is enabled and an INVITE is received with only SRTP stream. This scenario was not handled and the SBC was not rejecting call.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Feature flag "multipleAudioStreamsSupport" under trunk group media is enabled for both ingress and egress. 2. A disallowSrtpStream flag under trunk group should be enabled on the ingress side. 3. Make a call with two SRTP stream. 	<p>The code is modified so that when the disallowSrtpStream is enabled and an INVITE is received with only an SRTP stream, the call is rejected with 488.</p> <p>Workaround: None.</p>
26	SBX-109966 SBX-110950	2	<p>Portfix SBX-109966: The SBC incorrectly accepts an SDP offer in an ACK</p> <p>Impact: The SBC is designed to ignore an SDP offer in an ACK, but is not doing so.</p> <p>Root Cause: In an early media call, the SIP stack was accepting the SDP in ACK as an offer. This should never happen.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. In an early media call, add SDP in ACK. 2. Send a re-INVITE with SDP from the same side. 3. SIPS stack should ignore the SDP in an ACK and accept the SDP from re-INVITE as an offer. 	<p>The code is modified to not forward the SDP to the application.</p> <p>Workaround: None.</p>
27	SBX-109298 SBX-109569	2	<p>PortFix SBX-109298: SBC: The DSP fails to modify ptime.</p> <p>Impact: When the SBC receives a re-INVITE with a change in the ptime for EVS, the DSP Modify command fails.</p> <p>Root Cause: Handling of DSP Modify Command for ptime change for EVS was not present.</p> <p>Steps to Replicate: Run a EVS to G711 call, send a re-invite on the EVS leg with change in ptime. The change in ptime should be put into effect through a DSP Modify.</p>	<p>The code is modified to support for ptime change for EVS is added. Allowable ptime changes include 20, 40, 60, 80 and 100ms.</p> <p>Workaround: None.</p>
28	SBX-106127 SBX-110221	2	<p>Portfix SBX-106127: The SBC product name is incorrect in a 10.0 SBC.</p> <p>Impact: The sbcDiagnostic incorrectly prints product name as "ConnexIP5000" instead of "AWS".</p> <p>Root Cause: Platform type is determined by querying platform data using dmidecode command, due a bug in the query platform type is returned as unknown. As a result of this bug, the sbcDaignostic shows generic product name ("ConnexIP5000").</p> <p>Steps to Replicate: Run sbcDiagnostic command from the Linux shell of the SBC. It shows the following: <pre>*****SBC Information ***** SBC Product Name: ConnexIP5000</pre> After a fix, it prints the following: <pre>*****SBC Information ***** SBC Product Name: AWS</pre> </p>	<p>The code is modified to return correct platform type.</p> <p>Workaround: No workaround.</p>
29	SBX-109993 SBX-110007	2	<p>PortFix SBX-109993: The PRS Process is coring with a pkt switchover with a loopback call.</p> <p>Impact: The PRS Process core dumped when testing a pkt port switchover.</p> <p>Root Cause: The statement to log the debug message was missing a string parameter.</p> <p>Steps to Replicate: Force pkt port switchovers. This is intermittent so may be difficult to reproduce.</p>	<p>The code is modified to address the issue.</p> <p>Workaround: No workaround.</p>

30	SBX-108956 SBX-110136	2	<p>Portfix SBX-108956: SBC: Default bitrate for G722 codec should be 64kbps in SBC but it is 48 kbps.</p> <p>Impact: The SBC is using 48kbps bit rate for G722 codec instead of 64kbps when setting up G722 calls.</p> <p>Root Cause: The 48kbps bitrate is incorrectly chosen for G722 codec, resulted in the media packets being generated with 48kbps bitrate.</p> <p>Steps to Replicate: Make one g722 to g711 transcode call and observing the RTP stream for g722 codec.</p>	<p>Changed bitrate to 64kbps when setting up G722 calls to address the issue.</p> <p>Workaround: None.</p>
31	SBX-107975 SBX-110176	2	<p>Portfix SBX-107975: A Serf event processor is unable to restart because restart check marker file is not getting removed.</p> <p>Impact: The Serf event processor is unable to restart because the restart check marker file is not getting removed.</p> <p>Root Cause: The restart check marker was only being removed if serfeventProcessor starts successfully, so if it fails to start, any attempts to restart it would be prevented because the marker is not removed.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Make serf event processor fail at some point. 2. It should try to restart up to 5 times if it keeps failing to come up. 	<p>The code is modified so that the user can attempt a restart.</p> <p>Workaround: None.</p>
32	SBX-105391 SBX-110222	2	<p>Portfix SBX-105391: SBC: The SM process leaks memory on an OAM active for an SBC switchover.</p> <p>Impact: While carrying out operations like a configuration upload to EMS, the memory is leaked.</p> <p>Root Cause: The Python/C APIs cause a memory leak while using functions to the upload/download configuration from the EMS.</p> <p>Steps to Replicate: Create a direct single instance(ASAN build) registered with EMS, carry out operations saveAndActivate/restoreRevision and change glog /sanitizer_SmProcess* and verify no leaks due to above operation.</p>	<p>Change the implementation from using Python/C APIs to libcurl.</p> <p>Workaround: None.</p>
33	SBX-109167 SBX-110150	2	<p>Portfix SBX-109167: The AddressSanitizer: SEGV on unknown address 0x0000000000028.</p> <p>Impact: During the pre-parsing the Messagebody, the SEGV on an unknown address is observed in SipsPreParseMsgBody(). This could result in an SCM coredump.</p> <p>Root Cause: When the Content-Type is NULL the code was performing string comparisons to see if its an expected Content-Type and the code did not verify that the header pointer was not NULL before accessing it.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Make an A to B call. 2. Send a re-INVITE with the MessageBody Content-Type as empty. 	<p>The code is modified to check the Content-Type string pointer is not NULL before accessing it.</p> <p>Workaround: None.</p>
34	SBX-105644 SBX-110151	2	<p>Portfix SBX-105644: The sonusSbxCertificateExpireSoonWarningNotification trap does not display for all certificates in the current and history alarm lists.</p> <p>Impact: The sonusSbxCertificateExpireSoonWarningNotification trap does not display the individual alarms for separate certificates.</p> <p>Root Cause: The SBC alarm for sonusSbxCertificateExpireSoonWarningNotification was just using trap ID as key identifier.</p> <p>Steps to Replicate: Add multiple certificates to a 1:1 HA system that will expire and configure certificate expiry date. Confirm:</p> <ul style="list-style-type: none"> • Only one trap is triggered per certificate. • One alarm exists for certificate. 	<p>The code is modified to use certificate name as well as trap ID as key identifiers to show the alarm.</p> <p>Workaround: None.</p>
35	SBX-109407 SBX-110061	2	<p>PortFix SBX-109407: MicroFlows are failing in the REGISTER Performance of 2400 RPS (256000 REGISTRATIONS) with an error code 0xf9.</p> <p>Impact: Unable to support the 256k concurrent subscriber registrations in Default memory profile for the SBC SWe.</p> <p>Root Cause: Existing logic for determination of flow hash causes increased number of hash collisions leading to increased depth of hash buckets, leading buffer starvation.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Spawn a SBC SWe instance (of default memory config profile). 2. Test the 256K registrations. 	<p>The code is modified to minimize hash collisions.</p> <p>Workaround: There is no workaround for this.</p>

36	SBX-109424 SBX-109716	2	<p>Portfix SBX-109424: The SBC sends a 420 Bad Extension response when an INVITE with both supported and require 100rel is sent.</p> <p>Impact: The SBC sends 420 Bad extension when Require:100Rel is received in initial Invite and e2e Prack flag is enabled</p> <p>Root Cause: Incorrect code was added to reject an INVITE with Require:100Rel when rel100Support flag is disabled and e2e prack flag is enabled.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. The rel100Support must be disabled. 2. The e2e Prack is enabled. 3. An INVITE is received with a Require:100Rel Header. 	<p>The code is modified for rejection in the require:100rel scenarios.</p> <p>Workaround: Use SMM to remove Require:100Rel Header.</p>
37	SBX-105437 SBX-109737	2	<p>Portfix SBX-105437: The CDR issues for non-INVITE messages when blacklisting is involved.</p> <p>Impact: In case of handling of failure responses for non-INVITE messages, before writing the CDR for a current failure cause code, the SBC was finding out next route and sending a message on network.</p> <p>This worked fine in normal cases as after sending out a request, the response was processed later, after writing a CDR for current failure response.</p> <p>However in a blacklisted entry case, no actual message is sent out, so the blacklisted entry CDR was written before the previous CDR response code.</p> <p>Root Cause: Whenever a blacklisted entry was involved, the CDR entries were in accurate for this blacklisted entry and the previous entry.</p> <p>Steps to Replicate: Configure Routes for non-INVITE as follows: R1 R2 -> Blacklisted R3 R4 -> Blacklisted</p> <p>The CDR's should be printed in order R1, R2, R3, R4 after a fix.</p>	<p>The code is modified to write the CDR for the current failure response code, and later find next route and send a request on the new route.</p> <p>Workaround: None.</p>
38	SBX-105175 SBX-108184	2	<p>PortFix SBX-105175: The SBC sends a re-INVITE while media is played on the ingress side in a GW-GW early media call.</p> <p>Impact: In a GW-GW call scenario, while the media is played on the ingress Gateway, the egress SBC is sending a re-INVITE to the UAS.</p> <p>Root Cause: Before the ingress GW completes its end to end activation, it received a MODIFY OFFER request from the egress GW due to the change in SDP received in 200 OK for lockdown INVITE. This caused the ingress GW to process modify offer first and then end to end activation later that triggered a re-INVITE towards UAS.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Send INV from UAC to GW1. 2. GW2 sends INV to UAS. 3. UAS sends 180 SDP. 4. GW1 sends 180 SDP to UAC and starts to play tone. 5. UAS sends 200 OK with out SDP. 6. GW1 sends 200 OK to UAC. 7. GW2 sends ACK and triggers a lock down INVITE. 8. UAS sends 200 OK with change in SDP. 	<p>The code is modified so that while modify offer-answer is handled properly during end to end activation.</p> <p>Workaround: None.</p>
39	SBX-109418 SBX-110173	2	<p>Portfix SBX-109418: The LeakSanitizer detected memory leaks Direct leak of 883820 byte(s) in 413 object(s).</p> <p>Impact: The SBC was not freeing memory in one of the failure cases.</p> <p>Root Cause: The SBC was not freeing memory in few cases where incoming INVITE handling fails.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Send an INVITE with Replaces having call id, to in the Replaces header that does not match to any existing leg. 2. After the call ends, the ASAN should not show memory leak. 	<p>The code is modified to free up memory allocated, in all cases when the INVITE handling fails.</p> <p>Workaround: None.</p>

40	SBX-107960 SBX-109734	2	<p>Portfix SBX-107960: The AWS IPs remain assigned to the standby SBC. Unnecessary dual restarts.</p> <p>Impact: If communication between active and standby is broken (over ha0 interface) both assume active roles (split brain). When this happens, the standby node that becomes active calls AWS APIs to move IP address to self. Once ha0 link is restored, the machine that becomes active does not call AWS API to move IP addresses, and this might cause issue when node that has IP does not come up as active. In this case, IPs are assigned on standby and another node becomes active.</p> <p>Root Cause: Root cause of this issue is not calling AWS switchover API (move IPs) during split brain recovery.</p> <p>Steps to Replicate: Perform a Split brain test and recovery of the SBC HA, and verify that API query is send by an Active SBC after recovery.</p>	<p>The code is modified to call AWS APIs to move IP to current active machine during split brain recovery path.</p> <p>Workaround: Manually move all secondary IPs to current active machine to restore calls.</p>
41	SBX-108434 SBX-109170	2	<p>PortFix SBX-108434: SBC: A Standby OAM fails to come up after the restart of active and standby.</p> <p>Impact: After restarting the active and standby or after a fault that causes the SBCs to go down, the standby OAM waits for active to come up first and never recovers if active is down.</p> <p>Root Cause: On the Standby OAM startup sequence, there is a makeSureActivesUp() loop that exits only after active is up. This results in standby to be in a hung state forever if active is down.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Start active OAM SBC. 2. Start the standby OAM SBC after 1 minute. 3. Stop the active OAM SBC before it becomes active. 4. The Standby OAM SBC should reboot and recover after 10 minutes. 	<p>The code is modified to reboot the instance that is starting as standby, if peer does not become active in 10 minutes.</p> <p>Workaround: Reboot the standby OAM SBC to fix the issue.</p>
42	SBX-109017 SBX-109198	2	<p>PortFix SBX-109017: SBC: Observed MAJOR logs for BRM "BrmAsyncCmdErrHdlr" on T140 load.</p> <p>Impact: Occasionally enable-RID errors are seen in NP when the system is subjected to large number of call flows that employ the RTCP termination.</p> <p>Root Cause: The message drops in internal queues due to momentary congestion.</p> <p>Steps to Replicate: Run an SBC AMRWB<->T140 full load with the RTCP enabled.</p>	<p>The code is modified to ensure the control messages experience a larger queue depth while traversing internal queues.</p> <p>Workaround: No workaround.</p>
43	SBX-107646 SBX-110230	2	<p>Portfix SBX-107646: For a revision not present in the OAM or EMS, upon doing a restore, an error should be thrown.</p> <p>Impact: Not showing the proper failed message when failing to download from the EMS.</p> <p>Root Cause: The confd was not waiting for the EMS download error as the EMS download was done in a different thread context.</p> <p>Steps to Replicate: Case 1:</p> <ol style="list-style-type: none"> 1. Create 3 revisions on the EMS. 2. Delete the second revision on the OAM and EMS. 3. Request the system admin vsbcSystem restoreRevision revision 2. o/p admin@Rahul-OAM1-192.168.20.13% request system admin vsbcSystem restoreRevision revision 2. This command will restart all nodes unless the target revision is for a previous version of software. Do you want to continue [yes,no] yes result failure reason bundle not found locally or on EMS, unable to view changes for this revision. 4. See the above failure message and no restart of nodes. <p>Case 2:</p> <ol style="list-style-type: none"> 1. Create 3 revisions on the EMS. 2. Delete the second revision on the OAM. 3. Request the system admin vsbcSystem restoreRevision revision 2. 4. Will observe a restart in all nodes. <p>Case 3:</p> <ol style="list-style-type: none"> 1. Create 3 revisions on the EMS. 2. Delete the second revision on the EMS. 3. Request the system admin vsbcSystem restoreRevision revision 2. 4. Will observe a restart in all nodes. <p>Any other failure case.</p>	<p>The code is modified to run the restoreRevision procedure on the same thread context. This helps to display proper error in case of failure.</p> <p>Workaround: None.</p>

44	SBX-110178 SBX-110919	2	<p>Portfix SBX-110178: The PRS Process gave heap use after free on address on latest 10.0 build.</p> <p>Impact: The PRS Process gave a "heap use after free on address" error while running the HA suite on ASAN build.</p> <p>Root Cause: Interface number was being returned after typecasting the main structure to packet LIF structure, and not management LIF structure.</p> <p>Steps to Replicate: Run SBX_504_HA suite on ASAN build.</p>	<p>The code is modified to correct the typecasting.</p> <p>Workaround: None.</p>
45	SBX-110128 SBX-110921	2	<p>Portfix SBX-110128: AWS CFN: The SBC auto-registration in the EMS is not working with EMS FQDN.</p> <p>Impact: The SBC auto-registration in the EMS is not working when using EMS FQDN.</p> <p>Root Cause: A service discovery is unable to resolve the EMS FQDN using a system DNS because there is no ACL rule to allow DNS query to go through.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Configure the SBC with EMS FQDN. 2. Ensure that it is registering to EMS successfully. 	<p>The code is modified to add an ACL rule to allow DNS query to go to the configured nameserver.</p> <p>Workaround: None.</p>
46	SBX-104319 SBX-110138	2	<p>Portfix SBX-104319: There is a conflicted design between two features "sdpRelayAttribute" and "Send RTCP Bandwidth Info".</p> <p>Impact: The SBC sends duplicate b=RR and b=RS attributes when "sdpRelayAttribute" and "Send RTCP Bandwidth Info" are enabled together.</p> <p>Root Cause: The SBC skips adding RR/RS to SDP due to "sendRTCPBandwidthInfo" only when "sdpRelayAttribute" is enabled with TFT.</p> <p>Due to this, b=RR and b=RS attributes are added twice in outgoing SDP, when "sdpRelayAttribute" was enabled without TFT.</p> <p>Steps to Replicate: Configuration:</p> <ol style="list-style-type: none"> 1. Enabled sdpAttributesSelectiveRelay flag under TG. set addressContext default zone ZONE_V6 sipTrunkGroup TRUNK_V6 media sdpAttributesSelectiveRelay enabled set addressContext default zone ZONE_V4 sipTrunkGroup TRUNK_V4 media sdpAttributesSelectiveRelay enabled 2. Enabled RTCP with below settings. set profiles media packetServiceProfile DEFAULT rtpOptions rtpc enable rrBandwidth 500 rsBandwidth 150 <p>Procedure:</p> <ol style="list-style-type: none"> 1. Make a Basic call such the the bandwidth parameters are received from the endpoint as b=RR:1125, b=RS:775 <p>Expected Result:</p> <ol style="list-style-type: none"> 1. The SBC should transparently send the b=RR & b=RS parameters in SDP and should not honor the bandwidth configured in the PSP. <p>Also, the SBC should not add duplicate b=RR & b=RS parameters in SDP honoring the configured bandwidth values in PSP.</p>	<p>If "sdpRelayAttribute" is enabled without TFT, and "sendRTCPBandwidthInfo" is also enabled, the SBC does not add the RR/RS due to "sendRTCPBandwidthInfo" since attributes are relayed.</p> <p>Workaround: None.</p>

47	SBX-110201 SBX-110383	2	<p>Portfix SBX-110201: A late offer call resulting in the SBC not sending DTMF for AMR-WB in initial offer towards ingress.</p> <p>Impact: The SBC is not sending 16k 2833 Payload type in the initial offer towards ingress during a Late media "convert" call.</p> <p>Root Cause: Answer received from egress contained both 8k and 16k 2833 Payload type and that resulted in the SBC incorrectly assigning the 8k PT value to 16k as well while generating offer towards ingress. As a result, the 16k PT get dropped by SIP stack.</p> <p>Steps to Replicate: Configuration:</p> <p>Transcode conditional rel100Support enabled honorSdpClockRate enabled DLRBT enabled Configure multiple codecs on both Routes</p> <p>To re-create the issue:</p> <ol style="list-style-type: none"> 1. The UE initiates Late media convert call. 2. The SBC sends out 180 answer with the below SDP: <pre>m=audio 1024 RTP/AVP 96 8 97 0 18 9 101 a=rtpmap:96 AMR-WB/16000 a=fmtp:96 mode-set=0,1,2; mode-change-capability=2; max-red=0 a=rtpmap:8 PCMA/8000 a=rtpmap:97 AMR-WB/16000 a=fmtp:97 mode-change-capability=2; max-red=0 a=rtpmap:0 PCMU/8000 a=rtpmap:18 G729/8000 a=fmtp:18 annexb=no a=rtpmap:9 G722/8000 a=rtpmap:101 telephone-event/8000 a=fmtp:101 0-15</pre> <p>Test Result without a fix:</p> <p>The SBC drops the telephone-event/16000 payload type when generating offer towards ingress in 180.</p>	<p>The code is modified to prevent the same PT value getting assigned to both 8k and 16k DTMF in this call flow.</p> <p>Workaround: None.</p>
48	SBX-106693 SBX-110228	2	<p>Portfix SBX-106693: SBC: There was a wrong warning message on the Direct-SBC while doing restore.</p> <p>Impact: The wrong warning message was displayed during a restoreRevision.</p> <p>Root Cause: Rewording of the text message required for when restorerevision was invoked.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. A new image was created after changes. 2. The new csar file was loaded to VNF Manager. 3. Created revision 2 on the EMS. 4. Run the command "request system admin vsbcSystem restoreRevision revision 1". <pre>test output admin@Rah-OAM1-192.168.20.3% request system admin vsbcSystem restoreRevision revision 1</pre> <p>This command will restart all nodes unless the target revision is for a previous version of software. Do you want to continue [yes,no]</p>	<p>The code is modified by rewording the warning text message.</p> <p>Workaround: None.</p>
49	SBX-110362 SBX-110363	2	<p>PortFix SBX-110362: The SBC generates RTP inactivity alarms when the policer mode is set to bypass.</p> <p>Impact: The SBC generates false RTP inactivity alarms when the policer mode is set to bypass.</p> <p>Root Cause: A bug in the media policing logic of SWe_NP code skips updating the timestamp of last RTP packet of a media flow, resulting in raising RTP inactivity notifications for every 10 seconds.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Configure a ISBC instance. 2. Set the policer mode to bypass (Unit tested by hacking the code). 3. Configure RTP inactivity. 4. Run the call for more than the rtp inactivity duration. 5. Observe that RTP inactivity alarms being raised, even if traffic flow is happening. 	<p>The code is modified to update the timestamp after every arrival of RTP packet for the media flow.</p> <p>Workaround: Ensure the policer mode is set values to other than 'bypass'.</p>

50	SBX-110054 SBX-110432	2	<p>Portfix SBX-110054: The encrypted packets not being sent towards racoon in case of IPsec tunnel mode.</p> <p>Impact: The lawful Interception fails when IPsec is enabled for LI media over UDP.</p> <p>Root Cause: An earlier fix assumed the LIG information to be stored in the selector structure to send notification to XRM, indicating if a IPsec tunnel is for LI media.</p> <p>The field was LIG was present in the selector but was never updated.</p> <p>This resulted in failure to send notification to XRM to indicate IPsec tunnel for LI media, for it always failed the match with the LIG configured under the CDC and as a result, the LI failed.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Generate Local and remote Certificates which has FQDN as a SAN value and signed by an Intermediate CA. 2. Copy Intermediate CA's and Racoon's .crt files to /opt/sonus/external on SBC. 3. Covert Remote Files to .der and local files to .p12. 4. Configure IPSEC related configuration on SBC. 5. Execute a basic call over IPSEC interface. <p>Expected Results:</p> <ol style="list-style-type: none"> 1. CA, Local(SBC) and Remote(Racoon) certificate generated successfully. 2. Copy of certificate files to SBC is successful. 3. Conversion of remote and local certificate is successful. 4. All configuration issuccessful. 5. IMS LI is configured successfully. 6. Basic call must run successfully over ipsec enabled interface. Signaling and media are captured on IMS LI 	<p>The LIG check is removed to ensure that the notification is sent to XRM to indicate the IPSEC Tunnel is for LI media.</p> <p>Workaround: No workaround.</p>
51	SBX-102618 SBX-109302	2	<p>PortFix SBX-102618: The SBC needs to handle the scenario where member-join event handling is missed at the serf level.</p> <p>Impact: A race condition in RGM event handling can show sync to be inProgress in N:1.</p> <p>Root Cause: A race condition in handling member-join event causes some member-join events to get discarded at serf level.</p> <p>Steps to Replicate: Perform multiple sbxrestarts but it is not readily reproducible.</p>	<p>The code is modified to ensure the Redundancy Group is up-to-date.</p> <p>Workaround: None.</p>
52	SBX-106589 SBX-110493	2	<p>Portfix SBX-106589: The SBC SWe does not forward the SIP Info and drops the call.</p> <p>Impact: After receiving 32763 SIP indialog INFO messages, the SBC sends a 500 error for further INFO messages instead of forwarding the message.</p> <p>Root Cause: The SBC maintains the callHandle structure with refcount, which will be incremented for every reference and decremented once reference is released. But in this case, the refcount is not decremented properly, for every indialog msg (INFO) request received, the count kept increased.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Establish a call. 2. Send indialog INFO messages more than 32763. 	<p>The code is modified to address the issue.</p> <p>Workaround: No workaround.</p>
53	SBX-106897 SBX-110224	2	<p>Portfix SBX-106897: The SBC does not send notify for successful trace activation.</p> <p>Impact: The SBC does not send a calltrace notify to C3 in a call setup with calltrace ON request for ADD commands of both first termination and second termination of the call, but the ADD command failed for the second termination.</p> <p>Root Cause: The SBC code logic was to send a calltrace NOTIFY after the ADD commands of both termination have been processed, but this result in calltrace NOTIFY not sent for the first termination in failure of ADD in the second termination.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Enable calltrace in the SBC. 2. Send ADD request with CALLTRACE/TRACEACTIVITYREQUEST=ON on both the terminations. The codec should be PCMU on the first termination and G726 on the second termination, such that ADD of second termination will fail (not supported codec). <p>Expect Result:</p> <ol style="list-style-type: none"> 1. The call fails. 2. The SBC sends calltrace NOTIFY with RES=success for the first termination. 	<p>The code is modified to send calltrace NOTIFY for successfully ADDED termination.</p> <p>Workaround: None.</p>

54	SBX-109439 SBX-110914	2	<p>Portfix SBX-109439: SBC: An activeRevision failure when state for even type audit is set to on/off.</p> <p>Impact: The configuration Playback on managed VMs fails on command.</p> <pre>set oam eventLog filterAdmin vsbc1 audit audit level major state off</pre> <p>Root Cause: The playback engine does not use user context.</p> <p>Steps to Replicate: Run the command on the OAM.</p> <pre>set oam eventLog filterAdmin vsbc1 audit audit level major state off commit</pre> <p>Perform a saveAndActivate.</p>	<p>The code is modified to address the issue.</p> <p>Workaround: Reboot on managed VMs to get the configuration at startup.</p>
55	SBX-109412 SBX-110910	2	<p>Portfix SBX-109412: The destination buffer was too small during the snprintf function.</p> <p>Impact: During the snprintf function call, the destination buffer was not big enough to copy the source string.</p> <p>Root Cause: The size of the destination buffer was smaller than the source buffer because the destination buffer length check was not present.</p> <p>The buffer file size should be 256 bytes. File: call/sonus/p4/ws/jenkinsbuild/sbxAsan100/marlin/SIPS/sipsParse.c</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. When the callld length was more then the defined max callld length this error was seen. 2. Run a call with callld length greater than 256 bytes. This error should not come. 	<p>The code is modified to address the issue.</p> <p>Workaround: None.</p>
56	SBX-109453 SBX-110622	2	<p>Portfix SBX-109453: The SBC is closing the socket for DNS query over TCP frequently.</p> <p>Impact: The SBC is closing the socket for DNS query over TCP frequently.</p> <p>Root Cause: Whenever the TCP connection is closed, the FIN packet is sent towards the SBC from DNS server. But the application was not closing the connection by sending the FIN and Even after a connection is closed by DNS server, the application is still holding socket information. This details were used in future queries has cause TCP connection to reset.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Have a v6 Cloud SBC with the below build SBC: V10.00.00-A006. 2. Configure the DNS server to send queries using TCP connection. 3. Make a call. 4. Wait for some time and run one more call here. <p>Actual behavior:</p> <p>In step 3, the SBC will query DNS server to resolve NAPTR records over TCP connection.</p> <p>In step 4, the SBC will try to send a NAPTR query and then immediately close TCP connection.</p> <p>Expected result:</p> <p>The NAPTR query should be successful in step.</p>	<p>The code is modified to close TCP connection (Sending FIN to close connection). Also, the connection information in application is freed.</p> <p>Workaround: No workaround.</p>
57	SBX-106520 SBX-106881	2	<p>PortFix SBX-106520: Among the exported setting values, there were setting values that were not set in other environments.</p> <p>Impact: The "comment" parameter is missing from some configuration related to AAA rules/cmdRules.</p> <p>Root Cause: The "comment" parameter is not being restored during a LSWU.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Spawned a SBC, exported the configuration and observed comment parameters are present. 2. Upgraded the SBC software. exported the configuration and observed comment parameters are not present. <p>Revised the steps above with fix, and the issue was not observed.</p>	<p>The code is modified to restore "comment" parameter during a LSWU.</p> <p>Workaround: None. The "comment" parameter is used only for description. It doesn't have functional impact.</p>

58	SBX-106869 SBX-110296	2	<p>Portfix SBX-106869: The SBC Node branch information is inconsistent after a Cleanstart/Restore Revision is performed in OAM's.</p> <p>Impact: The SBC Node branch information is inconsistent after the Cleanstart /Restore revision is performed for OAMs.</p> <p>Root Cause: Managed VMs registration failed with active and standby OAMs as OAMs were not up. VMs did not retry to register again.</p> <p>Steps to Replicate: After the issue is fixed, create a setup OAM and S/M/T SBC.</p> <ol style="list-style-type: none"> 1. Log in to the CLI as admin and execute the following: show configuration node o/p: All nodes should be listed. 2. Configure the SBC (create AddressContext, zone, TG) saveAndActivate Revision o/p: new configRevision should be propagated to all managed VMs. 	<p>The code is modified to keep retrying to register with active and standby OAMs.</p> <p>Workaround: None.</p>
59	SBX-108112 SBX-109576	2	<p>PortFix SBX-108112: MS TEAMS - The Media Stats are not correct when the DLRBT profile is removed from TGs and ICE is enabled.</p> <p>Impact: On an SBC was configured for MS Teams LMO centralized mode with ICE, if DLRBT is not correctly configured, this can lead to media not flowing correctly for a call even after ice learning gets completed.</p> <p>Root Cause: The early ICE learning logic for non DLRBT mode in the SBC was not fully deactivating ICE media resources on receipt of 200 OK . As a result, the resources were unable to be re-activated for end to end media flow.</p> <p>Steps to Replicate: On an SBC configured as MS Teams LMO centralized mode with ICE on MS Teams TG:</p> <ol style="list-style-type: none"> 1. Disable DLRBT on PSTN and MS Teams TG's. 2. Establish a PSTN endpoint to MS Teams call that uses primary (Internal) LIF address towards Teams. 3. Once call has established and ice learning has completed, send media (voice) in both directions between PSTN and Teams endpoints. 4. Media should flow as expected in both directions and call Media status should correctly show the number of media packets sent and received for ingress and egress. 	<p>The code is modified to correctly deactivate early ice learning media resources on receiving an SDP from MS Teams.</p> <p>Workaround: The DLRBT should be enabled.</p>
60	SBX-109220 SBX-110005	2	<p>PortFix SBX-109220: The monitor.c "runtime error: signed integer overflow: 0 - -9223372036854775808 cannot be represented in type 'long int'" error in the np.log.</p> <p>Impact: Internal ASAN builds report runtime signed integer overflow error in SWe for standard deviation metric for jitter meant for consumption by Ribbon Protect.</p> <p>Root Cause: Existing algorithm of standard deviation metric did not handle integer overflow issues.</p> <p>Steps to Replicate: Stream the custom RTP packet where expected standard deviation and verify via Ribbon Protect metric that computed standard deviation values are in range.</p>	<p>The code is modified to handle the overflow.</p> <p>Workaround: No workaround. This metric is only exposed to Ribbon Protect. This is not exposed as a metric in any SBC display or CDR.</p>
61	SBX-70800 SBX-109606	2	<p>PortFix SBX-70800: AWS:Observing that metaVariable table is getting modified on the loading the backup configuration file of one instance in other instance.</p> <p>Impact: When loading the backup configuration from one SBC instance to another, the metavar table is getting populated with the list of metavars from both instances. This by itself does not cause any issues so long as the metavars are unique, it is incorrect to see the metavars for another instance in the table.</p> <p>Root Cause: The code was not removing the existing metavars prior to loading the configuration of another instance. This meant the metavar table had both sets of information.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Create a backup configuration of one cloud instance (instance1). 2. Loaded the backup configuration file into a new cloud instance (instance2). 3. Check the metaVariables table and see if it contains the metavars from both instances. 	<p>The code is modified to flush the metavars for the existing instance prior to loading the configuration from another instance.</p> <p>Workaround: None.</p>
62	SBX-108577 SBX-109166	2	<p>PortFix SBX-108577: The CE_2N_Comp_SmProcess==11649==ERROR: AddressSanitizer: SEGV on an unknown address 0x000000000000.</p> <p>Impact: The SBC was performing a write operation on one of the un-allocated memory space while restoring NTP server configuration. As a result, SEGV on unknown address was reported.</p> <p>Root Cause: This issue was caused because a flag variable was not initialized. As a result, if condition was evaluated true instead of false, a write operation would be performed.</p> <p>Steps to Replicate: Configuring the NTP server and restoring the configuration by switchover or restart this error should not come up.</p>	<p>The code is modified to address the issue.</p> <p>Workaround: None.</p>

63	SBX-106759 SBX-110175	2	<p>Portfix SBX-106759: In a N:1 mode, the SBC performs a switchover for a different node than the switchover request is honored for.</p> <p>Impact: In N:1 during few scenarios, SBC switchovers to node that was not request honored. When the issue occurred, a switchover will be processed to other node instead of processing to request honored node.</p> <p>Root Cause: During switchover process SBC was not checking the node which was requested honored. It was switching over to a node that comes first/is available while processing in the SBC.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Bring up a 2:1 SBC. 2. Try to reboot the both active server. 3. While switchover process delay the up of request honored node and allow the another node to come up first. 	<p>The code is modified to check the service id of the request honored node before a switchover.</p> <p>Workaround: None.</p>
64	SBX-106613 SBX-110297	2	<p>PortFix SBX-106613: The SBC adds duplicate header on new INVITE upon 422 response, when a transparency profile is attached to egress TG.</p> <p>Impact: The SBC adds a Duplicate Header on a new INVITE after the 422 Response is received.</p> <p>Root Cause: The SIP Stack adds a semi-known twice (Header is unknown to SIP Parser but configured in transparency Profile) in this 422 scenario.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Header has to be unknown to the parser. 2. The same header name needs to be configured in the Transparency profile. 	<p>The code is modified to Skip formatting the second instance of same header based on header type.</p> <p>Workaround: SMM can be added to remove the duplicate header.</p>
65	SBX-110246 SBX-111002	2	<p>PortFix SBX-110246: The OAM and SBC nodes would not start after revert to 9.1 from 10.0.</p> <p>Impact: The SBC node fails to start after reverting from 10.0 to 9.1.</p> <p>Root Cause: The start-up failure after the SBC node revert from 10.0.0 to 9.1 is due to incorrect permissions of logs in /var/log/postgresql/ directory.</p> <p>Steps to Replicate: Perform revert from 10.0.0 to 9.1 on the SBC node and verify the SBC comes up fine post revert.</p>	<p>The code is modified to update the right permissions on files in /var/log/postgresql/ directory after revert.</p> <p>Workaround:</p>
66	SBX-107973 SBX-110017	2	<p>Portfix SBX-107973: The SBC adding RR and RS attributes twice in the egress INVITE when multiple m lines present in SDP.</p> <p>Impact: The SBC was adding RR and RS attributes twice in the egress INVITE when multiple m lines present in a SDP.</p> <p>Root Cause: The SBC is not setting the default value of M lines present in SDP when number of lines is greater than 1 and as a result, the SIPS value is getting added.</p> <p>Steps to Replicate: Configure these:</p> <pre>set profiles media packetServiceProfile DEFAULT rtcpOptions rtcp enable set profiles signaling ipSignalingProfile DEFAULT_SIP commonIpAttributes flags sendRTCPBandwidthInfo enable set profiles signaling ipSignalingProfile DEFAULT_SIP commonIpAttributes flags sendRtcpPortInSdp enable comm set addressContext default zone ZONE_EGRESS sipTrunkGroup TG__EGRESS media multipleAudioStreamsSupport enabled set addressContext default zone ZONE_INGRESS sipTrunkGroup TG_INGRESS media multipleAudioStreamsSupport enabled comm</pre> <p>An incoming INVITE has multiple Audio line with:</p> <pre>b=RS:250 b=RR:250</pre>	<p>The code is modified so that if value is set to default and it is not getting modified, the SIP does not send the RR and RS value twice.</p> <p>Workaround: When the sdpAttributesSelectiveRelay is enabled, the SBC does not send RR and Rs twice.</p>
67	SBX-108008 SBX-110137	2	<p>Portfix SBX-108008: SBC: Observed MAJOR logs with MegacoSendAmmsResp failure after a switchover during an EVS and T140 -> Mulaw Load.</p> <p>Impact: There was a MegacoSendAmmsResp: MegacoSendAmmsResp: Failure while encoding the response, and Error code = 197 Major logs observed after a switchover of the SBC during a load run of call setup with audio and text stream.</p> <p>Root Cause: The text stream is not properly synchronized to the standby node and triggers the h248 message encoding error in reply to a H248 MODIFY command.</p> <p>Steps to Replicate: Perform failover during call load run with EVS and ToIP on ingress side and the G711U+Baudot on egress side.</p>	<p>The code is modified to address the issue.</p> <p>Workaround: None.</p>

68	SBX-107142 SBX-108456	2	<p>PortFix SBX-107142: The SBC VM was unable to power on a post power off procedure.</p> <p>Impact: The VMware GPU SWe instance does not come up post reboot or GPU is not visible in the instance.</p> <p>Root Cause: When the SWe instance runs in GPU mode, the persistence mode of the NVIDIA driver is enabled using the nvidia-smi to prevent the GPU state from being unloaded. This is a kernel-level solution and creating problem when the hypervisor is VMware.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Set sweActiveProfile to use GPU. 2. Reboot the instance after the SBC application is up. <p>RESULT:</p> <p>Instance should come up and GPU should be visible in the instance.</p>	<p>The code is modified to address the issue</p> <p>Workaround: No workaround.</p>
69	SBX-105457 SBX-110171	2	<p>Portfix SBX-105457: An error is thrown on the EMA while configuring the SMM Profile having messageBody criteria with regex.</p> <p>Impact: An error is thrown on the EMA while configuring the SMM Profile having messageBody criteria with a regex.</p> <p>Root Cause: The EMA assumes that the Num Match was a mandatory field but it is an optional field (user may enter that value or he may not).</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Log into the EMA. 2. Navigate profile -> signaling -> sip adaptor profile 3. Create the profile that should have a Num match value and after that delete the Num Match value from the CLI. 4. As a result, you cannot see the issue. 	<p>The code is modified to make the Num Match field optional.</p> <p>Workaround: None.</p>
70	SBX-107639 SBX-110135	2	<p>Portfix SBX-107639: The CA CMR with offset 2 and priority LOW is not being processed or rejected.</p> <p>Impact: When a Channel Aware (CA) Mode CMR for priority LOW and offset 2 was the first CA CMR received in a call, the CMR was not processed.</p> <p>Root Cause: The present default value of curFeclndOffset in the code was set to 0 that corresponds to offset 2 and priority LOW.</p> <p>Due to this default value when a CMR for offset 2 and priority low is received as the first CA CMR, it is not getting processed.</p> <p>Steps to Replicate: Test 1:</p> <ol style="list-style-type: none"> 1. Enable IR9.2cmr in the SBC. 2. Send ADD request with A1 profile having ch-aw-recv=-1 parameter in T1 termination. 3. Pump pcap with 13.2br with cmr byte of CA-L-O2 followed by CA-H-07 <p>Expected Result:</p> <ol style="list-style-type: none"> 1. The SBC should accept the ADD request. 2. The SBC should not accept the cmr request. InvalidCMRCount should be incremented <p>Actual Result:</p> <ol style="list-style-type: none"> 1. The call is successful. 2. invalidCMRCount gets incremented as per the number of invalid CMRs received <p>Test 2:</p> <ol style="list-style-type: none"> 1. Enable IR9.2cmr in the SBC. 2. Send ADD request with A1 profile having ch-aw-recv=0 parameter in T1 termination. 3. Pump pcap with 13.2br with cmr byte of CA-L-O2 followed by CA-L-03 <p>Expected Result:</p> <ol style="list-style-type: none"> 1. The SBC should accept the ADD request. 2. The SBC should accept the cmr requests. codecModeChangeRxCnt should be incremented to 2. 	<p>The code is modified so that the default value does not match any of the valid curFeclndOffset that is 0-7.</p> <p>Workaround: None</p>

71	SBX-108232 SBX-110929	2	<p>Portfix SBX-108232: SBC: The AddressSanitizer: detected a heap-use-after-free on address 0x6180000a5a8 at pc 0x559e8989c4ac bp 0x7f4411fde290 sp 0x7f4411fde288 READ of size 8 at 0x6180000a5a8 thread T9.</p> <p>Impact: While the SBC node is shutting down it can access memory after its been freed, this could result in unexpected behaviour and in the worst case a core dump. But would have limited impact as it only happens when shutting down.</p> <p>Root Cause: During a sbxstop/sbxrestart or switchover because of race-condition, when the SBC is in deactivation the oamNodeRegisterRetry can access already deallocated resource leading to a core dump.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Setup build with HA SBC using OAM. 2. Perform a sbxrestart/switchover of active instance. 	<p>The code is modified to handle this race condition.</p> <p>Workaround: None.</p>
72	SBX-109443 SBX-109675	2	<p>PortFix SBX-109443: ERROR: The AddressSanitizer: detected heap-use-after-free on address 0x61900412bbce at pc 0x55d0b8c48037 bp 0x7fb50a457b00 sp 0x7fb50a4572b0 READ of size 2 at 0x61900412bbce thread T11.</p> <p>Impact: The ASAN reported "AddressSanitizer: heap-use-after-free" error for subscribe message received with Proxy-Authorization header having auts parameter.</p> <p>Ex: Proxy-Authorization: Digest auts*="0x01P+20"</p> <p>Root Cause: An invalid access of the freed memory occurred in this case. Accessing memory after it is freed can cause unexpected behavior that may result in core dumps.</p> <p>Steps to Replicate: Send a SUBSCRIBE message received with Proxy-Authorization header having auts parameter.</p>	<p>The code is modified to address the issue.</p> <p>Workaround: None.</p>
73	SBX-109186 SBX-110422	2	<p>Portfix SBX-109186: [ASAN]: sanitizer.CE_2N_Comp_SamProcess.32180: /localstore/ws/jenkinsbuild/sbxmainasan/marlin/TRM/trmTgPsStat.c:1427:74: runtime error: signed integer overflow: -2147061904 - 552880 cannot be represented in type 'int'</p> <p>Impact: There was a runtime error: signed integer overflow.</p> <p>Root Cause: A runtime error occurred due to wrong type casting from ULONG to LONG.</p> <p>Steps to Replicate: Run the codenomicon INVITE response with outgoing bye codenomicon UAC suite.</p>	<p>The code is modified to address the issue.</p> <p>Workaround: None.</p>
74	SBX-104817 SBX-109818	2	<p>Portfix SBX-104817: The SBC does not answer the SDP correctly when the remote SDP contains both session and media port state attributes.</p> <p>Impact: The SBC sets incorrect media port state in the SDP in reply of a MODIFY command when media port state is present in both session level and media level.</p> <p>Root Cause: The SBC parsed the media port state incorrectly when it is present in both session level and media level, and put them both at media level in the reply SDP.</p> <p>Steps to Replicate: Create a 3pcc call, then send a re-invite that triggers C3 sending a MODIFY with media port state present in both session level and media level. The SBC should keep them the same session level and media level in the SDP in the reply.</p>	<p>The code is modified when parsing media port state attribute when it is in both session level and media level.</p> <p>Workaround: None.</p>
75	SBX-109209 SBX-110139	2	<p>Portfix SBX-109209: SBC: Observed MAJOR logs for SIPFE "/localstore/ws/jenkinsbuild/sbxMain/marlin/SIPFE/sipFeSigPortCsv.c, SipFeGetSipSigPortStatisticsGetNextReqMsg, 1111] Another Query in process" on T140 load.</p> <p>Impact: Observed MAJOR logs for SIPFE "/localstore/ws/jenkinsbuild/sbxMain/marlin/SIPFE/sipFeSigPortCsv.c, SipFeGetSipSigPortStatisticsGetNextReqMsg, 1111] Another Query in process" on T140 load.</p> <p>Root Cause: The addressContext zone sipSigPortStatistics table is not supported on the SBC, so the code used to return the information was sometimes not returning, causing the error messages above.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Run the following command: snmpwalk -c admin -v2c -mgmt address of sbx> 1.3.6.1.4.1.2879.2.10.2.121 Repeat this command 4 times. 2. Verify the following error: "/localstore/ws/jenkinsbuild/sbxMain/marlin/SIPFE/sipFeSigPortCsv.c, SipFeGetSipSigPortStatisticsGetNextReqMsg, 1111] Another Query in process" on T140 load Does not occur in the DBG log. 	<p>The code is modified so SipFeGetSipSigPortStatistics routines return immediately.</p> <p>Workaround: Do not query that table.</p>

76	SBX-109187 SBX-109266	2	<p>PortFix SBX-109187: [ASAN]: sanitizer.CE_2N_Comp_ScmProcess_0.32472: /localstore/ws/jenkinsbuild/sbxmainasan/marlin/SIPS/sipsParseActions.c:16242: 70: runtime error: null pointer passed as argument 2, which is declared to never be NULL.</p> <p>Impact: The NULL pointer was accessed during the codenomonicon test with an ASAN SBC build.</p> <p>Root Cause: The Codec Attribute has been accessed in the SDP without a proper NULL check.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Install the ASAN build on the SBC. 2. Run Codenomonicon INVITE response with the outgoing Bye suite. 	<p>The code is modified to add the defensive NULL check.</p> <p>Workaround: Not applicable.</p>
77	SBX-108619 SBX-110145	2	<p>Portfix SBX-108619: [ASAN]: /localstore/ws/jenkinsbuild/sbxmainasan/marlin/SIPS/sipsPreParse.c:1742:46: runtime error: signed integer overflow: 2002002002 * 10 cannot be represented in type 'int'.</p> <p>Impact: A runtime error is seen in the SAM process. When an INVITE is sent out and response code received is very large, we will see the following issue: Runtime error: signed integer overflow: cannot be represented in type 'int'</p> <p>Root Cause: When the SIP Response code is very large, there is a signed integer overflow during the processing of the SIP PDU.</p> <p>Steps to Replicate: An INVITE is sent out to the egress. If the response code received is very large, the issue is seen.</p>	<p>The code is modified so if the SIP response is greater than or equal to max response code, the SBC throws an error.</p> <p>Workaround: None.</p>
78	SBX-108572 SBX-110148	2	<p>Portfix SBX-108572: [ASAN]: /localstore/ws/jenkinsbuild/sbxmainasan/marlin/SIPS/sipsParseActions.c:12698:76: runtime error: index 20 out of bounds for type 'sip_nameval_str [20]'</p> <p>Impact: When the PUBLISH message is received with 20 params in the Contact Header, the SBC throws a runtime error: index 20 out of bounds for type 'sip_nameval_str'.</p> <p>Root Cause: The param check for boundary condition was missing while parsing the contact header. While it was checking the params, the number of params should be less than 20, but the condition to handle number of params is not specified.</p> <p>Steps to Replicate: Send a PUBLISH message with 20 or more params in the Contact header.</p>	<p>The code is modified so if the number of params is 20, the SBC should throw a parse error.</p> <p>Workaround: None.</p>
79	SBX-108411 SBX-110024	2	<p>PortFix SBX-108411 : [ASAN]: sanitizer.CE_2N_Comp_ScmProcess_2.30828: ==CE_2N_Comp_ScmProcess_2==30828==ERROR: AddressSanitizer: heap-buffer-overflow on address 0x61b00008dc88 at pc 0x5621c53c6946 bp 0x7f9854100bc0 sp 0x7f9854100bb8.</p> <p>Impact: While running a INVITE CANCEL Proxy Call, the UAS observed a Address Sanitizer leak in SCM process and the SBC services stopped.</p> <p>Root Cause: Issue in parsing string when string is blank and very large that lead to a heap buffer overflow. Example: SIP/2.0/UDP 10.xx.xx.xx:7009;sigcomp-id=" LWS "</p> <p>When the token length becomes large 100 say (which is the LWS length) and this token length point outside the PDU, we will see this issue.</p> <p>Steps to Replicate: Send an INVITE with a VIA header as the last header.</p>	<p>The code is modified to use temptokLen instead of the tokLen. As a result, the tokLen does not reach outside the PDU.</p> <p>Workaround: The VIA header should not be last header, and the empty token string should be small.</p>
80	SBX-109873 SBX-110385	2	<p>Portfix SBX-109873: The SBC includes the "text port = 0" in its response for the re-INVITE to insert the text at the mid call.</p> <p>Impact: The SBC rejects the t140 stream in handling a MODIFY command that changes the audio codec from PCMU to AMR-WB and a t140 stream, with "adaptive codec" enabled.</p> <p>Root Cause: The SBC delays the process of MODIFY command when "adaptive codec change" is enabled, and as a result the MODIFY reply is sent back to C3 without t140 stream media resource allocation (IP and RTP port, etc).</p> <p>Steps to Replicate: With an adaptive codec change enabled on a VMG, create a 3pcc call from EVS + t140 — PCMU + t140. Then MODIFY T2 side to AMR-WB + t140 and the SBC should reply with a valid port number for t140 stream.</p>	<p>The code is modified to address the issue.</p> <p>Workaround: None.</p>
81	SBX-108410 SBX-109379 SBX-110155	2	<p>PortFix SBX-108410: [ASAN]: sanitizer.CE_2N_Comp_ScmProcess_3.8866: ==CE_2N_Comp_ScmProcess_3==8866==ERROR: AddressSanitizer: heap-use-after-free on address 0x6190001c77dd at pc 0x558bcc9ff877 bp 0x7fea305f4e00 sp 0x7fea305f45b0.</p> <p>Impact: The ASAN reported a "AddressSanitizer: heap-use-after-free" error when Subscribe request having a NULL character in quoted string.</p> <p>Root Cause: Invalid access of the freed memory occurred. Accessing memory after it is freed can cause unexpected behavior that may result in core dumps.</p> <p>Steps to Replicate: The codenomonicon subscribe-notify suite.</p>	<p>The code is modified to now log a parser error.</p> <p>Workaround: None.</p>

82	SBX-108574 SBX-110908	2	<p>Portfix SBX-108574: [ASAN]: CE_Node2.log:snprintf buffer too small. Need 327 Have 128 File: /localstore/ws/jenkinsbuild/sbxmainasari/marlin/SIPSG/sipsgLibUtils.c Line: 2524.</p> <p>Impact: The length of the destination buffer was smaller than the source buffer while calling the snprintf function that why buffer too small error was seen.</p> <p>Root Cause: The destination buffer size was smaller than the source buffer while calling snprintf.</p> <p>Steps to Replicate: Run publish requisition codenomicon test suite. This error should not come.</p>	<p>The code is modified to upper limit of source buffer size.</p> <p>Workaround: None.</p>
83	CHOR-7443 SBX-107356	2	<p>PortFix CHOR-7443: Call failures are observed with a 503 error response due to UXPAD process utilization spiked to 20% more that causes overall CPU usage of media container reaches above 90% utilization.</p> <p>Impact: CPU congestion was observed for media core in 2 VCPU scenarios, leading to call failures.</p> <p>Root Cause: There was a scheduling issue between NP and UXPAD running on the same core that was causing random spikes of processes causing media congestion.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Create a 2 VCPU scenario in public cloud and run transcode heavy load. (This issue was seen specifically in container environment where we do not have much control on performance tuning parameters). 2. Run the load for 2-3 days. 3. Check if there are any call failures due to media core CPU congestion during the execution. 	<p>The code is modified to adjust the process priority of NP and UXPAD to have better scheduling and avoid spikes of processing.</p> <p>Workaround: None.</p>
84	SBX-105073	2	<p>The *XrmAsyncCmdErrHdr: ERROR NpMediaIntf cmd 2c gcid.</p> <p>Impact: The set grace command returned an error code 0xF0 from the NP, which triggered unsolicited call cleanups.</p> <p>Root Cause: The set grace command can only be performed on a disabled media flow in the NP. But, the NP had received some set grace commands issued on an already enabled media flow. Unfortunately, we are unable to determine from just the core WHY XRM and NP became out of sync.</p> <p>Steps to Replicate: The issue is not reproducible. Perform regular regression tests.</p>	<p>The code is modified to set after XRM has sent media flow enable command to NP and is cleared on media flow disable. The XRM checks this flag and only send set grace commands to NP when this flag is clear. If the flag is set, the XRM reads the media control block in NP and log a MAJOR level debug message to help future debugging efforts.</p> <p>Similarly introduced a set of new flags in the BRES structure for RTCP Gen enable/disable commands sent to NP. The BRM checks the new flags and only enables RTCP Gen in NP when the flag is clear. If the flag is set, BRM reads the rtcpGen control block in NP and log a MAJOR level debug message to help future debugging efforts.</p> <p>Workaround: No workaround.</p>
85	SBX-108369	2	<p>The FIPS mode was broken on the AWS SBC.</p> <p>Impact: After enabling the FIPS mode on the AWS SBC or OpenStack, the SBC functions such as 'certificate import' do not work.</p> <p>Root Cause: While enabling FIPS mode, the FIPS flag is set in sbx.conf. But on Openstack/AWS deployments, the sbx.conf gets reset on the reboot causing a reset of the FIPS flag.</p> <p>Steps to Verify Fix:</p> <ol style="list-style-type: none"> 1. Enable fipsMode from CLI. 2. After a reboot, check /opt/sonus/conf/sbx.conf or try a certificate import from the CLI. <p>The fipsMode in sbx.conf is set to 'enabled'.</p>	<p>The code is modified to retain FIPS flag in sbx.conf on reboot.</p> <p>Workaround: After enabling the FIPS mode, set fipsMode=enabled in /opt/sonus/conf/sbx.conf.reconfigHa.orig and reboot the SBC.</p> <p>After a reboot check, /opt/sonus/conf/sbx.conf</p>
86	SBX-109105	2	<p>OOD PUBLISH vs. OOD INFO: There are different NRM triggers/congestion profiling.</p> <p>Impact: OOD PUBLISH vs. OOD INFO: There are different NRM triggers used for congestion profiling.</p> <p>Root Cause: Some of the OOD message types were not being included for triggering NRM congestion.</p> <p>Steps to Replicate: Run with various OOD PUBLISH and INFO messages that should trigger NRM congestion and calculate CPS resource appropriately.</p>	<p>The code is modified to trigger the NRM congestion by including all OOD message types.</p> <p>Workaround: There is no workaround.</p>

87	SBX-109493	2	<p>Forwarding info was redirecting info for a customer.</p> <p>Impact: When testing with JJ90.30 call flows and sending in the following parameters, the RDI (redirection information) parameter created from the history header information had the redirecting indicator field set to "call diverted, all redirection information presentation restricted" even though there was no privacy parameter information in the history header.</p> <p>Root Cause: The SBC interworking code was following the 3GPP 29.163 specification table 7.4.6.3.2.3 that allows the standard SIP privacy header information to be used when setting the RDI parameter information. Due to the setting RDI parameter information "id", it resulted in the redirecting indicator parameter value of "call diverted. All redirection information presentation restricted" instead of it being set to "call diverted".</p> <p>Steps to Replicate: Send an INVITE to the SBC, with parameters similar to those in the impact statement and check in the policy request that the RDI redirecting indicator field is set to "call diverted". The SIP trunk group variant needs to be set to TTC and the signaling acceptHistoryInfo control needs to be set to enabled.</p>	<p>The code is modified to not consider the SIP Privacy header value when determining the RDI parameter redirecting indicator field value if the SIP trunk group variant is set to TTC.</p> <p>Workaround: None.</p>
88	SBX-109013	2	<p>Increased the healthcheck interval.</p> <p>Impact: Core dumps are occasionally seen in the field due to the health check timer expiring when the process gets stuck.</p> <p>Root Cause: In SWe environments, disk/CPU timing issues can lead to the process slowing down and hitting these health checks.</p> <p>Steps to Replicate: This issue was only reproduced using debug code.</p>	<p>The code is modified to be more forgiving to cope with short spikes. The health check ping interval is now 2 seconds and needs to have 15 consecutive non responses in order for the process to be declared deadlocked and a core dump initiated to recover.</p> <p>Workaround: None.</p>
89	SBX-106788	2	<p>TOD Routing broken for non-ALL timeRangeProfiles</p> <p>Impact: Time of Day Routing is broken for any configured time range profile other than the default ALL profile</p> <p>Root Cause: The time of the day values are stored in the DB as hexadecimal octets and A to F digits are stored as lower case. While matching the configured data, the stored values are compared against upper case A to F digits, thus the result of this match always failed.</p> <p>Steps to Replicate: Set timeRangeProfile other than ALL to test. set global callRouting route trunkGroup INTERNAL_IPTG99 VSBCSYSTEM standard Sonus_NULL 1 all TEST0001 none Sonus_NULL routingLabel TO_EXTERNAL_TG99</p> <p>Note, the TRP TEST0001 is defined as below, and is set to match all days and all times:</p> <pre>set profiles callRouting timeRangeProfile TEST0001 entry 7 timeZone psxLocal set profiles callRouting timeRangeProfile TEST0001 entry 7 dayMatching dayOfWeek monday,tuesday,wednesday,thursday,friday,saturday,sunday set profiles callRouting timeRangeProfile TEST0001 entry 7 dayMatching holidays disable set profiles callRouting timeRangeProfile TEST0001 entry 7 dayMatching specialDays range none set profiles callRouting timeRangeProfile TEST0001 entry 7 timeMatching range all</pre>	<p>Code is updated in the matching function to check for both lower case and upper case hexadecimal digits.</p> <p>Workaround: No workaround</p>
90	SBX-110719	2	<p>HA setup on ASAN is not coming up after call config - pathcheck & SMM</p> <p>Impact: While testing with the following configuration the code could potentially read of the end of an internal memory block while printing a debug log. In the worst case scenario this could result in an SCM core dump. set addressContext default zone <zonenumber> sipTrunkGroup <trunk group> signaling messageManipulation smmProfileExecution fixedOrder</p> <p>Root Cause: The internal structure was not always getting initialized correctly prior to trying to print the contents and this led to the code reading past the end of the memory block.</p> <p>Steps to Replicate: Configure SMM with the fixedOrder configuration and make some calls. This issue was highlighted while running with ASAN images in the engineering lab.</p>	<p>The code has been updated to correctly initialise the internal structure and avoid reading past the end of a memory block.</p> <p>Workaround: Avoid using the "signaling messageManipulation smmProfileExecution fixedOrder" configuration if possible.</p>
91	SBX-111059	2	<p>Memory leaks are observing on ASAN build when SMM is configured</p> <p>Impact: While assigning SMM or flexible policy profiles to the zone configuration small amounts of memory leaked.</p> <p>Root Cause: The code was allocating memory in order to read information from CDB as part of processing the assignment and the memory did not get freed up at the end of processing.</p> <p>Steps to Replicate: Make SMM configuration changes at the zone level.</p>	<p>The code has been updated to correctly free the memory.</p> <p>Workaround: None.</p>

92	SBX-108998 SBX-111014	2	<p>PortFix SBX-108998 to ##9.2.2## - concise view config file is not getting created through REST API.</p> <p>Impact: concise view config file is not getting created through REST API.</p> <p>Root Cause: Backend API was unable to read username from curl command to append in exportConfig.sh. which was causing the issue.</p> <p>Steps to Replicate:</p> <p>Export the concise config file using curl -k -i --user admin:Sonus@123 -X GET 'https://\$device_ip:444/pm/api /ImportExport/exportConciseStart&#39; To check the file name, use the below REST API curl -k -i --user admin:Sonus@123 -X GET 'https://\$device_ip:444/pm/api /ImportExport/getStatus</p>	<p>Modified code to read username from curl command and added in exportConfig.sh to generate the export directory and exportStatus file.</p> <p>Workaround: Need to update ImportExport.php.</p>
93	SBX-110640 SBX-110916	2	<p>Portfix SBX-110640 to ## 9.2.2 ## - PC 2.0: CDC config not replicated to standby during SWO when there are no active calls.</p> <p>Impact: In a PC2.0 setup with HA, when a switch over is performed without LI calls then the interception stops working on the new active.</p> <p>Root Cause: The standby SBC is not processing CDC config for realm2hash. Realm2hash is only created if SBC has active LI calls at the time of transition from Active to standby. In case when SBC don't have Active call realm2hash is not created on transition from Active to standby.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Configure SBC for LI PC2 flavor call. 2. Disabled the peer and realm. 3. Enable peer and realm. 4. Check the realm2hash entry on current active. 5. Perform switch over from active to standby. 6. Check the realm2hash entry on current active. 	<p>Standby code is updated to process CDC config and create realm on standby.</p> <p>Workaround: Disable/enable CDC Diameter Realm Route and Peer on the new-active.</p>
94	SBX-111162 SBX-111343	2	<p>Portfix SBX-111162: The SBC fails to defragment/assemble fragmented IPSEC ESP packets received from the customer.</p> <p>Impact: The SBC drops an in-coming 1500B IP fragments sent through an IPsec tunnel, where the ESP packet is also fragmented into approximately equal-sized IP fragment packets. This complex IP frag-IPsec-IP frag problem primarily affects large SIP INVITE packets in certain network designs.</p> <p>Root Cause: This problem affects large packets sent through IPsec gateways that (1) do no reassemble received IP fragmented packets before sending them through the IPsec tunnel (IP fragments themselves are sent through the tunnel), and (2) fragment resulting ESP packets into approximately equal-sized packets instead of a 1500B and 72B fragments.</p> <p>Steps to Replicate: Test should send SIP packets 1500B to SBC over an IPsec tunnel through a device/devices that:</p> <ol style="list-style-type: none"> 1. Fragments the SIP packet into 1500B + small IP fragments. 2. Sends the IP fragments through the IPsec tunnel. 3. Fragments (again) the larger ESP packet into approximately equal-sized IP fragment packets. 	<p>The code is modified to reassemble IP fragments up to 1580B divided any way into a single internal packet buffer, which the IPsec decryption code and subsequent IP frag reassembly code can handle.</p> <p>Workaround: Two workarounds are possible:</p> <ol style="list-style-type: none"> 1. Terminate the IPsec tunnel on a router in front of the SBC instead of directly on the SBC. 2. Do not send the IP fragments through an IPsec tunnel terminated on the SBC. Instead, reassemble IP packets before sending them through the tunnel towards the SBC, so that there is only one level of IP fragmentation (of the ESP packet itself).

Resolved Issues in 09.02.01R003 Release

The following Severity 1 issues are resolved in this release:

Table 25: Severity 1 Resolved Issues

Issue ID	Sev	Problem Description	Resolution
----------	-----	---------------------	------------

SBX-107397 SBX-110861	1	<p>Portfix SBX-107397: There was an SBC switchover DEADLOCK detected for sysID 62, task SIPSG.</p> <p>Impact: An SIP call load can trigger a healthcheck timeout in the SIPSG module.</p> <p>Root Cause: The problem was with the flags that is used for creation of shared memory between SCM process and fault avalanche handler process. As a result of incorrect setting of flags some of the writes to the shared memory was taking more time resulting in the SCM thread getting blocked and health check timeout.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Run a 50 cps SIP-SIP call load for an extended period of time. 2. Ensure that the fault avalanche control feature is turned on. <p>"show system faultAvalancheControl facState" should be enabled.</p>	<p>The code is modified to disable fault avalanche control functionality by default.</p> <p>Workaround: Disable the Fault Avalance control functionality using the following configuration control.</p> <p>set system faultAvalancheControl facState disabled</p>
SBX-110248	1	<p>The SBC SWe crashes when making a video call with a certain type of phone.</p> <p>Impact: The SBC SWe NP crashes when IPsec fragmented signaling packets scenario calls made by using IPsec crypto as NULL encryption.</p> <p>Root Cause: During the IPsec crypto null processing of fragments (chained mbuf segments) by the SWe NP, an incorrect first segment length corrupted adjacent buffers, which resulted in a crash of the SBC SWe.</p> <p>Steps to Replicate: Establish IPsec session with null encryption and make calls.</p>	<p>The code is modified to address the issue.</p> <p>Workaround: Use IPsec non-null encryption suites such AES/3DES.</p>
SBX-111966	1	<p>The SBC/GSX fail to decode trigger response received from the PSX and leads to call failure.</p> <p>Impact: Call failures are seen at the SBC when the STI Service is executed for a Trigger Request on the PSX.</p> <p>Root Cause: For a two-staged call, the SBC performs a trigger request following an initial policy request to the PSX. If any STI service like signing or verification is executed for this trigger request, the SBC fails to decode a trigger response received from the PSX and leads to a call failure.</p> <p>Compiled diameter Read methods for STI AVPs that are executed on the GSX and SBC are incorrectly looking for TRG_REQ(trigger request) tag while decoding TRG_RESP (trigger response), which is leading to decode errors on the GSX/SBC and resulting in a call failure.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Perform a two stage call flow with the STI service enabled. 2. Ensure that the SBC sends a trigger request to the PSX for a two stage call call flow and perform the STI service. 3. Observe that there are no diameter decode errors on the SBC that can lead to a call failure. 	<p>The code is modified to look for the correct TRG_RESP tag. The SBC builds updates to use new diameter read methods.</p> <p>Note: This fix only stops the DIAMETER decoding errors and a further fix is being tracked under another JIRA for a future release to have the SBC pass STI information in the trigger request, as well as process the information in the trigger response to STI service works in conjunction with 2-stage calls.</p> <p>Workaround: As a workaround, STI service escapes for a trigger request.</p>

The following Severity 2-3 issues are resolved in this release:

Table 26: Severity 2-3 Resolved Issues

Issue ID	Sev	Problem Description	Resolution
SBX-109013 SBX-110784	2	<p>Portfix SBX-109013: Increase the healthcheck interval.</p> <p>Impact: Coredumps are occasionally seen in the field due to the health check timer expiring when the process gets stuck.</p> <p>Root Cause: In SWe environments, disk/CPU timing issues can lead to the process slowing down and hitting these health check issues.</p> <p>Steps to Replicate: This issue was only reproduced using debug code.</p>	<p>The code is modified to be more forgiving to cope with short spikes. The health check ping interval is now 2 seconds and must have 15 consecutive non responses in order for the process to be declared deadlocked and a coredump initiated to recover.</p> <p>Workaround: None.</p>

SBX-109153 SBX-110776	2	<p>Portfix SBX-109153: DTLS: A 503 Service Unavailable with GCM Ciphers.</p> <p>Impact: The following ciphers could not be used with DTLS:</p> <ol style="list-style-type: none"> 1. ECDHE_RSA_WITH_AES_128_GCM_SHA256 2. ECDHE_RSA_WITH_AES_256_GCM_SHA384 3. RSA_WITH_AES_128_GCM_SHA256 4. RSA_WITH_AES_256_GCM_SHA384 5. ECDHE_ECDSA_WITH_AES_256_GCM_SHA384 <p>Root Cause: The TLS profile had been updated with additional ciphers listed below but the DTLS code was not updated to support these.</p> <p>Steps to Replicate: Use aDTLS client that support these ciphers make a connection to the SBC.</p>	<p>The code is modified to include support for these ciphers.</p> <p>Workaround: None.</p>
SBX-110205	2	<p>D-SBC: The calling party cannot hear disconnect treatment announcement if M-SBC is a N:1 cluster.</p> <p>Impact: The connection IP is not updated when a different M-SBC is selected in clustered D-SBC deployment during disconnect treatment. The SBC was using the old connection IP which was causing the Ingress Peer endpoint not being able to hear the Disconnect Announcement.</p> <p>Root Cause: The newly allocated M-SBC media IP during Disconnect Treatment is not used by NRMA due to incorrect validation check and the previous one is used which caused the one way audio issue.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Configure Disconnect Tone Treatment in the SBC with an M-SBC cluster. 2. From the Ingress endpoint, send SIP INVITE to the SBC. 3. From Egress side send SIP 404 USER NOT FOUND in response to the INVITE. 4. Ensure Ingress side endpoint hears "Disconnect Tone TreatMent". 	<p>The code is modified to not use the newly allocated M-SBC media IP during Disconnect Treatment by the NRMA (due to incorrect validation check), and instead use the previous treatment that caused the one-way audio issue.</p> <p>Workaround: None.</p>
SBX-110558	3	<p>A call hold from both ends causes a re-INVITE handling issue.</p> <p>Impact: Call Holds received from both ends results in a re-INVITE handling issue.</p> <p>Root Cause: UserA to UserB call is connected. UserA puts call on hold. Later, UserA triggers a latemedia re-INVITE, and the SBC sends 2xx with offer SDP with a=sendrecv, the local dpm mode changes to sendrecv. The ACK received from UserA with a=inactive. call is still on hold.</p> <p>When userB also initiates the hold (ie., hold from other direction), it should be relayed to UserA. But the SBC not relaying the re-INVITE. Due to improper late media re-INVITE handling, the SBC media state changed, which blocks the relaying of hold re-INVITE from other direction.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. UserA to UserB call is connected. 2. UserA puts call on hold. 3. UserA triggers latemedia re-INVITE, the SBC sends 2xx with offer sdp with a=sendrecv. The ACK received from UserA with a=inactive. 4. UserB sends re-INVITE to put call on hold (ie., hold from other direction). 5. Nothing will be sent to UserA. 	<p>The code is modified to relay the hold re-INVITE from the other direction.</p> <p>Workaround: None.</p>

Resolved Issues in 09.02.01R002 Release

The following Severity 1 issues are resolved in this release:

Table 27: Severity 1 Resolved Issues

Issue ID	Sev	Problem Description	Resolution
SBX-109465 SBX-109686	1	<p>Portfix SBX-109465: The Leadership algorithm workaround for the openclovis issue can cause a core dump.</p> <p>Impact: The safplus_gms process crashes when coming out of split brain.</p> <p>Root Cause: Incorrect/inconsistent data results in the code asserting.</p> <p>Steps to Replicate: This issue is not easily reproducible and is caused by HA link instability /flapping.</p>	<p>The code is modified so that the data is consistent.</p> <p>Workaround: Run HA link stability.</p>

SBX-109893 SBX-109914	1	<p>Portfix SBX-109893: The SBC frequently switches over with a coredump after 9.2.1R0 upgrade.</p> <p>Impact: An SCM core occurred in code that is executed only when the STI feature is enabled.</p> <p>The core was the result of the code in SipSgStiCopyDisplayNametoCPC() attempting to de-reference a NULL pointer. The fix is to add a check for NULL before attempting de-reference the pointer.</p> <p>Root Cause: The root cause is that there is code in an STI specific function that attempted to de-reference a NULL pointer. A NULL pointer check is missing in this code.</p> <p>Steps to Replicate: Specific steps to reproduce this issue are not known. The root cause and the fix were found by code inspection and core analysis.</p>	<p>The code is modified to check for NULL before attempting de-reference the pointer.</p> <p>Workaround: The only known workaround is to disable the STI.</p>
-------------------------	---	---	--

The following Severity 2-3 issues are resolved in this release:

Table 28: Severity 1 Resolved Issues

Issue ID	Sev	Problem Description	Resolution
SBX-108469 SBX-109085	2	<p>A registration issue with a switchover case.</p> <p>Impact: Security Mechanism of Registration is set to the TLS in the RCB with two different scenarios. The scenarios are of basic registration, which does not have any security profile.</p> <p>Root Cause: Cause is when Reconstruction of RCB happens during switchover by default security is set to TLS without verifying Digest structure. Also whenever the Digest structure is deleted for any reason the code is not setting security back to NONE.</p> <p>Steps to Replicate: Test requires a HA setup.</p> <p>Scenario 1:</p> <ol style="list-style-type: none"> 1. In Active Node make a successful registration. Send a Fake registration such that it gets rejected with 403 error from server side. 2. Now perform a switchover and when standby node becomes active make another fake registration that gets rejected by 403 again. 3. Verify the Security Mechanism in CLI using "show status addressContext default sipActiveRegisterNameStatus" and also try to make a call. <p>Scenario2: (This is for pre-present TLS security before upgrade)</p> <ol style="list-style-type: none"> 1. In Active node, make a successful registration with response code other than 200 ex: 202 Accepted. [send 202 instead of 200 in server script] 2. Now, send a refresh register, it will be internally rejected with 403 from the SBC. In logs, you will see these statements "invalid state auth-rcvd" and "Refresh register did not meet security requirements". If you verify the security Mechanism it should show TLS. 3. Upgrade the Standby node to 824R2 build. Perform a switchover and send a refresh register to verify the CLI for security mechanism if set to NONE. Try making a call or send a refresh register again both should be successful. 	<p>The code is modified so when RCB reconstruction occurs, it verifies the Digest structure whether to set security to TLS or NONE based of the DigestWithoutTLS variable. Whenever the Digest structure is deleted for any reason, the code sets security back to NONE.</p> <p>Workaround: Try performing switchover twice.</p>
SBX-109083 SBX-109820	2	<p>Portfix SBX-109083: The SCM process core dumps during a SipSgAORHashRemove.</p> <p>Impact: The SCM Process core dumps due to an entry corruption in the Registration hash table post switchover. This corruption is rare and occurs infrequently.</p> <p>Root Cause: The result below is likely to core as a result of the core dump and SYS error logs.</p> <p>The corrupted AOR entry in hash table was allocated during the reconstruction for an Active RCB from standby context after switchover. The SYS Err logs indicate the presence of duplicate AOR entry. This could potential lead to corruption.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Run a Basic Registration test with switchover. 2. Test with application server sending more than one P-Associated URIs. 	<p>The code is modified to ensure that only one AOR entry exists in the hash table after a switchover on the Active SBC's cache.</p> <p>If the AOR entry is not found during remove operation, manually remove the entry to avoid the corruption later.</p> <p>Workaround: None.</p>
SBX-107960	2	<p>The AWS IPs remain assigned to the standby SBC, causing unnecessary dual restarts.</p> <p>Impact: If the communication between the active and standby is broken (over ha0 interface), both assume active roles (split brain). When this occurs, the standby node that becomes active, calls AWS APIs to move IP address to self. Once ha0 link is restored, the machine, which becomes active, does not call AWS API to move IP addresses, this might cause an issue when the node that has IP does not come up as active, in this case IPs are assigned on standby and another node becomes active.</p> <p>Root Cause: The root cause of this issue is not calling a AWS switchover API (move IPs) during split brain recovery.</p> <p>Steps to Replicate: Run a split brain test and recovery of SBC HA, and verify that API query is send by the active SBC after recovery.</p>	<p>The code is modified to call AWS APIs to move IP to current active machine during split brain recovery path.</p> <p>Workaround: Manually move all secondary IPs to current active machine to restore calls.</p>

SBX-70800 SBX-109822	2	<p>Portfix SBX-70800: Observing that the metaVariable table is getting modified on loading the backup configuration file of one instance in other instance.</p> <p>Impact: When loading the backup configuration from one SBC instance to another, the metavar table is getting populated with the list of metavars from both instances. This alone does not cause any issues so long as the metavars are unique, its just confusing to see the metavars for another instance in the table.</p> <p>Root Cause: The code was not removing the existing metavars prior to loading the configuration of another instance. This meant the metavar table had both sets of information.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Create a backup configuration of one cloud instance (instance1). 2. Loaded the backup configuration file into a new cloud instance (instance2). 3. Check the metaVariables table and see if it contains the metavars from both instances. 	<p>The code is modified to flush the metavars for the existing instance prior to loading the configuration from another instance.</p> <p>Workaround: None.</p>
SBX-109694 SBX-109823	3	<p>Portfix SBX-109694: The /node/actualCeName is not updated with new name when there is a change.</p> <p>Impact: If a cloud instance is completely rebuilt with a new HA IP value and actualCeName, the new actualCeName is not getting updated into the CDB. The CDB is left with the original actualCeName value. This leads to the SBC being unable to process the new configuration data as it is still trying to read the metavar information based on the original actualCeName value. This in turn leads to the SBC application being unable to start.</p> <p>This problem happens when the textual part of the actualCeName matches but the IP address is different.</p> <p>Root Cause: The code was only using the textual part of the name to check if the data in CDB already existed and was not updating the full actualCeName if there was a match on the textual part of the name.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Rebuilt a cloud instance using the same node name but new IP address and then try to apply this configuration to an instance that is already up and running with the same node name. 2. Check that the SBC application is up and running following the application of the new configuration. 	<p>The code is modified to correctly update the actualCeName, even if the node name (textual part of the name) already exists in the CDB.</p> <p>Workaround: None.</p>
SBX-108173 SBX-109685	3	<p>Portfix SBX-108173: The openclovis split-brain recovery data was not always correct.</p> <p>Impact: On recovery from split brain, a leader may not be properly chosen, and both machines could stay running for multiple minutes.</p> <p>Root Cause: The data passed to the leader election algorithm does not properly indicate that both nodes are leaders.</p> <p>Steps to Replicate: Repeatedly break and reconnect the HA connection, checking that the nodes realize they are coming out of split brain (through logs) and one node will properly restart to again become standby.</p>	<p>The code is modified so that the "isLeader" field is properly set and the leader election algorithm can properly detect we are coming out of split brain.</p> <p>Workaround: None.</p>

Resolved Issues in 09.02.01R001 Release

The following Severity 1 issues are resolved in this release:

Table 29: Severity 1 Resolved Issues

Issue ID	Sev	Problem Description	Resolution
SBX-108557	1	<p>The SBC continuously core dumps for SCM process since upgrade to V09.02.01R000</p> <p>Impact: The SCM Process may coredump due to memory corruption.</p> <p>Root Cause: There is code that is using an invalid pointer when writing to a buffer. This code was only added recently in 9.2.1R0.</p> <p>Steps to Replicate: This problem is triggered by the receipt of an invalid PDU and/or an SMM rule to reject the incoming Invite and an early ATTEMPT record was attempted to be written. The issue is random and depends on what info is not available when trying to write the accounting record, therefore the issue may not reproduce all the time.</p>	<p>The code is modified to address the issue.</p> <p>Workaround: If there is SMM rule (ignore/reject the Invite), then the SMM rule needs to be disabled until patch is applied.</p>

SBX-107690	1	<p>There was call failures observed on T140 load with various MAJOR logs in DBG.</p> <p>Impact: Call fails due to RID Enable errors. (where RID = receiver ID and is mapped to an allocated resource) DBG log shows many BrmAsynCnCmdErrHdlr logs with cmd 0x30 (RID Enable): MAJOR .BRM: *BrmAsynCnCmdErrHdlr: ERROR NpMediaIntf cmd 0x30 gcid 0x2128915e</p> <p>Root Cause: When RTCP Generation is disabled, RTCP RID for the call is expected to be disabled by Network Processor (NP).</p> <p>With introduction of SBX-86241 Streaming RTCP packets to Protect Server, there is now a case where RTCP Generation is enabled to stream RTCP packets to Protect Server but RTCP packets are not generated for the call and RTCP RID for the call is not enabled.</p> <p>When RTCP Generation is disabled NP uses rtcpMode=RTCP Terminate to indicate RTCP RID needs to be disabled. This is incorrect, since rtcpMode=RTCP Relay Monitor also has RTCP RID enabled and NP is expected to disable the RTCP RID.</p> <p>If a call has rtcpMode=RTCP Relay Monitor and RTCP Generation is disabled, leak this particular RTCP RID resource.</p> <p>The next time the particular RTCP RID is allocated, the NP returns an error indicating RTCP RID is already allocated.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Enable the Packet Service Profile //Resources/profiles/media/packetServiceProfile /tcpOptions/generateRtcpForT140IfNotReceivedFromOtherLeg. 2. On the SBC, set //System/media/mediaRtcpControl/t140RtcpMonitorInterval to 20 seconds. 3. Start a T140 call but do not send RTCP packet. Terminate the call in 10 seconds. 4. If you keep making this type of call, you will eventually see the Enable RID error. The DBG log will have the BrmAsynCnCmdErrHdlr entry. 	<p>The code is modified so that the BRM indicates the RTCP RID needs to be disabled or not. When the NP receives the command, it uses this new parameter to decide if RTCP RID should be disabled or not.</p> <p>Workaround: Disable the RTCP and disable RTCP termination.</p>
SBX-108516	1	<p>A call outage led to DSP errors.</p> <p>Impact: The call fails due to the RID Enable errors.</p> <p>The DBG log shows many BrmAsynCnCmdErrHdlr logs with cmd 0x30 (RID Enable): MAJOR .BRM: *BrmAsynCnCmdErrHdlr: ERROR NpMediaIntf cmd 0x30 gcid 0x2128915e</p> <p>Root Cause:</p> <ol style="list-style-type: none"> 1. A defect was found in the XRM redundancy processing code where xres->options field was not updated on the standby XRM. 2. When modifying a media flow in NP if RTCP relay monitor is being modified, the XRM did not provide rtcpMode in the media flow modify command. And NP is assuming the RTCP relay monitor is enabled by default. <p>Steps to Replicate: Test with calls that use RTCP terminations and will set rtcp-xr-relay-drop to TRUE.</p>	<p>When processing media flow modify request, the code is modified to set the RTCP relay monitor state based on the value in NP_MEDIA_RTCP_REL_MON_STR.</p> <p>Workaround: Disable RTCP and RTCP termination in the PSP.</p>

Resolved Issues in 09.02.01R000 Release

The following Severity 1 issue is resolved in this release:

Table 30: Severity 1 Resolved Issues

Issue ID	Sev	Problem Description	Resolution
SBX-101155 SBX-104620	1	<p>Portfix SBX-101155: A DSP coredump occurred on the SBC.</p> <p>Impact: The DSP faults with a memory fault in area that pulls buffers out of EMAC port. Suspicion is on buffer size which is reported by EMAC buffer. These errors result in a DSP reload, which is essentially ensures call continuation. Any calls which are using this DSP could have a small media outage while the DSP is reloaded.</p> <p>Root Cause: Root cause is a speculation that EMAC system is occasionally is a incorrect bufSize, similar to a bit flip.</p> <p>Steps to Replicate: This issue cannot be reproduced, it has been fixed based on code review and coredump backtrace analysis.</p>	<p>The code is modified to look for consistency in packet size and EMAC buffer size reported. In case an inconsistency is found, the packet is rejected, avoiding illegal memory access.</p> <p>Workaround: There is no workaround.</p>

SBX-104334 SBX-106056	1	<p>PortFix SBX-104334: Call Drop when being placed on hold - IPv6.</p> <p>Impact: The "anonymous.invalid" in IPv6 media is not considered as a call hold, and is rejected.</p> <p>Root Cause: There is no handling for "anonymous.invalid" while handling calls on hold.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Establish a normal IPv6 media call. 2. Send a relinvite with "anonymous.invalid" in the c line of the media. 3. Check that the relinvite is not rejected and handled as call hold. 	<p>Check for "anonymous.invalid" and if present, consider it as call hold and avoid going for DNS resolution.</p> <p>Workaround: Use the SMM to modify "anonymous.invalid" to ":" on the incoming side.</p>
SBX-106452 SBX-107256	1	<p>PortFix SBX-106452: The Standby SBC SWE on AWS is not coming up.</p> <p>Impact: There are hard-coded references to the admin user in SBC LCA code whereby if the admin user is removed, bringing up the SBC may fail on virtual platforms.</p> <p>Root Cause: The lack of exception handling to account for admin user usage caused the startup process to error out.</p> <p>Steps to Replicate: Create a new Administrator user from CLI/EMA on the Cloud SBC and delete default admin user. Reboot the SBC to check if the SBC comes back up as active or standby.</p>	<p>The code is modified to address the issue.</p> <p>Workaround: Follow the MOP:</p> <ol style="list-style-type: none"> 1. From Active's CLI: Create another Administrator user (e.g. newadmin) <p>CLI commands:</p> <pre>set oam localAuth user emsadmin accountRemovalState disabled set oam localAuth user emsadmin passwordAgingState disabled accountAgingState disabled set oam localAuth user newadmin group Administrator</pre> <ol style="list-style-type: none"> 2. From the Active's CLI, delete emsadmin user. <p>CLI commands:</p> <pre>delete oam localAuth user emsadmin</pre> <ol style="list-style-type: none"> 3. From shell of both instances, create admin user: <p>Shell command:</p> <pre>useradd -p Sonus@123 -G Confd,sftponly,upload,Administrator -s /bin/sh -d /home/sftproot/Administrator /admin admin</pre> <ol style="list-style-type: none"> 4. Reboot the standby instance. 5. Delete the admin user and group from both instances using shell command. <p>Shell commands:</p> <pre>userdel admin groupdel admin</pre> <ol style="list-style-type: none"> 6. Login as 'newadmin' user and add 'admin' user <p>CLI commands:</p> <pre>set oam localAuth user newadmin accountRemovalState disabled set oam localAuth user newadmin passwordAgingState disabled accountAgingState disabled set oam localAuth user admin group Administrator set oam localAuth user admin passwordLoginSupport disabled passwordAgingState disabled accountAgingState disabled accountRemovalState disabled commit request oam localAuth userStatus admin setRsaKey keyName adminKey rsaKey "<YOUR-PUBLIC-KEY-HERE>"</pre> <p>Execute the shell command below on both boxes to reset keep admin user value in cdb.</p> <pre>/opt/sonus/sbx/taif/bin/confd_cmd -o -c "set /SYS:system /admin[0]/sftpUserPassword/keepAdminUser true"</pre> <ol style="list-style-type: none"> 7. Now, the admin user present in /etc/passwd file of both instances.
SBX-105687 SBX-106176	1	<p>PortFix SBX-106175: An SBC 5400 post upgrade to 9.1, the SMM rule for options stopped working.</p> <p>Impact: The SMM rules were not applied after an upgrade to 9.1.</p> <p>Root Cause: The logic was modified to pick the same TG for request and response. In the Options ping scenario, request and respond with defaultiptg.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Upgrade from 8.1 to any later release. 2. Attach the SMM at zone. 	<p>The code is modified to address the issue.</p> <p>Workaround: NA</p>

SBX-105961 SBX-106480	1	<p>PortFix SBX-105961: The SBC reinvite with sendonly was causing a one-way audio.</p> <p>Impact: The SBC unexpectedly sends out a 'reinvite sendonly' when the relay DPM is disabled.</p> <p>Root Cause: Internal logical error to queue the SDP on ingress when egress send 18x (sendonly).</p> <p>Steps to Replicate: Enable minimize and disable relay DPM on ingress. A call B, B answer 183 (sendrecv), 183 (sendonly), 183 (sendrecv)...200 (sendrecv)</p>	<p>The code is modified so when subsequent 183 (sendonly) received on egress, the SBC updates the queue sdp (recvonly). Since the relay DPM is disabled, the SBC does not change the direction of datapath on ingress.</p> <p>Workaround: n/a. This is an application bug per rfc, and the SDP should not change in subsequent 18x.</p>
SBX-106178 SBX-106183	1	<p>PortFix SBX-106178: There was a switchover // sonusCpSystemProcessCoredumpGeneratedNotification.</p> <p>Impact: The SCM process has cored.</p> <p>Root Cause: The SCM has cored when attempting to write a log message because the code is attempting to de-reference a NULL pointer to access information for the log message.</p> <p>Steps to Replicate: There are no specific steps for reproducing this issue. This core will only occur in a multi-transfer scenario.</p>	<p>The code is modified to check that the pointer is non-NULL before de-referencing it.</p> <p>Workaround: There is no workaround. This core will only occur in a multi-transfer scenario.</p>
SBX-106611 SBX-107113	1	<p>PortFix SBX-106611: The SBC sends a BYE message within dialog sometimes.</p> <p>Impact: After a call is connected, if the SBC triggers internal re-INVITE due to minimizeRelayingOfMediaChangesFromOtherCallLegAll flag on one leg and at the same time. If the SBC receives late media re-INVITE on other leg, the SBC clears the call.</p> <p>Root Cause: The internal offer answer state of call at the SBC SIP subsystem becomes invalid.</p> <p>Steps to Replicate: Config: minimizeRelayingOfMediaChangesFromOtherCallLegAll enabled. lateMediaSupport -->passthrough E2E Ack Enabled. Egress PSP crypto and SRTP Enabled.</p> <p>Call Flow: Ingress (RTP), Egress (RTP).</p> <ol style="list-style-type: none"> 1. After call is connected, the SBC triggers internal re-INVITE to egress with one crypto. 2. At the same time , Ingress peer sends late media re-INVITE. 3. LateMedia will get queued up and processed after internal reinvite. 4. But as SBC receives 200OK from egress , it generates 200OK to ingress as a response to late media re-INVITE from ingress. <p>This is wrong. The SBC also sends re-INVITE to egress (late media). Due to this internal state gets messed up and after receiving ACK with SDP from Ingress, the SBC clears up the call.</p> <p>With a fix, verified that the SBC correctly sends 491 Request Pending to ingress when handling internal re-INVITE transaction on egress leg.</p> <p>When the SBC receives another late media re-INVITE on ingress leg , SBC sends it to egress and this second re-INVITE transaction completes successfully.</p>	<p>The code is modified to ensure if the SBC is handling internal re-INVITE on egress leg, late media re-INVITE on ingress leg is responded with 491 Request Pending with RetyAfter header.</p> <p>After the egress re-INVITE transaction is completed, if the ingress peer send another late media re-INVITE it is sent to egress leg correctly and offer answer transaction succeeds for this second re-INVITE.</p> <p>Workaround: Suppress the internal re-INVITE on egress leg.</p>
SBX-106691 SBX-106918	1	<p>PortFix SBX-106691: Adding IPSP fails with the application error.</p> <p>Impact: Adding IPSP fails with application error when ipSignalingProfile destinationTrunkGroupOptions is includeDtG, but the error message is not descriptive.</p> <p>Root Cause: ERE validates to ensure the check box 'Include DTG' is not selected in the IP signaling profile. The validation was done guiserver but error message was not set and return for why instead of meaningful error message, a very generic error message was coming.</p> <p>Steps to Replicate: admin@SBXPLTF1H% set addressContext ADDR_CONTEXT_1 zone ZONE_IAD sipTrunkGroup TG_SIPART_IAD policy signaling ipSignalingProfile AANN [ok][2021-02-02 19:22:08]</p> <p>[edit] admin@SBXPLTF1H% co Aborted: 'sipServiceGroupExt TG_SIPART_IAD ipSignalingProfile': IP Signaling Profile Id 'AANN' cannot be assigned as it is DTG(Destination TrunkGroup) option is selected.</p>	<p>The code is modified to set the error message and and return to provide error message during configurations.</p> <p>Workaround: N/A</p>

SBX-106206 SBX-106710	1	<p>PortFix SBX-106206: An existing hairpin call gets silenced after a switchover.</p> <p>Impact: A media loopback is not happening in after switchover.</p> <p>Root Cause: There was extra check performed during NP incoming path for the combination of MAC address and loopback flag.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. In the SWe SBC Active-Standby HA setup. 2. Using AS/3GPP call flow scripts/setup establish a call that creates an SBC internal loopback media flow. 3. Verify the media flow and then issue a switchover. 4. Verify the media path on this existing active call. 	<p>The code is modified to address the loopback media issue and also saves the NP cpu cycles.</p> <p>Workaround: None.</p>
SBX-106968 SBX-108087	1	<p>PortFix SBX-106968: There was an abnormal switchover on a cluster SBC.</p> <p>Impact: The PRS process cored due to accessing NULL destination ICM handle provided by DRM.</p> <p>Root Cause: The DRM does not mirror txcn handle in DRES data structure to standby and the code was not validating the pointer in one specific code path. The code intentionally crashed to identify the bad code.</p> <p>Steps to Replicate: The problem could not be reproduced and the solution was found based on code review. There was Edge case timing issues where internal message sent before a switchover and response is processed after a switchover.</p>	<p>The code is modified to validate the txcn pointer that is not null before trying to use it and thereby avoids the intentional system error crash.</p> <p>Workaround: No workaround</p>
SBX-106722	1	<p>The S-SBC application goes down with a the MESSAGE call load of 11 cps or more.</p> <p>Impact: Run the message call flow with the SMM configured with Dialog Scope variables on ingress and egress sides.</p> <p>Root Cause: Memory corruption is occurring when we store ingress dialog scope variable id in the Relay CB as it is already freed.</p> <p>Steps to Replicate: Run the message call flow with SMM configured with Dialog Scope variables on ingress and egress sides.</p> <p>Run 10 CPS load.</p>	<p>Delaying Relay CB free until we send 200 OK for Message</p> <p>Workaround: Disable the Ingress dialog Statefull variable for the SMM rules.</p>
SBX-107586	1	<p>A SAM Process core dump occurred while executing DoS on an SBC SWe SIP signaling port with a malformed Register Message containing 80 multiple unique and spoofed IPs.</p> <p>Impact: SIPSG is sending an update to SIPFE for deleting the RCB details with Username, Hostname and PhoneContext as 0 when Register is received with malformed PDU.</p> <p>Root Cause: In this malformed Register PDU scenario SIPSG should not send a Delete update to SIPFE.</p> <p>Steps to Replicate: Run a load with Register Malformed PDU from IXia or any test tool.</p>	<p>The code is modified for non zero phone context key in SIPSG for sending delete operation update to the SIPFE.</p> <p>Workaround: Run in normal mode instead of sensitive mode.</p>
SBX-105764	1	<p>A customer Teams core dump occurred in a TEAMS call flow.</p> <p>Impact: A SCM Process core dump resulted for TEAMS call flow.</p> <p>Root Cause: Null Pointer exception lead to a SCM Process core dump.</p> <p>Steps to Replicate: The steps cannot be reproduced.</p>	<p>The code is modified to prevent the crash.</p> <p>Workaround: None.</p>
SBX-107853	1	<p>The SCM process core dump was observed on the standby SBC while running a load with the SBC generated subscribe.</p> <p>Impact: The SCm process cores on the standby while running the load for a Reg-Event subscription.</p> <p>Root Cause: There is a leak on the standby for SIP dialog data structure, which is causing a SYS_ERROR on the standby.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Enable the below configuration for Reg-Event feature. set profiles signaling ipSignalingProfile DEFAULT_SIP commonIpAttributes subscriptionPackageSupport supportRegEvent enable 2. Run the load of REGISTRATION for Reg-Event subscription feature. 	<p>The code is modified to fix the leak of SIP dialog structure on the standby SBC.</p> <p>Workaround: No workaround</p>

SBX-106058	1	<p>Multiple core dumps were observed during a Load run for generated Subscription call Flow tested during Registration Display Enhancement feature.</p> <p>Impact: Multiple core dumps were observed when Reg-Event Subscription feature flag gets enabled.</p> <p>Root Cause: Core 1: Incorrect type cast of a structure led to invalid memory access and core was observed.</p> <p>Core 2: During a switchover, the list "pendingCcDslcmList" was never initialized; however, it was incorrectly getting freed.</p> <p>Core 3: Accessing of a NULL pointer led to this core.</p> <p>Steps to Replicate: Load test for Reg-Event Subscription initiation.</p>	<p>Core 1: The code is modified to point to valid memory.</p> <p>Core 2: In case of a switchover, the code is modified to initialize the list correctly.</p> <p>Core 3: Before accessing the pointer, the code is modified to resolve the issue.</p> <p>Workaround: N/A</p>
SBX-105905 SBX-106000	1	<p>PortFix SBX-105905: Observed the NP crash during an overload test.</p> <p>Impact: Possible memory corruption in NP KNI request queue when KNI requests are not processed in time.</p> <p>Root Cause: An older KNI request buffer in the KNI request queue could be concurrently overwritten by a new request while user-space NP logic is still processing it, leading to memory corruption. This issue is more frequent with Mellanox NICs, particularly with multicast mac programming flow.</p> <p>Steps to Replicate: Perform an sbxrestart on the SBC with port redundancy enabled, and using a Mellanox-based NIC.</p>	<p>The code is modified to ensure currently handled requests are not overwritten by a new request from the KNI module.</p> <p>Workaround: None.</p>
SBX-105269 SBX-107045	1	<p>PortFix SBX-105269: An SBC crash during a call transfer produced a core dump.</p> <p>Impact: The SCM Process core dumped due to NULL pointer access for a pointer that was freed due to BYE and HOLD race condition in a transfer call scenario.</p> <p>Root Cause: There was an illegal memory access, absence of NULL check, and exposed due to race condition.</p> <p>Steps to Replicate: A calls B, B REFERS to C and now A and C talk. After sometime, A sends BYE and C sends re-INVITE with a=inactive at the same time.</p>	<p>The core is modified to address the issue.</p> <p>Workaround: None.</p>
SBX-105888 SBX-106060	1	<p>PortFix SBX-105888: The SBC is sending the wrong TO tag to AS side.</p> <p>Impact: The response 200OK for Multi Notify has invalid ToTag.</p> <p>Root Cause: Internal logical error when process the 2nd 200OK Notify from AS. The application misinterpretation of the direction of the relay message as IAD. And passing down the wrong data down to SIPS for relaying 200OK to AS.</p> <p>Steps to Replicate: IAD Subscribe to AS. the AS sends multiple Notify immediately one after another. The 2nd 200OK of Notify sends to AS has wrong ToTag.</p>	<p>The code is modified so the correct data passes down to the SIPS.</p> <p>Workaround: Peer can send 1 Notify at a time. N/A on the SBC.</p>
SBX-108081 SBX-108199	1	<p>PortFix SBX-108081: Unable to see Video from Cisco 8865/9971 on ConnectMe BYOD client.</p> <p>Impact: The SBC drops RTCP packets for video stream if audio stream is transcoded by SBC for a multi-stream call.</p> <p>Root Cause: Non-RTCP binding resource was incorrectly picked by the SBC for video stream if audio is transcoded in a multi-stream call. The reason for picking this resource is, natively the SBC did not support audio transcode for a multi-stream call. As a result, the code assumed audio to be in pass-thru mode.</p> <p>Steps to Replicate: Configuration: ----- 1. Configure the SBC for Audio, Video Call. 2. Enable the allowAudioTranscodeForMultiStreamCall in PSP. 3. Enable the rtcp 4. Configure thisLeg and otherLeg, so that audio call will be transcoded.</p> <p>Procedure: ----- 1. Place an Audio transcoded and Video Relay call through the SBC.</p> <p>Without a Fix: ----- The SBC drops RTCP packets for video stream.</p> <p>With a Fix: ----- The SBC relays RTCP packets for video stream.</p>	<p>The code is modified to pick the RTCP binding resource for video stream irrespective of audio being pass-through or transcoded.</p> <p>Workaround: Disable the allowAudioTranscodeForMultiStreamCall flag at Route PSP which would result in Audio pass-through call.</p>

SBX-106329 SBX-106862	1	<p>PortFix SBX-106329: The SBC disconnects a call every few seconds after recovering a DNS server.</p> <p>Impact: The SBC disconnects a call every few seconds after recovering a DNS server.</p> <p>Root Cause: The DNS client sends a failed lookup response to a DNS agent (SCM ID 03,01 ...) when a probe query fails to get a response for a blacklisted DNS server because a probe query and regular query were used the same FQDN, record type and zone id.</p> <p>The timeout response for probe query in the DNS client process as triggered DNS failed response towards SCM process.</p> <p>Steps to Replicate: DNS settings: Three DNS servers are configured. Each weight is set to 50.</p> <p>DNS#1 10.xxx.x.xxx has RRs with ttl=5 DNS#2 10.xxx.x.xxx has RRs with ttl=0 DNS#3 10.xxx.x.xxx has RRs with ttl=0</p> <p><Time series> DNS#1 10.xxx.x.xxx process was down (Pkt No.29615). DNS#2 and DNS#3 were available. After few Seconds DNS#2 10.xxx.x.xxx process was down. Only DNS#3 was available.</p>	<p>The code is modified so the DNS probe query does not trigger any response towards a DNS agent from DNS client.</p> <p>So do not send any failure response to the DNS agent in case of probe query failure.</p> <p>Workaround: None.</p>
SBX-105439 SBX-106049	1	<p>PortFix SBX-105439: The SBC does not check if "serverCertName" exists while configuring it as part of the TLS profile.</p> <p>Impact: The SBC allows the configuration of the certificate, "serverCertName", as part of the TLS profile even though that certificate is not present in the SBC.</p> <p>Root Cause: The SBC is not validating that a certificate is present during configuration of "serverCertName" as part of TLS profile.</p> <p>Steps to Replicate: Verify the SBC does not configure a certificate which is not installed.</p> <p>Steps to verify issue:</p> <ol style="list-style-type: none"> 1. Bring up the instance. 2. Configure the certificate that is not present in the below path. show system security pki certificate. set profiles security tlsProfile defaultTlsProfile serverCertName <Name> <p>Expected result:</p> <p>The SBC does not allow configuring "serverCertName" as part of the TLS profile.</p> <p>Steps to verify fix:</p> <ol style="list-style-type: none"> 1. Bring up the instance. 2. Configure certificate that is present in the below path. show system security pki certificate. set profiles security tlsProfile defaultTlsProfile serverCertName <Name> <p>Expected result:</p> <p>The SBC allows configuring "serverCertName" as part of the TLS profile.</p>	<p>The code is modified for the SBC to validate whether a certificate being configured is present or not on the SBC.</p> <p>Workaround: None.</p>
SBX-106224 SBX-106481	1	<p>PortFix SBX-106224: An SBC failover with the SCM process core dumps due to memory corruption.</p> <p>Impact: When Prack is enabled on the ingress and advanced SMM for "variableScopeValue message" is enable, the SBC may core when the egress responds with 18x/2xx in quick succession, and the Prack is still pending.</p> <p>Root Cause: When 18x/2xx is queuing on the ingress, the SBC did not allocate proper resources for "variableScope" data. By the time the SBC sends the data out, it accesses invalid data.</p> <p>Steps to Replicate: This is random issue that is not easily reproducible.</p> <ol style="list-style-type: none"> 1. Egress: inbound config variableScopeValue message and send to ingress. For example: store "99" in var1. 2. Ingress: prack enable, outbound read the var1 and append to the userName of To Header. 3. Egress: send multiple 18x/200 fast to trigger some delay on ingress for sending. 4. Run a load. 	<p>The code is modified to properly allocate resource for "variableScope" data.</p> <p>Workaround: Disable prack.</p>

SBX-106476 SBX-107115	1	<p>PortFix SBX-106476: The sipAdaptiveTransparencyProfile is not working for a re-INVITE coming from Egress.</p> <p>Impact: When the sipAdaptiveTransparencyProfile is configured for Egress TG and re-INVITE comes from egress peer with change in P-ASSERTED-ID, the SBC does not relay the invite from egress to ingress.</p> <p>Root Cause: The SBC code does not set service bit for reINVITE transparency for re-INVITEs coming from egress peer.</p> <p>Steps to Replicate: Test case 1.</p> <p>Configure the sipAdaptiveTransparencyProfile for Egress TG for re-INVITE.</p> <pre>config set profiles services sipAdaptiveTransparencyProfile ADP2 sipMethod INVITE triggerHeader P-ASSERTED-ID action new-value trigger value-change set profiles services sipAdaptiveTransparencyProfile ADP2 state enabled set addressContext ADDR_CONTEXT_1 zone ZONE4 sipTrunkGroup SBXSUS7_LABSIP2 services sipAdaptiveTransparencyProfile ADP2 commit</pre> <p>When egress sends a re-INVITE with modified PAI after call is connected, no re-INVITE is generated towards the ingress side.</p> <p>Test case 2</p> <p>Verify the issue.</p> <p>With a fix, the re-INVITE is relayed to ingress for both late and early media cases.</p>	<p>The code is modified to ensure the SBC sets the service bit for egress properly when the sipAdaptiveTransparencyProfile is configured for egress TG.</p> <p>Workaround: None.</p>
SBX-106920 SBX-107730	1	<p>PortFix SBX-106920: The FmMasterProcess dumped core after a switchover with stable call.</p> <p>Impact: The FmMasterProcess core dumps during shutdown.</p> <p>Root Cause: The FmMasterProcess may deadlock during shutdown due to receiving an event while trying to shutdown/finalize the event handling.</p> <p>Steps to Replicate: This is extremely time-sensitive and requires publishing an event from AMF to FM at just the right time in the shutdown sequence. There is no way to consciously reproduce the issue.</p>	<p>The code is modified to avoid the possibility of the deadlock.</p> <p>Workaround: No workaround.</p>
SBX-107041 SBX-107802	1	<p>PortFix SBX-107041: The OAM and MRFP goes down due to component error time reset for SreqProcess,SLwredProcess, and when the FmMasterProcess is going down.</p> <p>Impact: The OAM and MRFP goes down due to component error time reset.</p> <p>Root Cause: The NTP is not synced before the SBC comes up. As a result, when the sync occurs after the SBC comes up and the new time is older than the current time, the SBC goes down.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Launch the SBC with a time zone and the NTP in sbx.conf. 2. Check if the NTP server was properly synced and there was no issue of "time getting reset to past" in the SBC logs. 	<p>Run NTP sync operations before starting serf to avoid any issue when the time sync occurs and time gets reset to past.</p> <p>Workaround: No workaround.</p>
SBX-106732	1	<p>The CE_Node1.log is filling up quickly.</p> <p>Impact: The CE_Node1.log file size grows to an excessive size causing a hard disk space alarm.</p> <p>Root Cause: An excessive number of log prints were coming from stack trace dumps on a SYS_ERR (EVLOG).</p> <p>Steps to Replicate: Run a suite of REFER call scenarios and see that NrmaCpcCallVerifyTimerFunc is not being written in CE_Node log.</p>	<p>Converted the SYS_ERR(EVLOG) event for the concerned log event to NrmaDlog, so the CE_Node.log file no longer contains the back trace.</p> <p>Workaround: echo > CE_Node.log when it grows too big.</p>
SBX-107439	1	<p>The SWe_NP - KNI was out of memory</p> <p>Impact: On the Azure platform with Mellanox ConnectX-3 accelerated NICs, the packet port interfaces may lose connectivity.</p> <p>Root Cause: The buffer pool associated with Packet interfaces was getting exhausted due to lazy free logic in DPDK's driver for Mellanox ConnectX-3 NIC.</p> <p>Steps to Replicate: Route few calls to the same packet interface as incoming and some calls to the other packet interface.</p>	<p>The code is modified in order to alleviate stress on the buffer pool.</p> <p>Workaround: None.</p>

SBX-107397	1	<p>After an SBC switchover, a DEADLOCK was detected for sysID 62, task SIPSG.</p> <p>Impact: The SIP call load can trigger healthcheck timeout in SIPSG module.</p> <p>Root Cause: The problem was with the flags that is used for creation of shared memory between SCM process and Fault Avalanche handler process. As a result of incorrect setting of flags, some of the writes to the shared memory was taking more time resulting in the SCM thread getting blocked and health check timeout.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Run a 50 cps SIP-SIP call load for an extended period of time. 2. Make sure that the fault avalanche control feature is turned on. "show system faultAvalancheControl facState" should be enabled. 	<p>The code is modified to disable Fault Avalanche control functionality by default.</p> <p>Workaround: Disable the Fault Avalance control functionality using the following configuration control: set system faultAvalancheControl facState disabled</p>
------------	---	--	---

The following Severity 2-3 issues are resolved in this release:

Table 31: Severity 2-3 Resolved Issues

Issue ID	Sev	Problem Description	Resolution
SBX-106815 SBX-107963	2	<p>PortFix SBX-106815: The PES process was leaking memory.</p> <p>Impact: In certain circumstance with high enough call rate, the SBC may experience PES memory leak.</p> <p>Root Cause: The newly ported Postgres code mishandled Postgres DATABASE cursor and counter.</p> <p>Steps to Replicate: This problem has been fund and reproduced in Comcast lab, when they were testing their call load.</p>	<p>The colde is modified in cursor and counter area.</p> <p>Workaround: Lower call rate should lower the risk.</p>
SBX-107593 SBX-107595	2	<p>PortFix SBX-107593: There is DTLS support for version 1.2.</p> <p>Impact: The SBC did not support DTLS clients which only supported DTLS version 1.2 to connect.</p> <p>Root Cause: The SBC was hardcoded to only support DTLS version 1.0.</p> <p>Steps to Replicate: Make a call from the DTLS client that only supports DTLS version 1.2 and check that the SBC is able to establish the call.</p>	<p>The code is modified to allow support for up to DTLS version 1.2.</p> <p>Workaround: None.</p>
SBX-105483 SBX-105739	2	<p>PortFix SBX-105483: There was a malformed P-charging vector.</p> <p>Impact: P-Charging-Vector was not passed transparently to egress on SIP-I to SIP call, when Create P-Charging-Vector is selected on egress IP profile and Store P-Charging vector is selected on ingress IP profile.</p> <p>Root Cause: The SIP-I INVITE contains IAM with a bad parameter and no parameter compatibility, causing the SBC to send 183 with CFN message and not transit the P-Charging-Vector.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Make SIP-I to SIP call. 2. SIP-I IAM contains an unrecognized parameter with no parameter compatibility information. 3. Ingress IP profile has Store P-Charging-Vector. 4. Egress IP profile has Create P-Charging-Vector. 	<p>Correct code to ensure P-Charging-Vector is passed to egress in this case</p> <p>Workaround: Disable support for the CFN message in ISUP signaling profile.</p>
SBX-105674 SBX-105892	2	<p>PortFix SBX-105674: Turk Telecom coverity issues part2.</p> <p>Impact: While processing SUBSCRIBE messages the coverity tool has highlighted that the code could dereference a pointer that is potentially null. Although no bad behaviour has been observed during testing there is a small chance that it could result in core dumps if the pointer really was null.</p> <p>Root Cause: Based on other validation in the code coverity highlighted that some legs of code could result in accessing a pointer that might be null. Dereferencing null pointers can cause unexpected behaviour and in the worst case core dumps.</p> <p>Steps to Replicate: Run various SUBSCRIBE related test cases.</p>	<p>The code is modified to validate that the pointer is not null before using it to avoid any potential issues/core dumps.</p> <p>Workaround: None.</p>
SBX-105497 SBX-107250	2	<p>PortFix SBX-105497: The wrong format of the SIP 400 BAD REQUEST.</p> <p>Impact: When the SBC received a message that is unable to find the required headers, the SBC responses 400 with syntax errors itself.</p> <p>Root Cause: The peer intentional sending garbage message.</p> <p>Steps to Replicate: Incoming request with have required headers as part of content body section.</p>	<p>The code is modified so the SBC discards the message.</p> <p>Workaround: Use the SMM to drop the message.</p>

SBX-102277 SBX-106516	2	<p>PortFix SBX-102277: The debug command caused ScmP to core dump</p> <p>Impact: We are hitting a Healthcheck timeout when displaying the output of the following debug command: "request sbx sipsg debug command "svcgrp -ce 0 -scm <x> -1""</p> <p>Root Cause: The Healthcheck timeout is because it is taking too long to display the data for all of the service groups when there are very large number of trunk groups configured.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Configure the 1000 TGs 2. Issue the following command: "request sbx sipsg debug command "svcgrp -ce 0 -scm <x> -1"" 	<p>The code is modified to only display up to 50 service groups in order to prevent this Healthcheck timeout.</p> <p>Workaround: To prevent this issue: Avoid using the following command if there are large number of TGs configured: "request sbx sipsg debug command "svcgrp -ce 0 -scm <x> -1""</p>
SBX-105160 SBX-107902	2	<p>PortFix SBX-105160: There was a CHM process core dump on standby OAM after reboot.</p> <p>Impact: The SBC application continues to come up even when asp.py zap (zapAsp()) call in sbxCleanup.sh) is invoked from sbxCleanup.sh script. When we try to access ceNum in one of the places, since it is not properly set due to failure in sbxCleanup.sh, it results in CHM core dump</p> <p>Root Cause: The SBC application continues to come up even when asp zap is called.</p> <p>Steps to Replicate: Fail sbxcleanup in one of the places and ensure the SBC is not being brought up.</p>	<p>The code is modified to avoid the SBC bring up if any errors are reported in sbxCleanup.</p> <p>Workaround: The issue is fixed, no workaround is required.</p>
CHOR-7376 SBX-107023	2	<p>PortFix CHOR-7376: There was a Swe_UxPad core dump observed during Teams Testing.</p> <p>Impact: The SWE_UXPAD process crashes due to the health check failure in extra-low memory SWE configuration on Microsoft Azure with presence of outgoing STUN/DTLS traffic.</p> <p>Root Cause: The root cause was identified to exhaustion of an internal buffer pool during the presence of STUN/DTLS traffic.</p> <p>Steps to Replicate: Launch a 6GB SWE instance in azure using accelerated NIC. Run a Teams test suite for multiple iterations which would trigger STUN/DTLS traffic to go out from SWE.</p>	<p>The code is modified to adjust the internal buffer pool size in order to account for additional requirements of STUN/DTLS traffic.</p> <p>Workaround: No workaround.</p>
SBX-107420 SBX-107880	2	<p>PortFix SBX-107420: Failed to download the Call Diagnostics file.</p> <p>Impact: Unable to download a large size call diagnostics log file.</p> <p>Root Cause: When we tried to download a large size call diagnostics log file. It was failing with JAVA heap error because IOUtils.toByteArray API was created internally ByteArrayOutputStream to store file data into byte format. Due to java memory restriction, we were getting a Heap error from ByteArrayOutputStream.</p> <p>Steps to Replicate: Steps to reproduce the issue:</p> <ol style="list-style-type: none"> 1. Login into EMA. 2. Troubleshooting > Troubleshooting Tools > Call Diagnostics. 3. Click on Save Call Diagnostics button to generate the log file. 4. Go to the Call Diagnostics Data Files section and try to download the tar file. 5. If the file size is more than 400MB, it should get failed with a proxy error. 	<p>The code is modified to internally copy the data from inputStream into OutputStream.</p> <p>Workaround: NA</p>
SBX-105290 SBX-108092	2	<p>PortFix SBX-105290: The SBC CDR: Redirecting digits captured incorrectly.</p> <p>Impact: The redirecting number recorded in the CDR was not correct if the first policy dip route was used for the call but there was an updated redirecting number from a later route in the policy dip.</p> <p>Root Cause: The code was incorrectly using the last per route redirecting number from the policy response to overwrite the redirecting number from the received INVITE message and using this to populate the redirection information field in the CDR record.</p> <p>Steps to Replicate: Configure the PSX routing label with two different route with two different ingress trunkgroups, and in each trunk add a DM rule for RDN. Now, make a call check what is the RDN from the CDR logs. If the call gets success with route 1 then it should have RDN according to DM rule in route1.</p> <p>The variations in tests:</p> <p>Test1 => route1 has DM/PM rule to remove CC in RDN and OCN, route 2 has DM/PM rule to modify RDN and OCN. CALL SUCCESS with Route1</p> <p>TEST2 => route1 has DM/PM rule to remove CC in RDN No rule for RDN, route 2 has DM/PM rule to modify RDN and OCN. CALL SUCCESS with Route1</p> <p>TEST3 => route1 has DM/PM rule to remove CC in RDN No rule for OCN, route 2 has DM/PM rule to modify RDN and OCN. CALL SUCCESS with Route1</p>	<p>The code is modified to populate the CDR using the per route redirecting number digits or the contain received from the INVITE.</p> <p>Workaround: Not Applicable.</p>
SBX-105901 SBX-106712	2	<p>PortFix SBX-105901: There is a heap-buffer-overflow in the ASAN build for SIPS module.</p> <p>Impact: Observing the heap buffer overflow in SCM process for info level log while decrypting the ROUTE header. Reading memory beyond the end of the allocated buffer can result in memory access faults and core dumps.</p> <p>Root Cause: A heap overflow occurs because the debug statement is trying to print from a string variable that is not null terminated.</p> <p>Steps to Replicate: Execute a test case where INVITE message contains encrypted route-header and verify that there are no failures.</p>	<p>The code is modified to create a local variable of type character array, which gets dynamically created and is always null terminated. This variable is used in the info log.</p> <p>Workaround: Not Applicable.</p>

<p>SBX-103183 SBX-106171</p>	<p>2</p>	<p>PortFix SBX-103183: The SBC incorrectly formatted the rport in the VIA header.</p> <p>Impact: The SBC incorrectly formats the rport value in the header of the response message when rport has a valid port number and is at the end of VIA header of request message.</p> <p>Root Cause: The code does not verify that rport has a valid port number. If rport is the last parameter in VIA header, the SBC appends "=<source port>".</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Setup: SIPp - SBX - SIPp 2. Send Register message with "rport=1111" in VIA header from client SIPp script with to the SBC. "rport" is at the end of VIA header 3. Run the client script with -p 30333 4. Verify that the SBC replaces rport port number (1111) with its own rport (30333) in VIA header of 200 OK response. 	<p>The code is modified so that it checks for any port number in rport parameter (rport is at the end of through the header of request message), it replaces the port number and appends it's own rport.</p> <p>Workaround: Use the following SMM as a workaround:</p> <pre> set profiles signaling sipAdaptorProfile HeaderModifications rule 1 applyMatchHeader one set profiles signaling sipAdaptorProfile HeaderModifications rule 1 criterion 1 type message message messageTypes response statusCode 200 set profiles signaling sipAdaptorProfile HeaderModifications rule 1 criterion 2 type header header name Via condition exist set profiles signaling sipAdaptorProfile HeaderModifications rule 1 action 1 type header operation store set profiles signaling sipAdaptorProfile HeaderModifications rule 1 action 1 headerInfo headerValue set profiles signaling sipAdaptorProfile HeaderModifications rule 1 action 1 from type header value Via set profiles signaling sipAdaptorProfile HeaderModifications rule 1 action 1 to type variable variableValue var4 set profiles signaling sipAdaptorProfile HeaderModifications rule 1 action 2 type variable operation regsub set profiles signaling sipAdaptorProfile HeaderModifications rule 1 action 2 from type value value "rport=" set profiles signaling sipAdaptorProfile HeaderModifications rule 1 action 2 to type variable variableValue var4 set profiles signaling sipAdaptorProfile HeaderModifications rule 1 action 2 regex string "rport=[0-9]"= matchInstance one set profiles signaling sipAdaptorProfile HeaderModifications rule 1 action 3 type header headerInfo headerValue set profiles signaling sipAdaptorProfile HeaderModifications rule 1 action 3 operation modify set profiles signaling sipAdaptorProfile HeaderModifications rule 1 action 3 from type variable variableValue var4 set profiles signaling sipAdaptorProfile HeaderModifications rule 1 action 3 to type header value Via </pre>
--------------------------------	----------	---	--

<p>SBX-104975 SBX-106101</p>	<p>2</p>	<p>PortFix SBX-104975: The MCF id sent back to the ingress as MPS during T.38 Fax to G.711 transmission.</p> <p>Impact: Some fax T.38 endpoints send an entire DCS V.21 signal in one packet as opposed to 1 octet per packet. This can causes fax failures.</p> <p>Root Cause: A burst of octets in a single packet is essentially a burst, and causes potential problems as these packets get queued for modulation and cause delay and later TCF (high speed) signal some packets get dropped.</p> <p>Steps to Replicate: Refer to unit test description in the JIRA.</p>	<p>The code is modified to accommodate such larger bursts.</p> <p>Workaround: This bug is a specific interoperability problem with such endpoints which exhibit burst single packet DCS signals. For such endpoints, a workaround is not available unless some configurations are available on those devices to modify this behavior.</p>
<p>SBX-106269 SBX-106800</p>	<p>2</p>	<p>PortFix SBX-106269: RTT-TTY uppercase behavior is inconsistent.</p> <p>Impact: Sometimes ASCII letters received in a T140 packet are sent as numbers and figures characters on g711 baudot side.</p> <p>Root Cause: Baudot has two modes - LTRS mode and FIGS mode. After transmitting 72 characters continuously in one mode, the recommendation requires to repeat LTRS or FIGS baudot tone to reassert the current mode.</p> <p>LTRS and FIGS mode have some common baudot codes such as 0 (BKSP),2(LF), 4(SPACE) and 8 (CR). There is no need to change mode when any of these common characters need to be transmitted.</p> <p>The issue is that when in LTRS mode 72nd character is received as one of common characters mentioned above, then the SBC incorrectly sends FIGS mode baudot tone. As a result, all subsequent LTRS characters are interpreted as FIGS and show incorrectly on screen of receiving phone</p> <p>Steps to Replicate: Create a t140 pcap with characters in each packet (per line) such as this. Note there is LF character after each line.</p>	<p>After 72 characters of same mode (LTRS or FIGS) the stack sends LTRS/FIGS baudot code. If 72nd characters is common code such as BKSP, SPACE, CR or LF and the if previous mode was LTRS, stack sends FIGS baudot tone.</p> <p>Fix to check current LTRS/FIGSmode and send the LTRS/FIGS baudot code.</p> <p>Workaround: There is no workaround as such.</p>

AAAAAAA
BBBBBBB
CCCCCCC
DDDDDDD
EEEEEEE
FFFFFFF
GGGGGGG
HHHHHHH
0000000
1111111
BBBBBBB
CCCCCCC
DDDDDDD
EEEEEEE
FFFFFFF
GGGGGGG
HHHHHHH
IIIIIII
AAAAAAA
BBBBBBB
CCCCCCC
DDDDDDD
EEEEEEE
FFFFFFF
GGGGGGG
HHHHHHH
IIIIIII
AAAAAAA
BBBBBBB
CCCCCCC
DDDDDDD
EEEEEEE
FFFFFFF
GGGGGGG
HHHHHHH
IIIIIII

Before the fix, the following appeared on g711 baudot side:

AAAAAAA
BBBBBBB
CCCCCCC
DDDDDDD
EEEEEEE
FFFFFFF
GGGGGGG
HHHHHHH
0000000
1111111
BBBBBBB
CCCCCCC
DDDDDDD
EEEEEEE
FFFFFFF
GGGGGGG
HHHHHHH
IIIIIII

???????

.....

\$\$\$\$\$\$\$

3333333

!!!!!!!

+++++++

=====

8888888

???????

.....

\$\$\$\$\$\$\$

3333333

!!!!!!!

+++++++

=====

8888888

After fix the T140 characters are displayed correctly.

SBX-106822 SBX-107399	2	<p>PortFix SBX-106822: There was a crash observed in Four GPU Codec Scenario.</p> <p>Impact: The G.729AB GPU codec may crash when packets are lost in the incoming G.729AB stream which, in turn, leads to SWe_UXPAD crash and an SBC application restart.</p> <p>Root Cause: An uninitialized stack variable in GPU G.729AB decoder code was identified as the root cause.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Set sweActiveProfile to use GPU, and sweCodecMixProfile to use G.729. 2. Make a G.729AB to G.711U call, and don't send the media. RESULT: SWe_UXPAD may crash and the SBC application restarts. <p>NOTE: The issue is not always reproducible.</p>	<p>The code is modified to initialize the stack variable appropriately.</p> <p>Workaround: No workaround.</p>
SBX-105413 SBX-105956	2	<p>PortFix SBX-105413: Support for 3072-bit long RSA keys is missing.</p> <p>Impact: The SBC needs to support 3072-bit RSA keys in certificates.</p> <p>Root Cause: A feature change is required to support 3072-bit RSA keys in certificates in addition to 2048-bit and 4096-bit RSA keys.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Test 3072-bit RSA key certificate on client. Make SIP-TLS call from the client to the SBC acting as TLS server. 2. Configure remote certificates with 3072-bit RSA key. 3. Configure local certificates with 3072-bit RSA key. 4. Test 3072-bit RSA key certificate on SBC acting as SIP-TLS server. 	<p>The code is modified to provide support for 3072-bit RSA keys in certificates.</p> <p>Workaround: None.</p>
SBX-106149 SBX-106641	2	<p>PortFix SBX-106149: The PSP QoS non-zero value of msrpDscp cause PES to crash.</p> <p>Impact: The PES process crashes when apps starting up if msrpDscp in PSP QoS values configured with non-zero.</p> <p>Root Cause: A debug statement tried to print a integer as a string, causing memory problem.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Configure the non-zero msrpDscp value. 2. Restart the SBC. <p>Before the code change, the PES will crash.</p> <p>After the code change, the PES will not crash.</p>	<p>The code is modified to print integer as integer.</p> <p>Workaround: Do not configure a non zero msrpDscp.</p>
SBX-107348 SBX-107360	2	<p>PortFix SBX-107348: The SIP LM call re-Invite sent in the 18x-PRACK stage on ingress.</p> <p>Impact: The SBC sends a reInvite to the late media while the call is still in prack pending state.</p> <p>Root Cause: Logical error that fail to detect the state of call not connect yet and the SBC tries to send the reInvite out.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Configure the LRBT. 2. Incoming late media with prack support. 3. Egress 180 trigger play tone, 180 again again, and 200OK(sdp). 4. Ingress delay sending prack for 5secs. 	<p>The code is modified to validate the state of the call before reinvite out.</p> <p>Workaround: Disable prack or disable LRBT.</p>
SBX-104113 SBX-105070	2	<p>PortFix SBX-104113: The mediationServer signaling channel was online even interface went down.</p> <p>Impact: Without this fix, the LI - mediationServer signaling channel stays online even after ipInterfacegroup attached to call data channel(CDC) is disabled post a switch-over.</p> <p>Root Cause: The SBC code, handling the LI - mediation server signalling socket functionality did not register for ipInterfaceGroup operational status in standby mode. As a result, post switch-over it does not get any notification when ipInterfaceGroup attached to CDC toggles. The signaling connection towards mediation server is not closed.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Create CDC and configure it for IMSLI and validate that the LI signaling channel is inService. 2. Perform a switch-over. 3. Disable the ipInterfaceGroup attached to the CDC. 	<p>The code is modified to register for the ipInterfaceGroup that is attached in CDC in standby mode as well so that if and once it becomes active due to switch-over it would get operational status of the ipInterfaceGroup attached to CDC.</p> <p>Workaround: N/A</p>

SBX-105711 SBX-105924	2	<p>PortFix SBX-105711: There was a ASAN MRFP: CE_2N_Comp_CpxAppProc leak during disable and enable of pkt port</p> <p>Impact: Creating an H248 signaling port results in a small memory leak.</p> <p>Root Cause: While processing the H248 signaling port creation command when metavaris are defined for the IP value, the SBC allocated memory to hold the metavar value from CDB internally for processing, but never freed up this memory block at the end of the port creation action resulting in a small leak.</p> <p>Steps to Replicate: Create an MRFP instance where metavaris are used to define the IP value for the H248 signaling port.</p>	<p>The code is modified to correctly free the memory block at the end of processing the signaling port creation.</p> <p>Workaround: None.</p>
SBX-103306 SBX-107647	2	<p>PortFix SBX-103306: Allocated bandwidth for opus call is 1032kb and 1028kb for single call.</p> <p>Impact: If "transcoderFreeTransparency(TFT)" is enabled at Route PSP, then extra bandwidth is allocated for opus call in case maxaveragebitrate attribute is not received in SDP from the endpoints.</p> <p>Root Cause: If maxaveragebitrate is not received in SDP, then the SBC was using the max value of 510kpbs as the default value (which was not as per RFC). Later, this bitrate value gets intersected with Route PSP configured value. However, for TFT calls, this intersection with route PSP does not happen. As a result, the SBC continues to maintain this value as 510kpbs, which results in extra bandwidth allocation.</p> <p>Steps to Replicate: Configuration:</p> <ol style="list-style-type: none"> 1. set profiles media packetServiceProfile DEFAULT packetToPacketControl transcode transcoderFreeTransparency 2. set profiles media packetServiceProfile DEFAULT secureRtpRtcp flags enableSrtp enable allowFallback enable (Ingress) <p>To re-create the issue:</p> <ol style="list-style-type: none"> 1. UAC sends INVITE with OPUS codec and SDP does not contain "maxaveragebitrate" attribute. 2. Since "maxaveragebitrate" is not received from UAC, the SBC defaults to max value of 510kpbs and sends the same to UAS. 3. UAS responds 200OK with OPUS codec and SDP does not contain "maxaveragebitrate". 4. The SBC sends out 200OK with maxaveragebitrate=510kpbs to UAC. <p>Test Result without a fix:</p> <p>Since maxaveragebitrate defaults to 510kpbs, the SBC ends up allocating more bandwidth than expected.</p>	<p>If "maxaveragebitrate" is not received in SDP, then derive the default value according to RFC using "maxPlayBackRate" and mode(mono/stereo).</p> <p>Workaround: None.</p>
SBX-107611 SBX-107612	3	<p>PortFix SBX-107611: The AMRWB bit-exactness problem in case of mixed mode test.</p> <p>Impact: Degraded audio in AMRWB stream when AMRWB (GPU) call load is running, particularly when each call uses different mode.</p> <p>Root Cause: Root cause was identified to an issue with context rearrangement code of GPU AMRWB codec.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Set sweActiveProfile to use GPU and sweCodecMixProfile to use AMRWB. 2. Make AMRWB to G711U call load using multiple AMRWB clients, each client using different mode. <p>RESULT: Some of the calls may have degraded audio in AMRWB stream.</p>	<p>The code is modified to address the issue.</p> <p>Workaround: No workaround.</p>
SBX-107924 SBX-107931	2	<p>PortFix SBX-107924: The OAM services are not coming up with 6GB Ram with network interface mapping error.</p> <p>Impact: The SWe_NP process fails to come up in a 6GB OAM instance configured with only 2 interfaces (ha0 and mgt0).</p> <p>Root Cause: Lack of sufficient huge pages in this particular scenario causes SWe_NP initialization failure.</p> <p>Steps to Replicate: Bring up a OAM node with 6GB RAM instance and configure two interfaces (without any pkt interfaces). The OAM instance would fail to come up.</p>	<p>The code is modified to reserve enough huge pages required for SWe_NP.</p> <p>Workaround: Configure OAM instance with >= 16GB.</p> <p>This is the officially supported configuration for OAM instance.</p>

SBX-60855 SBX-106922	2	<p>PortFix SBX-60855: The OPTIMA FTTH Stat for TG-A and TG-C Stat mismatching.</p> <p>Impact: The active register count on ingress TG is not decremented when a bad refresh REGISTER is received. Due to this there is difference in count of total stable registrations across zones.</p> <p>Root Cause: When the load of bad Refresh or initial REGISTERs received, at SBC for some of the registers not getting userNPhoneContextKey due to this the code that is responsible for reducing activeRegCount is not executed.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. First make one active Registration. 2. After active registration send one bad refresh REGISTER. <p>Expected Result:</p> <p>After bad refresh REGISTER, the SBC sends 400 Bad to Refresh REGISTER and also reduces the activeRegCount on ingress Side but it will not reduce count on egress side,</p>	<p>The code is modified to address the issue.</p> <p>Workaround: None.</p>
SBX-102700 SBX-106306	2	<p>PortFix SBX-102700: The number mapping (translated, RDI, To_hdr, R-URI) is unexpected.</p> <p>Impact: The SBC is not adding translated number received from PSX to RequestURI when Diversion header is present in the ingress Invite.</p> <p>Root Cause: If the ingress Invite does not contain Diversion header, then the SBC is adding translated number to RURI. So same result is expected when diversion header is present.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Make a SIP call by sending an Invite with Diversion header from UAC. 2. Configure PSX to return "translated number" for the called party number in policy response. 3. The SBC includes the translated number in Request URI and routes the call. 	<p>The fix is to populate the RequestURI with the translated number even when the Diversion header is present in the initial Invite.</p> <p>Workaround: None.</p>
SBX-104507 SBX-106552	2	<p>PortFix SBX-104507: The SBC is not passing URI parameter while History to Diversion interworking</p> <p>Impact: The SBC is not passing URI parameter "user=phone" while interworking History-Info header to Diversion header.</p> <p>Root Cause: The SBC was not adding "user=phone" during interworking History-Info header to Diversion header for the SIP uri.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. On the ingress leg, enable acceptHistoryInfo. 2. On the egress leg, enable diversionHistoryInfoInterworking. 3. Run a basic call with History-Info header having SIP uri and "user=phone" parameter. 	<p>The code is modified to add "user=phone" during interworking History-Info header to Diversion header for the SIP uri.</p> <p>Workaround: None.</p>
SBX-106042 SBX-108135	2	<p>PortFix SBX-106042: The RCB(s) can be hijacked and effect on registered users/EPs.</p> <p>Impact: Issue 1: When the register is sent to the registrar, the SBC creates a RCB for that particular User. When a fake/hijacked register is rejected with 403, some of the parameters in RCB are being modified and users were unable to make a call due to security mechanism being set to TLS.</p> <p>Issue 2: If a timeout on the RCB leads to it moving to a terminated state and a refresh register arrives, subsequent calls are rejected.</p> <p>Root Cause: Issue 1: When a Register request is rejected with a 403, it moves to the terminated state. Also, when a register is received, the RCB details are modified irrespective of the response received from the registrar.</p> <p>Note: A register is considered fake when we receive a 403 response for that.</p> <p>Issue 2: When the RCB was moved to the terminated state, the code was partially deleting internal memory that led to reporting the RCB mechanism as TLS, and not rejecting the TLS calls.</p> <p>Steps to Replicate: Use the conas mentioned in the description and use the config file attached for reference:</p> <ol style="list-style-type: none"> 1. After config make a successful registration, later make a register so that it gets rejected with 403 error. 2. Verify the parameters in RCB using this command: show status addressContext default sipActiveRegisterNameStatus 3. Check for all the displayed parameters and also check for sipSigPort and other essential parameters in Debug logs. 	<p>The code is modified for:</p> <ol style="list-style-type: none"> 1. When we receive a fake register and 403 response the RCB should remain in previous state with the previous information. The RCB contents are backed up so they can be restored on failure. 2. Remove the security mechanism of TLS when the RCB is in terminated state. <p>Workaround: No Workaround.</p>

SBX-102726 SBX-106883	2	<p>PortFix SBX-102726: The Diameter RX had the wrong data in media-component AVP post T.38-488.</p> <p>Impact: Sending the wrong data in media-component AVP (codec-data) when T.38 fallback fails.</p> <p>Root Cause: Instead of sending last negotiated codec, the SBC was sending previous offer-answer data in the media-component AVP of AAR message.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Configure PSPs for faxfallback and enable Rx feature. 2. Run transcode call (A(PCMA) and B(PCMU)). 3. The B sends fax tone, on detecting fax tone the SBC sends T.38 Re-INVITE towards A and A rejects with 488. 4. When a fallback happens, the SBC sends G711 Re-INVITE towards A and gets 200 OK. On getting 200 OK. The SBC sends final AAR that should contain last negotiated SDP information. 	<p>The code is modified to address the issue.</p> <p>Workaround: Not Applicable.</p>
SBX-105450 SBX-106573	2	<p>PortFix SBX-105450: A failure to delete remote server configured with FQDN for multiple commits.</p> <p>Impact: Failure to delete remote server configured with FQDN for multiple commits.</p> <p>Root Cause: The confd iter stops on returning failure for not finding any child remote servers for the first remote server and we fail to process the delete request for the second remote server that leaves the child remote servers of the second remote server as is.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Configure multiple FQDN remote servers such that first server fails to resolve FQDN and the second server creates child remote servers. 2. Disable both the remote servers. 3. Delete the first remote server and then the second in a single commit. 4. The remote servers created by second remote server fails to delete. 	<p>Return success from delete remote server function so that next CDB operation is processed. Failure to find a server is anyways logged and is not shown on the CLI as an error.</p> <p>Workaround: Delete the remote servers in separate commits.</p>
SBX-106913 SBX-107027	2	<p>PortFix SBX-106913: There was a SCM process coredump for register with timeout from Application server with no response.</p> <p>Impact: The coredump observed when the SBC logs account info incase of register timeout.</p> <p>Root Cause: There was a NULL check miss in the code so coredump observed when we try to access account information.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Send a REGISTER from UAC, the SBC sends towards AS. 2. The AS sends 302 to REGISTER. 3. The SBC sends redirected REGISTER to next AS and tries for 5 times since there is reply from AS. 	<p>The code is modified to address the issue.</p> <p>Workaround: Not applicable.</p>
SBX-105050 SBX-105675	2	<p>PortFix SBX-105050: The SBC DRBD mount not visible on active the SBC.</p> <p>Impact: The SBC DRBD mount was not visible on the active SBC.</p> <p>Root Cause: When the sbxrestart is issued from platform manager, DRBD gets mounted on apache's mountspace.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Bring up both the nodes. 2. Do an SBC restart from pm on active. 3. The new active must have drbd mounted. 	<p>Modify the apache service file and set the responsible parameter (PrivateTmp) to false.</p> <p>Workaround: None.</p>
SBX-105736 SBX-106560	2	<p>PortFix SBX-105736: Sending the m=application 0 UDP/BFCP (null) in case of UPDATE.</p> <p>Impact: The SBC sends m=application line in incorrect format in SIP UPDATE message.</p> <p>Root Cause: The SBC does not format m=application line for UDP/BFCP when sending UPDATE message.</p> <p>Steps to Replicate: Test case</p> <ol style="list-style-type: none"> 1. Reproduce the issue. <p>The SBC sends a INVITE with SDP(with audio and video lines and m=application) and receives rel 18x with SDP(with audio and video lines only) followed by final 200OK with SDP (with audio and video lines and m=application) if flag is enabled then SDP of 200OK should be ignored by the SBC.</p> <p>UPDATE message content: m=application 0 UDP/BFCP (null).</p> <ol style="list-style-type: none"> 2. Verify the issue. Repeat step 1. <p>UPDATE message content: m=application 0 UDP/BFCP *</p>	<p>The code is modified to ensure the SBC sends m=application line in correct format.</p> <p>Workaround: None.</p>

SBX-104851 SBX-105142	2	<p>PortFix SBX-104851: The SBCb is down and the standby registration with active failed, error 160004.</p> <p>Impact: The standby is not allowed to join cluster and fails to start</p> <p>Root Cause: The safplus checkpoint file is corrupt and the section needing to be overwritten is not found.</p> <p>Steps to Replicate: The root cause of the checkpoint corruption is unknown/cherckpoint corruption cannot be forced and therefore, directly testing this fix is not possible.</p>	<p>The code is modified to re-add the missing section if it is not found.</p> <p>Workaround: A complete outage is required as the active must be restarted.</p>
SBX-102469 SBX-106294	2	<p>PortFix SBX-102469: Time zone is defaulting to EST/EDT on the SBC instances while using image based instantiation.</p> <p>Impact: Time zone is defaulting to EST/EDT on the SBC instances while using image based instantiation because the timezone details passed during launch are not being applied when the SBC is brought up.</p> <p>Root Cause: The timezone details passed during launch are not being applied when the SBC is brought up.</p> <p>Steps to Replicate: Launch the SBC on the SWE. Provide values for timezone and NTP in sbx.conf. Check if proper timezone values and NTP configuration is applied on launch.</p>	<p>The code is modified to update timezone with value in sbx.conf.</p> <p>Workaround: No workaround other than setting timezone manually in cli.</p>
SBX-106231 SBX-107398	3	<p>PortFix SBX-106231: The Metallic noise present in the stream coming out from G722 side.</p> <p>Impact: Metallic noise in G722 stream when GPU G722 to Narrow band codec call is made.</p> <p>Root Cause: Resampled output is incorrectly copied in GPU G722 Encoder upsampling (narrowband to wideband) code.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Set sweActiveProfile to use GPU and sweCodecMixProfile to use G722. 2. Make G722 to G711U call. <p>RESULT: G722 stream contains metallic noise.</p>	<p>The code is modified to properly copy the resampled output.</p> <p>Workaround: No workaround.</p>
SBX-106343 SBX-107400	2	<p>PortFix SBX-106343: The GPU EVRCB decoder asserts in Erasure frames simulation test.</p> <p>Impact:The GPU EVRCB decoder may crash when there are lost packets in the incoming EVRCB stream, which in turn leads to SWe_UXPAD crash and the SBC application restart.</p> <p>Root Cause: Precision errors in the floating point comparison in EVRCB decoder code was identified as the root cause.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Set sweActiveProfile to use GPU and sweCodecMixProfile to use EVRCB. 2. Make EVRCB to G711U call and do not send the media. <p>RESULT: SWe_UXPAD will crash and the SBC application restarts.</p> <p>NOTE: The issue is not always reproducible.</p>	<p>The code is modified to handle precision errors.</p> <p>Workaround: No workaround.</p>
SBX-107613 SBX-107618	2	<p>PortFix SBX-107613: The AMRWB encoder produces corrupted output when channel is reused by lower mode.</p> <p>Impact: Degraded audio in AMRWB stream when AMRWB (GPU) call load is running, particularly when each call uses different bitrate.</p> <p>Root Cause: Problem occurs when the same codec context is reused and the previous used was for higher bitrate, the root cause was identified due to reinitialization logic for an internal buffer.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Set sweActiveProfile to use GPU and sweCodecMixProfile to use AMRWB. 2. Make AMRWB to G711U call load using multiple AMRWB clients, each client using different mode. <p>RESULT: Some of the calls may have degraded audio in AMRWB stream.</p>	<p>The code is modified to appropriately reinitialize the internal buffer.</p> <p>Workaround: Problem does not occur when all channels use the same bitrate.</p>
SBX-106004 SBX-106239	2	<p>PortFix SBX-106004: The EMA display error when SMM deleted through the CLI.</p> <p>Impact: When a rule is deleted from SMM through CLI and if we try to load the SMM from EMA, an error is displayed in the UI.</p> <p>Root Cause: EMA checks the ordering of the rules and if the Order is not continuous then an error is shown in the UI.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Delete any SMM rule (other than the last) from CLI. 2. Login to EMA. 3. Navigate to Profiles > Signaling > SIP Adaptor Profile. 4. Select the SMM whose rule was deleted from CLI. 5. SMM is shown without any error. 	<p>The code is modified to address the issue.</p> <p>Workaround: N/A</p>

SBX-105387 SBX-105936	2	<p>PortFix SBX-105387: LeakSanitizer:SCMP_3 gave memory leaks at MemAlloc2 and same gave Heap overflow issue, both from the same caseID.</p> <p>Impact: Heap use after free detected in ASAN for a downstream forking call flow. Accessing memory after it has been freed can cause unexpected behavior and in the worst case potentially coredumps.</p> <p>Root Cause: When copying multiple contacts from different downstream forking response, the username of the contact header was not updated from the call block to sip message handle. The Sip Message handle was holding a address that was already freed.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Run Downstream Forking Call Scenario. 2. UE sends Initial INVITE towards UAS through the SBC. 3. The SBC receives multiple 18x with different to tag and different username in the Contact header. 4. The SBC receives 200 OK for any of the downstream forking dialog. 	<p>The code is modified with the correct username from the call block for every downstream forking response.</p> <p>Workaround: None.</p>
SBX-106167 SBX-106808	2	<p>PortFix SBX-106167: The call ends up in one-way audio after the called party puts the call on hold and off hold twice.</p> <p>Impact: Call ends up in one-way audio after the called party puts the call on hold and off hold twice</p> <p>Root Cause: As a result of call-modify a couple of times, RTCP NAPT learning completes before RTP NAPT learning.</p> <p>This results in RTCP Remote Address being updated, which has remote RTCP Port.</p> <p>Due to incorrect code in RTCP modify flow, remote RTCP port, gets assigned to RTP port. This results in one-way media.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Run basic SIP to SIP call with NAPT and RTCP enabled. 2. Hold/Unhold the call a few times to check for proper 2-way audio. 	<p>Set the correct RTP Port as part of RTCP modify flow.</p> <p>Workaround: Since the issue is caused to RTCP NAPT learning completed before RTP during multiple hold/unhold scenario. Work around could be:</p> <ol style="list-style-type: none"> 1. Disable RTCP or 2. Disable NAPT.
SBX-105804 SBX-107918	2	<p>PortFix SBX-105804: The CDR field is not populated even though the SBC writes the value to CDR field for '200 Ok of BYE' received/sent.</p> <p>Impact: The CDR field added by SMM is not present in the 'STOP' record.</p> <p>Root Cause: The SBC was not adding CDR field in the STOP record that was added during the process of 200 response of BYE method.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Attach the SMM profile which will add the CDR field for 200 response of BYE method. 2. Enable endToEndBye flag. 3. Run a basic A-B call. 	<p>The code is modified for updating the CDR field in the STOP record during the processing of response of BYE method.</p> <p>Workaround: Add the CDR field in the SMM for BYE method instead of 200 response of BYE method.</p>
SBX-103570 SBX-106546	2	<p>PortFix SBX-103570: Observed major logs flooding for "MAJOR .SIPSG: sipsgMsgProc.c (-19034) 49487. SipSgRemoveDbiTrackingEntry, Hash entry not found"</p> <p>Impact: Flood of Error logs are seen during load testing of TLS/TCP Registration and certain Register messages are under the scanner of DBL.</p> <p>Root Cause: The DBL has limitations on the offenders list. When this limitation is reached, due to certain bug in the code, the DBL was sending unintended messages to other modules. Due to this message, the flood of logs were observed.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Execute a 600K TLS/TCP Registration at 1000 rps. 2. Let all the endpoints register. 3. Start the Supported call load (here Approx 20K, 254cps * 90cht= 22860). 50% from the registered Endpoints and 50% from Peering EPs. 4. Ext to Ext IntraSBC Call load of 5000, 50cps*100cht from SIPP. 5. The DBL applied/removed for Registered endpoints should deny entries to 10% of registration /call load. Repeat this for 10 times. 6. Let the load run for 3 hours. 7. Do operator initiated CLI switch-over and revert. 	<p>The code is modified to address the issue.</p> <p>Workaround: Not Applicable.</p>

SBX-99258 SBX-108094	2	<p>PortFix SBX-99258: For OOD NOTIFY, SBC sending 481-Call Leg/Transaction Does Not Exist when same SBC acting as P-CSCF and IBCF.</p> <p>Impact: SBC fails to send NOTIFY if the User part is not present in the To Header.</p> <p>Root Cause: SBC is not able to find the entry in the hash-table as to-tag is not parsed properly, because of user part being absent in the To header.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Make a basic Subscribe-Notify call, ensure that in From header of subscribe and To header of NOTIFY does not have User part. eg: To: sip:10.xx.x.xx;tag = abcd-1234-012 2. See that for Notify, the SBC will respond with 481 call leg does not exist. 3. Try the same with Fixed build, notify should reach the subscriber. <p>The examples below are To headers that were tested in the Notify:</p> <ol style="list-style-type: none"> 1. To: sip:10.xx.x.xx;tag = abcd-1234-012 2. To: <sip:10.xx.x.xx>;tag = abcd-1234-012 3. To: sip:xxx@10.xx.x.xx;tag = aabcd-1234-012 4. To: <sip:xxx@10.xx.x.xx>;tag = aabcd-1234-012 	<p>The code is modified to fetch the tag properly even if the user part is not present in To header.</p> <p>Workaround: No Workaround.</p>
SBX-88007 SBX-106099	2	<p>PortFix SBX-88007: The call flow should work with 5 video and 1 audio streams, which is not working, but when one video stream is removed then the callflow is working.</p> <p>Impact: The call fails if peer SDP contains six streams (5 video and 1 audio) and directMedia along with ICE is enabled at the SBC.</p> <p>Root Cause: Memory allocated for this SDP was not enough during inter process communication resulted in the failure.</p> <p>Steps to Replicate: Configuration: Test requires ingress TG in Zone1, egress TG in Zone2 and AS TG on the SBC, with call to be routed as follows:</p> <p>UE1 -> ingress TG -> AS TG -> AS peer (sipp) -> AS TG - egress TG -> UE2 Enable Direct Media on the ingress and egress TG's (but not on AS TG): TG - media directMediaAllowed enable PSP for all TG's should have DM flag enabled PSP - flags useDirectMedia enable Enable ICE (webrtc) on the ingress and egress TG's: TG - services natTraversal iceSupport iceWebrtc Enable DTLS on the ingress and egress. TG's and PSP's TG - media dtlsProfileName defaultDtlsProfile PSP - dtlsCryptoSuiteProfile DEFAULT enableDtlsSrtp enable Set the directMediaGroupID on ingress and egress TG's to be same e.g. 200 and AS to be different e.g. 400 e.g. TG - media directMediaGroupID 200 Enable Direct Media Anti Trombone on AS TG: TG - media directMediaAntiTrombone enable</p> <p>Steps:</p> <ol style="list-style-type: none"> 1. From UE1 send an INVITE with ICE and DTLS in the SDP with 5 video and 1 audio stream to ingress TG and route towards the AS TG 2. From the AS send back an INVITE to AS TG of the SBC, the C line of the sent SDP must match the C line of the received SDP. 3. From UE2, send 180 with no SDP followed by 200 OK with ICE and DTLS in the SDP (UE2-sdp 1 audio and 5 video). 4. From the AS send the 180 followed by 200 OK back towards the SBC. <p>Expected Result: The call succeeds.</p> <p>Actual Result: The call fails.</p>	<p>The code is modified so the memory buffer size is increased to accommodate six streams.</p> <p>Workaround: None.</p>
SBX-103539 SBX-105934	2	<p>PortFix SBX-103539: Observed flooded Logs for refresh registration in the Standby SBC restart stop for sometime "sipsgRegSecurity.c (-2745) 1752. SipRaDigestTlsProcessRefreshRegister: No Digest TLS negotiation in progress"</p> <p>Impact: The DBG logs were filling up with the following MAJOR level log. 167 09142020 143342.220452:1.01.05.41210.MAJOR .SIPSG: sipsgRegSecurity.c (-2745) 1753. SipRaDigestTlsProcessRefreshRegister: No Digest TLS negotiation in progress</p> <p>Root Cause: This is an information message and should not have been getting generated at MAJOR level.</p> <p>Steps to Replicate: Run registration call flows.</p>	<p>The code is modified to address the issue.</p> <p>Workaround: None.</p>

SBX-106170 SBX-106189	2	<p>PortFix SBX-106170: The AddressSanitizer: heap-use-after-free on address 0x60f000047d38 at pc 0x562fcd2973a5 bp 0x7f3963983f40 sp 0x7f3963983f38 READ of size 20 at 0x60f000047d38 thread T7.</p> <p>Impact: The ASAN reported of accessing a structure pointer that is already freed.</p> <p>Root Cause: The SBC is trying to access structure pointer to get socket address, but that structure pointer is already freed in other function.</p> <p>Steps to Replicate: Use ASAN build for testing.</p> <ol style="list-style-type: none"> 1. Send INVITE from UserA, respond from the DNS1 with RCODE error 4 for A query and respond from DNS2 with RCODE 0 with proper DNS answer for A query and check dnsServerStatistics. 2. Run a show command to check the dnsFallback flag and ednsFailures stats: show addressContext <addressContext_Name> dnsGroup <dnsGroup_Name> dnsFallback show table addressContext default dnsGroup <dnsGroup_Name> dnsServerStatistics 	<p>The code is modified so now getting socket address from pstSrcAddr structure.</p> <p>Workaround: No workaround.</p>
SBX-102115 SBX-106504	2	<p>PortFix SBX-102115: The SBC modifies the Replaces parameter with wrong Call-ID and tags if the Replaced call was looped back to the SBC through a SIP Proxy that does not modify SIP Call-ID and tags</p> <p>Impact: Relay refer with replaces for loopback call, the SBC picks up the wrong call leg.</p> <p>Root Cause: The SBC query for the replaces callId, and found matching both legs (loopback). The SBC pick up the wrong leg.</p> <p>Steps to Replicate: This is specific to customer call flow.</p>	<p>If loopback is detect, only pick the one with to-tag match local-tag. This is per RFC-3891, section 3.</p> <p>Workaround: None.</p>
SBX-106200	2	<p>The DSP is coring and the system unstable</p> <p>Impact: Some Fax T.38 IAD calls (T.38 protocols version 3) result in a DSP crash and reload and calls are not successful.</p> <p>Root Cause: This condition is caused because this specific V3 IAD is sending CM messages with incrementing UDPTL seq numbers. Typically, repeated messages carry same UDPTL seq number to indicate that they are redundant. Packets with unique seq numbers are assumed to be independent and data from these is causing an unchecked buffer overflow in 3rd party T.38 stack code.</p> <p>Steps to Replicate: The main cause of the crash based on core inspection seems to be multiple CM messages with incrementing UDPTL seq numbers Rx by stack.</p> <p>As a result, the test case involves a setting up G711 to V3 call with such CM packets.</p> <p>UAC: start call in G711 and reinvoke to V3 and send CM g711v3t38_cmseq.xml UAS: start in g711 and accept a re-invite to silsup-off uas_reinvite_g711.xml PCAP: cmt38seq.pcap</p> <p>Delete beforefix.PKT and afterfix.PKT entries because they refer to files attached to the JIRA, which customers don't have access to.</p> <p>Before fix: The call is set up, and a DSP coredump occurs with a single DSP reload.</p> <p>After fix: The call continues without crash.</p>	<p>The code is modified to address the issue.</p> <p>Workaround: Use Version 0 for T.38 calls.</p>
SBX-106429	3	<p>The EMA Configuration Script and Template import page file list doubles each time the search magnifier is used</p> <p>Impact: When the search magnifier is hit, the entries are doubled.</p> <p>Root Cause: Table with old data is not cleared before starting a new search. As a result, entries were doubled in dataTable.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Have some entries in the table. 2. Apply a filter and click search button. 3. Entries are not doubled. 	<p>The code is modified to clear the data in the table before starting a new search.</p> <p>Workaround: No Workaround.</p>

SBX-106583	2	<p>With the Q1912 in a 302 Redirect msg, the SBC sends rn and npdi even though the Contact Header in 302 does not have rn and npdi.</p> <p>Impact: When making a call where the INVITE contains NPDI/RN parameters and the SIP variant on the ingress trunk group is set to Q1912 when the call goes through 3xx processing the wrong called party number is sent in the subsequent INVITE.</p> <p>e.g. The initial INVITE contains INVITE sip:11111111;npdi;rn=222222@<IP>:<PORT></p> <p>The policy dip performs an ENUM query and gets a new called party number of 33333333</p> <p>The initial INVITE goes out with INVITE sip:33333333@<IP>:<PORT></p> <p>The end point responds with 3xx and the subsequent INVITE then contains INVITE sip:11111111@<IP>:<PORT> because when using the Q1912 the TOA_PORTED_NUMBER of 11111111 is passed up to in the second policy dip and confuses the routing logic.</p> <p>Root Cause: When processing the 3xx and making a second policy dip, the code was meant to remove the generic number parameter of type TOA_PORTED_NUMBER. The code assumed there would only ever be one generic number parameter present. However, when the ingress SIP variant is set to Q1912 the SIP code internally creates a generic number parameter of type additional calling party number based on the contents of the FROM header. In this scenario, the code was unable to delete the generic number of type TOA_PORTED_NUMBER and this was passed back to the PSX in the second policy dip and resulted in routing confusion and unexpected parameter content in the INVITE following the 3xx.</p> <p>Steps to Replicate: Perform a call where the SIP variant on the ingress trunk group is set to Q1912 and the INVITE contains the npdi/rn parameters e.g. INVITE sip:11111111;npdi;rn=222222@<IP>:<PORT> The initial INVITE goes out with INVITE sip:33333333@<IP>:<PORT> The end point responds with 3xx and the subsequent INVITE then contains INVITE sip:33333333@<IP>:<PORT></p>	<p>The code is modified to correctly delete the generic number parameter of type TOA_PORTED_NUMBER associated with the number translation information from the first policy dip to avoid routing confusion on the second policy dip.</p> <p>Workaround: If possible change the SIP variant on the ingress trunk group to be "sonus" instead of "Q1912" to avoid the creation of the generic number parameter of type additional calling party number based on the FROM header contents.</p>
SBX-103950	2	<p>There are PAI differences between the SBC and GSX.</p> <p>Impact: The SBC was using the ingress calling URI information to create the egress PAI header contents, even when the calling party number digits are all removed using DM/PM rules in PSX.</p> <p>Root Cause: Because the SBC supports additional functionality that is not present on the GSX, the SBC can use the contents of the calling URI from the ingress side of the call even when the PSX does not set the username mask flag in the calling URI parameter in the policy response.</p> <p>Steps to Replicate: Configure the PSX to generate the PAI header on the egress INVITE, but remove all the calling party digits, all the calling URI username digits and the generic name parameter. Then make a basic call and check that the SBC does not include PAI in the egress INVITE.</p>	<p>The code is modified to check that the PAI header contains username and/or displayname parameters before adding it into the egress INVITE. On the PSX, when the customer wants to delete the calling party number digits, they also need to remove the calling URI username and the generic name information.</p> <p>Workaround: Need to use SMM to remove the PAI header.</p>
SBX-106197	2	<p>The PKT port manualSwitch triggered twice.</p> <p>Impact: Executing port switchover command from CLI triggers port switchover twice.</p> <p>Root Cause: Link detection sub system on standby node subscribes to port switchover action command with CDB (Confd) twice:</p> <ol style="list-style-type: none"> 1. When the node comes up as standby. 2. When it initiates switchover to become active. <p>When standby node becomes active, executing port switchover command from CLI results in notifying link detection sub system twice due to duplicate subscription. This results in initiating port switchover twice.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Bring up nodes with port redundancy in HA mode. 2. Initiate CE switchover to make sure standby node becomes active. 3. Execute port switchover command on new active node and validate the behavior. 	<p>The code is modified to not subscribe for action command when the node is in standby mode and subscribe only upon switchover when it transitions from standby to active.</p> <p>Workaround: None.</p>
SBX-107372	2	<p>Observing lot of errors in DGB files since upgrading to 9.2.0R001.</p> <p>Impact: Recently, a fix was added in the SBC for updating the subscription timer on getting a NOTIFY message with the value of expires param in Subscription-State header. But as per RFC 3265 this param is a "should" and is not mandatory.</p> <p>Root Cause: When the SBC updated the subscription timer on receiving a NOTIFY message, the case for the expired param missing was not considered.</p> <p>Steps to Replicate: Create a subscription on the SBC, with SUBSCRIBE relay.</p> <ol style="list-style-type: none"> 1. From UAC send SUBSCRIBE to the SBC, which it relays to UAS. 2. Send a NOTIFY without expires param to the SBC from UAS. 3. The log "SipSgStartUpdatedSubscriptionTimer: Invalid subscription timer 0 second for callId" should not be seen after fix. 	<p>The code is modified to address the issue.</p> <p>Workaround: If it is really required, SMM could also be used to send some expires param in NOTIFY Subscription-State header.</p>

SBX-107999	2	<p>The LeakSanitizer: CpxDnsCreate leaks observed during SBX-85432/<redacted> E2E C-SBC automation run.</p> <p>Impact: There is a small memory leak in the Cpx process when DNS elements are added or deleted.</p> <p>Root Cause: As part of adding/deleting DNS entries the Cpx process reads in the interface group name from CDB and stores it in a local buffer but does not free it when finished processing.</p> <p>Steps to Replicate: Add DNS server entries.</p>	<p>The code is modified to correctly free up the local buffer at the end of the configuration logic.</p> <p>Workaround: None.</p>
SBX-107978	2	<p>There are coverity issue in the SIPSG.</p> <p>Impact: There are coverity issues for NULL check.</p> <p>Root Cause: The null pointer dereferences while handling a session refresh.</p> <p>Steps to Replicate: Run a call where re-INVITE does not contain SDP.</p>	<p>The code is modified to add the NULL check before accessing the pointer.</p> <p>Workaround: No workaround.</p>
SBX-105105	2	<p>Adjust the MAX_LEN_REALM in camRfAppRedund.h to correct length.</p> <p>Impact: If the realm string under realmRoute of a diamNode configuration is configured with more than 129 characters, there is a possibility of the string value not being retrieved correctly from the cdb after the SBC restarts or not being made redundant correctly on a HA system.</p> <p>Root Cause: The maximum size for the realm string was incorrectly defined as 129.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. On a HA SBC system, Configure diamNode with a realmRoute that has realm string greater than 129 chars. For example: set addressContext default diamNode DIAMNODE realmRoute RX.EXAMPLE.COM realm ims111111111122222222223333333333444444444455555555556666666666 777777777888888888899999999900000000011111111122222222223333333333 mnc094.mcc235.3gppnetwork.org peer RX.EXAMPLE.COM appld rx 2. Show addressContext default diamNode to verify configuration has applied correctly and realmRoute has realm value string as configured. 3. Perform a switchover the SBC. 4. On the newly active SBC, show the addressContext default diamNode to verify configuration is the same and realmRoute has realm value string as configured. 	<p>The definition of maximum size for the realm string is modified to 257.</p> <p>Workaround: The realm string should be limited to 129 characters.</p>
SBX-106687	2	<p>The Platform Interface Status was incomplete.</p> <p>Impact: In the Network Tools, the Platform Interface Status section can be used to display the status of Interface selected from the left side section. For certain interfaces like pkt0, pkt1, mgt0, mgt1 the status message was incomplete.</p> <p>Root Cause: For interfaces like pkt0, pkt1, mgt0, mgt1 the status message displayed in textarea element, had more than one < and > characters and so the text within them was considered as HTML block by the browser and it was hidden.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Login to EMA, open Administration > System Administration > Network Tools. 2. Select interfaces in the "Platform Interface Status" section and ensure that the displayed status message matches with the status message from CLI. 	<p>The code is modified from <textarea> to <pre> to address the issue.</p> <p>Workaround: Not Available.</p>
SBX-104439	2	<p>The SBC is not generating the UPDATE message towards the UAC for the 183 Dialog-2 when there is delay in PRACK message.</p> <p>Impact: A call failure was observed during the processing of second dialog provisional response with SDP for downstream forking scenarios when PRACK is delayed at ingress.</p> <p>Root Cause: When the SBC needs to send UPDATE at ingress during downstream forking for codec re-negotiation while PRACK is pending, it tries to formulate an internal auto answer. If the UPDATE is offered with a new codec other than previously locked, the SBC while forming such local answer adds the previously negotiated codec. This created a mismatch in offer-answer state machine resulting in call failure.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Enable customer configurations. 2. Send 183 Dialog-2 from UAS. 3. Once 183 Dialog-2 received to UAC, Send PRACK for that 183 after 2 seconds delay. 4. Check if the SBC sends UPDATE for the second dialog 18x. 	<p>The code is modified to consider presence of "doNotAutoAnswer" TrunkGroup flag to decide if a local answer needs to be formed or wait until PRACK comes to release the UPDATE.</p> <p>Workaround: Not Applicable.</p>

SBX-104591	2	<p>The SBC is releasing the call during customer-1-1-4 call flow when 200OK answered with dialog-1 and delay in ACK from UAC</p> <p>Impact: Call failure observed during processing of second dialog provisional response with SDP for downstream forking scenarios when PRACK is delayed at ingress</p> <p>Root Cause: When the SBC need to send UPDATE at ingress during downstream forking for codec re-negotiation while PRACK is pending, it tries to formulate an internal auto answer. If the UPDATE is offered with a new codec other than previously locked, the SBC while forming such local answer adds the previously negotiated codec. This created a mismatch in offer-answer state machine resulting in call failure.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Enable customer configurations. 2. Send 183 Dialog-2 from UAS. 3. Once 183 Dialog-2 received to UAC, Send PRACK for that 183 after 2 seconds delay. 4. Check if the SBC sends UPDATE for the second dialog 18x. 	<p>The code is modified so consider the presence of "doNotAutoAnswer" TrunkGroup flag to decide if such local answer need to be formed or wait until PRACK comes to release the UPDATE.</p> <p>Workaround: Not Applicable</p>
SBX-108066	2	<p>The /opt/sonus/sbx/bin/opensslSelftest cmd failed in the SBC 7000.</p> <p>Impact: When the FIPS mode is enabled on hwType SBC7000 (7k), services will fail to come up and the SBC becomes inaccessible.</p> <p>Root Cause: After enabling the FIPS mode when FIPS selfTests run, opensslSelfTests fails due to using a deprecated SSL engine on SBC 7000.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Enable FIPS mode from CLI. 2. Ensure the SBC services are accessible after reboot. 	<p>The code is modified to remove the deprecated SSL engine from opensslSelfTest and this allows FIPS mode to function as expected.</p> <p>Workaround: None.</p>
SBX-107486	2	<p>The EMA display error when the SMM is having criteria with reg-exp like "cpm.msg cpm.largemsg".</p> <p>Impact: The EMA display error when SMM is having criteria with reg-exp like "cpm.msg cpm.largemsg".</p> <p>Root Cause: The EMA is not checking whether nonmatch value is present in database or not.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Log into the SBC from CLI. 2. Configure the profile rules provided in the Description in the DB. 3. Log into the EMA GUI. 4. Navigate Profiles->Signaling ->sipAdaptorProfile. 5. Edit the Configured profile that is done from the CLI. 6. There we can see the issue is fixed. 	<p>The code is modified to address the issue.</p> <p>Workaround: Not Applicable.</p>
SBX-104685	2	<p>The SBC is not triggering UPDATE to ingress and egress for 183 Dialog-2 during E2E precondition interworking. Issue with ISBC(SIP_FW)</p> <p>Impact: The SBC was not sending UPDATE towards ingress while process downstream forking call.</p> <p>Root Cause: Precondition values are not properly updated in the data structures, resulting in failure to send the UPDATE.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Bring up customer E2E setup. 2. Transparency precondition scenario. 3. Send 18x dialog-2 with same SDP that of 1st dialog 18x. 	<p>The code is modified so now the SBC processes the precondition values and triggers an UPDATE to ingress.</p> <p>Workaround: None.</p>
SBX-104624	2	<p>The SBC is not generating UPDATE message for 183 Dialog-2 when 183 Dialog-1 and 183 Dialog-2 having same SDP.</p> <p>Impact: The SBC was not sending UPDATE towards ingress while process downstream forking call.</p> <p>Root Cause: Precondition values are not properly compared at egress, resulting in failure to send the UPDATE.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Bring up customer E2E setup. 2. Transparency precondition scenario. 3. Send 18x dialog-2 with same SDP that of 1st dialog 18x. 	<p>The code is modified to compare the precondition values and if there is change an indication message is sent to ingress leg which further triggers an UPDATE.</p> <p>Workaround: None.</p>

SBX-108029	2	<p>The SIPSG_MSG_PARAM_SIPSG_REDUND_ACT_STR is not registered for serialization correctly before 10.0 release</p> <p>Impact: After LSWU, some of the event record fields are not serialized properly for the calls related to Register and OOD messages.</p> <p>Root Cause: Event record accounting details are not registered for Serialization.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. The SBC is an older release(< 9.2.1). 2. Configure to generate event record for Register. 3. Make a Register call. 4. Send a Refresh register. 5. LSWU to 9.2.1. 6. Complete the registration call. 7. Check the details of generated event record. 	<p>The code is modified to address the issue.</p> <p>Workaround: None.</p>
SBX-108025	2	<p>Receiving an unexpected CANCEL in the SBX-SBX GW early media case.</p> <p>Impact: 503 is sent for Update because of Server Request not present for Update Request</p> <p>Root Cause: Update Server Request is removed from SIP Dialog Structure because we are unable to send a 200 OK for the Update.</p> <p>Steps to Replicate: In this scenario Update is received from Egress and it is sent to ingress. In the Ingress, the 200 OK for this Update is delayed and not sent immediately.</p> <p>In the mean time, there is one more Update is received from egress. This is rejected with 500 Internal Error stating request already pending.</p> <p>Now, the 200 OK for the First Update is received from ingress and it is being sent to egress</p>	<p>The code is modified to not remove an Update request from Server request list in the SIP dialog request.</p> <p>Workaround: None.</p>
SBX-107243	2	<p>The mid call tracing is collecting only MGSG logs.</p> <p>Impact: If the C3 sends MODIFY command with only CallTraceActReq=ON and without any media parameters, then the media logs are not captured in TRC file even if media parameters are changed in subsequent MODIFY command from C3.</p> <p>Root Cause: Whenever the MRFP/MGSG receives MOD command with only CallTraceActReq=ON, the MGSG does not send any MODIFY command to NRMA and as a result call tracing info is not updated in NRMA and other media related modules.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Establish a MRFP call. 2. Send MODIFY command from C3 with only callTraceActReq=ON, no change in media parameters. 3. Send MODIFY command from C3 with changes in media parameters. 4. Verify that .TRC file contains logs from NRMA, XRM and other media related modules. 	<p>The code is modified to send modify command to NRMA whenever MGSG receives MOD command with only callTraceActReq=ON.</p> <p>Workaround: No Workaround.</p>
SBX-107642	2	<p>The SBC terminates a call as 491 does not get relayed.</p> <p>Impact: Run a Scenario with Re-invite without SDP and send 491 request pending for that re-invite. 491 is not being relayed even though relay4xx-6xx is enabled and E2E Re-invite is enabled.</p> <p>Root Cause: Even though Re-invite is received from the Network the following bit bisE2ENtwkReInviteReq in the SIP call structure is not set to true. This is causing 491 to be handled locally.</p> <p>Steps to Replicate: Run a Re-invite without SDP scenario and send 491 for that re-Invite.</p>	<p>The code is modified to address the issues.</p> <p>Workaround: None.</p>
SBX-104266	2	<p>The SBC is rejecting the call when 200OK of INVITE is answered with Dialog-1 .Dialog-1 is having the AMR-WB codec and Dialog-2 having the EVS codec</p> <p>Impact: Call failure observed during certain downstream forking scenarios when 200 OK of initially negotiated dialog comes.</p> <p>Root Cause: When the SBC need to send Re-Invite at ingress during downstream forking for codec re-negotiation while ACK is pending, it tries to formulate an internal auto answer. If the INVITE is offered with a new codec other than previously locked, the SBC while forming such local answer adds the previously negotiated codec. This created a mismatch in offer-answer state machine resulting in call failure.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Enable customer configurations. 2. Send 183 Dialog-1 with the AMR-WB codec. 3. Send 183 Dialog-2 with the EVS codec. 4. 200OK of INVITE sent with the dialog-1. 	<p>The code is modified to consider presence of "doNotAutoAnswer" TrunkGroup flag to decide if a local answer needs to be formed or wait until ACK comes to release the Re-Invite.</p> <p>Workaround: Not Available.</p>

SBX-93922	2	<p>The SBC is sending Min-Expires header twice to endpoint for 423 Interval Too brief response.</p> <p>Impact: The Min-Expires header is sent twice in the REGISTER response.</p> <p>Root Cause: The Min-Expires header was sent twice as it was getting added by the application and the transparency profile.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Enable Min-Expires headers in the Transparency profile. 2. Send a REGISTER message to SBC from UAC. 3. UAS sends 423 response to REGISTER message received from SBC. 	<p>The code is modified to address the issue.</p> <p>Workaround: Remove the Min-Expires headers from the Transparency profile.</p>
SBX-106995	2	<p>The SBC is not relaying the 183 Dialog-2 towards UAC when it receives UPDATE with preconditions from UAS for Dialog-1.</p> <p>Impact: The SBC was not relaying the second dialog 18x to ingress after the completion of preconditions UPDATE for first dialog during downstream forking scenario</p> <p>Root Cause: The SBC was consuming the second dialog 18x at egress and skipped further processing. When P-Early-Media value received as part of second dialog 18x had less precedence than that of previous dialog, the SBC did not process the second dialog 18x.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Bring up the customer E2E setup. 2. From UAS sent 183 Dialog-1 with preconditions. 3. The SBC would send the UPDATE message for the preconditions met for dialog-1 4. UAS sends 183 Dialog-2 with P-E-M sendonly. 	<p>The code is modified to follow second dialog 18x and not to consider P-Early-Media value precedence during this process.</p> <p>Workaround: Modify the value of P-E-M header to sendrecv using SMM.</p>
SBX-104253	2	<p>The "Deleting the a=maxtime" SMM is not working after an SBC restart.</p> <p>Impact: SMM with sdpContent not getting applied after restart.</p> <p>Root Cause: At time of restart CPX trying to read SMM with sdpContent from the wrong path.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Spawned an HA setup. 2. Configured SMM rule with sdpContent and run the scenarios. 3. Switchover the system and re-run the scenario. 4. Switchover again and re-run the scenario. 	<p>Fixed the path from where sdpContent will be read for a SMM action.</p> <p>Workaround: Delete the SMM rule and configure again.</p>
SBX-105544	2	<p>Dialog scope variable is not present after SMM reject on re-invite message.</p> <p>Impact: Dialog scope variable is not present when it is saved in the request which is rejected by SMM reject Operation.</p> <p>Root Cause: Dialog scope saved in case of rejected request seems to be not inserted in hash and freed.</p> <p>Steps to Replicate: Run a Re-invite call flow and reject the re-invite with 488 and have a SMM Rule to store dialog scope variables for that re-invite,</p>	<p>The code is modified to address the issue.</p> <p>Workaround: None.</p>
SBX-108278	2	<p>Enabling the facstate flag can cause healthcheck timeout in SBC.</p> <p>Impact: The 'facState' flag if enabled, may result in some potential crash that is caused by health check timeouts.</p> <p>Root Cause: The thread that is writing per call key elements in to the file is taking more time and resulting in health check failures.</p> <p>Steps to Replicate: Run the calls at 100cps rate for around 5 to 6 hours with flag enabled may hit the health check failure problem.</p>	<p>The code is modified to address the issue.</p> <p>Workaround: Disable the facState flag.</p>
SBX-107388	3	<p>Multiple threads are writing to same socket and causing issues.</p> <p>Impact: In case of CE_switchover, active went down and standby did not come up as new ACTIVE which led to unstable state until the active becomes stable again.</p> <p>Root Cause: To track a healthcheck, the VNFR used a ZMQ mechanism to send/receive messages from VNFCs. As a result, multiple threads tried to write at same time, which caused failure to send messages from VNFR to VNFCs.</p> <p>Steps to Replicate: Run multiple CE switchovers with fix.</p>	<p>The code is modified to use the mutex to avoid concurrent writing at same time with its lock/unlock mechanism.</p> <p>If one thread is writing then other threads need to wait in queue until the previous one is not done.</p> <p>Workaround: None.</p>

SBX-106850	3	<p>The eventLog combined log size calculation does not consider the disk size properly.</p> <p>Impact: For the SBC5400 product, the user is allowed to configure the maximum accounting file size and the maximum number of accounting files so that those files can consume up to 150GB of diskspace. However, the maximum amount of disk space allowed should be 80GB.</p> <p>Root Cause: The code that calculates the limit did not properly compute the disk limit if the product was the SBC 5400.</p> <p>Steps to Replicate:</p> <p>On an SBC 5400, set oam eventLog typeAdmin acct fileSize 65535 fileCount 2047 [ok][2021-02-17 15:02:03]</p> <p>commit Aborted: 'oam eventLog typeAdmin': The fileSize and fileCount values configured for the event logs would result in a potential combined log file usage of 128.4GB, which exceeds the maximum of 80.0GB allowed [error][2021-02-17 15:02:06]</p> <p>The preceding error message should be displayed indicating that the maximum log file usage is 80GB.</p>	<p>The code is modified to properly compute the limit for the disk limit if the product was the SBC 5400. The disk limit is 80GB.</p> <p>Workaround: Ensure that the sum of the fileSize*fileCount for each of the SBC log files does not exceed 83886080.</p>
SBX-107162	3	<p>There was an error when dumping System Diagnostics md5 files on the EMA.</p> <p>Impact: The checksum files were not shown while generating a system dump through the EMA/PM.</p> <p>Root Cause: The EMA/PM only searches for sha256 checksum files now, but the checksum generated was md5, so it was not visible in the list present in EMA GUI.</p> <p>Steps to Replicate: Use the EMA/PM to generate System Dump, and verified the presence of sha256 checksum file for every tarball generated.</p>	<p>The code is modified so the System Dump program now generates a sha256 checksum, instead of md5. This is visible to EMA/PM.</p> <p>Workaround: None.</p>

Resolved Issues in 09.02.00R002 Release

The following Severity 1 issue is resolved in this release:

Table 32: Severity 1 Resolved Issues

Issue	Sev	Problem Description	Resolution
SBX-106989 SBX-107017	1	<p>Portfix SBX-106989: Multiple Switchovers occurred because of a DSP core after an Upgrade to V09.02.00R000.</p> <p>Scenario: A transcoded call with any Codec<=>G711 and RFC2833 DTMF relay enabled on both legs of a call, and the media probe is disabled.</p> <p>Impact: If a peer device sends a RFC2833 without EOP packet (end-of-digit marker), then a DSP crash occurs.</p> <p>Root Cause: The code tried to add an error condition (LOST_EOP flag) to media probe data structure without checking whether a media probe was enabled or not. This resulted in a NULL pointer access and a subsequent memory protection fault resulting in a dsp core dump.</p> <p>Steps to Replicate: The customer case had a G711A<=>G711A transcoded call with RFC2833 enabled on both legs</p> <p>Use the following setup:</p> <p>SipP UAC=>SBX 5x10 => SipP UAS G729AB(RFC2833) <=> G711(RFC2833) show configuration system media mediaProbe state state disabled; From UAC send dtmf_2833_1_noEop.pcap with '1' digit w/o EOP packets</p> <p>This resulted in dsp core dump before the fix and no core dump was observed after the fix and digit '1' was correctly sent to g711 side.</p>	<p>The code is modified so if the media probe is enabled and not write to media probe structure if media probe is disabled.</p> <p>Workaround: Use RFC2833<=>Inband digits or RFC2833<=>SIP Info instead of RFC2833<=>RFC2833.</p>
SBX-103881 SBX-105740	1	<p>PortFix SBX-103881 to 9.2.x - MEO: Qseries CDR field not matching Qseries format</p> <p>Impact: When rn parameter is received as part of SIP R-URI, this value is being used to populate field 10 of the QSBC format CDR record, rather than the received Called Number</p> <p>Root Cause: In case of rn parameter received, source for field 10 of QSBC CDR is incorrect</p> <p>Steps to Replicate: Send SIP INVITE with R-URI including rn parameter.</p>	<p>Change code such that if rn is received, field 10 of QSBC is populated with the "Called Number Before Translation" (field 24 of Ribbon format CDR). For other scenarios, field 10 is populated as before.</p> <p>Workaround: None.</p>

SBX-106206 SBX-107018	1	<p>Portfix SBX-106206: An existing hairpin call gets silenced after a switchover.</p> <p>Impact: When switchover occurs in the SWe SBC Active-Standby HA with the SBC internal loopback media calls flows, the media loopback is not working in existing calls. The looped back media is dropped internally in the SWe NP, because packet skb was not including the internal loopback flag set condition in the NP incoming path validation checks from the internally looped back media for the existing calls.</p> <p>Root Cause: When a switchover occurs in the SWe SBC for existing medial calls, a loopback occurs in the SWe NP based on the matching SA.</p> <p>The DA IP address was configured in BRES by setting an internal loopback flag for packet work /skb, even though MAC DA is not updated in BRES flows to the new active instance interface MAC Address to match SA MAC.</p> <p>However, this internal skb loopback flag settings was removed in the release 9.0 release NP refactoring code which caused the packet drops in incoming packet checks for MACs at NP from the looped back media of existing calls. The new calls will not have this issue post-switchover.</p> <p>Steps to Replicate: With the SWe SBC Active-Standby HA setup, using AS/3GPP call flow scripts /setup, establish a call that creates the SBC internal loopback media flow. Verify the media flow and then Issue a switchover and verify the media again on this existing active call.</p>	<p>The code is modified to save NP the CPU cycles also while addressing the loopback media issue after a switchover issue for existing calls.</p> <p>Workaround: None.</p>
SBX-105269 SBX-107044	1	<p>Portfix SBX-105269: The SBC crashed and a core dump created.</p> <p>Impact: The SCM processes coredump due to NULL pointer access on a pointer that has been freed due to BYE and HOLD REINVITE in a transfer call scenario.</p> <p>Root Cause: There was an illegal memory access, absence of NULL check, exposed due to REINVITE.</p> <p>Steps to Replicate: A calls B, B REFERS to C and now A and C talk. After sometime, A sends BYE and C sends re-INVITE with a=inactive at the same time.</p>	<p>The code is modified to address the issue.</p> <p>Workaround: None.</p>

Resolved Issues in 09.02.00R001 Release

The following Severity 1 issue is resolved in this release:

Table 33: Severity 1 Resolved Issues

Issue	Problem Description	Resolution
SBX-106080	<p>Release 9.2.0 should have CAM version of "00610000".</p> <p>New CDR fields were added in the 9.2.0R0 release, but the CAM version in the top of the ACT header file was not updated to "00610000". It was still set to "00600000", which was used for release 9.1.0R0.</p> <p>Impact: Incorrect CAM version in the top of the ACT header file.</p> <p>Root Cause: The code to increment the CAM version was missed.</p> <p>Steps to Verify Fix: Run a basic call and check the ACT log header to confirm the CAM version is set to "00610000"</p>	<p>The code is updated to correctly set the CAM version to "00610000".</p>

Resolved Issues in 09.02.00R000 Release

The following Severity 1 issues are resolved in this release:

Table 34: Severity 1 Resolved Issues

Issue	Problem Description	Resolution
-------	---------------------	------------

<p>SBX-101451</p>	<p>The DSP Threshold setting is not generating a trap on SBC 5400.</p> <p>Impact: The g711PacketThreshold, g729Threshold, and g726Threshold onset and abate traps are not sent.</p> <p>Root Cause: The NRM did not receive CLI updates to the g711PacketThreshold, g729Threshold, and g726Threshold, and the trap generation code used wrong trap names.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Provision the threshold levels: set oam traps dspAdmin dspAvailabilityTrap g729Threshold 40 set oam traps dspAdmin dspAvailabilityTrap g726Threshold 60 set oam traps dspAdmin dspAvailabilityTrap g711PacketThreshold 80 commit set oam traps admin sonusSbxDspAvailG729OnSetCrossThresholdNotification state enabled commit set oam traps admin sonusSbxDspAvailG729AbateCrossThresholdNotification state enabled commit set oam traps admin sonusSbxDspAvailG726AbateCrossThresholdNotification state enabled commit set oam traps admin sonusSbxDspAvailG726OnSetCrossThresholdNotification state enabled commit set oam traps admin sonusSbxDspAvailG711OnSetCrossThresholdNotification state enabled commit set oam traps admin sonusSbxDspAvailG711PacketAbateCrossThresholdNotification state enabled commit 2. Configure the trap target: set oam snmp trapTarget EMS160 ipAddress 10.xxx.xx.xxx port 162 trapType v2 state enabled commit. 3. Limit the compression resources to make the issue readily occur: set system mediaProfile tone 98 compression 2 commit. 4. Perform a bunch of transcoded (e.g. G711 to G729) calls that exceed threshold limits, and see onset traps are sent to the trap target. 5. Clear the transcoded calls, and see abate traps are sent to the trap target. 	<p>The code is modified to support g711PacketThreshold, g729Threshold, and g726Threshold values.</p> <p>Deprecated the support for allThreshold, as it was never implemented in the SBC.</p> <p>Workaround: None.</p>
<p>SBX-102833</p>	<p>Calls are dropping when placed on hold through MS Teams.</p> <p>Impact: When MS Teams puts a call on hold and then attempts to retrieve it while the microphone is disabled, the action results in the call being cleared.</p> <p>Root Cause: MS Teams performs the call retrieve by sending an INVITE with a replaces message to the SBC and when the microphone is disabled this message includes a=recvonly.</p> <p>The SBC send 200 Ok response to the INVITE with replaces with a=sendrecv while Teams is expecting a=sendonly that causes the call to be cleared by MS Teams.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Establish a PSTN to Teams call through the SBC. 2. Place call on hold from Teams. 3. Disable the microphone in browser running Teams and then retrieve (place off hold) the call. 4. The call should retrieve should be successful. 	<p>The code is modified to respond to an INVITE with replaces containing a=recvonly, with 200 OK containing a=sendonly.</p> <p>Workaround: None.</p>
<p>SBX-103974</p>	<p>The STI sends 2 Identity fields in Contact header, but SBC only passes one</p> <p>Impact: In the case where multiple "Identity Headers" are present as embedded headers in a 3xx "Contact Header", the SBC sends out only 1 Identity Header for the re-directed INVITE.</p> <p>Root Cause: The SBC did not support handling of the scenario where SIP headers are repeated in embedded contact headers in 3xx. The SBC would end up the sending only the last of the repeated headers in the re-directed Invite.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. The SBC is configured with ERE, set the config required for handling embedded headers in 3xx as shown below. "set profiles signaling ipSignalingProfile DEFAULT_SIP egressIpAttributes redirect flags forceRequeryForRedirection enable honorEmbeddedHeadersIn3xx enable" 2. Run a UAS script with 302 redirection and add multiple "Identity headers" as shown below. Ex- Contact: 9710622482<sip:9710622482@10.xx.xx.xx:xxxx?Identity=test1&Identity=test2&Identity=test3> 3. Run a 2nd UAS Script to handle the Redirected INVITE. 4. Run SIPP call, trigger UAC script. <p>Expected Result: The redirect INVITE shall contain both the Identity headers received in 3xx as embedded contact header.</p>	<p>The code is modified to support repeated SIP headers.</p> <p>Workaround: None.</p>
<p>SBX-103645</p>	<p>A customer SBC upgrade had a spiltbrain SBC boot issue</p> <p>Impact: One of the customer SBCs in N:1 mode is not coming up after upgrade.</p> <p>Root Cause: The customer SBC node went for multiple reboots due to config profile change and split brains after upgrade exceeding max limit for reboot.</p> <p>Steps to Replicate: Upgrade all customer SBC nodes in one shot so that they all come up together.</p>	<p>The code is modified to avoid split brain due to nodeld/serviceId collision by informing the peer nodes about the self node's nodeld/serviceid as soon as its allocated.</p> <p>Workaround: Restart the customer SBC application manually on the node which has exceeded max reboot limit.</p>

SBX-99850	<p>The BYE message was sent to the wrong port.</p> <p>Impact: When endpoint is registered over TLS initiates a calls and after call is established, endpoint sends refresh register from modified port and also sends refresh INVITE with no SDP change, the SBC does not send BYE to modified port upon call disconnect.</p> <p>Root Cause: The SBC does not update remote connection address when refresh INVITE is received from different port.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Register endpoint over TLS. 2. Endpoint initiates a call, call gets connected. 3. Endpoint sends refresh register from modified port. 4. Endpoint sends refresh re-INVITE from modified port. 5. Called party disconnects the call. 6. Verify SBC sends BYE to modified port towards caller. 	<p>The code is modified to ensure the SBC update remote address when the refresh INVITE is received from modified port.</p> <p>Workaround: None.</p>
SBX-86293	<p>PreInstall Check improvements for file permissions.</p> <p>Impact: PreUpgrade checks failure due to permission issues on external directory.</p> <p>Root Cause: Pre-checks script failed to run commands on peer due to permission issues.</p> <p>Steps to Replicate: Perform upgrade to the fix version and verify that upgrade is successful.</p>	<p>The code is modified to ensure permissions/ownership are set properly before running through pre-checks /upgrade.</p> <p>Workaround: Set the right ownership for external directory as: <code>chgrp -R upload /opt/sonus /external/</code></p>
SBX-103057	<p>The IAC:terraform apply output is displaying only mgt Active and Standby IP's. I</p> <p>Impact: When orchestrating a HFE 2.0 or HFE 2.1 setup in Azure using Ribbon RAF modules, the Terraform output does not show any information about HFE node(s) created, even though the HFE nodes are started correctly.</p> <p>Root Cause: Terraform was missing the configuration to tell it to output the information about the HFE node(s).</p> <p>Steps to Replicate: Create a HFE 2.0 or HFE 2.1 setup using Ribbon RAF modules.</p>	<p>The code is modified so that Terraform outputs the management IP and instance name for each HFE node when the orchestration is successful.</p> <p>Workaround: None.</p>
SBX-103571	<p>There are call failures due connection toggling between the SBC and SLB.</p> <p>Impact: The SBC's can lose connectivity to the HFE nodes when there is high amounts of traffic on the network in Azure</p> <p>Root Cause: The health check timeout from the HFE to SBC is too small, therefore was not receiving the replies back from the SBC, causing HFE to switchover</p> <p>Steps to Replicate:</p>	<p>Increased the timeout from 20ms to 100ms</p> <p>Workaround: Manually edit the following line in HFE_AZ.sh from: <code>\$FPING -c 3 -t 20-p 200</code> <code>\$(ACTIVE_SBC_IP_ARR[0]) &> /dev /null</code> to: <code>\$FPING -c 3 -t 100 -p 200</code> <code>\$(ACTIVE_SBC_IP_ARR[0]) &> /dev /null</code></p>
SBX-103058	<p>The IaC is creating same NSG for all the network interfaces created. Recommended SGs should be created for each interface.</p> <p>Impact: IaC is creating single NSG for all the 4 sbc interface.</p> <p>Root Cause: lac is creating one NSG and using same for all sbc interfaces created.</p> <p>Steps to Replicate: Provided option to create different NSG for different interface created for the SBC.</p> <p>Option available in terraform.vars Ex: <code>sbc_security_group_names = ["TF-SVT-SG-MGT0", "TF-SVT-SG-HA", "TF-SVT-SG-PKT0", "TF-SVT-SG-PKT1"]</code></p>	<p>The code is modified to create different NSG for different interface of the SBC.</p> <p>Workaround: None</p>
SBX-103277	<p>The Active SLB failed in the LCA to coming up "getInterfaceInfoFromMDS: Could not get the interface details in the metadata information!"</p> <p>Impact: If the Azure fails on the request when verifying the authorization of the managed identity, then the cloud-init immediately fails.</p> <p>Root Cause: The cloud-init is missing tolerance when testing the authorization.</p> <p>Steps to Replicate: TBD</p>	<p>Add retries if the connection is reset or times out to address the issue.</p> <p>Workaround: Reboot the instance after cloud-init has failed.</p>
SBX-103781	<p>The LeakSanitizer detected memory leaks on various processes for a t140 Call.</p> <p>Impact: Small memory leak while configuring SNMP trap targets.</p> <p>Root Cause: While processing the configuration requests the SBC code was reading content from CDB into local memory blocks but failed to release the memory blocks at the end of the configuration action.</p> <p>Steps to Replicate: Configure the SNMP trap targets.</p>	<p>The code is modified to correctly free the internal memory blocks used to hold the temporary CDB configuration data.</p> <p>Workaround: None</p>

SBX-104494	<p>The PES Process cored while testing Flex AD support with the ERE (SBX-72926).</p> <p>Impact: PES process dumps core while executing AD service.</p> <p>Root Cause: While executing the AD service, PES fetches the AD profile from cache. Since there is no AD profile was configured in this case, cache shall return NULL. The code was missing to check for NULL for the AD profile pointer.</p> <p>Steps to Replicate: Configure a AD number translation criteria with ingress TG as the trigger criteria type. Make a call on the ingress TG. While executing the AD service, PES process will core dump.</p>	<p>The code is modified to check for a NULL for AD profile and AD attribute profile. If there is no profiles configured, the control returns and the service is marked as failed.</p> <p>Workaround: Configure an AD Profile and AD Attribute profile so that the existing code can be executed further.</p>
SBX-104769	<p>The SCM Process cored while testing SBX H323 Conformance and Interworking.</p> <p>Impact: The SCM Process cored while running SIP-H323 interworking call to verify rfc2833 to inband DTMF.</p> <p>Root Cause: Accessing of a NULL pointer caused the core dump.</p> <p>Steps to Replicate: Run a SIP-H323 interworking call with SIP side RFC 2833 and H.323 side inband DTMF being enabled.</p>	<p>The code is modified to check for NULL pointer before accessing it.</p> <p>Workaround: None.</p>
SBX-105263	<p>Observed a PRS Process core followed by a "DsPr, Ssre, Cpx" core on both active and standby box, for the Openstack S-SBC HA pair and "Ccsp" process core on the TSBC active while running calls on the 2CPS with the ASAN build.</p> <p>Impact: There was a memory leak detected by the ASAN in the active S-SBC while running the G.711 to G.729 transcoding load on the D-SBC.</p> <p>Root Cause: A code analysis of the ASAN reported function revealed that in some corner scenarios, the SBC may end up leaking the memory related to the place holder used to store the remote DSP resource.</p> <p>Steps to Replicate: Re-run the same load in ASAN and ensure no memory leaks are detected</p>	<p>The code is modified to free this memory.</p> <p>Workaround: None.</p>
SBX-102082	<p>The SCM Process core dump was observed when REGISTER is received and useRandomUserInfolnContactHdr is enabled.</p> <p>Impact: Invalid memory access issue when both embeddedRegInfolnUserPart and useRandomUserInfolnContactHdr flag were enabled</p> <p>Root Cause: In the SBC, the appropriate memory was not allocated for contactUrl puchUsername in case both embeddedRegInfolnUserPart and useRandomUserInfolnContactHdr flag were enabled, due to this invalid memory access was happening.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. On SLB - Enable remoteDeviceType to access 2. On SBCs enable the embeddedRegInfolnUserPart flag as:- set addressContext <ADDRESS_CONTEXT> zone <ZONE> sipTrunkGroup <TG> signaling embeddedRegInfolnUserPart enabled 3. On SBCs enable useRandomUserInfolnContactHeader as :- set addressContext <ADDRESS_CONTEXT> zone <ZONE> sipTrunkGroup <TG> signaling useRandomUserInfolnContactHdr enabled 	<p>The code is modified to allocate appropriate memory for puchUsername based on embeddedRegInfolnUserPart or useRandomUserInfolnContactHdr enable case.</p> <p>Workaround: As per the SBC design, any of the flags below should be enable not both. embeddedRegInfolnUserPart or useRandomUserInfolnContactHdr flag</p>
SBX-98177 SBX-104816	<p>PortFix SBX-98177: The SBC 7000 TCP window size.</p> <p>Impact: Add CLI support to set the kernel parameter net.ipv4.tcp_window_scaling.</p> <p>Root Cause: In certain SBC deployment scenarios (a large number of devices running SIP-TLS), the tcp_window_scaling caused small TCP window size. In such cases, the customer had to use Linux command to disable tcp_window_scaling on each SBC node.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Show initial setting for tcp_window_scaling on both active (SBCSWE01a) and standby (SBCSWE01b) [root@SBCSWE01a ~]# cat /proc/sys/net/ipv4/tcp_window_scaling=1 [root@SBCSWE01b ~]# cat /proc/sys/net/ipv4/tcp_window_scaling=1 2. Use CLI to change the kernel param value and display the changed value on both active and standby. admin@SBCSWE01a% set system admin SBCSWE01 kernelParams tcpWindowScaling disable admin@SBCSWE01a% commit Commit complete. [root@SBCSWE01a ~]# cat /proc/sys/net/ipv4/tcp_window_scaling=0 [root@SBCSWE01b ~]# cat /proc/sys/net/ipv4/tcp_window_scaling=0 	<p>The code is modified for setting the kernel parameter net.ipv4.tcp_window_scaling. The Linux command is no longer needed.</p> <p>Workaround: Use Linux command to set the kernel parameter on both active and standby. [root@SBCSWE01a ~]# sysctl net.ipv4.tcp_window_scaling=1 [root@SBCSWE01a ~]# sysctl -w net.ipv4.tcp_window_scaling=0 [root@SBCSWE01b ~]# sysctl net.ipv4.tcp_window_scaling=1 [root@SBCSWE01b ~]# sysctl -w net.ipv4.tcp_window_scaling=0</p> <p>To be persistent across SBC restart and system reboot, use Linux the following shell command on both active and standby. [root@SBCSWE01a ~]# echo "net.ipv4.tcp_window_scaling=0" >> /etc/sysctl.conf [root@SBCSWE01b ~]# echo "net.ipv4.tcp_window_scaling=0" >> /etc/sysctl.conf</p>
SBX-105351 SBX-105366	<p>Portfix SBX-105351: The 823R0 SCM Process cored.</p> <p>Impact:The SCM cores after a long extensive load of the OOD relay.</p> <p>Root Cause: The SBC accesses an invalid pointer of link list when the cleanup relay control block.</p> <p>Steps to Replicate:System required HW_TYPE_SBS5000 and max support Relay per SCM is 4194304. Run the OOD relay load (Notify). After a long period time (at least 4m relay per SCM).</p>	<p>The code is modified to properly initialize the link list.</p> <p>Workaround: None.</p>

<p>SBX-105019 SBX-105184</p>	<p>Portfix SBX-105019: A crosstalk issue seen on the SBC 5400.</p> <p>Impact: When there is a transcoded call with Opus codec and Opus termination does not send any packets to the SBC, then other termination of opus call, may hear audio from another unrelated opus transcoded call or it may hear white noise.</p> <p>This occurs on SBC5x10/5400/7K but not on SWE platforms. This problem only happens if no opus packet is received. When the decoder receives first opus packet, subsequently problem does not happen for that call.</p> <p>Root Cause: Initially, when no packet is received, Opus decoder is not called. As a result, the decoder's output buffer is uninitialized and it happened to be stale buffer of another opus call or a leftover buffer of a previous Opus call on that DSP. That buffer is used to encode data for other leg of the Opus call and that results in cross talk audio or distorted noise to other termination of Opus call.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Make two sipp Opus<=>g711 calls, say call 1 and call 2 on SBC5x10/5400/7K 2. These two calls have to land on same DSP core. So issue a unhide debug command 'request sbx drm debug command "loadbalance disable" before making two calls. Without this command, both calls may not land on same DSP core and you may not see same problem. 3. Call 1 sends audio packet (.pcap) from Opus termination but call 2 does not send any packets to the SBC from it's Opus termination. 4. Call 2's g711 termination will hear audio from call 1's opus termination. 	<p>The fix is to clear the buffer (with silence) when the Opus decoder is not primed with any packet.</p> <p>Workaround:</p> <ol style="list-style-type: none"> 1. Use a different codec other than Opus. 2. Use SWE platform instead of h/w SBC5x10/5400/7000.
<p>SBX-102122 SBX-104540</p>	<p>Portfix SBX-102122: SAM Process core dumps with v06.02.01 F012.</p> <p>Impact: The SAM and PRS process cored in the Standby node.</p> <p>Root Cause: The DRBD split brain lead to DRBD GI reset. This leads to a DRBD full sync, increasing the i/o congestion and ultimately causing a healthcheck timeout.</p> <p>Steps to Replicate: To test, run "drbdadm disconnect mirror". Once we get "standby split-brain" log in DBG logs, run "drbdadm disconnect mirror" again so that DRBD does not connect automatically and we get into the block which we need to test. Run "drbd-overview" and check if drbd does not went for full sync.</p>	<p>The code is modified to perform the DRBD checks after DRBD split-brain recovery.</p> <p>Workaround: None.</p>
<p>SBX-102625 SBX-104065</p>	<p>PortFix SBX-102625: The SIP calls stop in receiving 183 (Precondition).</p> <p>Impact: Call stopped working after 183 received without SDP.</p> <p>Root Cause: As part of egress precondition interworking, to negotiate the SBC should send an UPDATE to the egress endpoint. When 183 received without SDP, SBC should not try to send UPDATE. SIPSG call state assumed UPDATE is sent. Therefore, when the next 183 is received it tried to queue the 183 assuming UPDATE is in progress.</p> <p>Steps to Replicate: As mentioned in the JIRA</p>	<p>The code is modified to avoid this state transition when 183 received without SDP.</p> <p>Workaround: Remove the 183 without a SDP message using SMM.</p>
<p>SBX-102837 SBX-104152</p>	<p>PortFix SBX-102837: The fra-tdg-sonus-01 SipSignalingPorts became OutOfService after a switchover.</p> <p>Impact:The SIP sigPorts stuck in the OOS after a switchover.</p> <p>Root Cause: The customer had the network/pkt port issue, on one of their HA node, which caused pkt port(s) to bounce randomly, i.e. pkt port went DOWN and came back UP within 2 to 3 seconds. So they have been keeping that node as standby node. They also have link detection enabled for pkt port(s) and have around 100 LIFs per pkt port, one per SIP SIGPORT. When the pkt port went down, NRS delays the port down event processing for 2 second to allow link failure detection to be ready and also to avoid the race condition between NRS and LVM. When a 2 second delay timer is up, the NRS starts to take down affected LIFs and notifies local SIPCM and SIPFE so they can take down affected SIP SIGPORTs.</p> <p>In SIPCM, all the sockets on the affected sigPort are put in a delete pending table and starts a 1 tick timer. Then the socket(s) is being deleted after 1 tick timer is up.</p> <p>When pkt port came back up, NRS processes the event with no delay and notifies SIPCM and SIPFE as well. Since there are around 100 LIFs, there were many messages exchanges between NRS/XRM/SIPCM/SIPFE. NRS LIF FSM has the mechanism in place to handle the timing issue and LIFs were all back in service. But SIPCM failed to activate some SIP SIGPORT(s) while binding the socket(s). These error messages indicated that SIPCM tried to activate the SIGPORT while it was still pending delete. Therefore SIGPORT got stuck in OOS(broken state) in both SIPCM and SIPFE on standby node. If there was a switchover happened later, user would then noticed one or more SIGPORTs were OOS. They have to manually bounce those SIGPORTs to bring them back in service.</p> <p>Steps to Replicate: The nature of the problem was the timing caused race condition. There is no good way to re-create/verify the fixes.</p> <p>Suggest to run regular SIP related regression tests.</p>	<p>The code is modified to:</p> <ol style="list-style-type: none"> 1. Introduce a new 1 tick timer in SIPCM_DATA_STR, activateRetryTimerId. 2. The deletePendingSocketTable and if the sigPort is found in the table, then start the 1 tick timer. When the timer is up, SipCmActivateCallSigPort() is invoked again. Once the sigPort is activated successfully, SIPCM notifies SIPFE as usual. <p>Workaround: No workaround.</p>
<p>SBX-101161 SBX-103677</p>	<p>PortFix SBX-101161: A memory leak was observed in the SAM process.</p> <p>Impact: A memory leak was observed in the SAM process.</p> <p>Root Cause: There is race condition in the code that handles CALL_AUDITs and CALL_CLEANUPs that can cause a memory leak.</p> <p>If the response to a CALL_AUDIT takes too long, the code that handles the "late" response has allocates memory for a structure that may never get freed.</p> <p>Steps to Replicate: This issue is caused by a race condition that cannot be forced - therefore we cannot specify steps to reproduce this issue.</p>	<p>The code is modified to add a timer every time a fault structure is allocated. When the timer expires, if the the structure still exists it will be freed.</p> <p>Workaround: There is no workaround.</p>

<p>SBX-104802 SBX-105186</p>	<p>PortFix SBX-104802: Enabling the Dialog Transparency causes calls to fail.</p> <p>Impact: The call fails for the dialog transparency and forking.</p> <p>Root Cause: When the ingress support PRACK and multiple 18x fork, the SBC fails to send 200OK to ingress.</p> <p>Steps to Replicate: Configure dialog transparency and downstream forking on egress. Ingress support PRACK, egress not support PRACK.</p> <p>The Egress answers 180fork1, 180fork2, 180fork3, and 200OKfork2 simultaneously.</p>	<p>The code is modified so the SBC resets the PRACK pending status properly, causing 200OK stuck in the queue.</p> <p>Workaround: Disable the dialog transparency.</p>
<p>SBX-104761 SBX-105065</p>	<p>PortFix SBX-104761: A SM Process coredump occurred on the server.</p> <p>Impact: The SM Process crashed while executing the "show table system syncStatus" command.</p> <p>Root Cause: The shell script used to get the oracle sync status - PolicyDBSyncStatus.sh - did not return within 10 seconds, causing a healthcheck timeout that caused the coredump.</p> <p>Steps to Replicate: This problem is not reproducible.</p>	<p>The code is modified to disable healthchecks while fetching the syncStatus.</p> <p>Workaround: There is no workaround.</p>
<p>SBX-103553 SBX-105277</p>	<p>PortFix SBX-103553: The SBC does not proceed multiple 183.</p> <p>Impact: The GSX-GW-SBX call erroneously queues 183 message at SBX when X-Service-Type precondition is received and downstream forking enabled.</p> <p>Root Cause: An interaction between downstream forking and GW-GW code.</p> <p>Steps to Replicate: Make GSX -> SBX SIP-GW-SIP call. Egress trunk group has X-Headers supported. Egress trunk group has downstream forking enabled.</p> <p>INVITE --> <- 183 with X-Service-Type: cf,precondition and SDP <- 183 with X Service-Type: cf,precondition and different SDP <- 180 with X Service-Type: cf and SDP <- 183 with X Service-Type: cf and no SDP This final 183 is queued.</p>	<p>The code is modified to not perform the queuing in GW-GW scenarios.</p> <p>Workaround: None.</p>
<p>SBX-104443 SBX-105071</p>	<p>PortFix SBX-104443: The SM Process and Cpx cores prevented fra-tdg-sonus-01-1 to come back up as standby after switchover</p> <p>Impact: The SM Process and Cpx processes cored on slot 1 after LDG triggered switchover to slot 2.</p> <p>The SM Process was deadlocked while retrieving peer NTP status.</p> <p>The Cpx cores happened after SM Process coredump, when writing SMM profiles into shared memory.</p> <p>Root Cause: The root cause was SM Process deadlock.</p> <p>Steps to Replicate: This is time sensitive issue. The steps cannot be re-produced or verified.</p> <p>Suggest to run full regression test with switchovers.</p>	<p>The code is modified to release smNtpLock_ after issuing a NTP restart command because we certainly do not need the lock while doing the restart.</p> <p>Workaround: No workaround.</p>

SBX-103254
| SBX-104839

PortFix SBX-103254: The call setup delay with the Option Ping 40+ endpoints SBC SWE.

Impact: The configuration support to disable Path MTU Discovery by setting kernel parameter net.ipv4.ip_no_pmtu_disc (ipNoPmtuDisc). When ipNoPmtuDisc is set to 2, the DF bit in IP packet header will be set 0.

Root Cause: The secure network does not support Path MTU Discovery per RFC-1191.

Steps to Replicate: Testing on a Standalone SBC SWE:

1. Start packet capture, and make SIP call (SIP-TLS). Verify Path MTU Discovery is enabled by checking outgoing IP packet's DF bit is 1.

- Test to show the issue

```
[root@SBCSWE03 ~]# date; tshark -t a -i pkt0 -f "ip host 10.x.xx.xx" -w t_tshark_pkt0_10.x.xx.1127_pmtu_1027_01.pcap
Tue Oct 27 16:21:11 EDT 2020
Capturing on 'pkt0'
35 ^C
[root@SBCSWE03 ~]# tshark -V -r t_tshark_pkt0_10.7.20.29_1127_pmtu_1027_01.pcap | more
```

```
Internet Protocol Version 4, Src: 10.x.xx.xx (10.x.xx.xx), Dst: 10.xxx.xx.xxx (10.xxx.xx.xxx)
Version: 4
Header Length: 20 bytes
Differentiated Services Field: 0x00 (DSCP 0x00: Default; ECN: 0x00: Not-ECT (Not ECN-Capable Transport))
0000 00.. = Differentiated Services Codepoint: Default (0x00)
.... 00 = Explicit Congestion Notification: Not-ECT (Not ECN-Capable Transport) (0x00)
Total Length: 60
Identification: 0x0000 (0)
Flags: 0x02 (Don't Fragment)
0... .. = Reserved bit: Not set
.1... .. = Don't fragment: Set
..0... .. = More fragments: Not set
Fragment offset: 0
```

```
[root@SBCSWE03 ~]#
--> The DF bit was set to 1.
```

2. Disable Path MTU Discovery.

```
admin@SBCSWE03% set system admin sbcswe03 kernelParams ipNoPmtuDisc ?
```

Description:

Configure /proc/sys/net/ipv4/ip_no_pmtu_disc (0 (default; disable); 1,2,3 (enabled modes))

1: when frag-required ICMP is received, PMTU to this destination is set to min_pmtu 552.

2: implicitly setting IP_PMTUDISC_DONT on every created socket - outgoing IP packet DF is set to 0.

3: hardened pmtu discover mode. Please refer to Linux kernel document.

Possible completions:

```
<int, 0 .. 3>[0]
```

```
admin@SBCSWE03% set system admin sbcswe03 kernelParams ipNoPmtuDisc 2
```

```
admin@SBCSWE03% commit
```

Commit complete.

```
[ok][2020-10-27 16:36:27]
```

3. Restart the SBC. (If using HA, restart both nodes)

```
[root@SBCSWE03 ~]# date; sbxrestart
```

```
Tue Oct 27 16:36:56 EDT 2020
```

4. Start packet capture, and make SIP call (SIP-TLS). Verify Path MTU Discovery is disabled by checking outgoing IP packet's DF bit is 0.

```
[root@SBCSWE03 ~]# date; tshark -t a -i pkt0 -f "ip host 10.x.xx.xx" -w t_tshark_pkt0_10.x.xx.1127_noPmtu_1027_02.pcap
Tue Oct 27 16:42:32 EDT 2020
Capturing on 'pkt0'
26 ^C
[root@SBCSWE03 ~]# tshark -V -r t_tshark_pkt0_10.x.xx.xx_1127_noPmtu_1027_02.pcap | more
```

```
...
Frame 2: 74 bytes on wire (592 bits), 74 bytes captured (592 bits) on interface 0
```

```
...
```

```
Internet Protocol Version 4, Src: 10.x.xx.xx (10.x.xx.xx), Dst: 10.xxx.xx.xxx (10.xxx.xx.xxx)
```

```
Version: 4
```

```
Header Length: 20 bytes
```

```
Differentiated Services Field: 0x00 (DSCP 0x00: Default; ECN: 0x00: Not-ECT (Not ECN-Capable Transport))
```

```
0000 00.. = Differentiated Services Codepoint: Default (0x00)
```

```
.... 00 = Explicit Congestion Notification: Not-ECT (Not ECN-Capable Transport) (0x00)
```

```
Total Length: 60
```

```
Identification: 0xe881 (59521)
```

```
Flags: 0x00
```

```
0... .. = Reserved bit: Not set
```

```
.0... .. = Don't fragment: Not set
```

```
..0... .. = More fragments: Not set
```

```
Fragment offset: 0
```

```
...
```

```
[root@SBCSWE03 ~]#
```

```
--> The DF bit was set to 0, and Path MTU Discovery is disabled.
```

Disable the Path MTU Discovery by setting kernel parameter net.ipv4.ip_no_pmtu_disc to 2. When set to 2, the DF bit will be set to 0 for the transmitted packets on every created socket. The configuration parameter is "system admin <name of system> kernelParams ipNoPmtuDisc".

Workaround: The workaround is a hack. It can be used one time test only. The hack will not survive Linux reboots.

1. Make changes on the current Standby (assuming SBC1 is currently active, and SBC2 is currently standby)

```
[root@SBC2 ~]# date; cat /proc/sys/net/ipv4/ip_no_pmtu_disc
[root@SBC2 ~]# date; echo "2" > /proc/sys/net/ipv4/ip_no_pmtu_disc
[root@SBC2 ~]# date; sbxrestart
```
2. Wait until SBC2 is up and synched. Ensure the p_no_pmtu_disc is set to 2.

```
[root@SBC2 ~]# date; sbxstatus | tail -4
[root@SBC2 ~]# cat /proc/sys/net/ipv4/ip_no_pmtu_disc 2
[root@SBC2 ~]#
```
3. Make changes on the current Active.

```
[root@SBC1 ~]# date; cat /proc/sys/net/ipv4/ip_no_pmtu_disc
[root@SBC1 ~]# date; echo "2" > /proc/sys/net/ipv4/ip_no_pmtu_disc
[root@SBC1 ~]# date; sbxrestart
```
4. Wait until the SBC1 is up and synched. Ensure the p_no_pmtu_disc is set to 2.

```
[root@SBC1 ~]# date; sbxstatus | tail -4
[root@SBC1 ~]# cat /proc/sys/net/ipv4/ip_no_pmtu_disc 2
[root@SBC1 ~]#
```
5. May perform a switchover to make the SBC1 to be active.
6. Run your tests.

SBX-95982 SBX-104938	<p>PortFix SBX-95982: The SBC was running version 6.2.2 and cannot capture media with MCT.</p> <p>Impact: The MCT reload of configuration (sbx restart or switchover) will not succeed.</p> <p>Root Cause: Missing small logic during the restart in order to properly reload the MCT configuration.</p> <p>Steps to Replicate: Retest MCT reload and switchover scenarios.</p>	<p>The code is modified to properly reload the MCT configuration.</p> <p>Workaround: User can delete the MCT config and reconfigure it after any restart/switchover.</p>
SBX-103948 SBX-103977	<p>PortFix SBX-103948: The SBC Application is crashing when CPaaS to PSTN call is made.</p> <p>Impact: The SBC cores after a reboot/upgrade.</p> <p>Root Cause: The issue was introduced when the SBC try to support 10000 SMM rules in 8.2.0.</p> <p>Steps to Replicate: Configure the customer rules (delete SDP line). Major logs trigger but the rule still success. After a reboot/upgrade, the core occurs.</p>	<p>The code is modified to address the issue.</p> <p>Workaround: Disable the delete SDP rule.</p>
SBX-105156 SBX-105239	<p>PortFix SBX-105156: The SBC was sending m=application 0 UDP/BFCP (null).</p> <p>Impact: The SBC generates the syntax error m line UDP/BFCP in 183.</p> <p>Root Cause: The SBC accesses the initialize data when format m line for UDP/BFCP</p> <p>Steps to Replicate: Incoming call with m-line UDP/BFCP. ERE/PSX using script NO_ROUTE and trigger announcement. The internal 183 has m-line UDP/BFCP with an invalid syntax.</p>	<p>Make the m-line initialize properly to address the issue.</p> <p>Workaround: Use the SMM to correct m-line syntax error.</p>
SBX-96239 SBX-105068	<p>PortFix SBX-96239: The display-name in the Diversion header or History-info header is deleted during the interworking.</p> <p>Impact: The SBC does not send display name present in ingress INVITE's diversion header in history info header in egress INVITE.</p> <p>The SBC does not send a display name present in ingress INVITE's history info header in egress INVITE's diversion header.</p> <p>Root Cause: The SBC does not consider display name when interworking between diversion header and history info header.</p> <p>Steps to Replicate: Ensure the display name is sent in history info header of egress INVITE when the SBC is converting diversion header to history info header.</p> <p>Also, the SBC should send the display name in diversion header of egress INVITE when the SBC is converting history info to diversion header.</p>	<p>The code is modified to ensure the display name is sent when the SBC is interworking history info and diversion header.</p> <p>Workaround: None.</p>
SBX-104225 SBX-105154	<p>PortFix SBX-104225: Both servers are registered but offline, and the SBCs cored after routing the changes from the customer.</p> <p>Impact: The S-SBC Core dump with PEM enabled call-flow.</p> <p>Root Cause: With PEM enabled, M-SBC performs NAPT learning for 1st RTP Packets. Upon receiving the Napt Indication for egress-leg from the M-SBC, the NRMA on the S-SBC performs flowChange/modification on both ingress and egress leg.</p> <p>As part of this flowChange/modification in ingress-leg, the ingress is NULL, resulting in coredump.</p> <p>The cktInfo on ingress should never be NULL, unless the call is being teared down.</p> <p>In this case, the ingress-leg has received a CANCEL, resulting in the call being taken down, as part of this call tear-down, ingress-leg cktInfo is reset to NULL, and then the de-alloc indication for ingress-leg is sent to the M-SBC, while the S-SBC is waiting for response of this tear-down, it processes the NAPT indication on egress-leg from the M-SBC. This results in the NULL cktInfo pointer getting accesses for ingress-leg.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Setup a D-SBC transcoded call with MRF. 2. Enable PEM on egress. 3. Ensure PRACK is enabled. 4. Ensure 180 without SDP, followed by 183 with SDP is received from UAS. 5. Ensure CANCEL is received from UAC, while 200 OK from MRF is received, such that the timing of RTP NAPT learning processing and CANCEL processing match. <p>NOTE: This is a very tricky issue to recreate, as the timing may vary every time we attempt, it is a matter is few milliseconds.</p>	<p>The code is modified to address the issue.</p> <p>Workaround: Disable NAPT or PEM. This workaround will ensure that we do not get to the race condition.</p>

<p>SBX-104561 SBX-104936</p>	<p>Portfix SBX-104561: The 8.1R5 customer SBC OAMs reaches 90% memory util and then restarts.</p> <p>Impact: The memory leak was observed on the OAM node.</p> <p>Root Cause:</p> <ul style="list-style-type: none"> The XRM and MRM are sending messages through the ICM to NRM. The NRM is not present/running on OAM. Since these are non-discardable messages, they will queue forever until we run out of memory. <p>Steps to Replicate:</p> <ol style="list-style-type: none"> Deploy the fix. Monitor memory usage of PRS process. Validate that PRS memory is not increasing. 	<p>The code is modified so that these messages are not sent on OAM node.</p> <p>Workaround: Before this fix, the workaround is to restart OAM app to recover the memory.</p>
<p>SBX-104484 SBX-104614</p>	<p>PortFix SBX-104484: The SBC is rejecting the second Update message from the Egress as 500 result into call failure.</p> <p>Impact: The SBC rejects 2nd Update from egress due to DLRB feature.</p> <p>Root Cause: The first Update SIPS response locally 200OK but not clear the server request message. As result the second Update, trigger the SIPS answer 500.</p> <p>Steps to Replicate: Configure the DLRB, Egress response 183 without the SDP, first Update, and the second Update.</p>	<p>The code is modified so after a 200OK response, the SIPS clears the server request so it can handle the subsequent one.</p> <p>Workaround: Disable the DLRB.</p>
<p>SBX-104146 SBX-104303</p>	<p>PortFix SBX-104146: The DNS Process observed a coredump when responding with a RCODE 1 error from the DNS server and deleting a DNS Group in the SBC.</p> <p>Impact: The DNS Process dumps core after getting RCODE errors 1,2,4 for a EDNS query and the DNS Group is deleted immediately on the SBC in the monitoring Interval specified in the "ednsRetryAfter" configuration has values ranging from 60 to 180 seconds with 180 being default.</p> <p>Root Cause: The EDNS failure timer callback is invoked, upon encountering an EDNS query failure based on the "ednsRetryAfter" configuration. The timer call back function is missing NULL check for the DNS Group</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> The SBC configured with EDNS support. The DNS server is configured to respond with the RCODE 1/2/4 error for EDNS query. Make a SIP Call that triggers the EDNS query. Disable the DNS Server and delete the DNS group within 180 seconds. 	<p>The code is modified to cancel the EDNS-failure timer, whenever the DNS Group is deleted and the NULL pointer check is added for the timer callback for robustness.</p> <p>Workaround: None.</p>
<p>SBX-103207 SBX-103580</p>	<p>Portfix SBX-103207: The SBC SWe duplicates the audio from call B to call A.</p> <p>Impact: If the CN is not negotiated for G711 codec and remote peer still sends CN packets that match default CN payload type (13) then user on other end may hear cross talk audio of completely unrelated channel or hear own reflect audio.</p> <p>Root Cause: When a g711 side does not negotiate CN in signaling, the DSP does not initialize Comfort Noise Generation object. However, if remote peer still sends a CN packet that matches the g711 SID payload type configured in PSP (default 13), DSP accepts the packet. It processes that CN packet incorrectly and as a result uses stale voice buffer that happens to be of another channel. This continues until next voice packet for that channel arrives. As a result, we see cross talk audio from the other call during silence period. In some cases, the user may hear own audio also and that is different manifestation of the same problem. This issue is specific to SWe.</p> <p>Steps to Replicate: Follow the step described in problem statement of JIRA.</p>	<p>The code is modified to initialize the comfort noise object even though the CN is not negotiated and process CN packets correctly.</p> <p>Workaround:</p> <p>Work around 1:</p> <p>Change the default payload type of comfort noise from 13 to something else (say 15) in PSP of peer that is sending CN packets. This will make DSP drop CN packets because payload type will not match.</p> <p>Work around 2:</p> <p>Enable silence suppression on PSP of peer that is sending CN packets and keep CN payload type that matches with CN payload type.</p>
<p>SBX-102704 SBX-103869</p>	<p>PortFix SBX-102704: Unable to Add or Edit route with # in Destination National, CLI or EMA.</p> <p>Impact: User is not able to store few special characters for the destination number field in the route entity.</p> <p>Root Cause: The pattern defined for the destination number field did not allow special characters.</p> <p>Steps to Replicate: Create a route with special character like #123 for the destination number field in the route entity.</p>	<p>The code is modified to allow only few specific special characters.</p> <p>So, reverted back the pattern to same as before 8.1 release, so that no customers face similar issue again who have configured special characters prior to 8.1.</p> <p>Workaround: None.</p>
<p>SBX-102290 SBX-104913</p>	<p>PortFix SBX-102290: The DBG file was filling up with messages "SIPCM: *ThreadPool: messageSequence".</p> <p>Impact: The DBG logs can be overrun with "SIPCM: *ThreadPool: messageSequence" messages.</p> <p>Root Cause: The DBG logs can be overrun with "SIPCM: *ThreadPool: messageSequence" messages, when Double CRLF "pings" are received by the SBC over UDP transport.</p> <p>Steps to Replicate: Send Double CRLF "pings" over UDP to the SBC.</p>	<p>The code is modified to properly dispose of Double CRLF "pings" received by the SBC over UDP.</p> <p>Workaround: Inhibit the transmission (or reception) of Double CRLF "pings" over UDP.</p>

<p>SBX-104934 SBX-105024</p>	<p>PortFix SBX-104934: The SBC generates the same ICID for more than 1 call, porting the fix for SBX-101887 into 8.2.3F1 did not help</p> <p>Impact: Duplicate the ICID generated for two different calls.</p> <p>Root Cause: The instance specific value is set to 0 every other call causing the ICID generation to always use 0 in all instances. This results in the same ICID generated in more than one SCM instance when the calls land within the same microsecond</p> <p>Steps to Replicate: Configure the SBC to generate ICID on the egress side. Run call load and check for duplicate ICIDs.</p>	<p>The code is modified to manipulate the last octet of the MAC address used in ICID generation.</p> <p>Workaround: None.</p>
<p>SBX-105417 SBX-105583</p>	<p>Portfix SBX-105417: During an S-SBC and M-SBC cyclic switchover, the S-SBC app did not come up after few switchovers.</p> <p>Impact: The service did not come up after a switchover due to inconsistent encryptedStore.</p> <p>Root Cause: The store became inconsistent after the oam_config got applied, the time taken during deletePeerEntries should also be reduced to allow a ChmClearEncryptedStoreOfPeers to finish faster during startup.</p> <p>Steps to Replicate: Perform multiple switchovers and ensure that:</p> <ol style="list-style-type: none"> 1. ChmClearEncryptedStoreOfPeers does not take too long during startup. 2. The deletePeerEntries runs in background and does not take too long either. 3. The service comes up, role is assigned successfully and store is consistent. 	<p>The code is modified so:</p> <ol style="list-style-type: none"> 1. deletePeerEntries run in background 2. decryptStore validates if decryption was successful to ensure the store is not inconsistent/corrupt. <p>Workaround: None.</p>
<p>SBX-103629 SBX-103654</p>	<p>PortFix SBX-103629: The SIP registrations are failing, and the SBC reporting REGISTER_PARSE_ERROR and replies with SIP 500 when it receives retransmitted initial REGISTER (CSEQ 1) followed by REGISTER with Authorization header (CSEQ 1).</p> <p>Impact: When the SBC recv rexmit of registration (cseq=1) after sending 401/407, the SBC relay to AS. At the same time, the IAD sends new registration(cseq=2) for authentication. It triggers race condition to the SBC and response 500.</p> <p>Root Cause: The SBC should relay rexmit (cseq=1) registration to application server.</p> <p>Steps to Replicate: The IAD send registration(cseq1) to AS, AS response 401 to IAD, IAD rexmit the same sseq1, and send a new sseq2 with auth header.</p>	<p>The code is modified to response rexmit request (cseq=1). So that a subsequent one cseq=2 can handle properly relay to AS.</p> <p>Workaround: None.</p>
<p>SBX-105486</p>	<p>The SBC 9.1 MD5 check sum.</p> <p>Impact: The SBC qcow2 image is shipped along with a sha256 checksum but Openstack glance supports only md5 checksum. Unable to validate the image in an automated way.</p> <p>Root Cause: Checksum was updated from md5 to sha256 as md5 is weaker compared to sha256 algorithm and there are reported attacks on md5.</p> <p>Steps to Replicate: With the fix build, ensure qcow2.md5 is one of the build artifacts.</p>	<p>The code is modified to facilitate openstack glance image verification.</p> <p>Workaround: User can manually create md5 checksum by running below command: md5sum sbc.qcow2 > sbc.qcow2.md5</p>
<p>SBX-104284</p>	<p>The SNMPv3 traps from the MRFP are not displayed in EMS Fault Management.</p> <p>Impact: The traps were not getting displayed in EMS.</p> <p>Root Cause: The SBC generates Engine ID at the time of ISO installation. Due to this, all cloud instances share same engine id from the qcow2 generation time.</p> <p>When same engine id is returned by two SBCs in different cluster, the EMS is unable to differentiate between them.</p> <p>Steps to Replicate: Launch instances from qcow2 in different clusters. Traps generated should be shown correctly now.</p>	<p>Engine ID is created at the time cloud instance launch now to address the issue.</p> <p>Workaround: None.</p>
<p>SBX-104688</p>	<p>Configure the import issue.</p> <p>Impact: The fields that supports \r like regex fields, changes it after a configuration Import(XML based).</p> <p>Root Cause: The \r is not xml encoded when parsing intermediate xml file.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Create SMM Profile with regex that has \r export configuration as xml-tar clear DB. 2. Import the configuration. 3. Verify SMM profile regex field. 	<p>The code is modified to handle \r as a special case.</p> <p>Workaround: None.</p>

<p>SBX-104463</p>	<p>Observed a SCM Process core dump in the goored while running the 30 SWOs on 7000 Platform.</p> <p>Impact: In case of switchover and switchback while running load, we expect some stale entries in SIPSG and CC control blocks. Such entries are cleared on the audit, but while clearing the SIPSG is sending accounting message to CC.</p> <p>This SIPSG was reconstructed, but in CC the index corresponding to this GCID was allocated to some other GCID.</p> <p>When the CC receives such message, the mismatch occurs and it logs a crash in sensitive mode.</p> <p>Root Cause: During load run with multiple switchovers, various call modules get reconstructed when standby transitions to active. Due to this some stale entries are expected. If one of these stale entries sends some message (e.g accounting message) and if index for that call was allocated to another call, this error can happen. In case of such mismatch the SBC logs a crash in sensitive mode.</p> <p>Steps to Replicate: Load run with multiple switchovers.</p>	<p>If the GCID sent from SIPSG module in accounting message is different from the GCID received in CC module, for switchover case do not do a sensitive crash. Instead add a major log to address the issue.</p> <p>Workaround: Run load scenarios in normal mode instead of sensitive mode.</p>
<p>SBX-104526</p>	<p>The emssftp fails in the 9.1R1.</p> <p>Impact: The emssftp login was failing.</p> <p>Root Cause: The sequence of asking for a password and keep credentials in sync caused the emssftp password failed.</p> <p>Steps to Replicate: Checking login, after multiple reboots, rebuild with 1:1 and N: 1, along with 1-2 EMS registered.</p>	<p>The code is modified on the mode(HA vs standalone), included a new script to handle password request and updated for some HA cases.</p> <p>Workaround: None.</p>

The following Severity 2 and 3 issues are resolved in this release:

Table 35: Severity 2-3 Resolved Issues

Issue	Sev	Problem Description	Resolution
<p>SBX-104962 SBX-104989</p>	<p>2</p>	<p>Portfix SBX-104962: The HFE setInterfaceMap() can be incorrect.</p> <p>Impact: The HFE_AZ.sh script can configure the HFE node incorrectly if the IP of an interface is a contained with in another.</p> <p>Root Cause: The logic used to determine which interface associates while the NIC can find the incorrect interface.</p> <p>Steps to Replicate: Create a HFE where the eth0 address contains either eth1 or eth2 IP address.</p>	<p>The code is modified to verify it the only the correct IP address is used.</p> <p>Workaround: Request an updated HFE_AZ.sh script from Ribbon. Replace the HFE_AZ.sh in the storage account and reboot the HFE node (s).</p>
<p>SBX-104156 SBX-104571</p>	<p>2</p>	<p>Portfix SBX-104156: The RTT-TTY support verifies the call flow in the SBC.</p> <p>Impact: Lower case characters from t140 packet were not getting converted correctly to Baudot (tty). Also, some other characters were not getting converted as per V.18 A2 table.</p> <p>Root Cause: The translation table from T140 to Baudot(TTY) only accounted for the subset of characters that were present in the Baudot (TTY) character set.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Setup a T140=>Baudot (AMRNB=>G711) call. 2. Create t140 pcap file with all UTF8 characters (128) and validate the baudot generation. 3. Create a t140 pcap file with multi-byte sequence and validate the baudot. 	<p>The code is modified to:</p> <ul style="list-style-type: none"> • The t140=>Baudot translation table as per V.18 Annx2. • Validate multi-byte (2,3 and 4 bytes). • Handle the Byte Order Mark (BOM). • Ignore an invalid byte sequence and substitute the valid but unknown multi-byte seq with an apostrophe. <p>Workaround: None. Instead of lower case letters use upper case letters.</p>
<p>SBX-96783 SBX-104926</p>	<p>2</p>	<p>PortFix SBX-96783: Some counts in the callCurrentStatistics keep incrementing.</p> <p>Impact: The "activeRegs" counter provided by the CLI zone callCurrentStatistics/ callIntervalStatistics command may continuously increase.</p> <p>Root Cause: Badly formed SIP REGISTER messages received by the SBC (i.e., a REGISTER message that the SIP parser determines is bad) will increment the "activeRegs" counter and never decrement it.</p> <p>Steps to Replicate: Send bad SIP REGISTER message(s) to the SBC, and see that the "activeRegs" counter provided by the CLI zone callCurrentStatistics / callIntervalStatistics command increases (and does not decrease).</p>	<p>Decrement the "activeRegs" counter when the SIP REGISTER message fails the SIP parser to address the issue.</p> <p>Workaround: None.</p>

<p>SBX-102364 SBX-104657</p>	<p>2</p>	<p>Portfix SBX-102364: The SBC behavior is inconsistent with the isfocus parameter handling.</p> <p>Impact:</p> <p>Issue 1: The SBC is not transparently passing the complete Contact header in 200 OK of SUBSCRIBE when there is a isfocus parameter.</p> <p>Issue 2: The SBC is not adding the Record-Route in any message when doing a full Contact header transparency</p> <p>Root Cause:</p> <p>Issue 1: The SBC passes contact header transparently only when the isfocus is present in request as well as response and would not send the contact header transparently only when the contact header in response has isfocus parameter.</p> <p>Issue 2: The code to add a record route for notify and its response is not present.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Configure the SBC for A to B call. 2. Enable the flag contactTransparencyForIsFocusMediaTag on both TGs. 3. Make a SUBSCRIBE-NOTIFY call flow. 	<p>Issue 1: The code is modified to check if the service bit is not set (for the response) and check if the contact header has the isfocus parameter, if the parameter is present, set:</p> <p>SIP_SERVICE_TYPE_CONTACT_TRANSPARENCY_FOR_ISFOCUS_PARAM</p> <p>for a single contact scenario.</p> <p>Issue 2: The code is modified to add the record route for notify and its response if isfocus parameter is present.</p> <p>Workaround: None</p>
<p>SBX-103988 SBX-104547</p>	<p>2</p>	<p>Portfix SBX-103988: The ICE not working after upgrade to V08.02.03A950.</p> <p>Impact: When a call with ICE changes from direct media to passthrough, ICE learning does not get enabled on the call.</p> <p>Root Cause: The code was not re-enabling the ICE FSM to start ICE learning when changing from direct media to passthrough.</p> <p>Steps to Replicate:</p> <p>Test 1 -----</p> <ol style="list-style-type: none"> 1. Establish a call with ICE on ingress and xdm on egress. 2. Send re-invite from egress without xdm and complete the re-invite related signaling. 3. Send stuns to ingress media to complete ICE learning and verify call state changes to established and media can be exchanged between ingress and egress. 	<p>The code is modified to re-enable the ICE FSM when changing from direct media to passthrough to allow the receipt and processing of stun messages to complete ICE learning.</p> <p>Workaround: None.</p>
<p>SBX-103496 SBX-103806</p>	<p>2</p>	<p>The MGT0 cannot reach the GW.</p> <p>Impact: The mgt0 interface appears to be inactive. The user cannot ping the gateway from the SBC 5400.</p> <p>The ethtool -S mgt0 shows the tx_pause count rising rapidly.</p> <p>Root Cause: This problem is specific to the SBX 5400 fix.</p> <p>The SBC 5400 has a backpressure mechanism for management ports so that when the receive buffer exceeds a certain level, Design sends pause packets out to the interface to ask the switch/router to pause transmission momentarily.</p> <p>The receive buffer count is incremented by a hardware module and decremented by software.</p> <p>The notification was sent from Network Processor to application about sRTP Rollover Counter (ROC) reset is specifying the wrong interface so software decrements the wrong receive buffer count.</p> <p>This causes the receive buffer count to wrap around, resulting in a high value. This triggers the back-pressure mechanism and Design continuously send out pause packets.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Set up a pass-thru SRTP call. 2. The call should be a long duration call so that the Rollover Counter (ROC) in the sRTP packet becomes non-zero. 3. Modify the call so that the sRTP packet's the SSRC changes. <p>Update:</p> <ol style="list-style-type: none"> 1. Do a long duration sRTP call (about 30 minutes.) 2. Hold the call. 3. Unhold the call. 4. As a result, this causes an SSRC change and triggers the problem. 	<p>Set the correct interface information in the notification to address the issue.</p> <p>Workaround: Please contact Ribbon support for a script that will restore contact to the management interface.</p>

SBX-103922 SBX-104649	2	<p>Portfix SBX-103922: There was a PRS Process core dump on an active/A node in the middle of an upgrade.</p> <p>Impact: The PRS Process cored during the LSWU while syncing ICE call data.</p> <p>Root Cause: The ICE redundancy code is taking a long time to sync the call data with the standby during LSWU, causing a Healthcheck timeout and a core dump.</p> <p>Steps to Replicate: The issue cannot be reproduced easily.</p>	<p>The code is modified so the ICE syncs the call data to standby and addresses the issue.</p> <p>Workaround: None.</p>
SBX-104704 SBX-104907	2	<p>PortFix SBX-104704: Duplicate the audio entries in the route PSP.</p> <p>Impact: The duplicate codec entry is being sent in the policy response.</p> <p>Root Cause: The function popPSPCodecEntryHlp that populates the codec entry was called twice, and as a result the codec entry was getting populated twice in the TLV.</p> <p>Steps to Replicate: Add more than 4 codec entries in the ingress PSP and egress PSP. This will cause the issue to be reproduced.</p>	<p>The code is modified so that the duplicate codec entry is not passed on the d+ response.</p> <p>Workaround: None.</p>
SBX-101575 SBX-104914	2	<p>PortFix SBX-101575: There was an offline upgrade failure from 8.2.1 to 8.2.2</p> <p>Impact: An offline upgrade on a HA setup failed for the upgrade from 8.2.1 to 8.2.2.</p> <p>Root Cause: An offline upgrade failed due to presence of a model update marker files that should not get created in case of offline upgrades.</p> <p>Steps to Replicate: Perform an offline upgrade on a HA setup to the fix build and ensure upgrade is successful.</p>	<p>The code is modified to ensure nodeB is powered-off/shutdown (instead of the sbxstop) while upgrading a nodeA. This ensures that the model update related markers files do not get created.</p> <p>Workaround: None.</p>
SBX-104668 SBX-105128	2	<p>PortFix SBX-104668: The SNMP PM collection is not working for newly added KPIs.</p> <p>Impact: SNMP stats collection at EMS for G711_WITHOUT_XCODE_RES_PEAK_COUNT and G711_WITHOUT_XCODE_RES_AVG_COUNT of DspResDspUsageIntervalStats failed.</p> <p>Root Cause: Some of the KPI stats fields are made obsolete as they are no longer supported. Due to this change EMS is not able to collect the stats through SNMP.</p> <p>Steps to Replicate: Verify EMS SNMP walk should not fail when queried to the SBC.</p> <ol style="list-style-type: none"> 1. Bring up SBC. 2. Register SBC to EMS. 3. Enable the sonusDspResDspUsageIntervalStatistics. <p>Expected Result: The EMS fetches stats from the SBC.</p>	<p>The code is modified so the EMS is able to collect stats using SNMP.</p> <p>Workaround: None.</p>
SBX-103175 SBX-103652	2	<p>PortFix SBX-103175: The large microflow profile was not getting enabled in the Custom Traffic Profile.</p> <p>Impact: There are multiple issues:</p> <ol style="list-style-type: none"> 1. For a default profile, max subs was always set to 256K. 2. For the custom and standard traffic profiles, the micro flow count in the NP resource was not proper. 3. Limited support of the 2M micro flows for standard profiles. <p>Root Cause:</p> <ol style="list-style-type: none"> 1. The large micro flow support was not there for the custom traffic profiles. 2. For standard and default traffic profiles, there was restricted /incomplete support of large micro flow that is made generic as part of this issue. <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Create an instance with 10GB memory and 16vpcu with the default traffic profile. Check microflow limits using "/opt/sonus/bin/cpsi -d summary" command. 2. Increase the memory to 50GB memory and check the microflow limits using "/opt/sonus/bin/cpsi -d summary" command. 3. Create an instance with 20GB memory and 16vpcu. Create and activate a custom traffic profile with access enabled in call mix. Check microflow limits using "/opt/sonus/bin/cpsi -d summary" command. 	<p>The code is modified to:</p> <ol style="list-style-type: none"> 1. Introduce a new estimation parameter maxSubs based on the micro flow count is decided and the NP hugepages are reserved. 2. Enable the large micro flow support for all the standard/default profiles where mem > 48GB and vcpu_count > 10. <p>Workaround: No workaround available. Need to use the build after the fix.</p>

SBX-104347 SBX-104735	2	<p>PortFix SBX-104347: The resource mem congestion level 3 is approaching threshold 90 sample 82.</p> <p>Impact: The SIPCM is leaking a structure that is associated with the SIPMM functionality.</p> <p>Root Cause: The code to free this memory under certain error/edge scenarios is missing.</p> <p>Steps to Replicate: Since the exact call flow is not know that caused the customer to hit the error condition, no test steps can be used.</p>	<p>The code is modified to free the SIPMM related structure in error scenarios.</p> <p>Workaround: Since the exact call flow is not know that caused the customer to hit the error condition, we cannot suggest a workaround.</p>
SBX-102707 SBX-104911	2	<p>PortFix SBX-102707: The network interfaces are not renamed as expected after an upgrade from 820R2 to 821R0.</p> <p>Impact: Network interfaces are not properly mapped when the system is abruptly brought down after the first boot and rebooted.</p> <p>Root Cause: Cloud-init service runs in parallel to sonusudev service and can cause issues with the network interface mapping as it tries to rename the interfaces.</p> <p>Steps to Replicate: Reboot the system after sonusudev is run on first bootup.</p>	<p>The code is modified to start only after the sonusudev has run and mapped the interfaces properly.</p> <p>Workaround: Reboot the instance again to recover from the issue.</p>
SBX-99208 SBX-102666	2	<p>PortFix SBX-99208: The SBC has the same issue as SBX-86420, this time for an outgoing REGISTER messages.</p> <p>Impact: When the TLS port is not sent from the PSX, the SBC is not sending the DNS SRV query to the DNS server.</p> <p>Root Cause: The root cause was the SBC just increments the port number by 1 sent by the PSX even if peer port was Zero.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Configure the SBC and PSX to send a register and subscribe SIP request to relay. 2. Configure the serve FQDN and zero for the FQDN port. 3. Send a register/subscribe request. 4. The SBC should send SRV query to the DNS server. 	<p>If port number sent by PSX is zero then do not increment it by +1 by default for selecting TLS port to address the issue.</p> <p>Workaround: None.</p>
SBX-103561 SBX-103793	2	<p>PortFix SBX-103561: The "tgrp" parameter are not passing transparently when the SIP in the core is enabled.</p> <p>Impact: The SipCore calls is passing the wrong TGRP parameter value.</p> <p>Root Cause: There was missing logic to support the SipCore.</p> <p>Steps to Replicate: Configure the SipCore feature. Both the IPSP core and egress leg have "originating Trunk Group Options" set to "Include Tgrp with IP address" Make a SipCore call.</p>	<p>The code is modified to support the SipCore feature.</p> <p>Workaround: None.</p>
SBX-103852 SBX-103978	2	<p>PortFix SBX-103852: The SBC sent an unnecessary 481 for PRACK when the CANCEL is received.</p> <p>Impact: The SBC sends unnecessary 481 for PRACK when received CANCEL.</p> <p>Root Cause: The issue was introduced by intercept feature (SIP_EVENT_PRACK_INTERCEPT_RCVD).</p> <p>When the SBC received PRACK, it save the message into the server list. Later the CANCEL arrives, UasCancelRequestCmd, the UAS found the message in server list, and thought PRACK is not completed therefore it try to send 481.</p> <p>Steps to Replicate: A calls B, B response 180 and send 180 to A with PRACK required. After 32 seconds A sends a CANCEL, the SBC sends an unnecessary 481 PRACK.</p>	<p>Once the 200OK for PRACK has send out, the SIPS deletes the PRACK message in the server list to address the issue.</p> <p>Workaround: Enable the e2e PRACK if the egress also supports PRACK.</p>
SBX-96788 SBX-103979	2	<p>PortFix SBX-96788: The DTMF 2833 PT change (in offer) has incorrect OA SBX handling (wrong answer PT).</p> <p>Impact: When the peer offer with new PT, the SBC is unable to process and respond to the previous one.</p> <p>Root Cause: There was a logical error that checking the wrong field to detect PT change.</p> <p>Steps to Replicate: A call B and connect with "0 100". An INVITE with an offer new PT "0 101". The SBC responds with "0 100".</p>	<p>The code is modified to check the correct field for PT change.</p> <p>Workaround: None.</p>

SBX-105428 SBX-105440	2	<p>PortFix SBX-105428: Error when attempting to delete negative cache record upon a probe response.</p> <p>Impact: The DNS query was not out sent from the SBC if a call is made immediately after the DNS server is recovered from a blacklist.</p> <p>Root Cause: Due to record is in the negative cache entry from earlier DNS dummy query, the SBC was not sending a query.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Configure the SBC with a basic call and also configure the DNS group. 2. Ensure that DNS server does not respond to the DNS query or stop DNS server. 3. Now make call and SBC query for DNS server for resolving FQDN, but DNS server is down and there no response from DNS server. 4. After retransmission's timeout value. DNS server will be blacklisted. 5. Now the SBC keep probing DNS server. 6. Bring the DNS server up. 7. The SBC receives probe response, DNS server is removed from Blacklist. 8. Immediately make SIP call. <p>Expected Result:</p> <ul style="list-style-type: none"> • The SBC selects a DNS query. 	<p>As soon as the DNS server is removed from blacklist entry, those records involved with the DNS Dummy query is removed from the negative cache and resolves the issue.</p> <p>Workaround: None.</p>
SBX-105010 SBX-105155	2	<p>PortFix SBX-105010: Observed a congestion for the G711 to G711 transcode calls even before the CPU resource limit is exhausted.</p> <p>Impact: The G711-G711 transcoded capacity test shows a lower capacity. (Test is Pump make and break g711 to g711 transcode calls on one of the active instances with 50cps, 90 CHT and 4500 calls).</p> <p>Root Cause: The root cause is an incorrectly built libfmd.a that was built with a debug option and makes the FMTD module take more cycles. This bug was addressed originally in SBX-104122.</p> <p>Steps to Replicate: Perform a test indicated in SBX-105010.</p>	<p>The code is modified to remove the -g flag and using -O3 flag for compilation (Port of SBX-104122).</p> <p>Workaround: None.</p>
SBX-105072 SBX-105150	2	<p>PortFix SBX-105072: Password padding with random characters by the SBC causes the RADIUS server to reject the password.</p> <p>Impact: The radius password sent to the server has no zero characters at the end following the password and a NULL.</p> <p>Root Cause: The radius passwords are padded to 16 characters. The existing implementation did not set those padded characters to 0.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Configure a radius authentication server. 2. Use a password that is less than 15 characters long for the external radius user. 3. Set the externalAuthentication to true. 4. Run a tshark session. 5. Login to the CLI. 6. Stop the tshark. 7. View the radius password element in wireshark after configuring the shared secret in wireshark under protocol preferences. 	<p>The padded characters are now set to 0 to address the issue.</p> <p>Workaround: Use 15 or 16 character passwords.</p>

SBX-66831 SBX-103723	2	<p>PortFix SBX-66831: Need to update/add proper warning message when deleting the ipInterface.</p> <p>Impact: The warning message does not display in the cloud SBC when deleting the IP interface, but there is no problem with functionality.</p> <p>Root Cause: This is working fine in the hardware SBC. We have the problem in the cloud SBC because when displaying the warning message, we checked few values if that values are matching. Then, show the warning popup but in case the cloud SBC values index are different and check is failing because of that, we are unable to show the warning popup when we delete IP Interface.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Login to EMA. 2. Go to All -> Address Context -> IP Interface Group -> IP Interface. 3. Click on New IP Interface button. 4. Fill the correct info in text field, select state as enable, Mode as In Service and click on save button. 5. The creation should be successful. 6. Select the same entry from IP interface list section and click on delete button. 7. We should be able to see warning popup like "addressContext PPP_AC ipInterfaceGroup PPP_IG ipInterface': Deleting ipInterface while in dryUp action can cause problems with calls on the ipInterface. Use force action to clear existing calls on the ipInterface before deleting. Do you wish to continue?" 	<p>The code is modified to handle the Cloud SBC scenario.</p> <p>Workaround: None.</p>
SBX-100286 SBX-103102	2	<p>PortFix SBX-100286 to 9.2: The trace file containing SIP Rec PDU was not imported in Ribbon Protect properly.</p> <p>Impact: When a level 4 call trace is active, if a SIP PDU is too large to fit on a single trace line, it is split over multiple lines. However, the Ribbon Protect only reads the first line - there is no way for it to know subsequent lines are continuation lines.</p> <p>Root Cause: Design issue.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Enable level 4 call trace. 2. Send a SIP PDU to be traced towards the SBC of total size 2000 bytes. 	<p>The code is modified to put the same header onto each level 4 trace line. Include in the header two new fields, the MSG ID and PART to allow Ribbon Protect to recombine multiple trace lines to recover the original message. The equivalent Jira VIGIL-17137 is required for Ribbon Protect compatibility.</p> <p>Workaround: None.</p>
SBX-100590 SBX-104867	2	<p>PortFix SBX-100590: The wrong "Egress Zone Name" in ATTEMPT CDR for a GW-GW call.</p> <p>Impact: For a call with multiple routes that include one or more local routes followed by routes on another gateway (for GW-GW) scenario, if the local routes fail and then other gateway route is selected, the resulting ingress gateway CDR for the call has incorrect egress zone name and ID.</p> <p>Root Cause: When there are multiple routes, when a route fails the call cranks back to use the next route. When selecting the next route the code is not clearing the internal CDR information for the egress zone. So if the next route selected is on another gateway, the CDR information retains the egress zone information from the previous attempt rather than being blank.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. With the SBC GW-GW setup create two routes for a call such that: <ul style="list-style-type: none"> - route 1 is routed through local egress trunk group on SBC1. - route 2 is routed through egress trunk group on SBC2. 2. Set crankback profile to enable attemptRecordGeneration and add reason 41 (503 is mapped to 41 in default SIP to CPC mapping profile). 3. Send a valid INVITE to the SBC1 ingress that will first select route1 and route out through the egress trunk group on the SBC1. From the egress endpoint reject that call attempt with a 503 Service Unavailable. 4. The call will crank back and re-route using route2 through the SBC2 egress trunk group. From the egress endpoint reject that call attempt with a 503 Service Unavailable. 5. Check there are two Attempt CDR records for the call on SBC1. <ol style="list-style-type: none"> a. ATTEMPT1 - should have the "Egress Zone Name" and "Egress Zone ID" fields populated correctly with zone information for zone of SBC1 egress. b. ATTEMPT2 - should have the "Egress Zone Name" empty and "Egress Zone ID" as 0 because this is an attempt for GW-GW route. 	<p>The code is modified to clear the internal CDR information for the egress zone when one route fails and the next route is selected.</p> <p>Workaround: None.</p>

SBX-74991 SBX-104980	2	<p>PortFix SBX-74991: Saving changes to a SMM profile after doing an edit /update operation is taking long time[8 to 16 minutes].</p> <p>Impact: The SIP Adaptor Profile with 250+ rules takes lot of time to create a profile including the update as well.</p> <p>Root Cause: When the SIP Adapter Profile with more than one rule is created, each rule is committed individually. However, after each commit EMA internally retrieves all the SIP Adapter Profiles along with their rules, though this is not required it is causing slowness of both create and update operation.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Login to EMA. 2. Navigate to Sip Adapter Profile Screen. 3. Create a profile with 200+ Rules. 4. Click on Save button. 	<p>The code is modified to prevent retrieval of sip adapter profiles after each commit</p> <p>Workaround: The only workaround is to create the profile from the CLI.</p>
SBX-101769 SBX-104912	2	<p>PortFix SBX-101769: The SBC sends the wrong SDP in 200 OK for UPDATE received from the Ingress during tone play(LRBT).</p> <p>Impact: The SBC responds with wrong codec in the answer to UPDATE received on ingress/calling side while local ring back tone is playing. As a result, the UPDATE is being used to change the codec on ingress. In this case, the SBC continues to play the RBT using the previously agreed codec.</p> <p>Root Cause: There was a logical error in the code that was picking the currently active packet service profile to answer instead of honoring the modify.</p> <p>Steps to Replicate: Configuration:</p> <ol style="list-style-type: none"> 1. Configure SBC to make a Basic A-B call. 2. Configure PSX with G711 and G729 in codecEntry and this Leg /other Leg configuration. 3. Configure LRBT, and attach to Ingress TG. <p>Procedure:</p> <ol style="list-style-type: none"> 1. Send an INVITE from UAC with G711. 2. Send an 180 from UAS, with no SDP. 3. Send an UPDATE from UAC, with G729. 4. Receive the 200 OK from INVITE from UAS. 	<p>The code is modified to address this issue.</p> <p>Workaround: None.</p>
SBX-103771 SBX-104825	2	<p>PortFix SBX-103771: There are special characters in DN hinders EMA display.</p> <p>Impact: The EMA does not display route with destination national containing special characters.</p> <p>Root Cause: The support for special characters was not available in EMA.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. In the SBC 7.2.4 R0, create a Route with Destination National containing special characters. 2. Upgrade to 8.2.4 R0. 3. After the upgrade, the EMA should display the route. 	<p>The code is modified to support special characters in the EMA.</p> <p>Workaround: The CLI can be used to view route details.</p>
SBX-104247 SBX-104644	2	<p>PortFix SBX-104247: The resource memory congestion level 3 is approaching the threshold 90 sample 81.</p> <p>Impact: There was a SAM Process is leaking memory when AKA is being used.</p> <p>Root Cause: There was a SAM Process is leaking an AKA related structure when the code is handling an error case scenario.</p> <p>Steps to Replicate: The issue cannot be reproduced.</p>	<p>The code is modified to correctly free the AKA structure in all scenario.</p> <p>Workaround: There is no known workaround.</p>
SBX-93898 SBX-104516	2	<p>PortFix SBX-93898: The "request sbx xrm debug command sec -stat gcid <gcid>" was not showing ENC and DEC details on the SBC SWe.</p> <p>Impact: This debug command in unhide section is not showing all the required fields populated in the SWe SBC.</p> <p>Root Cause: NP response is not framed in expected order to application layer.</p> <p>Steps to Replicate: Run the SRTP call and issue this debug command for internal use.</p>	<p>The code is modified so the NP Response is correctly framed in expected order to application layer.</p> <p>Workaround: This debug command is to see SSN field value with the RoC, all other details can be seen from show call mediastatus.</p>

SBX-103599 SBX-103893	2	<p>PortFix SBX-103599; The SBC is sending multiple re-INVITEs to the ingress and egress leg during a 200OK of INVITE answered with Dialog-1.</p> <p>Impact: The SBC is sending multiple re-INVITEs to the ingress and egress side if the first forked leg is answering with a 200 OK.</p> <p>Root Cause: When the 200 OK is coming without SDP, we were taking last received SDP and that precondition attributes to further processing.</p> <p>Steps to Replicate: Precondition Transparency is set on both the Ingress and Egress Steps:</p> <p>UAC sends INVITE with Supported: precondition and SDP with precondition attributes.</p> <p>First forked UAS sends 183 with SDP with preconditions.</p> <p>Second forked UAS sends 183 with SDP with preconditions.</p> <p>First forked leg send 200 OK without SDP.</p> <p>Expected Behaviour:</p> <ol style="list-style-type: none"> 1. The SBC transparently passes precondition attributes in SDP. 2. The SBC gets first forked 183 with SDP with preconditions. 3. The SBC forwards the 183 to ingress with preconditions. 4. The SBC sends an UPDATE to ingress and egress to complete precondition state. 5. The SBC forward second the 183 to ingress with preconditions. 6. The SBC sends an UPDATE to ingress and egress to complete precondition state. 7. The SBC forwards 200 OK and get ACK. 8. The SBC completes reINVITE-200-ACK towards egress. 	<p>The code is modified to address the issue.</p> <p>Workaround: None.</p>
CHOR-6320 SBX-105286	2	<p>PortFix CHOR-6320: The UxPAD core dump observed on the media container while pumping 2X overload of G711U to G711U DSP (media transcoding disabled) call load.</p> <p>Impact: The SWe_UXPAD crash was seen during an overload transcode test in a 2 vcpu SWe deployment in the public cloud.</p> <p>Root Cause: Potential corruption in the packet buffers due to issues related to concurrency specifically in the 2 vcpu transcode overload scenarios.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Configure a 2 vcpu instance. 2. Run 700 G711-G711 transcode sessions for overnight/long duration. 3. Check if there is any UXPAD/NP crash seen. 	<p>The code is modified to fix concurrency issues in transcode call scenarios in the 2 vcpu SWe deployments.</p> <p>Workaround: No workaround apart from avoiding overload scenarios for transcoded calls.</p>
SBX-105182 SBX-105414	2	<p>PortFix SBX-105182: After the SBC receives RCODE = 2 from the DNS server, the SBC makes a strange query request to the DNS and subsequent new calls are rejected.</p> <p>Impact: When RCODE error received for the DNS query and no other servers were available, the record was never being deleted from cache and this remained as a RECORD in "RESOLVING STATE" that impacted further queries with this domain and resulted in call failures.</p> <p>Root Cause: This particular test scenario's was not covered during the testing feature.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Setup with EDNS supported and DNS fall back disabled. 2. The single DNS server configured for the DNS Group. 3. The DNS Server configured to return RCODE Error 2 for SRV Query. 4. Setup a call that trigger NAPTR, SRV and A Query and the DNS Server responds with RCODE 2 for SVR. 5. Modify the DNS Server to return successful RCODE for SRV. 6. Make another call for the same domain. The call fails because the SRV record is in "RESOLVING STATE" in the cache. 	<p>The code is modified to delete cache entry (which is in "RESOLVING STATE") when the RCODE error is received.</p> <p>Workaround: NA</p>

SBX-103255	2	<p>Enhance the sbxPerf on the SBC for lesser resource consumption.</p> <p>Impact: The SBC performance monitoring tools like top and top2 at times take 20% of a CPU core there by reducing total available CPU resources for management activities on the SBC.</p> <p>Root Cause: The sbxPerf that contains list of tools such as top, mpstat, and iostat currently runs periodically without any linux value configured. This can cause these processes to get prioritized and scheduled ahead of other management processes such as ConfD and SSH.</p> <p>Steps to Replicate: Install fix build and ensure processes such as top, top2, mpstat sbxPerf are running with corrects value of 15 using the top command.</p>	<p>The code is modified so that priority of these processes are less compared to other management processes on the SBC.</p> <p>Workaround: None.</p>
SBX-74155	2	<p>The ACL rule is missing on the SBC in LD triggered switch-overs and link recovered after.</p> <p>Impact: When an IPACL rule is created that references an IP interface, but that IP interface has been down since the system started up, the rule is not successfully created.</p> <p>Root Cause: When an IP interface is failed on system startup, the system does not add that interface to the kernel. When an IPACL rule is added that references that IP interface, it fails to be added to the NP.</p> <p>Steps to Replicate: Start an SBC with an IP interface failed, that also has one or more IPACL rules that have been configured that reference that IP interface. Bring up the IP interface and logs will show the IPACL rules have been successfully added.</p>	<p>The code is modified to detect this situation and maintain a retry list. Once the IP interface comes up, it will retry the associated IPACL rules that were previously failed.</p> <p>Workaround: This issue can be worked around by having IPACL rules that do not reference the IP interface.</p>
SBX-90854	2	<p>There was a customer call forwarding MRF interaction failure on the receipt of an UPDATE.</p> <p>Impact: The SBC sends INVITE towards MRF upon receiving UPDATE from Egress End Point</p> <p>Root Cause: In some situations when there is delay in incoming PRACK request (In this situation PRACK from Ingress Endpoint), some internal logic of the SBC queues the SDP on ingress leg and let's SBC answer to egress, thereby unblocking the egress peer which can send a new offer using SIP UPDATE.</p> <p>This causes the SBC to send INVITE towards MRF, even before sending an UPDATE out towards the Ingress.</p> <p>Steps to Replicate: Configure the PSPs accordingly:</p> <ol style="list-style-type: none"> 1. UAC sends Invite with AMR-WB, AMR, telephone-event. 2. UAS sends 183 with SDP with AMR-WB with 100rel. 3. UAC sends PRACK/200 OK. 4. UAS sends 180 without SDP without 100 rel. 5. PRACK for 180 is pending on Ingress. 6. UAS sends UPDATE with PCMU, telephone-event. <p>Expected Results: =====</p> <ol style="list-style-type: none"> 1. Upon receiving an UPDATE with PCMU, the SBC will not auto answer towards UAS. 2. The SBC will wait for PRACK to be received from UAC, then will relay UPDATE towards UAC. 3. Based on answer in 200 OK (Update) from UAC, the SBC will send INVITE towards MRF for reserving transcoding resources. 	<p>The code is modified to ensure that SBC does not auto answer to egress.</p> <p>Workaround: None.</p>
SBX-91111	2	<p>The No Way Video after A/V calls is resumed after hold.</p> <p>Impact: No way video after audio video call is resumed using late media re-INVITE.</p> <p>Root Cause: The SBC does not send "sendrecv" in 200 OK to late media re-INVITE, instead sends the same datapath mode that was negotiated last - that is - inactive.</p> <p>Steps to Replicate: Setup an audio and video call. Put the call on hold by sending a=inactive in audio and video SDP. Then, send a late media re-INVITE to take the call off-hold. The idea is that the SBC will send SendRecv for audio and video in 200 OK and let the remote end choose if it wants to remain in hold or come out.</p>	<p>The code is modified to send sendrecv in other streams too when call is being "may be" resumed after hold.</p> <p>Workaround: Do not use late media re-INVITE to come out of hold.</p>

SBX-92066	2	<p>Observed a PRS process core dump "systemerror healthcheck timeout"</p> <p>Impact: A PRS core dump was found due to healthcheck timeout upon executing a CLI debug command</p> <p>Root Cause: The CLI debug command was taking a long time due to too many ACL entries resulting in a health check timeout.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Add more than 1000 ACL entries 2. Execute below debug command: Run CLI debug command "request sbx xrm debug command acl -stat" <p>NOTE: The debug commands are not available on customer sites.</p>	<p>The code is modified to process only 200 ACL entries to allow CLI to recover sooner.</p> <p>Workaround: As we do not run debug commands on customer sites, this error should not occur.</p>
SBX-98843	2	<p>There was a problem with management of the maxptime and ptime in the Direct Media mode.</p> <p>Impact: There was a problem with management of maxptime and ptime in the Direct Media mode.</p> <p>Root Cause: Generally, the SBC ignores the ptime value if only maxptime attribute is present in incoming SDP. If "sendPtimeInSdp" IPSP flag is enabled and peer is sending both ptime and maxptime, SBC is preferring maxptime value while sending ptime, The existing IPSP flag "sendPtimeInSdp" makes the SBC to send out a=ptime irrespective of what we receive from the peer.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Do a basic call from A to B. 2. A sends ptime=20 and maxptime=40 towards the SBC. 3. An INVITE going out from the SBC should contain ptime=20 as we enabled both "sendPtimeInSdp" IPSP flag and "preferPtimeInSdp" TG flag. 	<p>A new flag "preferPtimeInSdp" is introduced under SBC CLI Trunk Group signaling. If this new flag "preferPtimeInSdp" is enabled, the SBC prefers ptime value received in a=ptime over a=maxptime in incoming SDP.</p> <p>This new flag "preferPtimeInSdp" needs be enabled in conjunction with existing flag "sendPtimeInSdp". If the SBC is configured to send a=ptime, then preferPtimeInSdp should be enabled on the incoming trunk and vice versa.</p> <p>Workaround: None.</p>
SBX-103704	2	<p>The RTP sourceAddressFiltering is not working (call leg dependency).</p> <p>Impact: The source address validation is not done if call flow does not involve NAPT learning. For calls that want source address validation but does not have NAPT enabled, the SBC would not validate the source address and end up forwarding the RTP packet to the other endpoint.</p> <p>Root Cause: There is no way for application to inform Network Processor to validate source if NAPT learning is not enabled.</p> <p>Steps to Replicate: Involves a CISCO MOH server but probably something else can be used.</p>	<p>The code is modified to validate source even when NAPT learning is not enabled to address the issue.</p> <p>Workaround: Enable NAPT learning for the call.</p>
SBX-102827	2	<p>The SBC5210 upgrade failure during Starting DB_RESTORE stage.</p> <p>Impact: An upgrade failure during Starting DB_RESTORE stage.</p> <p>Root Cause: Creation of foreign key fails on table packet_service_profile column CODEC_ENTRY_FK9, as it had "0" in one of its rows and dbimpl.codec_entry table did not have this value.</p> <p>Steps to Replicate: Upgrade a dump that has "0" in packet_service_profile.CODEC_ENTRY_FK9, it should be successful.</p>	<p>Before add the referential keys, the data is checked and ant value that is not there in parent table is nullified to address the issue.</p> <p>Workaround: None.</p>
SBX-103669	2	<p>Empty "Egress Local Signaling IP Addr" field in the CDR</p> <p>Impact: When an incoming call is routed out of the SBC but is then rejected by the SBC with cause 132 Module failure before a backwards response message has been received, the resulting attempt CDR has an empty "Egress Local Signaling IP Addr" field.</p> <p>Root Cause: The SBC currently populates the "Egress Local Signaling IP Addr" CDR field when it processes a backwards response message for the call, and as a result the field remains empty until a backwards response message is received and processed.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Initiate a call by sending a valid INVITE to the SBC ingress TG to route towards the egress TG. 2. Once the INVITE is sent to egress by SBC, put the signaling port to egress out of service. 3. The SBC will clear the call by sending a 503 message towards egress. The resulting attemptCDR should have valid address in "Egress Local Signaling IP Addr" field. 	<p>The code is modified to populate the "Egress Local Signaling IP Addr" CDR field on sending out the egress INVITE message.</p> <p>Workaround: None.</p>

SBX-104705	2	<p>There was a memory leak on the glusterfsd.</p> <p>Impact: The OAM application cores.</p> <p>Root Cause: Memory leak is observed for the glusterfsd process, that is the brick process. The use of the "gluster volume heal info" function causes the process memory usage to increase and not go down afterward.</p> <p>Steps to Replicate: Ensure both Active and Standby OAM are running. Monitor glusterfsd memory usage on both nodes.</p>	<p>Use the "gluster volume heal statistics heal-count" command to determine the gluster bricks heal state to address the issue.</p> <p>Workaround: Reboot the OAM node.</p>
SBX-103593	2	<p>The SBC is ignoring Use Max Bitrate Only flag.</p> <p>Impact: The SBC is ignoring Use Max Bitrate Only flag.</p> <p>Root Cause: During Re-Invite answer processing, NrmaAdjustPeerT38MaxBitRate() routine is doing T38MaxBitRate adjustment based on packet size that is changing T38MaxBitRate value from 14400 to 4800.</p> <p>Steps to Replicate: Call Flow:</p> <ol style="list-style-type: none"> 1. Do a basic call between UAC and UAS. 2. The UAS sends fax re-invite to UAC with T38MaxBitrate set to 14400. 3. The SBC sends Invite out with T38MaxBitrate set to 14400. 4. The UAC answers with T38MaxBitrate set to 14400 in 200OK. 5. The SBC sends out 200ok with T38MaxBitrate set to 4800. <p>Actual Result:</p> <pre> -----UAS(Re-Invite) -----> SBC -----> UAC T38MaxBitRate:14400 T38MaxBitRate:14400 -----UAC(200ok for Re-Invite) -----> SBC -----> UAS T38MaxBitRate:14400 T38MaxBitRate:4800 </pre> <p>Expected Result:</p> <pre> -----UAS(Re-Invite) -----> SBC -----> UAC T38MaxBitRate:14400 T38MaxBitRate:14400 -----UAC(200ok for Re-Invite) -----> SBC -----> UAS T38MaxBitRate:14400 T38MaxBitRate:14400 </pre>	<p>The code is modified to relay the T38MaxBitRate in SDP, which is received from peer for Direct Media calls.</p> <p>Workaround: None.</p>

SBX-104560	2	<p>The redirection NOA set wrong in the CDR.</p> <p>Impact: The "Redirecting Orig Cd Num - NOA" subfield of "Redirection Feature Spec Data" is not set correctly in CDR record if the the Redirecting Original Called Number - NOA value has been modified by PSX.</p> <p>Root Cause: The existing code was not updating the value for the CDR record based on route specific information returned from the PSX.</p> <p>Steps to Replicate: Test Setup =====</p> <ol style="list-style-type: none"> 1. Test requires call routing through the PSX. Routing on PSX, setup to route incoming call through routing label with two routes: - route 1, through local egress trunk group 1 on SBC to UE2. - route 2, through local egress trunk group 2 on SBC to UE3. 2. On egress trunk group 1 only, associate a DM/PM rule to change the "Number Type" to International for "redirecting Original Called Number". 3. Set the crankback profile to enable attemptRecordGeneration and add reason 41 (503 is mapped to 41 in default SIP to CPC mapping profile) set profiles callRouting crankbackProfile default attemptRecordGeneration enabled reason 41 <p>Test Procedure =====</p> <ol style="list-style-type: none"> 1. Send INVITE to the SBC ingress, INVITE should include a diversion header e.g. Diversion: < sip:+4969667740185@xxx.xxx.xxx.xxx>;privacy=off;screen=no;reason=unknown;counter=1 2. Call is routed using route 1, respond from first egress with 503 Service Unavailable 3. Call is re-routed using route 2, respond from second egress with 503 Service Unavailable 4. Check the CDR records for call on the SBC that should have two attempt records - <ol style="list-style-type: none"> a. ATTEMPT1 - should have "Redirection Feature Spec Data" that has subfield 22 "Redirecting Orig Cd Num - NOA" as 3 (Unique International Number), because the number was updated by PSX to international. b. ATTEMPT2 - should have "Redirection Feature Spec Data" that has subfield 22 "Redirecting Orig Cd Num - NOA" as 2 (Unique National Number), because the number was not updated by PSX 	<p>The code is modified to populate the CDR value from the route specific data returned from PSX.</p> <p>Workaround: None.</p>
SBX-103986	2	<p>There was an incorrect RTP time stamp in the SBC packet capture.</p> <p>Impact: After an SBC upgrade to V09.01.00R001, RTP/RTCP time stamp in the SBC packet capture shows the year 2036.</p> <p>Root Cause: Present logic was not considering the edge cases while checking the last modification time of the NTP log.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Configure a remote syslog server: set oam eventLog platformRsyslog linuxLogs ntpLog enabled set oam eventLog platformRsyslog servers server1 port 10514 protocolType udp remoteHost <remote host ip> com set oam eventLog platformRsyslog syslogState enabled com 2. Create a second dummy NTP server (just to trigger some logs to be written to ntp.log): set system ntp serverAdmin 1.2.3.4 com 3. Delete the dummy NTP server: del system ntp serverAdmin 1.2.3.4 com 4. Wait at least several seconds. 5. Repeat steps 2 and 3. 	<p>The code is modified for the edge case of NTP log last modification time.</p> <p>Workaround: None.</p>

SBX-94798	3	<p>The MS Teams with FLRBT enabled hangs the call.</p> <p>Impact: When the Force Local Ringback Tone is applied on a call that has been through SIP replace followed by SIP refer procedure, and the final trunk group has downstream forking enabled. The final 200 OK response may be queued indefinitely, meaning the call does not get completed.</p> <p>Root Cause: Interactions between Force Local Ringback Tone, multi-party and downstream forking.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Make a SIP to SIP call. 2. A calls B. 3. B' replaces B. 4. B' refers to C. 5. The trunk for A has Force Local Ringback Tone applied. 6. The trunk for C has downstream forking enabled. 7. C responds with 180 with no SDP, 183 with SDP, then 200 OK. 	<p>The code is modified to prevent the queuing.</p> <p>Workaround: None.</p>
SBX-99306	3	<p>The SIPCM (SAM) deadlock reported on SNMP request.</p> <p>Impact: The SAM Process cores when getTlsStatus command is executed.</p> <p>Root Cause: The issue is because when fetching the TLS Status, we were trying to access sessData that was no longer valid.</p> <p>Steps to Replicate: Execute getTlsStatus command when running TLS load. SAM process cores randomly.</p>	<p>The code is modified to fetch the status from socketPtr instead of sessData.</p> <p>Workaround: None.</p>
SBX-100086	3	<p>The sbxAutoBackup.sh changes for AWS SWE.</p> <p>Impact: The Cloud SBC uses default system name 'vsbcSystem' when creating system backups, making it difficult to distinguish between setups.</p> <p>Root Cause: The Cloud SBC configuration file backup logic uses the system name to create the back up file instead of the specific name.</p> <p>Steps to Replicate: Run /opt/sonus/sbx/scripts/sbxAutoBackup.sh on the SBC running in a cloud (e.g. AWS) that has SystemName configured too something specific in user-data.</p>	<p>Use the actual system name (configured in user-data) in the configuration backup filename to address the issue.</p> <p>Workaround: None.</p>
SBX-91016	3	<p>Unable to see the signaling messages on the PKT log file.</p> <p>Impact: The Packet capture does not work properly for pkt ports.</p> <p>Root Cause: The version of libpcap0.8 library in Debian9 had a defect that resulted in the capture not working.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Login to the SBC and run a basic call configuration. 2. Create a signaling packet capture filter: set global callTrace signalingPacketCapture state enable devices pkt0 0. 3. Run the SIPP calls. <p>Expected Results:</p> <ol style="list-style-type: none"> 1. The SBC should capture all the signaling messages on the pkt0 port in a PKT file under the gsxlog directory. 	<p>Update the libpcap0.8 to a newer version to address the issue.</p> <p>Workaround: None.</p>
SBX-74709	3	<p>The SBC REST API. pam phase authorization failed to login through a PAM: timeout</p> <p>Impact: Too many requests on the REST API interface causes 401 error.</p> <p>Root Cause: Issue with third party tool ConfD that is used for configuration management.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Send REST API requests for long duration. 2. Open some CLI session for some time to cause max session limit. 3. REST APIs start failing with 401 error due to session limit, as expected. 4. Close the CLI sessions. 5. REST APIs keep failing, even after keeping CLI sessions under max limit. 	<p>The code is modified to fix the issue.</p> <p>Workaround: None.</p>

SBX-101408	3	<p>The SIPREC was not working. No request sent to recorder.</p> <p>Impact: Without this fix, SBC is generating INVITE towards the SIPREC SRS with bad format on SDP header line 'label'. The 'label' header is missing the CR in the CRLF line break.</p> <p>Root Cause: The code that formed the SDP header line "label" does not add CR character for CRLF line break.</p> <p>Steps to Replicate: Make a call that would match the configured SIPRec criteria and thus triggering a SIPRec session for the corresponding Communication Session.</p>	<p>The code is modified to add CR character for CRLF line break, for SDP header line "label" .</p> <p>Workaround: None.</p>
SBX-104324	3	<p>The reenableOsAccount silently sets an account expiration to 30 days.</p> <p>Impact: Using the reenableOsAccount will add account aging to any OS account.</p> <p>Root Cause: The logic does not take into account the state of OS account aging or if OS account aging should be applied (e.g. root).</p> <p>Steps to Replicate: Run with an OS account aging enabled:</p> <ol style="list-style-type: none"> 1. Test root: <ol style="list-style-type: none"> a. ReenableOsAccount for user root. b. Check no aging has been applied at OS level: chage -l root. 2. Test the Confid user: <ol style="list-style-type: none"> a. Create user through CLI. b. ReenableOsAccount for this CLI user. c. Check no aging has been applied at OS level: chage -l <CLIUser>. 3. Test OS user with a password: <ol style="list-style-type: none"> a. Create OS user with a password. b. Disable and then Enable OS account aging state. c. ReenableOsAccount for this OS user. d. Check correct aging has been applied at OS level: chage -l <OSUser>. 	<p>The code is modified in the OS account aging is enabled and is applied to the account.</p> <p>Workaround: None.</p>
SBX-97555	2	<p>The active OAM node for the Signaling SBC generates header-only ACT files and pushing it to the billing server, where the OAM node should be responsible for the configuration only.</p> <p>Impact: The active OAM node generates header-only ACT files and pushing it to billing server.</p> <p>Root Cause: The SBC does not check whether it is OAM node or Managed VM while generating ACT files.</p> <p>Steps to Replicate: Spawned a Nto1 OAM node, check whether .ACT files exist or not. Restart active Node, Check if any .ACT file gets created. Restart the current Active again (assigned role Standby) and check whether any .ACT file is present on OAM node or not.</p>	<p>Check whether VM is OAM before creating the ACT files to address the issue.</p> <p>Workaround: None.</p>
SBX-102501	2	<p>The SBC fails to relay embedded header in Contact of 3xx with statusCode3xx relay enabled</p> <p>Impact: The SBC fails to relay embedded header in Contact of 3xx with statusCode3xx relay enabled</p> <p>Root Cause: The code required to send the embedded part in contact of 3xx was not present</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Configure the SBC for A to B call. 2. Enable the flag statusCode3xx on egress TG IPSP. 3. Send an INVITE from UAC. 4. Send 300 Multiple Choice from UAS with below Contact header that has embedded headers. 	<p>The code is modified so the contact header transparency is set in the relay 3xx scenario that is now being changed to full contact header transparency and set the SIP_HEADER_URI_HEADERS so that the embedded contact header is transparently passed in the 3xx. A new header type is introduced that is set only in the relay 3xx case to send the entire contact header transparently.</p> <p>Workaround: None.</p>
SBX-102234	2	<p>The SubsystemAdmin filter affects calltrace (TRC) logging showed useless logs.</p> <p>Impact: After creating then deleting an entry in the oam eventLog subsystemAdmin, the original behaviour is not restored. INFO level events for that subsystem are no longer logged, even if the oam eventLog typeAdmin filterLevel is info.</p> <p>Root Cause: Deleting an entry in oam eventLog subsystemAdmin does not restore settings back to default.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Create an entry for a subsystem in the oam eventLog subsystemAdmin table. 2. Delete the entry. 	<p>The code is modified to process to deleted the oam eventLog subsystemAdmin entry so that settings are put back to default values.</p> <p>Workaround: None.</p>

SBX-86090	2	<p>Our SIPREC implementation is prone to failure to record a stream and to tear down the SRS call leg when SRS (SIP Recording Server) sends a re-INVITE to SBC</p> <p>Impact: SIPREC handling had race conditions issues in the scenario where a RE-INVITE was received on the main call and also RE-INVITE was received from SRS server at the same time.</p> <p>This resulted in recording failure and the SBC sent out BYE towards SRS.</p> <p>Root Cause: This race condition of RE-INVITE from the main call end point and SRS server simultaneously was not handled on the SBC.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Make a SIPREC call. 2. Send RE-INVITE with codec change/hold on main call (either from UAC or UAS) this triggers a RE-INVITE towards SRS. 3. SRS also sends RE-INVITE at the same time when SBC receives 200 OK for main call RE-INVITE. 	<p>The RE-INVITE triggered/sent towards SRS due to the main call RE-INVITE is queued in the situation where the SIPREC leg is in middle of processing the RE-INVITE that was received from the SRS. The queued RE-INVITE would be sent once the RE-INVITE transaction received from the SRS was completed.</p> <p>Workaround: The SRS can avoid sending immediate hold RE-INVITE if no media is observed on SIPRec leg.</p>
SBX-104093	2	<p>The EMA codec entry screen was not usable.</p> <p>Impact: The EMA codec entry screen was not usable.</p> <p>Root Cause: The method call place mismatched.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Log into EMA. 2. Navigate to Profiles > Media > Codec Entry 3. See that the screen is working correctly. 	<p>The code is modified so the EMA codec entry screen is placed at the appropriate position.</p> <p>Workaround: None.</p>
SBX-74907	2	<p>A SIPREC Codec issue.</p> <p>Impact: Without this fix, the wrong codecs are negotiated with SRS when communication sessions is renegotiated.</p> <p>Root Cause: Because of wrong conditions in code, the codec information that is to be sent towards SRS is fetched from peer packet service profile rather than active service profile from the communication call leg.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. The SBC receives INVITE with G711A, G711U 2. Egress leg responds with A law and RTP stream is A law. 3. The SBC sends an INVITE to SIPRec server with G711A law and SIPRec servers responds with A law. 4. Egress leg sends re-INVITE and changes the order of codecs, G711U and G711A respectively. 5. The SBC relays the re-INVITE to ingress peer with G711U, G711A respectively and ingress peer responds with G711A. 6. The SBC was sending re-INVITE with G711U codec to SIPRec Server. With this fix, SBC now sends re-INVITE with G711A codec to SIPRec server. 	<p>The code is modified to fetch correct codec information from the communication session and the same is sent in recording session.</p> <p>Workaround: None</p>
SBX-102993	2	<p>The "numMatch notmatch" gets auto converted to "numMatch match" when SMM profile is modified/saved from GUI</p> <p>Impact: The "numMatch notmatch" gets auto converted to "numMatch match" when SMM profile is modified/saved from GUI</p> <p>Root Cause: The numMatch field was not there in the EMA.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Navigate SipAdaptorProfile screen. (ProfilesSignallingSipAdaptorProfile). 2. Create profile with the selection numMatch field. 3. Create profiles with all possible values. 4. Here the numMatch field was newly added. 5. Verify whether all field values are coming properly or not. 	<p>The code is modified to address the issue.</p> <p>Workaround: None.</p>
SBX-103692	2	<p>The SWE application will not start: The memory mismatch after a AWS instance restart.</p> <p>Impact: The standby SBC shuts down if the memory product capability mismatches with the active SBC.</p> <p>Root Cause: The SBC checks for a minimum required memory with the memory allocated to peer and it should be either equal or greater than the peer SBC.</p> <p>Steps to Replicate: Launch the SBC HA pair on AWS to see if the memory allocated to the standby SBC is lower than an active SBC.</p>	<p>The code is modified to change the memory from 'Minimum required' to 'Informational' to not shut down the app in case of a mismatch and only raises a trap.</p> <p>Workaround: Try rebooting in case memory allocated to standby the SBC is lower than the active SBC.</p>

SBX-104552	2	<p>After a switchover, the MRFP does not send a Goodbye packet for the Text stream.</p> <p>Impact: When a call is using T140/TTY interworking, after a switchover RTCP bye is not sent toward the T140 stream endpoint.</p> <p>Root Cause: In a T140/TTY interworking, the resource for T140 stream is not mirrored properly to standby causing deactivating does not work after a switchover.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Make call with T140/TTY interworking (for example, AMR-WB+T140 - PCMU). 2. Once call is established, switchover. 3. After a switchover was successful, end the call and see that RTCP bye is sent properly to T140 stream. 	<p>The code is modified to mirror resource correctly for T140 stream in T140 /TTY interworking.</p> <p>Workaround: None.</p>
SBX-104563	2	<p>The SBC is not updating the DNS servers order in DnsServerStatistics on deleting and creating a fresh DNS Group.</p> <p>Impact: The SBC was showing the incorrect DNS Servers Index in DnsServerStatistics command "show table addressContext default dnsGroup <dnsGroupName> dnsServerStatistics " when the DNS Servers were deleted and re-created in a different order with the same server IPs.</p> <p>After creating 128 DNS Servers with different unique IPs are and then deleting some of them such that the total number of DNS Servers is less than 128, no new DNS servers created post the delete operations shall have the corresponding statistics information, even though the total number of servers is less than 128 on the system.</p> <p>As an example, if on the SBC we create a DNS Server with IP1 and then delete it and then proceed on to repeat the create and delete operations with different IPs say IP2 to IP128, then from the IP128 onwards, no statistics shall be displayed, even though there is only one DNS Server on the system.</p> <p>Root Cause: The DNS Server Statistic blocks are not deleted when the DNS Servers are deleted and continues to remain hashed with the IP with that it was created.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Configure the DNSGROUP. 2. Add the 1st DNS server with IP1 to DNSGROUP. 3. Add the 2nd DNS server with IP2 to DNSGROUP. 4. Check the DnsServerStatistics. show table addressContext default dnsGroup <dnsGroupName> dnsServerStatistics 5. Delete the DNSGROUP. 6. Configure the DNSGROUP. 7. Add the 1st DNS server with IP2 to DNSGROUP. 8. Add the 2nd DNS server with IP1 to DNSGROUP with weight1. 9. Check the DnsServerStatistics. 10. The IP2 should have 3 as Index and IP1 should have 4 as index in dnsServerStatistics. 	<p>The code is modified to delete the DNS Server Statistics block and remove it from the IP hash when the DNS Server is deleted.</p> <p>Workaround: None.</p>
SBX-103732	2	<p>The AddressSanitizer: heap-buffer-overflow on address 0x61d001429c18 at pc 0x556c876cf74c bp 0x7fb14f64fab0 sp 0x7fb14f64f260 READ of size 3488 at 0x61d001429c18 thread T21 in the SAM Process.</p> <p>Impact: In a D-SBC environment while processing the remote media data command message, the SAM process could read off the end of the received message. Reading off the end of the allocated message block could, in the worst case, result in a coredump.</p> <p>Root Cause: The SAM process was assuming that the size of the received message was larger than it actually was and this resulted in reading off the end of the received message buffer. In most cases, this does not cause a problem but in could potentially result in a coredump if the associated buffer is at the end of the addressable memory space.</p> <p>Steps to Replicate: Run the RFC4117 test case for audio transcoding and video relay.</p>	<p>The code is modified to not read off the end of the received memory block.</p> <p>Workaround: None.</p>

SBX-102175	2	<p>The HA SBC Admin Connection failure issue within sometime for the ASAN build.</p> <p>Impact: In an HA setup, the standby SCM process can read off the end of a memory block. Reading off the end of allocated memory blocks can cause unexpected behaviour and in the worst case, it could result in coredumps. This problem only impacts the standby server so will not be service impacting.</p> <p>Root Cause: While making P-Asserted-Id header information redundant between active and standby, the standby was reading off the end of the allocated memory buffer because the active instance had passed over a bad data length value for P-Asserted-Id header.</p> <p>Steps to Replicate: Run a call load where the INVITE messages contain P-Asserted-Id header in an HA setup.</p>	<p>The code is modified to correctly format the P-Asserted-Id header information with valid length when sending it to the standby to avoid the standby accessing invalid memory.</p> <p>Workaround: None.</p>
SBX-105280	2	<p>The CPX Process had a memory leak for single call on the OAM node.</p> <p>Impact: The CPX process can leak small amounts of memory when creating/modifying the SNMP configuration.</p> <p>Root Cause: While processing the SNMP configuration commands, the SBC was creating the temporary internal memory blocks to read and write information from/to the CDB. But it was not freeing up this memory at the end of the configuration action.</p> <p>Steps to Replicate: Modify the SNMP configuration.</p>	<p>The code is modified to correctly free the temporary memory allocated during the configuration process.</p> <p>Workaround: None.</p>
SBX-105542	2	<p>Fix the Customer Telecom coverity issues.</p> <p>Impact: While processing the SUBSCRIBE messages, the coverity tool has highlighted that the code could dereference a pointer that is potentially null. Although no bad behaviour has been observed during testing, there is a small chance that it could result in coredumps if the pointer really was null.</p> <p>Root Cause: Based on other validation in the code, the coverity highlighted that some legs of code could result in accessing a pointer that might be null. Dereferencing null pointers can cause unexpected behaviour and in the worst case coredumps.</p> <p>Steps to Replicate: Run various SUBSCRIBE related test cases.</p>	<p>The code is modified to validate that the pointer is not null before using it to avoid any potential issues/coredumps.</p> <p>Workaround: None.</p>
SBX-91127	2	<p>The leaksanitizer had an CpxAaaGetNextEntry.</p> <p>Impact: While processing the requests to display the user information /status, the Cpx process was leaking memory.</p> <p>Root Cause: While processing the requests to display user information /status, the Cpx process had to allocate the internal memory blocks to collect information from CDB and was not freeing up this memory at the end of the request.</p> <p>Steps to Replicate: From the CLI, run commands to request user information.</p>	<p>The code is modified to free up the memory at the end of the processing request.</p> <p>Workaround: None.</p>
SBX-105389	2	<p>De-reference after a NULL check and de-reference the NULL return value.</p> <p>Impact: The coverity scanning tool identified a potential edge case scenario, where the lawful intercept code could potentially access a NULL pointer leading to unexpected behaviour and potentially a coredump.</p> <p>Root Cause: While the code was performing X2 peer status updates, it retrieved the peer information from a hash table but did not verify that the pointers inside the record it retrieved were valid. It would be an extreme error case for this to result in reading of a null pointer and would likely need to be due to another problem. But coverity highlights it as an edge case issue to fix up.</p> <p>Steps to Replicate: Run test cases for lawful intercept using packet cable V2.0 configuration.</p>	<p>The code is modified to validate the pointer is not null before using it and avoid any error scenarios.</p> <p>Workaround: None.</p>

SBX-100225	2	<p>The AddressSanitizer: detected heap-use-after-free in UasProcessMsgCmd on address 0x623000187198 at pc 0x5623cc7b3b42.</p> <p>Impact: The SCM process is accessing memory after it has been freed during the timer expiry handling when there is no response to an OPTIONS message that was triggered due to debug optionsPing using an FQDN. Accessing memory after it has been freed can cause unexpected behavior and in the worst case, this potentially causes core dumps.</p> <p>Ex: "request addressContext default cmds optionsPing peerFQDN bats3.gsxlabs.com peerPort 14090 sigPort 2 transport udp"</p> <p>Root Cause: The SIP dialog memory block was being removed in two places while handling the timeout for the debug optionsPing command, which could result in unexpected behaviour and potentially result in core dumps.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Run debug optionsPing command "request addressContext default cmds optionsPing peerFQDN bats3.gsxlabs.com peerPort 14090 sigPort 2 transport udp" 2. The OPTIONS message will be sent out. 3. Let the options timeout. (do not send any response). 	<p>The code is modified to address the issue.</p> <p>Workaround: Do not use this debug command, or use it with an IP instead of an FQDN.</p>
SBX-105591	2	<p>Observed that PRS process heap buffer overflow issue on 9.2 ASAN build 63</p> <p>Impact: The BRM process was accessing memory after it had been freed. As the memory is being read immediately after being freed, then it is unlikely this would cause a problem.</p> <p>Root Cause: If the BRM process received an unexpected message, then it was trying to print a debug message to indicate the message type. However, the message had already been freed up because it was unexpected. This was observed on the standby server and generated at the point of switchover.</p> <p>Steps to Replicate: Run a normal call load and perform switchovers.</p>	<p>The code is modified to collect the message type information before the message is freed so that the debug log content can be safely generated without accessing freed memory.</p> <p>Workaround: None.</p>
SBX-103650	2	<p>The SCM Process cores when the targets set for OPTIONS/MESSAGE.</p> <p>Impact: Without this fix, the SBC will core dump when the Out Of Dialogue messages like OPTIONS/MESSAGE are intercepted.</p> <p>Root Cause: Double free of a data structure is causing the code dump.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Provision packet cable 2.0 Targets corresponding to the out of dialogue OPTIONS that will be sent to the SBC. 2. Send the OPTIONS to the SBC. 	<p>The code is modified to avoid double freeing of the corresponding data structure. Also, addressed couple of memory leaks concerned with the same data structure in error cases such as ARS blacklisting.</p> <p>Workaround: None.</p>
SBX-103788	2	<p>The LeakSanitizer: detected memory leaks in CpxAppProcess.</p> <p>Impact: Small memory leak when making configuration changes to the GW signaling port.</p> <p>Root Cause: The configuration logic was reading the interface group name from CDB into a temporary local variable in order to perform validation logic but never freed up this memory at the end of the validation.</p> <p>Steps to Replicate: This issue is highlighted in the engineering lab when performing GW sig port configuration changes on an ASAN enabled build.</p>	<p>The code is modified to correctly free the temporary memory to avoid the leak.</p> <p>Workaround: None.</p>
SBX-105412	2	<p>The MRFP: CE_2N_Comp_CpxAppProc leak for port SWO (BFD).</p> <p>Impact: There was a small memory leak during the configuration of packet port with BFD configuration associated.</p> <p>Root Cause: The configuration code was reading information from the CDB and storing information in an internal memory block but not freeing it up.</p> <p>Steps to Replicate: With a BFD configuration on the SBC, disable and enable the pkt port by setting mode OOS and state disabled and then enable the pkt port again.</p>	<p>The code is modified to correctly release the internal memory block at the end of the configuration action.</p> <p>Workaround: None.</p>

SBX-105496	2	<p>Fixing the NRMA coverity issue CID:11149.</p> <p>Impact: The coverity scanning tool identified a potential code leg where a null pointer could be dereferenced and could potentially result in a core dump, although not observed in internal testing.</p> <p>Root Cause: While trying to allocate resources for PCMU to PCMA call where the ingress leg is being tapped, there is a possibility the code could access and invalid pointer resulting in unexpected behaviour and in potential core dumps.</p> <p>Steps to Replicate: This part of the code could be triggered for PCMU to PCMA call where the ingress leg is being tapped.</p>	<p>The code is modified to verify the pointer is not NULL before trying to use it to avoid potential core dumps.</p> <p>Workaround: None.</p>
SBX-105200	2	<p>The AddressSanitizer detected heap-use-after-free on address 0x6080002177a0 at pc 0x55a942cce5b1 bp 0x7fb90b2df190 sp 0x7fb90b2df188.</p> <p>Impact: There were two types of issues identified in this Jira:</p> <ol style="list-style-type: none"> 1. There was a memory leak when putting a virtual media gateway in service. 2. The MRFP is accessing memory after it is freed while cleaning up a call following an internal resource allocation failure. Accessing memory after it has been freed can have unexpected behavior and in the worst case result in core dumps. <p>Root Cause:</p> <ol style="list-style-type: none"> 1. The MRFP was allocating internal memory while performing validation of the request to put a virtual media gateway in service, but the memory was not being freed up at the end of the configuration action leading to a small memory leak. 2. In case of internal call failure, the call data was being deleted but the memory was being access later on in the clean up procedure. <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Configure a virtual media gateway in service. 2. Configure the MRFP without licenses and then try to make some calls. 	<ol style="list-style-type: none"> 1. The code is modified to free the memory at the end of the configuration action. 2. The code is modified to no longer access the call data after is free. <p>Workaround: None.</p>
SBX-104118	2	<p>The LeakSanitizer detected memory leaks in the SCM Process.</p> <p>Impact: While relaying REGISTER messages, the SBC may leak memory blocks allocated to hold the record-route header information.</p> <p>Root Cause: The SBC allocates memory blocks to hold the contents of the record-route headers in the REGISTER message. But if the record-route header information is associated with the SBC IP address, then the SBC does not correctly free up the memory blocks causing a leak.</p> <p>Steps to Replicate: Run test cases for stateless REGISTER relay call scenarios where the IP information in the record-route header is the SBC's IP.</p>	<p>The code is modified to correctly freed up memory allocated to store the record route header information.</p> <p>Workaround: None.</p>
SBX-105850	2	<p>The MRFP failed to come up after an sbxstart.</p> <p>Impact: At startup of the SBC, there is a possible race condition where the SBC processes may go for an additional restart before they come up and restore the configuration. data.</p> <p>Root Cause: The race condition in the code between threads.</p> <p>Steps to Replicate: Installation and startup on the various platforms.</p>	<p>The race condition in setting some common variables is avoided by using static path of binaries to address the issue.</p> <p>Workaround: None.</p>

SBX-103626	2	<p>Due to the Adaptive Codec Change, the MRFP does not disable the text stream although there is conflict between the payload types that the non audio and text streams use.</p> <p>Impact: In a audio and T140 call setup, with Adaptive Codec Change enabled and the MRFP does not disable the text stream when there is conflict between the payload types that the audio and text streams use.</p> <p>Root Cause: In the Adaptive Codec Change handling, the code does not validate if the payload type used in the T140 stream are in conflict with payload types being used in the audio stream codecs. As a result, the T140 stream is created in the media plan that should be disabled in this scenario.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Enable Adaptive Codec Change in the VMG configuration. 2. Setup a call audio and T140 media streams from both call terminations: A- side AMRWB and text 140 to B side, AMR, text T140, and Telephone event. 3. At the mid call, send a Re-Invite to change B's side to AMRWB, T140 and Telephone Event with its payload type is the same as the T140's payload type. 4. Verify that MRFP should disable the T140 stream by setting the t140 stream port number to zero in the reply to MGC(C3). 	<p>The code is modified to check the conflict for T140 payload types against the payload types being used in an audio stream, and reject T140 stream when a conflict is found.</p> <p>Workaround: None.</p>
SBX-105039	2	<p>The SBC fails to perform Alert-Info to PEM interworking when the relayPemState is disabled on the ingress.</p> <p>Impact: The SBC fails to perform Alert-Info to PEM interworking when the relayPemState flag is disabled on the ingress TG.</p> <p>Root Cause: The relayPemState flag implementation was not considered for Alert-Info to PEM Interworking scenarios.</p> <p>Steps to Replicate: Configuration:</p> <ol style="list-style-type: none"> 1. The relayPemState flag is disabled on the Ingress TG. 2. The iToPemInterworking flag is enabled on the Egress IPSP. 3. The acceptAlertInfo flag is enabled on the Egress IPSP. <ol style="list-style-type: none"> 1. The UAC sends an Invite with the SDP with PEM: supported to the SBC. 2. The UAS sends 180 without the SDP with an alert-info as rt to the SBC. 3. The SBC should interwork the alert-Info to PEM and insert PEM: inactive while sending 180 responses towards UAC. 	<p>The code is modified to consider the Alert-Info to PEM Interworking scenarios when the relayPemState flag is disabled.</p> <p>Workaround: None.</p>
SBX-65509	2	<p>The preferred payload number was not used for either dtmf or amrwb when the "different2833PayloadType" transcode option is enabled.</p> <p>Impact: The preferred payload number was not used for either dtmf or amrwb when the "different2833PayloadType" transcode option is enabled.</p> <p>Root Cause: Both the AMR-WB and dtmf do not use 101, as 101 was being neglected and thinking the other one was using it. So at the end, it was not being considered and was not being populated.</p> <p>Steps to Replicate: Make a SIPP call using AMR-WB codec with dtmf enable different2833palyloadtype flag. enable honorSdpClockRate flag. set preferredRtpPayloadType as 101 set preferredRtpPayloadTypeForDtmfRelay 101</p> <p>Check for the Preferred payload number.</p>	<p>The code is modified to scan even the workingDynamicPtSet and verify that 101 is not being used and is considered.</p> <p>Workaround: The workingDynamicPTSet represents a more accurate information w.r.t the PTs used. As a result, that condition is being added to scan workingDynamicPtSet along with peer PSP and Route PSP.</p>
SBX-104471	2	<p>There was a CE_node error "glusterSetup.sh Abort: Mount is in bad state!!"</p> <p>Impact: A glusterSetup.sh script error log showing in CE_Node log file.</p> <p>Root Cause: A glusterSetup.sh script error text is printed to the terminal.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Shut down the Standby OAM. 2. Ensure there is no glusterSetup.sh log appear in CE_Node log file on managed node. 	<p>The error text already captured in two separate log files, so remove the output to terminal to address the issue.</p> <p>Workaround: None.</p>

SBX-101406	2	<p>Implement SAN validation for cert authentication in VNFR REST Server</p> <p>Impact: The VNFM request did not have SAN validation in addition to a cert chain verification.</p> <p>Now, the SAN validation checks against IPv4 and IPv6 both if present in userData.</p> <p>Root Cause: The SAN validation was not implemented.</p> <p>The SAN validation used to work only in mode(either IPv4 or IPv6) depending upon the SBC instantiation mode.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Bring VNFR(running HTTPS) from VNFM GR. 2. Perform a successful reassignment. 3. Do a sbxrestart. 4. Again perform a successful reassignment. 5. Bring up a dual stack VNFM, and spawn IPv4 and IPv6 SBC. 6. Register VNFR successfully in a simplex VNFM. 	<p>Implemented SAN validation using GnuTLS that is also used for cert chain verification.</p> <p>Loading both IPv4 and IPv6 from userdata addresses the issue.</p> <p>Workaround: No workaround.</p>
SBX-94760	2	<p>The cinder volume detach/attach required a reboot from Openstack to bring up all the functioning related to the SBC.</p> <p>Impact: In the case of Openstack, if cinder volume is detached from the running the SBC instance, it does not cause the SBC system to shutdown properly.</p> <p>Root Cause: The shutdown triggered on the volume detach from a running instance was not going through successfully due to failure to run abnormal reboot command.</p> <p>Steps to Replicate: Test cinder volume detach from a running instance and ensure the node goes for a reboot and if running on active, should switchover to standby successfully.</p>	<p>The code is modified to update the command to forcefully reboot the node on volume detach from a running instance.</p> <p>Workaround: None.</p>
SBX-93790	2	<p>A TEAMS user was unable to resume the call when transfer is failed, when the call is initiated in beginning by PSTN user.</p> <p>Impact: During a Refer, if the transfer target sends early answer in 18x but then rejects the call, the SBC fails to resume the previously active call.</p> <p>Root Cause: During call transfer, after tone play, while processing early answer in 183, the SBC wrongly freed the previous cut-thru context and instead retains the previously activated tone context (for A-B call).</p> <p>As a result, after transfer target rejects the call, the SBC attempts to resumes the previously active call (A-B) which fails due to unavailability of correct context.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. PSTN1 to TEAMS call TEAMS transfer call to PSTN2 PSTN2 rejects the call TEAMS resume the call 2. PSTN1 To TEAMS call TEAMS transfer call to PSTN2 PSTN2 does not answer the call TEAMS resume the call 	<p>The code is modified to retain the previously activated cut-thru context and free the previously activated tone context if current tone context is more recent.</p> <p>Workaround: No</p>
SBX-104331	2	<p>Observed the "NrmUpdateLicensesForXferFeatureType cant get origGCID" MAJOR Logs on the 5400 Platform while running a call with Multiple INVITE Replace.</p> <p>Impact: The Call Pickup logs were incorrectly printing a major error event.</p> <p>Root Cause: The Call Pickup logs were incorrectly printing a major error event.</p> <p>Steps to Replicate: None.</p>	<p>The code is modified to address the issue.</p> <p>Workaround: None.</p>

SBX-104738	2	<p>The outputAdaptor profile rule gets applied after the clearDB.</p> <p>Impact: The SMM rule was not applied after the upgrade from 8.2 or 9.0 or 9.1 to 9.2, even the same issue is observed during restart or switchover.</p> <p>Root Cause: The SMM rule was not applied due to wrong path in the switchover scenario.</p> <p>Steps to Replicate: Configure the SMM profile in the SBC.</p> <p>Perform any of the of the below steps.</p> <ol style="list-style-type: none"> 1. Restart the SBC. 2. The SBC performs a switchover. 3. Upgrade from 8.2/9.0/9.1 to 9.2. 	<p>The code is modified to address the issue.</p> <p>Workaround: ClearDB and load the configuration.</p>
SBX-95677	2	<p>The SBC is not feeding the delayed RBT on monitoring a failure in a late media passthrough scenario in the CLOUD ISBC and SWE ISBC.</p> <p>Impact: The SBC is not feeding delayed RBT on monitoring the failure in a late media passthrough call.</p> <p>Root Cause: The fix for the SBC was not feeding delayed RBT on the monitoring failure in a late media passthrough scenario was applicable only when bToneAsAnnc flag is enabled.</p> <p>Steps to Replicate: Procedure:</p> <ol style="list-style-type: none"> 1. An INVITE is received without SDP and no "PEM: supported" from the UAC. 2. The UAS sends the 183 with SDP and no PEM having PCMU, AMR. 3. The UAC send PRACK with SDP PCMU,G729 and UAS sends 200 OK for PRACK. 4. The UAS sends an authorized RTP. 5. The UAS sends 180 without SDP and no PEM. 6. The UAC send PRACK and UAS sends 200 OK for PRACK. <p>After a RTP Monitoring failure, the SBC should play the tone.</p>	<p>The code is modified to address the issue.</p> <p>Workaround: None.</p>
SBX-104136	2	<p>Unable to login to the CLI session.</p> <p>Impact: The SWe Active profile configuration may cause a password change commit to hang.</p> <p>Root Cause: The Internal Configuration change related to sweActiveProfile leaves the changes in the candidate database. This causes another commit from the SM Process deadlock, as sweActiveProfile change subscription needs the SM Subscribers acknowledged.</p> <p>Steps to Replicate: This cannot be recreated through User Interfaces. This was an internal error condition.</p>	<p>The code is modified to revert the changes from the candidate database, if the sweActiveProfile commit fails.</p> <p>Workaround: None.</p>
SBX-102958	2	<p>The Audit Compliance issues are found in the MRFP Setup.</p> <p>Impact: The Nessus Scan with the CIS Plugin shows compliance failures on the SBC.</p> <p>Root Cause: The failures were due to a missing a user directory and bad file permissions.</p> <p>Steps to Replicate: Run the Nessus scan with the same policy and verify that the failed compliances do not exist.</p>	<p>Below changes are made to meet the CIS requirements.</p> <ol style="list-style-type: none"> 1. Set nologin shell for cwagent 2. Change remoteExecution.log file permission. <p>Workaround: None.</p>
SBX-103273	2	<p>Dual NUMA support in the SWe.</p> <p>Impact: The SWe does not come up in the VMs having multiple NUMA nodes except for the signaling traffic profile (SLB, S-SBC, SO-SBC).</p> <p>Root Cause: Due to a deliberate software restriction to not allow multiple NUMAs, the issue was shown.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Configure a KVM instance with the dual NUMA and install in Non-Gold/Non-GPU setup. 2. Let it come in default traffic profile. 3. Verify that the SBC would come up fine without any HostCheck errors. 	<p>The code is modified to allow multiple NUMA irrespective of the personality and profile.</p> <p>Workaround: None.</p>

SBX-103627	2	<p>The BFD status is not up for the MGMT port (the BFD packets are seen though the LDG state is disabled).</p> <p>Impact: Even after disabling Link Monitor on the MGMT port, the BFD packets are still seen on the MGMT port.</p> <p>Root Cause: The issue was that when state is disabled and the BFD session was deleted. A timer is started so the expiry initiates a BFD session.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Create BFD session for MGT port. 2. Check the status of BFD session. 3. Disable the BFD session. 4. Observed the BFD packets even when the BFD session is disabled. 	<p>The code is modified to check for the Link Monitor state before initiating the BFD session.</p> <p>Workaround: None.</p>
SBX-103489	2	<p>The LeakSanitizer: detected memory leaks at the confd_malloc.</p> <p>Impact: The small memory leak during the configuration of lawful intercept server.</p> <p>Root Cause: The configuration code was reading information from the CDB and storing information in an internal memory block but not freeing it up.</p> <p>Steps to Replicate: Create a lawful intercept server and make configuration changes.</p>	<p>The code is modified to correctly release the internal memory block at the end of the configuration action.</p> <p>Workaround: None.</p>
SBX-103387	2	<p>The Video NAT learning is failing on the D-SBC cloud.</p> <p>Impact: The Video NAT learning failed in the D-SBC.</p> <p>Root Cause: On the D-SBC, NAT learning for video stream was not processed successfully and as a result caused the video packets to be dropped by the SBC.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Configure the SBC for an Audio and Video call. 2. Enable the Signaling NAT and Media NAT in the Ingress and Egress Trunk Groups. 3. Make an Audio and Video call between the NATed EndPoints EP1 and EP2. 4. After a call gets established, the EP1 and EP2 sends Audio and Video packets. 5. NAT learning happens for Audio and Video and then the packets are sent and received from EP1 and EP2. 	<p>The code is modified to process the NAT learning for the video stream in the D-SBC.</p> <p>Workaround: None.</p>
SBX-103353	2	<p>The AddressSanitizer: detected heap-use-after-free</p> <pre data-bbox="282 1203 672 1283">/localstore/ws/jenkinsbuild/sbxmainasan/marlin/SIPSG /sipsgCallNotification.c:1123 in</pre> <pre data-bbox="282 1308 711 1350">SipSgSendCallNotificationApi(SIPSG_CCB_STR*, SIPSG_CALL_NOTIFY_TYPE_ENUM, sip_nameaddr_str*).</pre> <p>Impact: While freeing up a SIP call, the code is accessing the SIP call control memory block immediately after it has been freed up in the same processing loop.</p> <p>Root Cause: This issue was detected using ASAN images, it has not been proven to cause bad behaviour using regular production images, but accessing memory after it has been freed can cause unexpected processing to happen which might potentially result in coredumps.</p> <p>Steps to Replicate: Run regular call scenarios with ASAN images.</p>	<p>The code is modified to avoid using the memory block after it is free.</p> <p>Workaround: None.</p>

SBX-104016	2	<p>Anomalies were observed after decoding a ACT File by sbxCamDecoder.pl.</p>	<p>sbxCamDecoder was updated to support decoding REBOOT, SWITCHOVER records and the "Ingress Signaling Information" and "Egress Signaling Information" in the EVENT Records.</p>
		<p>Impact:</p> <p>Issue 1: The camDecoder script "sbxCamDecoder.pl" does not decode a reboot, the switchover records.</p> <p>Issue 2: It does not display the field names under the "Ingress Signaling Information" and "Egress Signaling Information" for an EVENT records as shown below.</p> <pre> 26. Ingress Signaling Information : 26.1 : "sip:7587339665@10.xx.xxx.xxx 26.2 : 1-23946@10.xx.xx.xxx 26.3 : <sip:sipp@10.xx.xx.xxx:xxxx>;tag=23946SIPpTag011 26.4 : <sip:7587339665@10.xx.xxx.xxx> 26.5 : 26.6 : sip:sipp@10.xx.xx.xxx:xxxx 26.7 : 26.8 : 26.9 : 26.10 : 0 26.11 : 26.12 : 26.13 : 26.14 : 1 26.15 : 26.16 : 26.17 : 26.18 : 26.19 : 26.20 : 26.21 : 26.22 : 26.23 : 26.24 : 26.25 : " </pre> <p>Root Cause: The sbxCamDecoder script was not updated to support the reboot, on switchover records.</p> <p>Also, the script was not updated to decode "Ingress Signaling Information" and "Egress Signaling Information" for EVENT records.</p> <p>Steps to Replicate:</p> <p>Issue 1: The switchover record issue. Setup: The SBC HA Procedure:</p> <ol style="list-style-type: none"> 1. Perform a switchover the SBC. 2. Decode the latest ACT log that contains "SWITCHOVER" record by using the sbxCamDecoder.pl. <p>Issue 2: REBOOT Record issue.</p> <ol style="list-style-type: none"> 1. Reboot the SBC. 2. Decode the latest ACT log that contains the "REBOOT" record, by using sbxCamDecoder.pl. <p>Issue 3:</p> <ol style="list-style-type: none"> 1. Send an OOD message (OPTIONS). 2. Decode the latest ACT log that has the event RECORD by using the sbxCamDecoder.pl. 	<p>Workaround: Decode the record manually.</p>
SBX-103894	2	<p>The AddressSanitizer: detected heap-use-after-free on address 0x61b000018f84 at pc 0x556c7afd4213 bp 0x7fcbf0905680 sp 0x7fcbf0905678.</p> <p>Impact: An invalid memory access of a termination object after it's been deleted. Accessing the memory after it has been freed can result in unexpected behaviour and in the worst case coredumps.</p> <p>Root Cause: In the call teardown processing, after the last termination is deleted, the calltracing function was accessing the already deleted call termination object to retrieve the context ID.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Setup a call in MRFP. 2. Teardown the call. 3. The previously mentioned invalid memory access occurs in handling the deletion of the last termination of the call. 	<p>In the same function, retrieve the context ID from the context object instead to address the issue.</p> <p>Workaround: None.</p>

SBX-103814	2	<p>The RECORDING CDR does not have correct value for the SRTP field during switchover scenario.</p> <p>Impact: For the SIPREC sessions with SRTP, after a switchover the "RECORDING" CDR generated had values as 0 in the SRTP statistics fields as shown below:</p> <p>24.7 ingress SRTP info1 : 0:0:0:0 24.12 egress SRTP info1 : 0:0:0:0</p> <p>Root Cause: The standby processing code did not support copying the SRTP information into the SIPREC Standby data blocks and as a result was lost during a switch over.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Execute a SIPREC call, with a SIPREC Session over the SRTP. 2. Perform a switchover. 3. Verify the RECORDING CDR for SRTP Statistics in the following fields: <p>24.7 ingress SRTP info1 : 0:0:0:0 24.12 egress SRTP info1 : 0:0:0:0</p>	<p>The code is modified to process the SIPREC SRTP information is added.</p> <p>Workaround: None.</p>
SBX-103807	2	<p>The SBC disables the SRTP for an audio when the EP disables and re-enables the audio.</p> <p>Impact: The SBC disables the SRTP for audio when EP disables and re-enables the audio.</p> <p>Root Cause: During the modify offer processing, the SRTP value is taken from previous Active security PSP without checking if the SRTP is valid or not.</p> <p>Steps to Replicate: Test Procedure:</p> <p>Setup Audio+Video RTP to SRTP pass-Thru call between UAC-UAS:</p> <ol style="list-style-type: none"> 1. UAC sends invite with Audio and video. 2. UAS responds with Audio and Video cryptoline. 3. UAC sends re-invite to disable Audio stream port=0 and inactive and Video with valid media port and sendrecv. 4. UAS responds 200OK with port=0, a=inactive and no crypto line for audio and video with valid port and crypto line. 5. UAC sends re-invite to add Audio stream back with valid media port and sendrecv and video with valid media port and sendrecv. 6. UAS responds 200OK with valid Audio and Video crypto lines. <p>Expected Result:</p> <ol style="list-style-type: none"> 1. RTP-SRTP A+V call gets established. 2. The SBC disables audio stream and sets up RTP-SRTP call with video stream. 3. The SBC re-establishes Audio+Video RTP-SRTP call. <p>Actual Result:</p> <ol style="list-style-type: none"> 1. RTP-SRTP A+V call gets established. 2. The SBC disables audio stream and sets up RTP-SRTP call with video stream. 3. The SBC re-establishes re-Invite without SRTP for audio stream. 	<p>The code is modified to copy the Active security PSP only if the SRTP admin state is enabled.</p> <p>Workaround: None.</p>
SBX-105310	2	<p>There was a SM Process leak for the MRFP call on the MRFP active.</p> <p>Impact: The redundancy group manager (RGM) that is used in conjunction with N:1 deployments has a memory leak.</p> <p>Root Cause: The RGM handles the switchover and fault messages in N:1 deployments and it was not freeing up the ICM message after processing that results in a memory leak.</p> <p>Steps to Replicate: This issue was observed in an MRFP configuration especially when restarted instances.</p>	<p>The code is modified to correctly free the ICM messages.</p> <p>Workaround: None</p>

SBX-103603	2	<p>The LeakSanitizer: detected memory leaks in the MrIRmProcessTrmAllocRpyMsg.</p> <p>Impact: While testing call scenarios for RTCP for T.140 Pass-through functionality, it was observed that the SCM process could leak memory for calls associated with an MRF (external media transcoder).</p> <p>Root Cause: The SBC was allocated memory while processing the SDP associated with this call but was not always freeing up the memory at the end of the call.</p> <p>Steps to Replicate: Run various call scenarios with MRF where the SBC is using the SIP to allocated media resources on an external media transcoder (MRF) or T-SBC.</p>	<p>The code is modified to correctly free all the memory allocated for the call.</p> <p>Workaround: None.</p>
SBX-103731	2	<p>The AddressSanitizer: detected heap-buffer-overflow on address 0x61a0000cb9d9 observed in the SAM Process while running OCSP feature.</p> <p>Impact: When the OCSP stapling feature is enabled on the SBC and the code was processing the response it writes to unallocated memory and in the worst case this could result in process core dumps.</p> <p>Root Cause: While processing OCSP response the code was allocating a memory buffer large enough to hold the response, but then incorrectly writing one byte off the end of the memory buffer while attempt to try and null terminate the string.</p> <p>Steps to Replicate: Setup - UAC > Dut<->Adapter -> UAS</p> <ol style="list-style-type: none"> 1. Create an OCSP profile by configuring the defaultResponder and stapling enabled. 2. Attach the OCSP profile to the TLS profile configured. 3. From Endpoint A, Intiate the TLS call with Client Hello from the SIPP UAC having OCSP parameter in it. 4. The SBC should send server hello certifiacte to the user client. 	<p>The code is modified to correctly terminate the OCSP response string without writing off the end of the memory buffer allocated to hold the response.</p> <p>Workaround: None.</p>
SBX-103801	2	<p>Observed the run time error in M-SBC "runtime error: load of value 42, which is not a valid value for type 'bool'"</p> <p>Impact: The M-SBC could potentially configure the wrong data path mode for a call.</p> <p>Root Cause: No issues were observed while running this functionality on a regular deployment image. But while testing with ASAN, it highlighted that some of the fields used in a call resource allocation message were not always initialized correctly, which could potentially lead to unexpected behaviour e.g. SYS_ERRS, wrong datapath mode setup.</p> <p>Steps to Replicate: This issue was highlighted while while running test suite related to RFC-4117 MRF Mid-Call modification Enhancement</p>	<p>The code is modified to correctly initialize the resource allocation request fields to avoid issues.</p> <p>Workaround: None.</p>
SBX-98024	2	<p>The contact header is not transparently passed when the Ingress and Egress had different transport types.</p> <p>Impact: Contact header is not passed transparently from Ingress, when egress side has different transport type, even when the IPSP flag 'passCompleteContactHeader' is enabled on both ingress/egress Trunk Groups.</p> <p>Root Cause: In API SipSgCheckAndSetContactHeaderTransparency(), irrespective of transparency control is enabled/disabled, if the egress Sig Zone is MS Teams, then contact header was not transparently passed to egress.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Configure, IPSP flag 'passCompleteContactHeader' to enable. 2. Attach the IPSP to both ingress/egress TG's. 3. Set Ingress transport to TCP / Egress to TL. 4. Create pathCheck profile with hostName 'sip.pstnhub.microsoft.com' and attach to the Teams side. 5. Make a successful call. 	<p>The code is modified so regardless of MS Teams zone, if the transparency control flag is enabled then pass the contact header transparently to the egress.</p> <p>Workaround: None.</p>
SBX-101937	2	<p>The one medium vulnerability found after the Qualys Scan.</p> <p>Impact: A medium vulnerability found after the Qualys Scan in JQuery 3.4.1 version.</p> <p>Root Cause: The jquery 3.4.1 version has security issue.</p> <p>Steps to Replicate: Run a Qualys Scan and vulnerability was not shown.</p>	<p>The code is modified to address the issue.</p> <p>Workaround: No workaround.</p>

SBX-102718	2	<p>The Confd updateAdmin user's configuration was failing.</p> <p>Impact: In cloud deployments, when a new user is created, it throws the error: "Cannot change accountAgingState to disabled while accountRemovalState is enabled" even though accountAgingState is set from disabled and accountRemovalState to disabled.</p> <p>Root Cause: The validation logic for user creation was done along with transformation callbacks. This makes the validation logic to be order dependent.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Launch a cloud standalone instance. 2. Create a new user with accountAgingState to disabled and accountRemovalState to disabled: commit fails with error "Cannot change accountAgingState to disabled while accountRemovalState is enabled" 	<p>The code is modified to make an order independent.</p> <p>Workaround: None.</p>
SBX-103634	2	<p>The LeakSanitizer: detected memory leaks in DcmRestoreAllMetaVarsToStandbyContext.</p> <p>Impact: A small memory leak was observed in the SAM process while performing a switchover from standby to active box.</p> <p>Root Cause: The standby instance stores information about the metavaris associated with signaling ports configured on the active instance. During the conversion from standby to active, the standby data is moved into active structures but the original standby memory blocks are not freed up correctly resulting in a memory leak.</p> <p>Steps to Replicate: Perform a switchover from standby to active while running a basic call, no memory leak will be observed in the SAM process.</p>	<p>The code is modified to correctly free the memory associated with the standby data when it gets transitioned to active to avoid the memory leak.</p> <p>Workaround: None.</p>
SBX-104360	2	<p>The SWITCHOVER ACT Records are not generated in the SBC with HA mode set as Nto1.</p> <p>Impact: On an N to 1 system, the SWITCHOVER record is not written to the accounting file on the new active node after a switchover.</p> <p>Root Cause: The SWITCHOVER record is sent to the ENM process before it has opened the accounting file, so the record is not written.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Perform a switchover a system. 2. Check that the switchover record is written to the accounting file. 	<p>The SWITCHOVER record is stored and then written out when the accounting file is opened to address the issue.</p> <p>Workaround: None.</p>
SBX-104826	2	<p>The SBC fails to relay to 3xx and picks the next route when EnhancedLocalRedirection and StatusCode3xx flags are enabled.</p> <p>Impact: The SBC fails to relay the received 3xx to the ingress side and picks the next route when the EnhancedLocalRedirection and StatusCode3xx flags are enabled with 2 routes.</p> <p>Root Cause: When both the StatusCode3xx and EnhancedLocalRedirection are enabled for the 2 routes scenario, the performCrankback is set to true which leads to this issue.</p> <p>Steps to Replicate: Procedure:</p> <ol style="list-style-type: none"> 1. Configure the SBC for A to B call over ERE. 2. Configure the ERE to return two routes for Initial call R1, R2. 3. Enable the flag statusCode3xx, EnhancedLocalRedirection on TG IPSP. 4. Send the INVITE from UAC. 5. Send the 300 Multiple Choice from UAS with Contact header and embedded headers. 	<p>The code is modified to ensure that performCrankback is not set to true when both the EnhancedLocalRedirection and StatusCode3xx are enabled.</p> <p>Workaround: None.</p>

SBX-91592	2	<p>The DRBD tuning is not optimal.</p> <p>Impact: On the non-cloud 1:1 SBCs, due to non-optimal tuning of the DRBD, the DRBD connection between active was getting disconnected leading to full sync and high i/o on the DRBD partition.</p> <p>Root Cause: Due to the peer not responding within a certain time (ko-count * timeout), ko-count feature of DRBD was disconnecting the peer leading to full sync.</p> <p>Steps to Replicate: The following are the steps to test the changes:</p> <ol style="list-style-type: none"> 1. Decrease the timeout value in the DRBD conf file. 2. Restart the DRBD service. 3. Resume the DRBD sync if not already enabled. 4. Check syslog, the DRBD must not get disconnection logs. 	<p>The DRBD ko-count feature is disabled on the SBC to address the issue.</p> <p>Workaround: None.</p>
SBX-101743	2	<p>The SBC start is showing some errors related to a serf file and config-version file.</p> <p>Impact: Missing the file error printed to terminal.</p> <p>Root Cause: The gluster setup script does not redirect output of 'cat' command.</p> <p>Steps to Replicate: Run the sbxstop; sbxstart</p>	<p>The 'cat' command error output is redirected to NULL to fix the issue.</p> <p>Workaround: None.</p>
SBX-91799	2	<p>The segfault in pamValidator during PM login with incorrect credentials.</p> <p>Impact: The segfault in pamValidator on a failed login for locked user.</p> <p>Root Cause: The pamValidator defines a conversation function that is called by pam modules to exchange values. The pam module was freeing a struct member in the first call and accessing the same freed value in another call to the conversation function that was causing segmentation fault.</p> <p>Steps to Replicate: Perform following steps on any of the SBC deployment and verify that segmentation fault does not happen:</p> <p>## TEST 1: Ensure success for correct username and password:</p> <ol style="list-style-type: none"> 1. Encode username and password to base 64 and set as environment variables USER and PSWD example: export USER="YWRtaW4=" ; export PSWD="U29udXNAMTlz" 2. Run pamValidator and verify it returns success. <p>##TEST 2: Authentication Failure for username and incorrect password:</p> <ol style="list-style-type: none"> 1. Encode the correct username and incorrect password to base64 and export as env variables USER and PSWD. 2. Run pamValidator and verify it returns "Authentication Failure". 3. Run pamValidator again multiple times (atleast 3 more) and verify the authentication failure and no segmentation fault. 4. Wait for 30 seconds and try TEST#1 with the correct password and verify it succeeds. 	<p>The code is modified so the pam_response struct properly in between calls to the conversation function.</p> <p>Workaround: None.</p>

SBX-105152	2	<p>While expecting a 200 OK to ingress endpoint, the SBC was sending BYE to the Egress endpoint.</p> <p>Impact: If the Media inactivity/activity monitoring is enabled on media leg, and if media is not received at all on that leg in initial 10 seconds, then the NP sends a media inactive notification to up layer.</p> <p>But later even if media starts coming to the SBC on that leg, the NP is not sending media is active now in notification to Application layer.</p> <p>As a result, the application layers were closing the call based on the action configured for media inactivity (peerAbsenceTrapAndDisconnect).</p> <p>Root Cause: In the call, if the media is preceded by no media for few seconds. The call can be terminated if the peerAbsenceAction is peerAbsenceTrapAndDisconnect.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Configure the peerAbsenceAction as peerAbsenceTrapAndDisconnect in the PSP. 2. After the call is established, do not stream the media for initial few seconds (this duration should be less than inactivityTimeout), start it only after few seconds. > set system media mediaPeerInactivity inactivityTimeout 20 > set profiles media packetServiceProfile DEFAULT peerAbsenceAction peerAbsenceTrapAndDisconnect 3. In this case, the pause should be less than 20 seconds. 4. With the fix, the calls will not be terminated. 	<p>The code is modified to fix the Inactive to Active detection functionality and address the issue.</p> <p>Workaround: If media is expected to come late, they should not configure the peerAbsenceAction action in PSP to peerAbsenceTrapAndDisconnect. The action can be selected as the peerAbsenceTrap or none.</p>
SBX-105679	2	<p>The ASAN leaksanitizer detected errors in the CPX and SCM.</p> <p>Impact: The SCM and CPX processes have a small memory leak while changing the IPsec and IP interface group configuration.</p> <p>Root Cause: While processing configuration related commands, the SBC was reading information from CDB into temporary memory blocks but failed to release the memory at the end of the configuration action, resulting in a small memory leak.</p> <p>Steps to Replicate: Make the configuration changes to IPsec and IP interface group.</p>	<p>The code is modified to ensure the memory is correctly freed up to avoid the leak.</p> <p>Workaround: None.</p>
SBX-105637	2	<p>The AddressSanitizer: detected heap-buffer-overflow on address 0x61b000013128 at pc 0x55e3eea0419e bp 0x7ff7a52768a0 sp 0x7ff7a5276898 in ScmProcess_0.</p> <p>Impact: There was invalid memory access when the SBC receives the 500 Internal Error to REGISTER.</p> <p>Root Cause: The root cause of this issue is accessing invalid memory while accessing callData, trying to read off the data from the end of memory block. It can potentially cause the coredumps if the memory block is at the edge of the accessible memory region.</p> <p>Steps to Replicate: Testcase: Description:</p> <ol style="list-style-type: none"> 1. Clean up SAs if unexpected response received <p>Procedure:</p> <ol style="list-style-type: none"> 1. Monitor the exchange. 2. Send an initial reg message. 3. Receive the 401 challenge. 4. S-CSCF responses with 500. <p>Expected Results:</p> <ol style="list-style-type: none"> 1. Verify IPsec SA's deleted (ip xfrm state). 	<p>The code is modified to access the respective application data (It can be call data, registration data or subscription data) based on type of SIP message.</p> <p>Workaround: None.</p>

SBX-104990	2	<p>In a secure call, the SBC does not increment the port number in R-URI after processing Refer.</p> <p>Impact: In a secure call with TLS configured, if the call is REFERed with a REFER-TO header containing a FQDN and port number, the SBC sends out a new INVITE to specified FQDN and port number. If that INVITE fails and the SBC then sends a subsequent INVITE on the next route, it does not correctly increment the RURI port number for the TLS.</p> <p>Root Cause: The SBC code does not take into account that a re-route after a REFER-TO with FQDN and port number target needs to increment the port number for TLS, if the target after the re-route is different to the original target specified in the REFER-TO.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. With the recommended SBC configuration for MS Teams with TLS enabled between the SBC and MS Teams, establish a call from the PSTN to MS Teams 2. Send a REFER from MS Teams that includes following header: REFER-TO: <sip:sip2.pstnhub.microsoft.com:xxxx;transport=tls> 3. Based on the REFER, the SBC should route the call and send an INVITE to sip2.pstnhub.microsoft.com using port xxxx 4. Reject this INVITE from MS Teams with a 503. 5. The SBC should then send an INVITE out using the second route to sip3.pstnhub.microsoft.com using port xxxx. 6. Complete the call signaling and verify referred call is established correctly. 	<p>The code is modified to increment the RURI port number for TLS if performing a re-route to a target FQDN and port number that is different to the original target specified in the REFER-TO.</p> <p>Workaround: None.</p>
SBX-104671	2	<p>The CpxAppProc leak for the MRFP call.</p> <p>Impact: During the SBC startup, the CPX process has a small memory leak.</p> <p>Root Cause: During the SBC startup processing, the CPX process reads various CDB configuration and performs DB schema upgrade validation logic. As part of this processing, it was creating temporary internal memory blocks but not releasing them at the end of initialization.</p> <p>Steps to Replicate: Restart the SBC after it has been configured.</p>	<p>The code is modified to correctly release the temporary memory blocks used for initialization processing.</p> <p>Workaround: None.</p>
SBX-105339	2	<p>The LeakSanitizer: detected memory leaks on the active OAM.</p> <p>Impact: The CPX process was leaking small amounts of memory while processing BFD configuration changes.</p> <p>Root Cause: The CPX process was allocating memory in order to interact with the CDB while process BFD configuration changes. But it was not freeing up the memory at the end of the configuration action, resulting in a memory leak.</p> <p>Steps to Replicate: Make configuration changes to the BFD profile.</p>	<p>The code is modified to correctly free the memory required to process the BFD configuration changes and avoid the memory leak.</p> <p>Workaround: None.</p>
SBX-105262	2	<p>The SBC is sending an unexpected re-INVITE to the Egress side in the SRTP early media scenario.</p> <p>Impact: The SBC is sending an unexpected re-INVITE to the Egress side in the SRTP early media scenario.</p> <p>Root Cause: This issue is caused as a side-effect of SBX-103807.</p> <p>Steps to Replicate: Run the following call flow:</p> <ol style="list-style-type: none"> 1. The UAC sends an INVITE with 100 rel required. 2. The UAS sends 18x with SDP and SRTP(SHA-1-32). 3. Ingress sends PRACK with SDP and SRTP(SHA-1-32). 	<p>Copy an active security PSP only if audio stream is present to address the issue.</p> <p>Workaround: None.</p>
SBX-104537	2	<p>Observed SIPFE MAJOR logs on the n1-standard-4 SBC_HA_HFE_SPLIT instance.</p> <p>Impact: The log was printed as major, when stale calls are present and whenever we start cleaning the stale calls, then the log is seen.</p> <p>Root Cause: This log is expected when clearing any stale calls.</p> <p>Steps to Replicate: Run call load and if any stale calls are seen, then this log issue is seen (only when the Minor logging is enabled).</p>	<p>The code is modified to address the issue.</p> <p>Workaround: None.</p>

SBX-103851	2	<p>Observed an error "runtime error: index -20 out of bounds for type 'uint8_t [16]'" on UXPAD when running T140 call with out-of-order of T.140 packets.</p> <p>Impact: If the T140 packet contains any T140block that has more than 16 characters, then a debug buffer to display the ASCII characters of T140 packet may write beyond its size.</p> <p>Root Cause: The debug buffer to display ASCII characters of T140 packet is 16 bytes in size. It saves the 16 bytes of last T140block. However, the T140 packet SBC accepts allows a maximum T140block size of 36 characters. The code did not properly limit the size of buffer to copy to 16 characters.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Make a SIPP call (AMRNB=>G711) with t140=>baudot enabled. 2. Send PCAP with the T140 from AMRNB termination having T140block in packet exceeding 16 bytes. 3. There may not be any observable effect, but the debug buffer that displays ASCII characters writes beyond its bounds and corrupts some other fields in the structure. 	<p>The code is modified to limit the size of ASCII buffer to copy to 16 characters.</p> <p>Workaround: None.</p>
SBX-105270	2	<p>There was a CpxAppProc leak for MRFP calls.</p> <p>Impact: The small memory leak occurs on the SBC/MRFP nodes when action/status/stats under the node branch is accessed through the OAM node CLI or EMS.</p> <p>Root Cause: Resource references are cleared mistakenly after serving the request from the OAM node.</p> <p>Steps to Replicate: Execute a node branch command repeatedly and monitor the CPX process size on the target node.</p>	<p>The resource references are cleared only when the system is shutting down. The resources are now getting reused by subsequent requests to address the issue.</p> <p>Workaround: Avoid using the node branch commands in automated /periodic operations. Manual use should work as the leak is small.</p>
SBX-105114	2	<p>The usage of a kill command output in the active and standby CE_node logs.</p> <p>Impact: The kill command usage gets printed in the CE_Node log file of the managed nodes.</p> <p>Root Cause: No glusterfs process is present when the kill command is executed by the glusterSetup.sh script called by sbxConfigUpdater.sh.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Reboot the Standby OAM. 2. Ensure that kill usage does not appear in the CE_Node log file of the managed nodes. 	<p>Redirect the kill command error output to /dev/null to address the issue.</p> <p>Workaround: None.</p>
SBX-105395	2	<p>There are coverity issues in the OAMNODE.</p> <p>Impact: When processing a show list command under the node branch from the OAM node, if the target node fails to read the command path in the request, the code will access memory immediately after freeing it. While in most cases this should not cause issues, accessing memory after it is freed is not good behaviour and could result in unexpected behaviour, potentially causing core dumps.</p> <p>Root Cause: The error handling flow in the code is incorrect. The handling code hits some code for the success flow.</p> <p>Steps to Replicate: Enter the CLI commands like "show table/status node SSBC-1 <path to some list> <partial key>" and hit "tab" for auto completion on a system with an unreliable HA network. Repeat until the target node showing this error in its DBG log:</p> <pre>CPX ConfdProxy::worker: could not deserialize parameter for PROXY_FIND_NEXT_REQ</pre>	<p>The code is modified to address the issue.</p> <p>Workaround: None.</p>
SBX-103103	2	<p>The BFD packet forwarded by the router is not received by the MRFP.</p> <p>Impact: Once the BFD session is established on a port (either primary or secondary), an immediate port switch over is followed due to the BFD packets dropped at NP for a brief amount of time (~1 second). This eventually triggers a node switchover.</p> <p>Root Cause: A race condition in ACL lookup is the cause of this issue.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Ensure the BFD session is up on a port (either primary or secondary). 2. Initiate a port switchover. 3. Monitor if the BFD session comes back on the switched-over port and an immediate port switch over is not followed. 	<p>The code is modified to address the issue.</p> <p>Workaround: A user ACL can be added as a workaround.</p>

SBX-103984	2	<p>For an existing call trace, when the call trace feature is disabled on the MRFP, the MRFP should reject the trace request (in MODIFY with CALLTRACE/TRACEACTIVITYREQUEST=ON) from the C3 by sending a NOTIFY to the C3 with RES=FAILURE.</p> <p>Impact: When the call trace feature is disabled in the MRFP (using CLI command: set global callTrace state disabled), and the C3 tries to enable call trace using the MODIFY command with CALLTRACE /TRACEACTIVITYREQUEST=ON, then the MRFP should reject the trace request from C3 by sending NOTIFY to C3 with {CALLTRACE/TRACACT {Stream=1,RES=FAILURE}}.</p> <p>However, the MRFP (9.1 R0) is not sending any NOTIFY message for the Call Trace request received in the MODIFY command.</p> <p>Root Cause: The issue was seen only when the MODIFY command does not have any SDP.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Disable the call trace feature, using the CLI command: set global callTrace state disabled. 2. Establish a MRFP call for the SBC from C3 by sending two ADD termination commands. 3. Send MOD termination command from C3 with: CALLTRACE /TRACEACTIVITYREQUEST=ON.termId ip/1/intf/3523215361{ Media { Stream = 1 { LocalControl { Mode=SendReceive, CALLTRACE/TRACEACTIVITYREQUEST=ON } } },Events = 1 { NT /NETFAIL , ADID/IPSTOP { DT=50 } , HANGTERM/THB { TIMERX=1800 } },Signals { } } 4. Since call trace is disabled at the SBC, verify that the SBC sends a NOTIFY with res=Failure after sending a MODIFY reply. 	<p>The code is modified to send a NOTIFY for TRACEACTIVITYREQUEST when the MODIFY command does not result into any change in media parameters. The MRFP now sends a NOTIFY with RES=SUCCESS /FAILURE (based on call trace configuration) after sending reply for the MODIFY command.</p> <p>Workaround: None.</p>
SBX-104325	2	<p>A SCM core dump was observed when multiple gateway TGs are created in a GW-GW call on a HA setup.</p> <p>Impact: On a HA setup, the Standby box SCM Process dumps core when the Standby starts to sync from Active post a switch over and as a result, the switchover occurs after a scenario where Gateway TG's are created and then a GW TG with a lower index (created earlier) is deleted.</p> <p>Example:</p> <ol style="list-style-type: none"> 1. Create GWTG1, GWTG2 and GWTG3. 2. Delete GWTG2 3. Switch over. <p>Root Cause: The coredump is caused due to difference in indexing the gateway TG's in active and standby boxes.</p> <p>The GW TG indexes were out of sync between the Active and Standby. Active SBC had holes in the indices of the GW TGs after deletion. The Standby SBC does not have holes for the GW TGs that are present post deletion and occupy a different index when compared to active.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. The HA setup with GW-GW configurations. 2. Create 3 Gateway Trunkgroups: GWTG1, GWTG2 and GWTG3. 3. Delete the GWTG2. 4. Perform a switchover. 	<p>The code is modified to ensure that the GWTG indexes on the Active and Standby are in sync.</p> <p>Workaround: None,</p>
SBX-104693	2	<p>When multiple codecs are received in descriptor, the call is getting rejected if license of the first preferred codec is not present in the license bundle.</p> <p>Impact: When the MRFP receives an ADD termination command with a list of codecs in the audio stream from the MGC, it rejects call if the license of first preferred codec is not present in the license bundle, even though MRFP could have succeeded the call with other codecs on the list.</p> <p>Root Cause: The MRFP's codec filtering function chooses the first allowed codec from the list and send it to media plane without check the codec license, so that the media plane returns error in case of codec license validation failure. The call failed as a result.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Ensure the MRFP is up and running 2. Generate a license xml with ONE unit of MRFP-RTU, MRFP-DSP-RTU MRFP-DSP-AMR and ZERO units MRFP-DSP-AMRWB. Install the license bundle b1 in MRFP. 3. Place a call from endpoint with AMRWB and AMRNB in SDP in same m line 4. Validate the call should be succeeded with AMRNB codec being used. 	<p>The code is modified so it always chooses a codec with a valid license to the media plane.</p> <p>Workaround: None.</p>

SBX-103281	2	<p>The route data was lost after offline PM upgrade from 625R0 to 823A13.</p> <p>Impact: The special character data (e.g. ?,*,#,\$) is not getting migrated from the Oracle version to postgres version.</p> <p>Root Cause: Special characters were causing issues with Postgres data loading.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Create 100+ routes that have special characters in the DN field (Domain Name) in the SBC 6.2.5.R0. 2. Upgrade to the SBC through the PM offline upgrade to 8.0 or above. 3. After upgrade all route data was lost while other data are unaltered. 	<p>The code is modified to address the issue.</p> <p>Workaround: No workaround.</p>
SBX-102445	2	<p>The media port range threshold alarm is not triggered after a switchover.</p> <p>Impact: The sonusMrfpRealmMediaPortRangeThresholdExceededNotification alarm does not get raised on a newly active box following a switchover, even if the conditions for the alarm are met.</p> <p>Root Cause: The realm status is not getting properly mirrored to the standby box.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Run calls such that 90% of the ports are used and alarm is triggered. 2. Trigger a switchover from the active MRFP. 3. Alarm should be triggered in new active MRFP. 	<ol style="list-style-type: none"> 1. The code is modified to check the UDP port usage, and to raise the alarm if appropriate. 2. The code is modified to mirror realm status so that, it can be used to raise alarm when transitioning to active or when new call on SBY arrives. <p>Workaround: None.</p>
SBX-104963	2	<p>The createConfigDrive.py --file option throws an error when executing.</p> <p>Impact: In the case of KVM/VMware deployments with qcow2/OVA/vmdk, the createConfigDrive.py script that creates the config-drive is not working with the '--file' option.</p> <p>Root Cause: There was an error in handling the '--file' option</p> <p>Steps to Replicate: Generate a configuration drive with the '--file' option.</p>	<p>The code is modified to address the issue.</p> <p>Workaround: Use the '--cli' option to generate a configuration drive.</p>
SBX-103761	2	<p>The createConfigDrive.py --cli throws an IndentationError.</p> <p>Impact: When deploying on the KVM/VMware using qcow2/OVA/VMDK and generating config-drive using createConfigDrive.py with '--cli' option, the script returns with an error.</p> <p>Root Cause: There was an indentation issue in the Python code.</p> <p>Steps to Replicate: Generate a config-drive with the '--cli' option and verified the generated config-drive.</p>	<p>The code is modified to address the issue.</p> <p>Workaround: None.</p>
SBX-103682	2	<p>The LeakSanitizer: detected memory leaks at confd_malloc</p> <p>Impact: The ASAN detected memory leaks while processing the E164 profile configuration changes.</p> <p>Root Cause: The code was allocating memory while processing E164Profile configuration changes but not releasing the memory at the end of the configuration action, resulting in a small memory leak.</p> <p>Steps to Replicate: Create and modify the E164Profile configuration.</p>	<p>The code is modified to release the internal memory block at the end of the configuration action.</p> <p>Workaround: None.</p>
SBX-103821	2	<p>The AddressSanitizer: detected heap-use-after-free on address 0x6180000ed180 at pc 0x5619a742719d bp 0x7f3b04960310 sp 0x7f3b04960308.</p> <p>Impact: While the MRFP node is shutting down, it can access memory after it has been freed, this could result in unexpected behaviour and in the worst case a core dump. But would have limited impact as it only occurs when shutting down.</p> <p>Root Cause: During the sbxstop/sbxrestart or switchover because of race-condition, when the SBC is in deactivation the oamNodeRegisterRetry can access already deallocated resource leading to a core dump.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Setup a build with HA MRFP using OAM. 2. Do the sbxrestart/switchover of an active instance. 	<p>The code is modified to handle this race condition.</p> <p>Workaround: None.</p>

SBX-105312	2	<p>The trunkgroups are not displayed while assigning the SMM profile to TG on the EMA.</p> <p>Impact: The trunkgroups are not displayed while assigning the SMM profile to TG on the EMA.</p> <p>Root Cause: The dropdown height is limited, and as the result the last entry is not visible.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Create more than one TG. 2. Create a SMM profile on EMA. 3. Click on Assign SIP Adaptor Profile. 4. Under 'Assign Message Manipulation Profile to TGs', check for Input Adaptor and Output Adaptor. <p>All the options should be available after performing the test steps.</p>	<p>The code is modified to show all the options to select.</p> <p>Workaround: None.</p>
SBX-98283	2	<p>The SBC is unable to find the TG for in-dialog NOTIFY message when received from different IPs and the dialogTransparency flag.</p> <p>Impact: When the indialog NOTIFY comes from different source IP, then the SBC is dropping the NOTIFY.</p> <p>Root Cause: The existing design for OOD does not support the indialog NOTIFYs when it comes from a different source.</p> <p>Steps to Replicate:</p> <ol style="list-style-type: none"> 1. Initiate a SUBSCRIBE dialog. 2. Send an indialog NOTIFY from a different source. 	<p>The code is modified to accept the indialog requests when the source is different.</p> <p>Workaround: None.</p>

Known Issues

Known Issues in Release 09.02.01R001 to 09.02.03R002

The following known issues exist in these releases.

Table 36: Known Issues

Issue ID	Sev	Problem Description	Impact/Workaround
SBX-113524	3	The SBC is not forwarding the received 200 OK message and fails the call.	Impact Statement: When using a GW-GW for direct media calls between signaling only SBCs, the 200 OK message is not being passed along and the call fails. Workaround: Disable the following control: set global signaling sigOnlyMode sigOnlyModeValue
SBX-99023	2	The CE_2N_Comp_Sm Process is dumping core.	Impact Statement: The SBC coredump directory shows a SM Process coredump, but services are not affected and they keep running. Workaround: You can remove the coredump from the SBC coredump directory. No other action is required since these types of cores do not cause the SBC services to go down.
SBX-111973	2	On a LRBT with lateCrankback, the call is torn down upon Hold.	Impact Statement: If a call route advances after lateCrankback and LRBT is enabled, the call is torn down upon egress connect (200 OK) after hold/resume request. Workaround: None.
SBX-111461	2	The SBC plays a LRBT using PCMU when the negotiation happened with AMRWB in the LM call.	Impact Statement: The tone played is with an old negotiated codec instead of codec selected by ingress endpoint in Prack. Issue is seen only when multiple codecs are sent in late media offer in 180 with sendSbcSupportedCodecsForLateMediaReinvite flag enabled. Workaround: None.
SBX-111460	2	The SBC does not offer the 16K dtmf in a LM call when the SDP is present in 2xx.	Impact Statement: The SBC is not sending 16k 2833 Payload type in the initial offer towards the ingress when the SDP is present in 2xx during a Late media "convert" call. Workaround: None.
SBX-111351	2	The CHM process is coring after recovering from a splitbrain.	Impact Statement: While successfully recovering from a split-brain in a specific scenario, it is sometimes noticed that a CHM coredump is created. Workaround: None. The system does recover correctly without any manual intervention.
SBX-110134	3	For a T140 call - one-way audio observed if NAPT enabled on egress.	Impact Statement: If NAPT for media is enabled, the dest media port (NAPT media) stays set to 0. Workaround: None.
SBX-111034	2	Unable to login as an admin user using keys after cleanDB.	Impact Statement: User is unable to log in to the Confd CLI using 'admin' user private SSH key after running clearDBs.sh script. Workaround: Do not run the clearDBs.sh script.
SBX-103724	2	The RECORDING CDR does not have Media data and stats field in a REFER scenario case.	Impact Statement: The Media Stats are not present in the Recording CDR when we record the C Leg. Workaround: None.
SBX-101226	2	The OAM should not configure the same IP and port for two different VMGs.	Impact Statement: The mis-configuration that was using the same IpVar for two different VMGs does not throw an error in the OAM CLI. Workaround: None. The mis-configuration recovery requires a reboot.

SBX-105824	2	The glusterfs had a core dump on the Active OAM for the T-SBC post upgrade.	<p>Impact Statement: The OAM node is using the third party package glusterfs and we are using version 3.8.8-1. During a scenario where both OAM nodes are starting up, we may encounter an intermittent core from the glusterfs process. The OAM functionality is not impacted by this core and it does not impact the shared directory that is mounted by the gluster process. The call processing nodes that are managed by the OAM are not impacted and their configurations are unaffected.</p> <p>Workaround: No workaround. The core maybe be produced during simultaneous reboot of OAM nodes but the functionality of the OAM nodes is not impacted.</p>
SBX-105921	3	Reduced the configuration limits that need to be used for certain fields.	<p>Impact Statement: The CDB schema supports larger strings than the SBC application code can currently support for the following configuration objects. This issue is due to the SBC application code not allowing for one additional character to include the string null terminator if the configuration in CDB actually contains the maximum number of characters.</p> <p>The maximum size the application can be supported, even though CDB schema would allow for one character more to be entered.</p> <p>addressContext/diamNode/realmRoute/realm - 128 characters</p> <p>global/genericCodec/audioEntry/name - 49 characters</p> <p>system/policyServer/remoteServer/fqdn - 255 characters</p> <p>global/signaling/srvcc/stnSr - 30 characters</p> <p>global/signaling/srvcc/eStnSr - 30 characters</p> <p>global/signaling/srvcc/pstopSti - 30 characters</p> <p>profiles/services/testCallNumberProfile/testCallNumber/number - 23 characters</p> <p>In a future release, the application code will either be extended to allow for one more character or validated to put in place to restrict the character size to one less.</p> <p>Workaround: Use the reduced size for the fields as mentioned in the Impact Statement.</p>

Known Limitations

The following limitations exist in this release:

1. The Access Control List (ACL) is not installed to configure SNMP traps for accepting traffic. A dynamic ACL is added to configure SNMP traps. An ACL must be installed for SNMP traps for accepting traffic.
2. The physical NIC connectivity must be in active state at the hypervisor level before starting the SWe instance on the SBC SWe platforms. In case of SWe instance with SR-IOV interfaces, manual restart of the SWe instance is required if physical NIC connectivity goes down while the instance is in progress.
3. The Antitrombone feature is not supported on the D-SBC.
4. EMS identifies the nodes based on the VNFC-ID. While instantiating SBC/PSX cloud nodes, ensure that you use a unique VNFC-ID only. If you reuse an existing VNFC-ID, EMS treats this as a re-registration request and overwrites the existing data on the cloud node.
5. While configuring the SBC SWe Cloud instances, the CLIs commits successfully even if any metaVariable provided is incorrect. The SBC SWe Cloud instance cannot validate the CLIs, as the CDB configuration file is stored in the OAM Node and is shared among all the other SBC SWe Cloud instances in the cluster.
6. Editing IP Interface is not reflected in the if configuration (ifConfig). This behavior is observed only on the S-SBC when action is set to "dryup" mode on the IP Interface. The IP address changes are not updated in the kernel and will not be displayed when ifconfig linux command is executed. In case of S-SBC, if the ipInterface configuration needs to be modified and if the action is set to "dryup" in ipInterface configuration, it must be set to "force" before disabling the ipInterface and making any changes.
7. A LSWU on an SBC 7000 should only be performed when the total number of active calls on the system is below 18,000. If the criteria is not met, a double failure during the upgrade may occur thereby losing all active calls. If such a failure occurs, both active and standby SBC services will go down. Contact Ribbon Support immediately.



The VLAN tagged SRIOV packet interfaces are unable to ping endpoint Gateway IPs in the VMware platform because of an issue with VMware.

Performing a Heat Stack Update when userdata is Updated with SSH Keys

When upgrading SBC SWe cloud instances to release 9.2.1, you must update your Heat template userdata section to include mandatory SSH key information. An issue in OpenStack requires that you use the stack-update process rather than re-launch after updating the template, which leads to a new UUID for the instance. As a result, you must regenerate and apply new license bundles to the upgraded instances during the upgrade.

Refer to [Upgrading SBC SWe N:1 HA Nodes on OpenStack using Heat Templates](#) for the relevant procedure.

MOP to increase vCPUs Prior to Upgrading SBC SWe on VMware or KVM Hypervisor (9.2.3R2)

The SBC Core release 9.2 includes a feature that extends the use of hyper-threading to SBC SWe when it is installed on either the VMware or KVM Hypervisor platform. To take advantage of the performance improvements provided by hyper-threading, you must increase (double) the number of vCPUs configured in the VM prior to software upgrade. When upgrading, modify the upgrade process to incorporate the following additional steps: If upgrading SBC SWe KVM Hypervisor or VMware from pre-07.01.00R000 release to 07.01.00R000 or higher:

If upgrading vCPUs from less than 10 to 10 or more, use the procedure below



Note

This procedure requires shutdown of both SBCs, Ribbon recommends that you perform this procedure during a maintenance window.

1. On the VMware platform only, before beginning the upgrade, disable the CPU reservation check inside the guest by renaming the file: `vmware-toolbox-cmd`. For example, issue the following command:

```
mv /usr/bin/vmware-toolbox-cmd /usr/bin/vmware-toolbox-cmd.bak
```
2. Stop standby SBC by issuing "sbxstop". Wait until the SBC processes stop.
3. Shutdown the standby SBC by issuing "poweroff" command.
4. Increase vCPU count on the powered off standby SBC.
5. Stop an active SBC by issuing "sbxstop". Wait until the SBC processes stop. Note that this is service affecting since both SBCs are stopped now.
6. Shutdown the active SBC by issuing "poweroff" command.
7. Increase vCPU count on the powered off active SBC.
8. Power on the active SBC. Wait for the SBC to fully start as active.
9. Power on the standby SBC. The SBCs should start as standby and get in sync with an active SBC.

For any other CPU upgrade combination, use the procedure below:

1. On the VMware platform only, before beginning the upgrade, disable the CPU reservation check inside the guest by renaming the file: `vmware-toolbox-cmd`. For example, issue the following command:

```
mv /usr/bin/vmware-toolbox-cmd /usr/bin/vmware-toolbox-cmd.bak
```
2. For either platform, before starting the upgrade procedure, shutdown the standby VM instance and double the number of vCPUs specified for the VM. Refer to the procedure you used creating a VM that is appropriate for your specific deployment:
(VMware) [Creating a New SBC SWe VM Instance with VMXNET3](#) (step 7)
(VMware) [Creating a New SBC SWe VM Instance with Direct IO Passthru](#) (step 7)
(KVM) [Creating a New SBC SWe Instance on KVM Hypervisor](#) (step 6)
(KVM) [Creating a New SBC SWe Instance with PCI Pass-Through Device](#) (step 6)
3. After increasing the number of vCPUs, power on the standby instance with the increased number of vCPUs in the same release.
4. Switch over the active instance, so that the standby becomes the new active instance. Don't wait for the standby instance to come up. It may not come up due to the resource mismatch.
5. Shutdown the new standby instance and double its number of vCPUs, as done earlier in step 2.
6. Bring up (power on) the new standby and wait for the active and standby nodes to sync.