

# Juniper Apstra Version 4.1.1 Release Notes

## Known Critical Apstra Issues

### Changes in Freeform Configuration Templates May Not Cause Device Configuration Deployments (AOS-35791)

In some instances, when a user makes changes to a configuration template in an Apstra Freeform blueprint, and there are no device model changes, even though a new configuration in Apstra is shown and the user commits a blueprint change, Apstra will not deploy the configuration change to devices.

#### Workaround

Please contact Juniper Support for a hotfix release.

## New Features

### Support for New Arista Hardware (RFE-2485)

**Feature Category:** Device Profiles

Apstra now supports the following Arista hardware: DCS-7050CX3M-32S and DCS-7280CR3-32P4

---

### Support for Cisco 93180YC-FX3 switch (RFE-2036)

**Feature Category:** Device Profiles

Support for Cisco 93180YC-FX3 has been added in Apstra 4.1.1.

---

### Support for Cisco 92348GC-X switch (RFE-2497)

**Feature Category:** Device Profiles

Support for Cisco 92348GC-X has been added in Apstra 4.1.1.

---

## **Support for Cisco 36180YC-R switch (RFE-2190)**

**Feature Category:** Device Profiles

Support for Cisco 36180YC-R has been added in Apstra 4.1.1.

---

## **Prevent predefined Device Profiles from being edited (RFE-2474)**

**Feature Category:** Device Profiles

Device Profiles which are predefined by Apstra are no longer editable in the UI to prevent issues during upgrade conflicting with user changes. The Device Profile must be cloned to insert user modifications.

---

## **Support for New Device Operating Systems (RFE-2381)**

**Feature Category:** Device Operating Systems

Apstra 4.1.1 now supports the following switch vendor software:  
Juniper Junos: 20.4R3-S3 and 21.4R2  
Juniper Junos Evolved: 20.4R3-S3-EVO and 21.4R2-EVO  
Arista EOS: 4.27.4M

---

## **Report a failure when device NOS version is not changed post upgrade (RFE-2165)**

**Feature Category:** Device Operating Systems

The post upgrade logic has been enhanced to verify that the device NOS version post upgrade is different to the one the upgrade started with.

---

## **Display image size for NOS images. (RFE-2414)**

**Feature Category:** Device Operating Systems

The device NOS image size is now displayed under Devices - OS Images page to provide better information to the user on the image size before deciding to upgrade devices.

---

## **Check memory and disk space before starting Apstra upgrade (RFE-1391)**

**Feature Category:** Device Operating Systems

Verify enough disk space and memory is available before starting the Apstra upgrade to prevent resource issues during upgrade.

---

## **Supported upgrade paths to 4.1.1 (RFE-2303)**

**Feature Category:** Design, Build, Operate

This release introduces upgrade paths from previous releases. See the user guide for the documented list of supported upgrade paths and upgrade methods.

---

## **Improve logging of platform level actions (RFE-2471)**

**Feature Category:** Design, Build, Operate

If a user modifies the allowed or banned IP list under the security platform settings, such as deleting a banned entry, it will now be logged in the platform event logs.

---

## **Freeform Reference Design: build any design (RFE-2609)**

**Feature Category:** Design, Build, Operate

Apstra has expanded its automation capability to any network design with the "Freeform Reference Design", allowing you to build the design you want and how you want it. You can leverage any feature, protocol, or architecture that fits your deployment scenario. Freeform presents you with an interactive canvas to visually design or model any arbitrary network topology. The configuration is administered via Configuration Templates that grant you complete control over the configuration on the devices. You can still leverage the same simple and powerful lifecycle management features from device operating system upgrades, simple device deployments, pre-deployment data center modeling, device telemetry, analytics dashboards, to even Apstra's powerful Intent-Based Analytics and Intent-Time Voyager.

---

## **Apstra upgrade to verify all cluster and agent credentials (RFE-2336)**

**Feature Category:** Design, Build, Operate

To avoid upgrade issues with incorrect credentials for cluster nodes and system agents, Apstra now verifies all configured credentials before starting the upgrade. This includes verifying all cluster worker node credentials, system agent credentials plus privileges for all devices, and device system agent access for all devices from the new controller plus worker nodes.

In addition, there are new command-line options to either "skip" or "dry-run" these checks  
--skip-connectivity-validation - Do not validate cluster and system agent credentials  
--dry-run-connectivity-validation - Only run and report errors from cluster and system agent credentials then exit without any other action

---

### **Apstra upgrade improvements for large deployments. (RFE-2451)**

**Feature Category:** Design, Build, Operate

Several improvements and optimizations in the Apstra upgrade process to better support large scale deployments.

---

### **Support of Access Switches in VMWare/NSX-T IBA probes (RFE-2263)**

**Feature Category:** Telemetry and Analytics

Support of Access Switches in VMWare/NSX-T IBA probes.

## **Changed Features**

### **UI Improvements for Device States (RFE-1966)**

**Feature Category:** Design, Build, Operate

Improved and enhanced the UI to better define device states to the user. Added tooltips for Ready, Deploy, and Drain state and renamed "Discovery (2) Config" to Ready in Dashboard.

---

### **Relaxed Route Target Validation (RFE-2469)**

**Feature Category:** Design, Build, Operate

Relaxed a validation that prevented user-defined VRF leaking of internally Apstra-managed route targets. Apstra users can now import and export Apstra-managed route targets. A tooltip was added to advise users to ensure importing and exporting Apstra-managed route targets do not have overlapping IPs and have proper IP management.

---

### **General Availability of ESI at Access layer - Phase 1 (RFE-2337)**

**Feature Category:** Design, Build, Operate

ESI at the access layer which was introduced as a Tech Preview in 4.1.0 is now delivered as General Availability feature. This feature allows attachment of dual-homed Servers/Generic Systems to the Access layer using ESI-LAG and is supported on 3-Stage, 5-Stage, and Collapsed Fabric Blueprints.

---

### **Changed Juniper DHCP reference design (RFE-2537)**

**Feature Category:** Design, Build, Operate

Changed the reference design for Juniper-based leafs from a DHCP stateful model to a stateless model. Apstra will now automatically render the DHCP `forward-only` command for non-default VRFs.

---

### **Change RSTP configuration of Junos leaf devices (RFE-2333)**

**Feature Category:** Design, Build, Operate

Made configuration change to default design for Junos leaf devices to render edge command on server-facing interfaces. This change protects users from introducing a data plane loop in the data center if someone inadvertently connected two leaf devices together.

---

### **Enable historical data retention for IBA probe "Device System Health" (RFE-2259)**

**Feature Category:** Telemetry and Analytics

IBA probe "Device System Health" has been augmented with 7 days of historical data retention for the following metrics:

- System cpu utilization
- System memory utilization

- Disk utilization
- Disk utilization partition

## **Tech Preview Features**

*Tech Previews give you the ability to test functionality and provide feedback during the development process of innovations that are not final production features. The goal of a Tech Preview is for the feature to gain wider exposure and potential full support in a future release. Customers are encouraged to provide feedback and functionality suggestions for a Technology Preview feature before it becomes fully supported.*

*Tech Previews may not be functionally complete, may have functional alterations in future releases, or may get dropped under changing markets or unexpected conditions, at Juniper's sole discretion. Juniper recommends that you use Tech Preview features in non-production environments only.*

*Juniper considers feedback to add and improve future iterations of the general availability of the innovations. Your feedback does not assert any intellectual property claim, and Juniper may implement your feedback without violating your or any other party's rights.*

*These features are "as is" and voluntary use. Support Services will attempt to resolve any issues that customers experience when using these features and create bug reports on behalf of support cases. However, Juniper may not provide comprehensive support services to Tech Preview features. Certain features may have reduced or modified security, accessibility, availability, and reliability standards relative to General Availability software. Tech Preview is not supported under existing service agreements, SLAs, or support service.*

*For additional details, please contact Juniper Support or your local account team.*

### **Tech Preview for 802.1x for Junos (RFE-2602)**

**Feature Category:** Design, Build, Operate

This is a tech preview feature, and we would love your feedback! You can manage 802.1X configuration on network devices with 802.1X server port authentication and a collection of interface policy settings for Junos devices with 802.1x. Juniper Evolved does not at this time

support this feature.

## Fixed Apstra General Issues

### "add Pod" fails with "Shorter than minimum length 1" (AOS-32030)

Apstra does not support PODs with zero rack types. However, it is possible to add a Pod and then delete all rack types within it. This results in an error when trying to add an embedded POD at a later point when this POD has no rack types. The following error is seen:

```
{
  "_schema": [
    "rack_based_templates:2:rack_type_counts: Shorter than minimum length 1.",
    "rack_based_templates:2:rack_types: Shorter than minimum length 1."
  ]
}
```

---

### Agent 'check-job' causes Operation Mode Status to TELEMETRY ONLY on NX-OS device with scp disabled (AOS-30958)

Agent 'check-job' causes Operation Mode Status to TELEMETRY ONLY on NX-OS device when scp is disabled in pristine config.

---

### Applied Connectivity Templates are not showing on bonded server interfaces (AOS-30153)

The applied Connectivity Templates (CT) are not displayed on the UI neighbor view for bonded server interfaces if this CT was assigned to the interface before updating the LAG Mode or forming a LAG.

---

### Apstra Cluster upgrade fails if ssh credentials are not identical on all of the VMs (AOS-27351)

When using different SSH credentials on the VMs in an Apstra Cluster, an Apstra 'distance' upgrade will fail with an authentication error.

### **Apstra controller upgrade fails with error "IOError: [Errno 28] No space left on device" (AOS-29175)**

During Apstra controller upgrade, a tarball of the metric db files is created on the source VM controller and this tar file is copied to the target VM controller. If the disk space on src VM is not sufficient, the upgrade procedure will not be able to create this tar ball and fails with error "IOError: [Errno 28] No space left on device"

#### **Resolution**

From 4.1.1 onwards, there will be improvement in the pre-upgrade validation workflow. A pre-upgrade validation will check the disk space available and stops the user from upgrade if enough disk space is not available.

---

### **Arista DCS-7280CR3-32P4 Device Profile Errors (AOS-31288)**

Errors in the Arista DCS-7280CR3-32P4 Device Profile will result in incorrect Interface Map port assignments for ports 17 and higher.

---

### **Arista Large Scale Deployments May Cause EOS eAPI Timeouts (AOS-30977)**

If an Apstra user is creating a "large scale" deployment (e.g. 100 VRF Routing Zones) with Apstra EOS devices, they may experience Arista EOS eAPI timeouts during deployment.

---

### **Arista Large Scale Deployments May Cause EOS eAPI Timeouts (AOS-30708)**

If an Apstra user is creating a "large scale" deployment (e.g. 100 VRF Routing Zones) with Apstra EOS devices, they may experience Arista EOS eAPI timeouts during deployment.

---

### **ARP Entries Show Last Modified Always as 52 Years for Junos (AOS-30728)**

ARP entries show last modified always as 52 years for Junos.

---

### **BGP flap probe should not be instantiated on Blueprints with Cisco and SONiC devices (AOS-31507)**

BGP Flapping probe is not supported on Cisco and SONiC devices. If it is instantiated on a



blueprint with Cisco or SONiC devices the probe will list all BGP sessions but the "Flap count" and "Flap count increment" will permanently show the value of "zero" which is misleading .

---

### **BuilderAgent crash seen when external generic system name is deleted (AOS-31115)**

Apstra BuilderAgent crash is seen when external generic system both "Name" and "Hostname" are deleted.

---

### **Cannot commit anything in BP with build error - Valid pristine config is not available (AOS-31021)**

Physical device was removed from the network while assigned to the BP then hit this condition.

---

### **Cannot Edit System in Managed Devices (AOS-30144)**

Editing a Managed Device by clicking the device IP to change its Device Profile results in a base64-encoded error.

---

### **Cisco NX-OS Agent May Hang Loading Paramiko Due to Lack of Entropy (AOS-28349)**

In very rare cases, the Apstra Cisco NX-OS device system agent is in the process of creating device driver objects which results in the import of the package Paramiko. In case the entropy on the system is low, Paramike import blocks until entropy is available. In this case, the Apstra agent couldn't recover within a minute and the agent failed.

---

### **Config deviation anomaly after NX-OS upgrade when 'feature scp-server' not configured (AOS-17492)**

When the NX-OS device does not have 'feature scp-server' configured in pristine, config deviation anomaly will be seen after the agent installation/upgrade.

---

### **Configlet Imports for Upgraded Apstra 3.2 and 3.3 Blueprints Will Display ID Instead of Name (AOS-32648)**

If a user has a blueprint with configlets that have been upgraded from Apstra 3.2 or 3.3, the

configlets will be imported into the blueprint catalog by ID (e.g. 3584b018-0f6b-11ed-861d-0242ac120002), not by label or hostname.

---

### **Configlet section condition string literals cannot contain a colon (":") character (AOS-31530)**

Configlet section condition string literals containing a colon (":") character will not be parsed correctly and will emit a "no available options" error.

#### **Resolution**

This bug is corrected in AOS 4.1.1.

---

### **Edits to AOS-provided predefined payloads are not recommended (AOS-30569)**

AOS has predefined payloads for device profiles, chassis profiles, linecard profiles, logical devices and many others. These json files are over-written with updated ones at the time of upgrades. Any change is therefore lost. To avoid said loss, if any edits are required, please first clone and then proceed to make required changes to this cloned payload.

---

### **Errors in Arista DCS-7280CR3-96 Device Profile (AOS-30915)**

In Apstra 4.1.0, there are errors in the Device Profile for the Arista DCS-7280CR3-96 which may cause deployment errors or eApi timeouts for users using the device.

---

### **Generic Systems' Deploy Mode are set to "Not assigned" when updating device assignments (AOS-29129)**

In Apstra 4.0.2, when updating the device assignments using the "change system IDs" dialog window, all Generic Systems are set to "Not assigned" after clicking "Update assignments".

---

### **IPv6 Extra Routes Defined in Routing Policy Not Rendered When IPv4-safi Isn't Enabled (AOS-31497)**

Any extra prefixes added in a routing policy are not rendered in the resultant device configuration of the BGP peering using that policy, unless the IPv4 SAFI is enabled in that peering. The peering details are typically defined through a Connectivity Template. If the IPv4 SAFI in the template isn't enabled, the extra prefixes are not added.

---

### **Max EVPN Routes Count Junos Config Applied to Family Inet Unicast Rather Than EVPN Signaling (AOS-31415)**

When setting a value for "Max EVPN Routes Count" in Virtual Network Policy in an Apstra Blueprint, the configuration deployed on the Junos devices sets a "prefix-limit" at group level for "family inet unicast" routes instead of "family evpn signaling".

---

### **NOS images are not transferred when upgrading (AOS-16886)**

When upgrading the Apstra VM using the 'distance' method, NOS images are not transferred. This applies to the URL based image as well.

---

### **On AOS upgrade changes made in DP and IM by user are not reflected and upgrade is successful with uncommitted changes and build errors (AOS-13125)**

During upgrade built in DPs and IMs are updated to the target builds shipped IM and DP. Hence user made changes to DP and IM are not reflected.

---

### **Property Sets Cannot Be Deleted From the Global Catalog (AOS-31382)**

If an Apstra user deletes a Blueprint with assigned Property Sets, if they try to delete the Property Sets from the Apstra global catalog, they will get an error "Failed to delete resource. Property set in use."

---

### **Rendered Config for Static LAG Changes (AOS-38375)**

In Apstra 4.0 and 4.1.0 versions, 'system-id' and 'force-up' are set in aggregated-ether-options incorrectly for Static LAG.

#### **Resolution**

This was corrected in Apstra 4.1.1, and later versions. Users upgrading will see configuration changes.

---

### **SONiC Device DeploymentProxyAgent Crashing With High CPU After Change (AOS-32198)**

After manual change on SONiC device DeploymentProxyAgent failing every minute causing highcpu. The large diff caused by the manual change causes timeout (60s) of DeploymentProxyAgent.

---

### **SONiC Device System Agent Logrotate Failure Causing Liveness Anomaly (AOS-30810)**

SONiC devices with an Apstra System Agent installed may experience a situation where the agent logrotate function cannot complete an operation causing the agent to repeatedly fail to result in a device agent liveness anomaly in Apstra.

---

### **SystemAgent NOS install task may never complete when AOS host filesystem has issues (AOS-31668)**

When the AOS host filesystem experiences issues, it is possible that a currently running SystemAgent task such as NOS upgrade/install never fails or completes.

---

### **The "Skip Revert to Pristine on Uninstall" function can not be used to retain the full Service Config (AOS-27524)**

When uninstalling a device agent, the checkbox "Skip Revert to Pristine on Uninstall" under "Manager Config" allows the user to retain the config currently on the device. However, it is currently not possible to uninstall an agent while it is still assigned to a Blueprint. This means the full Service Config can not be retained, and only the Discovery1 config can be retained when using this function.

---

### **Unable to connect a single homed generic connected to a leaf in mlag rack with /31 subnet (AOS-24809)**

If user assigns a /31 subnet to the VN, IP address is not allocated to generic on using the connectivity template with the VN and static route or BGP session.

---

### **Unexpected Anomalies From VxLAN Flood List Validation IBA Probe (AOS-30773)**

Due to the caching logic used by the collector for the Apstra VxLAN Flood List Validation IBA Probe, the probe may fail under certain conditions where the collection of VLAN to VNI mapping

failed but the last update was not reset. This will cause the collector to incorrectly believe that the cache is empty and therefore it did not post any data.

---

### **User Can Import the Same Configlet More Than Once (AOS-30263)**

In all current versions of Apstra, a user is able to import an Apstra Configlet from the Global Catalog into their Blueprint. This may cause issues if the user tries to delete them from the Blueprint due to the negation commands causing the state of the other Configlet to become invalid.

### **Fixed Third-Party Issues**

#### **BGP session from dual homed router with leaf's loopback IP may not come up (AOS-25326)**

BGP session from dual-homed router with leaf's loopback IP may not come up.

---

#### **Cisco NXOS 7.0.3.I7(9) (AOS-30786)**

Cisco NXOS 7.0.3.I7(9) does not support unnumbered BGP & Apstra IBA VxLAN Flood List Probe may encounter false anomalies because of which we strongly discourage using of this NXOS version.

---

#### **Removing BFD-enabled BGP neighbor on Cumulus 3.7 or SONiC up to 3.4.1 can cause a config deploy failure (AOS-24679)**

Due to an FRR bug on Cumulus 3.7.x and SONiC versions up to 3.4.1, when a configuration starts with a BFD-enabled generic BGP session, and that session is removed, fr-reload.py fails to apply the new configuration. This is an FRR bug.

### **Known Apstra General Issues**

#### **"management\_ip" field not visible in rendered config preview (AOS-24760)**

When previewing the device context or the device configuration, the 'management\_ip' field is null in the context and any config templates which make use of it may render None.

The management\_ip status is not visible to the config preview APIs. If a user creates a configlet referring to the management\_ip, the preview may show an empty value but the actual configuration pushed to the device will correctly include the proper management\_ip.

---

### **'Export existing' Logical Device functionality does not work (AOS-34514)**

Using the 'export existing' Logical Device functionality is not working when trying to export a Logical Device to the global catalog from the blueprint.

#### **Workaround**

If you need to export a Logical Device from the blueprint to the global catalog, you will need to use the 'export as new' option in the Logical Device and manually map the port references to the Interface Map to keep it aligned with the one used at the blueprint.

---

### **Adjusting the 'next-hop' and 'interfaces' setting under "forwarding-options vxlan-routing" (AOS-32272)**

On Juniper EX4400, QFX5110 and QFX5120 devices running EVPN deployments the default values for 'interfaces' and 'next-hop' under "forwarding-options vxlan-routing" may require adjustments.

#### **Workaround**

The following values are recommended to be configured, if possible, prior to Apstra device system agent installation:

EX4400:

```
set forwarding-options vxlan-routing next-hop 16384
set forwarding-options vxlan-routing interface-num 6144
```

QFX5110:

```
set forwarding-options vxlan-routing next-hop 32768
set forwarding-options vxlan-routing interface-num 8192
```

QFX5120:

```
set forwarding-options vxlan-routing next-hop 45056
set forwarding-options vxlan-routing interface-num 8192
```

If the Apstra device system agent installation and deployment in the Apstra blueprint is already done, the user can use Apstra Configlets to add the configuration. Note, "set" Configlets are only

supported in Apstra 4.0.2 and later. For Apstra 4.0.1 and earlier, the user will need to use a Configlet with a "hierarchical" Junos configuration.

**WARNING! The Juniper EX/QFX PFE will restart automatically when "forwarding-options vxlan-routing" configurations are changed on the device. Traffic will be interrupted!**

Refer to the following Knowledge Base article for details on how to apply these changes using a Configlet: <https://kb.juniper.net/KB69735>

---

### **All AOS Deployments Running a Specific Version Have the Same Set of Secret Keys (AOS-30511)**

All AOS deployments running a specific version have the same set of secret keys. This is potentially a security flaw as a user having access to an AOS VM of a version can get access to secret keys installed in a different VM as they are all the same.

---

### **Apstra 4.1.0 to 4.1.1 Upgrade Failure When Using Modular Device Profiles (AOS-33358)**

In Apstra 4.1.0 when a modular device profile is imported into a blueprint as a part of an interface map, only some properties are created as a part of the blueprint node. The missing properties, while not affecting the operation of the blueprint, will cause the upgrade to Apstra 4.1.1 to fail with a "TypeError: 'NoneType' object is not iterable" error.

#### **Workaround**

Contact Juniper Support for a hotfix script to correct the missing properties in Apstra 4.1.0 which will fix the upgrade.

---

### **Apstra Authentication Agent Crashes When More Than One LDAP Servers Timeout (AOS-42566)**

If the user adds more than one LDAP server and the server does not respond, Apstra will timeout and crash the Apstra authentication agent (Authagent), causing all new login attempts to fail until the agent reco:rvers.

#### **Workaround**

Edit the LDAP provider under Provider-specific Parameters, Advanced Config, set the Timeout(seconds) to 15 seconds or lower to prevent the provider timeout from crashing Authagent.

---

## **Apstra Sysdb Crash (AOS-34904)**

Underlying issue with Sysdb database service may cause a crash in certain conditions.

---

## **Apstra UI Crash When Interface Map View in Blueprint Differs Global Catalog (AOS-32811)**

The Apstra UI may crash (blank screen) if the user changes the Interface Map Global Catalog view to "card view", then if the user goes to the the Interface Map Blueprint view with "table view".

---

## **Apstra Upgrade Connectivity Validation Uses SSH to Check Connectivity to vCenter (AOS-44413)**

Normally, Apstra uses the VMware vCenter API for communication. If the user configures any vCenter servers in Apstra, the `aos_import_state` upgrade will check connectivity to vCenter using SSH, not API. If the SSH is blocked or the SSH credentials are different than the API credentials, this pre-upgrade check may fail, causing the upgrade process to fail.

### **Workaround**

Before upgrading, please check if SSH from the controller to vCenter can be successfully established using vCenter credentials entered in Apstra. If necessary, the user can skip Apstra upgrade connectivity validation with the `--skip-connectivity-validation` option when running `aos_import_state`.

---

## **Apstra upgrade to 4.1.1 adds RSTP configuration (AOS-34638)**

Apstra 4.1.1 upgrade pushes Junos "protocols rstp interface <\*> edge" configuration to all configured generic system facing links immediately after coming out of maintenance mode. This is the change we made in 4.1.1: Change RSTP configuration of Junos leaf devices. Made configuration change to default design for Junos leaf devices to render edge command on server-facing interfaces. This change protects users from introducing a data plane loop in the data center if someone inadvertently connected two leaf devices together, but this change may impact users who have downstream layer 2 connected devices.

### **Workaround**

The user may configure a Junos configlet before upgrade which will remove the RSTP configuration added during upgrade.

---



### **Apstra ZTP Duplicate Entries for Junos Devices (AOS-40023)**

When monitoring Apstra ZTP device status in the Apstra UI under "ZTP Status" / "Devices", there may be duplicate entries for Junos devices. Apstra ZTP will try to ensure the physical management interface for the Junos device is used instead of any virtual management interface (e.g. "vme" interface). Junos may use the virtual interface when ZTP starts but cannot be added to the required "mgmt\_junos" routing-instance. This is done as the first step in ZTP in order to ensure that the management IP address does not change during the rest of the steps involved in ZTP (especially those involving connectivity to Apstra). Enabling a different management interface will cause the DHCP server to give out a new lease. Also, the vendor class identifier for the new management interface is cleared so that the DHCP server does not give out vendor-specific options to this interface, which may re-trigger a new ZTP session while the current session is active.

---

### **Apstra ZTP Faliure During Junos Upgrade with Console Special Characters (AOS-43732)**

Apstra ZTP may fail due to device console issues messages (e.g. "Scheduler Oinker") with special characters during a Junos upgrade.

#### **Workaround**

Manually reboot the device to complete the Junos upgrade, then repeat ZTP.

---

### **Apstra-CLI "system-agents update" Command Resets System Agent Credentials (AOS-42921)**

The Apstra-CLI (a.k.a. AOS-CLI) "system-agents update" command is used to update an existing Apstra system agent. However, if the "username" and "password" options aren't used, any existing system agent credentials will be removed.

#### **Workaround**

The user must use the "username" and "password" options with proper credentials when updating a system agent with the Apstra-CLI "system-agents update" command.

---

### **Arsita EOS VXlan Floodlist Anomalies When Flood Map for Vteps Programmed Correctly (AOS-43128)**

Occasional race conditions may exist for the VXlan collector when cached VNI entries from the device, which will cause false positive IBA VXlan Floodlist probe anomalies even though the VXlan floodmap is programmed correctly in the devices.

## Workaround

The user needs to either restart the Apstra service on the device or initiate a config change for the device from Apstra.

---

## BGP Anomalies Are Unexpectedly Raised for External Generic BGP Sessions While Draining (AOS-32878)

BGP Telemetry continues to expect external generic BGP sessions to be up even if the leaf or spine is in deploy mode 'drain'. This is a cosmetic issue and will not impact the operation of the network.

---

## Build Error (caused by wrong API call) persists even if rollback or revert is executed (AOS-33842)

Build Error(e.g. invalid interface error: Interface with name "E3/1", speed "25G" and role "leaf" not found in interface map) persists even if rollback or revert is executed. The build errors are caused by wrong API calls, however they should not persist but should be cleared after rollback or revert operation.

---

## Workaround

```
sudo service aos restart  
or  
docker exec -it aos_controller_1 bash -c "ps -ef | grep BuilderAgent|grep python | head -n 1"  
docker exec -it aos_controller_1 bash -c "kill -9 "
```

---

## BuilderAgent crashes are seen after name for link is cleared in the UI (AOS-37065)

Link name for a link between system nodes are automatically created when link is created. However, when link name (link label) is accidentally cleared by clicking clear button in the UI, BuilderAgent will crash because it uses link label as one of keys for sorting purpose.

## Workaround

Make non-empty name(accepted to only 65 characters) in the UI or use Swagger REST-API call for Blueprint node patch with the non-empty label string(more than 65 chars are allowed). Recommend not clicking clear button in the link name change. if BuilderAgent restarting stops by continuous crash, please execute service aos restart after fixing.

---

## Can't set speed to 100m/10m in the interface section of JUNOS DP (AOS-35035)

Can't set speed to 100m/10m in the interface section of JUNOS DP

### **Workaround**

Please contact JTAC Apstra support.

---

### **Cannot use exclamation point in NX-OS Configlet for passwords (AOS-14084)**

When creating a configlet for custom username or SNMP3 password on an NX-OS device, "!" cannot be part of the password. This will result in an error as "CLI execution error", clierror: "% Ambiguous command".

### **Workaround**

For NX-OS password, use an encrypted password in AOS Configlet. For SNMP3 passwords, the user will need to use passwords without an exclamation point "!".

---

### **Changing the Apstra Controller IP Address Will Not Update Uploaded OS Image URLs (AOS-37170)**

If the user changes the Apstra Controller IP address either using netplan or aos\_config, the URLs for any uploaded OS Images will continue to use the previous IPs, causing NOS upgrades to fail.

### **Workaround**

After changing the Apstra Controller IP address, the user must either delete and re-upload the OS images or the user can manually edit the image JSON files in the /opt/aos/frontend/www/dos\_images/ directory on the controller.

```
sudo sh -c 'sed -i "s/\\/(old-ip-address)\\/\\/\\/(new-ip-address)\\/\\/g" /opt/aos/frontend/www/dos_images/*.json'
```

---

### **Contiguous Aggregate Routes Specified in Custom Routing Zone Policy Are Aggregated (AOS-38444)**

When contiguous routes within a custom policy applied to a RoutingZone are used, the Apstra rendering engine will incorrectly summarize routes when rendering the VRF config for border leafs. Policy for external BGP sessions does not summarize aggregate routes, which may cause a summarized route to not be announced externally. For example, defining two aggregates, '7.7.6.0/24' and '7.7.7.0/24' will result in a BGP aggregate of '7.7.6.0/23', but the RoutesToExt prefix-list will list both ['7.7.6.0/24', '7.7.7.0/24'], preventing the aggregate route from advertising.

## **Workaround**

Add only the summarized large aggregate. When attempting to aggregate ['7.7.6.0/24', '7.7.7.0/24'], specify the BGP aggregate in the routing policy as ['7.7.6.0/23'].

---

## **Controller CPU History May Fail After VM Hard Reset (AOS-31975)**

If you perform a hard reset of your Apstra Controller VM, querying the Controller CPU history may fail due to a truncated file.

## **Workaround**

Contact Juniper Support

---

## **Creating a Virtual Network and Assigning It With a Connectivity Template May Trigger EVPN Type-5 Anomalies (AOS-32439)**

If you create a Virtual Network and assign it with a CT, an EVPN Type-5 anomaly may raise.

## **Workaround**

Disable and re-enable the related IBA probe.

---

## **Duplicate BGP Neighbor IP Addresses in CT/Remote GW Not Validated (AOS-36684)**

Apstra will not validate BGP neighbors with duplicate remote gateways IP address configured in Connectivity Templates.

## **Workaround**

The user will need to remove any duplicates.

---

## **ESI MAC MSB Change When Enabling IPv6 (AOS-33718)**

When enabling IPv6 Applications in an Apstra blueprint, ESI MAC MSB will change for ESI leafs causing an unexpected, incremental configuration change for Junos devices.

---

## **EVPN VXLAN TYPE-5 Route Probe Anomaly (AOS-36190)**

EVPN VxLAN type 5 route anomaly is expected when using a single connected generic on an ESI-connected leafs. This is caused by an Apstra bug expecting routes that will not be present in

the table of ESI leaf.

### **Workaround**

Upgrade to Apstra 4.1.2 if EVPN VxLAN type-5 probe is needed. If not, the probe anomaly can be verified and ignored.

---

### **Fabric expansion operation or flexible fabric expansion operation can lead to the involved system being removed from its blueprint (AOS-33756)**

A variety of fabric expansion operations or flexible fabric expansion operations on a node can result in the node being unassigned from its blueprint and its `deploy_mode` cleared. Examples include:

- Addition of a new generic system to a managed system's port that does not belong to the default transformation.
- If two link speed changes are made to a managed system in the same commit, and the speeds are swapped symmetrically (e.g. 10G -> 1G for the first link and 1G -> 10G for the second link), the specific system can be unassigned from the blueprint.
- Other cases where changes have to be effected to the interface map blueprint node used by the system undergoing fe or ffe expansion.

### **Workaround**

After any fabric or flexible fabric expansion operation on a managed system, please do not immediately commit. Check the system's deploy mode and, if unset, please set it to "deploy" again, then also reassign the device that was originally assigned to the managed system.

---

### **FFE Operations in POD Based Blueprints Involve the Formation of Different rack\_type With the Same ID (AOS-31854)**

If a user performs an add rack operation in an existing POD, adding a new POD based on the same rack will fail with a "rack\_types": "Values are not unique" error.

### **Workaround**

Contact Juniper Support

---

### **Freeform Deploy Mode Not Reset When System ID Is Unassigned (AOS-32751)**

In an Apstra Freeform blueprint, when a System ID is unassigned for a node, the deploy mode for the node is reset. This will result in a build error with `System ID must be set when deploy mode is "ready"/"deploy"/"drain".`

## **Workaround**

The user will need to manually change the deploy mode for the node to "undeploy".

---

### **Freeform Unassigned System ID Isn't Available for Other Nodes (AOS-32758)**

In an Apstra Freeform blueprint, if the user has 2 internal nodes, assigns a system ID to one node, reopens the systems page, selects both nodes, clicks the "edit" button to manage assigned systems, unassigns the system ID from the first node, they will be unable to assign the system ID to the second node.

## **Workaround**

The user can work around this issue by closing the managed assigned systems window after unassigning the first node. When they reopen the window, they will be able to assign the system ID to the second node.

---

### **Full Filesystem Will Cause Incorrect Rollback on Revert (AOS-32966)**

If the Apstra controller server's `/var/lib/aos/db` filesystem becomes full, blueprint changes can continue to be made and deployed to the devices, but the changes will not be written to the Apstra controller server disk. If the user reverts any uncommitted changes, the last blueprint state successfully written to disk will be loaded and the user will not be able to restore any changes.

## **Workaround**

The user must not make any blueprint changes if there are any disk or memory utilization warnings about the state of the Apstra controller (e.g. "Some partitions are almost full"). Please refer to <https://kb.juniper.net/KB37699> for instructions to add disk space to filesystems on the Apstra controller server.

---

### **IBA "Critical Services Trending and Altering" Dashboard Does Not Show Graph (AOS-30162)**

In the Apstra IBA "Critical Services Trending and Altering" Dashboard, the "Individual interfaces bandwidth 1-day trending" Probe does not show the graph.

---

### **IBA -EVPN VXLAN Type-5 Route, Processor: EVPN Type 5 Routes â†’ EVPN Table, "Endpoint" search not returning expected results (AOS-32622)**

In EVPN VXLAN Type-5 Route, Processor: EVPN Type 5 Routes -> EVPN Table, if you search by "Endpoint", the result is an empty list.

---

### **In the VM Query, VM Doesn't Show the Connected Leaf Node and Interface (AOS-37913)**

When Leaf node, connected by VMware ESXi hosts, is configured with domain name and hostname, fully qualified hostname is reported to ESXi host via LLDP. When the fully qualified hostname is exactly matched against the leaf node's hostname (non-fully qualified name), it leads to match failure so that the connected leaf node can't be found.

#### **Workaround**

The user must not use the domain-name in the leaf node where VMware ESXi hosts are connected.

---

### **In-Place Upgrades From Versions Before Than Apstra 4.1.0 Are Not Supported (AOS-30442)**

In-place upgrades from Apstra 4.0 versions are not supported. Due to a change in the backend database version, in-place upgrade may break communication between Controller and Device agents.

#### **Workaround**

The user must use a VM-VM upgrade when upgrading from any Apstra 4.0 version.

---

### **Incorrect Selection of Items When Using Query on "Change Link Speed" (AOS-37589)**

Under Staged>Physical>Links>Change Link Speeds, the Query option displayed items according to the pagination of items. When selecting all for that page, it will select all the items comprised of other pages as well.

#### **Workaround**

The user can select individual links or contact Juniper Support for an Apstra UI hotfix patch for Apstra 4.1.1 (UI patch 4.1.1-394) or Apstra 4.1.2 (UI patch 4.1.2-212) to use the select all function.

---

### **Inter-VRF Routing Problem w/ Single Spine Path (AOS-38834)**

Due to an issue with the Apstra EVPN reference design, problems with inter-vrf routing via an external router can occur if there is a single spine (either by design, failure, undeploy, or drain)

and the route is originated on a non-border-leaf. The one spine will drop the route due to as-path loop.

### **Workaround**

Workarounds include originating the virtual network route on the border-leaf, adding a Junos set configlet "set protocols bgp group l3clos-s-evpn family evpn signaling loops 2" on all spines, or reconfiguring the external router to remove spine ASNs from the as-path.

---

### **Interfaces Field Is Empty When Editing an Interface Configlet Inside the Blueprint (AOS-32869)**

When editing an interface configlet inside the blueprint, the interfaces field will always be empty and will not contain previously configured interfaces.

### **Workaround**

When editing an interface configlet inside the blueprint, the user will need to always configure all necessary interfaces including previously configured interfaces.

---

### **Juniper QFX5120-48YM Device Profile error for port 50 and 52 (AOS-34374)**

The Juniper QFX5120-48YM Device Profile has an error for ports 50 and 52 which will cause deployment errors when the user tries to connect 40G ports.

### **Workaround**

Contact Juniper Support for an updated Device Profile.

---

### **Junos "statement has no contents" Warning Causes Deployment Failure (AOS-33355)**

If you use a Junos system configlet with an "empty stanza", when Apstra deploys this configuration to the Junos device, Junos will respond with a warning "warning: mgd: statement has no contents; ignored", however Apstra will treat this warning as an error causing a deployment failure, example "Apply config failed: ConfigLoadError(severity: warning, bad\_element: et-0/0/49, message: warning: mgd: statement has no contents; ignored".

### **Workaround**

You must remove the "empty stanza" from the Apstra Junos system configlet and re-import the configlet into the blueprint.

---



### **Junos 'device-count' for lag is incorrectly counting ae interfaces (AOS-40448)**

Junos 'device-count' configuration for defining the number of aggregated ethernet ports is incorrectly including layer3 port-channel subinterfaces in the total device count.

---

### **Junos EVPN\_IMPORT policy-statement config rendering change (AOS-33852)**

On Junos devices, the EVPN\_IMPORT policy-statement used for custom import & export route targets generates an `accept` action on every statement, leading to only the first statement being evaluated.

#### **Workaround**

Correct this behavior via configlet or upgrade to Apstra 4.1.2.

---

### **Link Tags Not Properly Associated With ESI/MLAG Interfaces (AOS-42414)**

Link tags applied to physical interface members are not associated with ESI/MLAG interfaces, but they are associated with non-ESI/MLAG LAG interfaces.

---

### **Missing Upgrade Plugin for node\_to\_node\_if\_counter Processor (AOS-40850)**

If the user has an Apstra blueprint created before Apstra 3.3.0 configured with the Headroom probe, upgrades to Apstra 4.1.1 and later may fail with error `AttributeError: 'NoneType' object has no attribute 'validate_config'` because the `node_to_node_if_counter` processor has been removed.

#### **Workaround**

Prior to starting the Apstra 4.1.x upgrade, the user must remove all probes with the `node_to_node_if_counter` processor.

---

### **Offbox System Agent and IBA Containers May Restart (AOS-33181)**

Certain conditions may cause Apstra offbox system agents and IBA containers to restart when it starts. This will only happen if a new container configuration is created while the agent is down. This should have no impact.

---

## **PODs May Not Show. in Drop-Down List When Trying to Add a New POD (AOS-32405)**

When trying to add a new POD from your blueprint, it may not show up in the drop-down list in the blueprint.

### **Workaround**

Clone the POD that is not showing in the dropdown menu from the Global Catalog.

---

## **Post-upgrade Virtual Network Policy Settings EVPN Type 5 Routes is Null (AOS-34376)**

For users running Apstra upgraded from versions prior to 4.0, after the upgrade, for Virtual Network Policy Settings, the "EVPN Type 5 Routes" default setting of "Disabled" is not retained and neither option will be selected.

### **Workaround**

After the upgrade, the user must modify Virtual Network Policy Settings in their blueprints, manually setting "EVPN Type 5 Routes" to "Disabled" and committing changes to their blueprint.

---

## **ref\_count of protocol\_session from dynamic BGP peering was not correctly set (carried with null) when upgrading from 4.0.0 to 4.x.x (AOS-33904)**

Un-assignment of switch's ports from CT with dynamic BGP peering always fails with "Server-side Validation Errors - Internal error: reference counter is not initialized" after upgrading 4.0.0 to 4.0.2/4.1.1. The ref\_count for protocol session is not properly set during upgrading process and validation logic in new upgraded version prevents CT from being unassigned.

### **Workaround**

Manually correct using curl or UI after finding out protocol\_session node from QE.

1. At first find out blueprint ID. You can get it when you click blueprint the UI.
2. After selecting the staged tab, launch graph explorer and execute query to find out protocol\_session nodes with ref\_count as null.

Graph Query: `node('protocol_session',name='protocol_session',ref_count=is_none())`

The output should show all the protocol sessions with ref\_count value as null (not set).

3. For the the collected IDs of protocol\_session nodes from step2, Choose Step 4 using UI or Step 5 using curl in the terminal
4. UI: Platform > Developers, Click REST API Documentation button, select PATCH `/api/blueprints/{blueprint_id}/nodes/{node_id}`
  - a. Update node.
  - b. fill out blueprint ID and node ID (from Step 2), body as {

```
"ref_count":1
}
```

Execute it. it should be done for all protocol\_session nodes collected from step 3.

5. Curl: `curl -X PATCH "https:///api/blueprints//nodes?type=config" -H "accept: application/json" -H "AUTHTOKEN: " -H "content-type: application/json" -d "[{ \"id\": \"\", \"ref_count\": 2 }]" --insecure`

a. from Step 1

b. from Step 2.

6. Execute Graph Query to check out whether protocol session has ref\_count value as 1.

Graph Query: `node('protocol_session',name='protocol_session')`

The result should show reference count value as 1.

Try again to remove CT from leafâ€™s interface.

After the whole operation is done, please commit the BP to make assure that the deployed one has correct ref count value.

---

### **Rendered Configuration Error for Juniper\_EX4400-48T\_EM-4Y and Juniper\_EX4400-48T (AOS-35037)**

Rendered configuration error is seen for Juniper\_EX4400-48T\_EM-4Y and Juniper\_EX4400-48T when the 100Mbps transformation is used.

---

### **RPC Timeout Error Message on Juniper QFX5100 Device (AOS>-36338)**

Apstra users with Juniper QFX5100 devices may receive the following rpc timeout error messages on the device "Rpc timed out: RpcTimeoutError(host: xx.xx.xx.xx, cmd: commit-configuration, timeout: 60)".

#### **Workaround**

Configure "Open options" "commit\_timeout" value to "120" seconds in the system agent or agent profile.

---

### **Running `sysctl --system` causes offbox agents to go offline (AOS-36918)**

If the user, on the Apstra VM CLI, runs `sysctl --system`, incorrect kernel forwarding parameters in `/etc/sysctl.d/60-aos_sysctl.conf` will be loaded causing off-box system agent Docker containers to go offline.

#### **Workaround**

Restart docker and aos services to fix kernel parameters and reloads offbox agents with the

following on the Apstra VM CLI:

```
sudo systemctl restart docker && sudo systemctl restart aos
```

---

### **Security policies are not supported in SONiC (AOS-34402)**

Security policies are not supported in SONiC. Any security policy configuration done on a SONiC device will be accepted, but will not be implemented.

---

### **Separate CTs Might Swap Assignments in Some Conditions (AOS-41465)**

Under specific conditions where the user has multiple, similar connectivity templates (CT), an incorrect graph database structure might occur. This may cause unexpected validation errors when the user makes subsequent CT changes.

#### **Workaround**

1. Identify these two binded CTs. They should be assigned to the same application points. The user can use the following graph query to check for any graph database errors.

```
match(
node('ep_group', name='corrupted_group')
.in_('ep_affected_by')
.node('ep_application_instance')
.out('ep_top_level')
.node('ep_endpoint_policy', name='ct', policy_type_name='batch')
.having(
node('ep_group', name='corrupted_group')
.in_('ep_affected_by')
.node('ep_application_instance')
.out('ep_top_level')
.node('ep_endpoint_policy', name='ct', policy_type_name='batch'),
at_most=2
),
node('ep_endpoint_policy', name='ct', policy_type_name='batch')
.in_('ep_top_level')
.node('ep_application_instance')
.out('ep_affected_by')
.node('ep_group', name='corrupted_group')
.in_('ep_affected_by')
.node('ep_application_instance')
.out('ep_top_level')
.node('ep_endpoint_policy', name='ct2', policy_type_name='batch')
.where(lambda ct, ct2: ct.label != ct2.label and ct.id > ct2.id)
)
.distinct()
```

2. Unassign all application points from **both** conflicting CTs. Both CTs should be reverted to "Ready" state when they have 0 application endpoints assigned.

3. Assign each CT to proper application endpoints.

---

### **Setting VXLAN for the First Time and Enabling DHCP Helper Addresses Simultaneously May Fail in SONiC 4.X Devices (AOS-38701)**

If vxlan vtep does not exist and is to be enabled in a blueprint as part of a day-2 operation, and in the same config apply any DHCP helper address is to be set in at least one vxlan-enabled vlan, the config apply operation will fail.

#### **Workaround**

- The first config apply of a blueprint is not a day-2 operation and does not apply to this case.
  - User can enable vxlan in a separate config apply and then apply necessary DHCP helper addresses in a subsequent config apply.
  - If the bug has already happened, please full config apply any device that failed deploying. This will clear any problem.
- 

### **Some local mac entries may not be reported by the aos mac telemetry in SONiC (AOS-32975)**

Some local mac bridge fdb entries might not be picked up by the SONiC aos mac telemetry.

---

### **SONiC DHCP Relay Towards Helper Goes Over the Default VRF (AOS-44242)**

The Apstra reference design implementation for SONiC, communication of the DHCPv4 and DHCPv6 relay always uses the default VRF. This means that the DHCP server must always be reachable over the default VRF, regardless of the VRF to which the DHCP client belongs. The DHCP relay process will not operate correctly if the DHCP server is not reachable over the default VRF.

#### **Workaround**

The user must ensure the DHCP server addresses is always reachable over the default VRF.

Alternatively, a full config apply has been observed to put the DHCPv6 and DHCPv4 relay in the correct VRF as well. Do note however, that any subsequent incremental manipulation of the

DHCP helper configuration will negate the correct VRF and reset it to default, necessitating another full config apply.

---

### **SONiC Non-"Admin" Device-User Not Added to sudo Group via Apstra ZTP (AOS-21365)**

When using Apstra ZTP to bring up a SONiC device, the non-"Admin" device-user added to SONiC by ztp.py is not added to the sudo group.

#### **Workaround**

Change ztp.py to allow additional administrative groups (sudo, docker) into non-default device-user. Contact Juniper Support.

---

### **Switching between speeds 10G and 1G on a Dell S5248F-ON or S5296FON may fail (AOS-33885)**

In case of a topology whose AOS has been upgraded to 4.1.2 from a previous version, setting a 10G port to 1G or vice versa may fail on a Dell S5248F-ON or S5296F-ON. The problem can happen only if the device has been previously configured by an AOS version earlier than 4.1.2.

#### **Workaround**

It is possible to proactively solve this problem by executing ``sonic-db-cli CONFIG_DB HSET PORT|EthernetX valid_speeds 10000,1000`` where EthernetX is the interface about to change speed. The command can be inserted in a configlet. Please contact support for a more generic script that will fix all interfaces at once.

---

### **The display\_id for a Cloned Device Profile Starts From 1 Instead of 0 (AOS-33656)**

When cloning a device profile, the display\_id starts from 1 instead of 0 and hence the interface numbering start from 1 instead of 0.

#### **Workaround**

The user will need to correct the display\_id in a cloned Device Profile.

---

### **Traffic Heat Layer Stays in Continuous Loading Status (AOS-35819)**

Traffic Heat layer in the active tab stays in continuous loading status without showing result after leaf or access switch with link to generic system is un-deployed. A request with malformed filter information, which includes interface from the un-deployed device, from frontend UI to the backend makes parsing of the request fail and UI not be updated with continuous retrials.

## **Workaround**

Re-deploy the device back into the blueprint or contact Juniper Support for a UI hotfix patch.

---

### **Unexpected Config Pushed When Neighbor Deploy Mode Changed (AOS-34051)**

If the user changes the deploy mode from "deploy" to "undeploy" for a node in an Apstra blueprint, the MTU for the neighbor's interface facing this node may unexpectedly change. For example, the MTU on the interface may change from 9216 to 9050.

## **Workaround**

There should be no impact as the BGP session on the link will expedectedly go down. If needed, the user may override this change with an Apstra configlet.

---

### **Unexpected Policy Changes And Potential BGP Flap When Draining One Leaf In Leaf Pair (AOS-37948)**

In Apstra, when the user sets one leaf in a leaf pair to drain mode, Apstra fabrics with Juniper Junos devices will deploy unexpected policy changes between the spines and the other leaf in the leaf pair. This may result in the network operating system resetting the BGP session between the spines and the other leaf. Apstra fabrics with other vendor devices (SONiC, Cisco, Arista) are not affected by this issue.

---

### **Upgrade Precondition Fails in Case Username Contains "@" (AOS-33738)**

The upgrade precondition checks fail in case the username has the char @.

## **Workaround**

The user can use the "--skip-connectivity-validation" option to skip this precondition validation.

---

### **User Can Add Virtual Network Endpoints on Invalid Devices (AOS-34306)**

Apstra is missing a validation that will allow the user to add Virtual Networks Endpoints on devices that do not host the Virtual Network.

## **Workaround**

The user will need to modify the Virtual Network so it is configured for the device.

---

### **vCenter VM Node's VNIC Node to VNET Node Relationships Are Not Updated on Network Adaptor Delete/Create (AOS-34281)**

New port groups were added, and there were VMs using these port groups, and there was no VN configuration on the fabric side. Apstra correctly flagged that the fabric is missing VLAN configurations. The IBA probe to detect the impacted VM's because of the missing Virtual Network configurations did not report the VMs.

---

### **Viewing Connectivity Template With Invalid JSON Data Causes Apstra UI Crash (AOS-35876)**

When Connectivity Template (CT) was created with invalid JSON data via API call, any following edit/view functions of the CT can lead into Apstra UI crash. The fix makes the Apstra UI parse the invalid JSON data safely not to cause the Apstra UI to crash.

#### **Workaround**

The user will need to delete the CTs with invalid JSON data via an API call.

---

### **When a System-Agent Job Is Currently in Progress, Starting Device Show Tech Collection Does Not Throw an Error (AOS-32807)**

If an Apstra user is running any System-Agent job for a device and the user starts a device Show Tech collection, the job does not start because there is already a system-agent job in progress, but the Apstra UI says "Successfully started collecting show tech" which is misleading.

---

### **While editing configlets with jinja device context the preview will render incorrectly (AOS-35141)**

In Apstra 4.1.2 and earlier, while editing configlets the config rendering preview may not be correct when using jinja device context as it requires a device id/model context to render properly.

#### **Workaround**

Please use the devices rendered config as a preview of the changes being made by the configlet.

### **Known Third-Party Issues**



## **All BGP peerings with password get restarted in every config apply on SONiC (AOS-34086)**

All BGP peerings with password get restarted in every config apply on SONiC. Any config apply on a SONiC device will cause all BGP peerings that use a BGP password to flap. It is noted that fabric links do not use BGP passwords.

### **Workaround**

Do not use BGP passwords in SONiC BGP peerings.

---

## **Apstra Sonic on-box agent installation failure on Sonic switches (AOS-32738)**

On switches running NOS versions earlier than Sonic 4.0, Apstra agent installation can fail if the SWSS service on switch is down.

Apstra agent installation workflow require that all core services from command output of "show system status" to be up (healthy) before Apstra device agents become operational.

If any of the system status is not UP(healthy), Apstra will stop agent installation and fail with error message.

### **Workaround**

The issue here is caused by a timeout occurring during boot up with swsswait.sh and wait\_for\_replay\_done.sh files so the workaround is to update timeout in swsswait.sh and wait\_for\_replay\_done.sh files.

Change the slice value from 96 to 300 in both the swsswait.sh and wait\_for\_replay\_done.sh in their respective containers:-

```
root@netst10mgmt14:~# docker exec -it swss bash
root@netst10mgmt14:/# vi /usr/bin/swsswait.sh
And make sure beginning of the file looks as: #!/usr/bin/env bash
```

```
function wait_until_config_replayed
{
slice=300 <<<<<<< Update this value to 300
VRFMGR=0
INTFMGR=0
VXLANMGR=0
ALLDONE=0
```

Follow same steps to update slice value in wait\_for\_replay\_done.sh file

```
root@netst10mgmt14:~# docker exec -it bgp bash
root@netst10mgmt14:/# vi /usr/bin/wait_for_replay_done.sh
```

Save the files and reboot the switches to fix the SWSS service issue

---

### **Apstra Using Cisco NX-OS Depreciated "soft-reconfiguration inbound always" Configuration Causing Packet Loss (AOS-35513)**

Apstra uses Cisco NX-OS depreciated "soft-reconfiguration inbound always" configuration in its EVPN reference design. Usually, this wouldn't be an issue however, there is a Cisco bug (CSCvz75734) where this can lead to potential packet loss.

#### **Workaround**

Contact Cisco TAC for a fix release.

---

### **Arista EOS Link-local BGP peering does not come online with older versions of EOS (AOS-30733)**

Arista EOS only supports neighbor-based link-local peering as of EOS 4.24+. Previous releases of EOS will not bring up a link-local (IP Unnumbered) BGP session. This configuration is accepted by EOS and no deployment failure is observed. The symptom will be under 'show bgp configuration unsupported', the 'neighbor interface ' bgp session will appear as unsupported, and Apstra may raise 0.0.0.0/0 routing anomalies if this is configured underneath the default VRF with an Ipv4 default route expectation. No BGP anomalies are raised for missing link-local peers.

#### **Workaround**

Upgrade to a more recent version of Arista EOS for link-local (IP Unnumbered) BGP peering, eg at least 4.24+.

---

### **BGP peering with password to an IPv6 peer over non-default VRF may fail to establish on SONiC Buzznik 3.5.3 or earlier and SONiC Cyrus 4.0.2 or earlier (AOS-34839)**

For a fully detailed description of bug, please refer to vendor issue SONIC-65999. Under specific circumstances, a password-using BGP peering with an IPv6 peer over a non-default VRF may fail to establish, if there are also IPv6 peerings over the default VRF. The issue does not consistently appear, but may stochastically appear when restarting the BGP service or the entire device.

#### **Workaround**

Please contact the vendor with reference to issue SONIC-65999 for details.

---

### **Cisco C9348GC NXOS Rollback Failures (AOS-34250)**

NXOS on the Cisco 9348GC platform apparently has a bug with the NXOS “rollback” command which Apstra uses to revert the device configuration to pristine. The verification process in this command may start failing without errors which will cause Apstra configuration deployment errors (NXAPI transaction timed out).

### **Workaround**

The user can workaround this NXOS bug by adding NXOS configuration `snmp-server enable traps lldp lldpRemTablesChange` to the device before installing the Apstra device system agent or editing the Apstra device pristine configuration after removing it from the blueprint.

---

### **DHCP Relay Not Working for Juniper QFX10000 Devices (AOS-27830)**

Juniper does not support having DHCP clients on QFX10000 devices (e.g. QFX10002) being used as border-leafs in an EVPN-VXLAN-based data center Edge-routed bridging (ERB) fabric.

### **Workaround**

The customer must disable dhcp-relay on QFX10000 devices (e.g. QFX10002) being used as border-leafs. Please contact Juniper Support. <hr>

### **DHCPv6 Relay Not Working for Juniper QFX10000 Devices (AOS-26987)**

Due to a Junos issue with DHCP relay packet handling with the expess ASIC, DHCPv4 and DHCPv6 relay fails for Juniper QFX10000 devices (e.g. QFX10002, QFX10008) used as a non-border-leaf when the DHCP server is behind an external router.

### **Workaround**

`forward-snooped-clients all-interfaces` needs to be enabled for applicable VRFs. An Apstra Configlet is available to add this. Please contact Juniper Support.

---

### **Juniper QFX10000 Devices Unable to Send Packets on Tagged Layer2 Interfaces (AOS-35096)**

Juniper QFX10000 devices cannot send packets on tagged layer2 interfaces for external router connections.

### **Workaround**

Use VLAN tagged layer3 sub-interfaces for external router connections on QFX10000 platforms

---

**Some non-qualified Junos releases in Juniper 10K switches may fail to provide a correct**

### **Serial Number after a NOS Upgrade (AOS-32421)**

Some non-qualified Junos release upgrades in Apstra may lead to an inconsistent state in the Juniper 10K platform due to the device providing a different chassis number after the reboot, this leads to a serial number mismatch in Apstra.

#### **Workaround**

Re-onboard the device if the issue is triggered. This is fixed in Apstra 4.1.2.

---

### **Unicast DHCP Packets Might Get Flooded When DHCP Relay Is Configured in Non-Default Routing-instance (AOS-31196)**

Due to Junos bug <https://prsearch.juniper.net/problemreport/PR1603444> , Apstra does not recommend using Junos version 20.2R2-S3.

---

### **VMware vCenter 7.0 error when deploying OVA from URL (AOS-32485)**

When deploying an OVA from a URL in the vCenter Server 7.0, due to a regression in vCenter, the deployment may fail with an error, "Unable to retrieve manifest or certificate file".

<https://kb.vmware.com/s/article/79986>