

# Juniper Apstra Version 4.1.2 Release Notes

## Known Critical Apstra Issues

### **NXOS BGP Incorrect local-as Command Negation (AOS-43313)**

When reverting a Cisco NXOS BGP configuration which uses the local-as command, Apstra uses just the "no local-as" configuration command which is invalid and will cause a deployment anomaly.

#### **Workaround**

The user can workaround this issue by manually issuing the "no local-as {asn}" command.

## Feature Configuration Rendering Changes

### **Route Tagging using BGP communities (RFE-1407)**

Apstra 4.1.2 introduces a new feature where all routes (IPv4 and IPv6) generated within the DC fabric, routes received from external generic systems, and routes received from remote EVPN gateways are tagged with BGP communities (RFC1997 - BGP Communities Attribute). These communities allow you to identify any BGP route within the DC fabric quickly and will be used for running more sophisticated route telemetry in future releases.

Introducing this new feature results in new lines of configuration on deployed network devices. These configuration changes will not impact the control or forwarding plane and thus will not be service-impacting. If you have existing route policies or custom BGP communities, we encourage you to carefully review the changes before upgrading.

---

### **Remove Soft-Reconfiguration Inbound on EVPN Address-Family Templates (RFE-2734)**

Due to Cisco bug CSCvz75734, it is recommended that 'soft-reconfiguration inbound' is removed from the EVPN address-family templates on NXOS.

Future releases of NXOS may begin to raise errors during deployment if this configuration is present.

Moreover, presence of the soft-reconfiguration inbound is known to cause issues on some NXOS software releases such as 9.3(6), 9.3(7), 9.3(8).

Upon upgrade to 4.1.2, Apstra will incrementally remove soft-reconfiguration inbound from l2 evpn address family templates.

This command is not required for telemetry or forwarding, and was mostly added as a convenience tool for debug purposes so network engineers can view unmodified BGP routing tables prior to RIB installation.

---

### **Alerting of configuration changes during Apstra upgrade (RFE-2698)**

Added a new capability that ensures you are aware of any configuration change Apstra may render to devices running in a blueprint, either to address issues or for enhancements, prior to performing the upgrade. During the upgrade process, you can review any changes and decide to proceed or abort the upgrade process. The install script presents a summary of the configuration changes, logically grouped into sections, that would be rendered on the devices.

## **New Features**

### **Support for Juniper QFX5700 Chassis Profile and 3 Line Cards (RFE-2206)**

**Feature Category:** Device Profiles

Support for Juniper QFX5700 chassis profile and 3 line cards in Apstra 4.1.2

---

### **Support for Juniper modular devices PTX10000 series (RFE-2000)**

**Feature Category:** Device Profiles

Apstra now supports Juniper modular chassis PTX10004, PTX10008 and PTX10016 with the following line cards PTX10K-LC1201-36CD (36 \* QSFP56-DD) and PTX10K-LC1202-36MR (32 \* QSFP28+ and 4 \* QSFP56-DD).

Note: Both line cards includes transformations for 1x200G and 2x200G as these are supported and compatible by hardware, although corresponding Juniper-supplied qualified transceiver is not yet available and these transformations are not yet supported by software.

---

### **Support for Juniper EX3400-48T as Access Switch (RFE-2102)**

**Feature Category:** Device Profiles

Support for Juniper EX3400-48T as Access Switch in Apstra 4.1.2

---

**Support for Juniper ACX7100-48L (RFE-2203)**

**Feature Category:** Device Profiles

Support for Juniper ACX7100-48L in Apstra 4.1.2

---

**Support for Juniper ACX7100-32C (RFE-2204)**

**Feature Category:** Device Profiles

Support for Juniper ACX7100-32C as part of Apstra 4.1.2

---

**Support for EX4400-24MP and EX4400-48MP (RFE-2335)**

**Feature Category:** Device Profiles

Apstra now supports Juniper EX4400-24MP and EX4400-48MP devices

---

**Support for Cisco Nexus 9508 (RFE-2524)**

**Feature Category:** Device Profiles

Apstra now supports Cisco Nexus 9508 modular chassis with N9K-X9732C-EX (32x100G) line card

---

**Support for Arista DCS-7280SR3-48YC8 (RFE-2411)**

**Feature Category:** Device Profiles

Apstra now supports Arista DCS-7280SR3-48YC8-F device profile.

---

**SONiC OS 4.0.5 (RFE-2880)**

**Feature Category:** Device Operating Systems

SONiC-OS-4.0.5-Enterprise\_Advanced SONiC-OS-4.0.5-Enterprise-Base and are now qualified with Apstra 4.1.2

---

### **Qualified EVPN Leaf support for Junos Evolved (RFE-2042)**

**Feature Category:** Device Operating Systems

Qualified EVPN Leaf support for Junos Evolved.

---

### **Junos OS version 22.2R3 is now qualified with Apstra 4.1.2 (RFE-2804)**

**Feature Category:** Device Operating Systems

Junos OS 22.2R3

---

### **Apstra CLI command to change device credentials (RFE-2517)**

**Feature Category:** Device Operating Systems

New Apstra CLI command to change device credentials

```
scenario change-device-password --blueprint (bp_id) --system (sys_id) --old-  
password (old_password) --new-password (new_password)
```

---

### **Apstra 4.1.2 NOS Version Support (RFE-2528)**

**Feature Category:** Device Operating Systems

Apstra 4.1.2 now supports the following switch vendor software:

Juniper Junos: 21.4R3

Juniper Junos Evolved: 21.4R3-EVO and 22.2R2-EVO

Enterprise SONiC: 3.5.4 and 4.0.5

Arista EOS: 4.27.6M

Cisco NXOS: 9.3(10) and 10.1.2

---

## **Usability improvements for Freeform designs (RFE-2590)**

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

Thanks to user feedback, various improvements were made to make your experience building Freeform blueprints even smoother. To name a few:

- the ability to clone config templates and property sets
  - allow editing the "assigned to" systems from the Config Template menu
  - when creating/editing config templates, you'll see a list of existing config templates that can be included
  - expose the error/warning labels to 2nd/3rd level navigation tabs, so you know exactly where the build errors/warnings are happening
- 

## **Support for "established" option in security policy (RFE-2433)**

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

You now have the ability to set permit/accept `established` options to TCP traffic. This ensures only sessions initiated from inside your network (has the Acknowledgment (ACK) or Reset (RST) bit set) will be allowed back. This option only exists with TCP traffic types. Any connections initiated from outside the Apstra fabric will be denied. This argument is relevant ONLY to the destination, not the source.

---

## **Node search capability in cabling-map editor for Freeform blueprints (RFE-2623)**

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

In Freeform Blueprints, Cabling-Map editor now exposes a node search field available in both View and Edit mode.

Use this field to search for nodes based on a key string. The matching results will be highlighted in the cabling map editor. You can also quickly select a node by clicking it or pressing Enter and then perform an action such as creating a link or editing node properties.

---

## **Granular Day-2 operations for ESI Access pairs (RFE-2314)**

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

You can now perform granular Day-2 changes on topologies with ESI-Access pairs, including:

- Adding an access-switch pair.
- Adding links to an access-switch pair.
- Adding peer links in an access-switch pair.
- Adding a generic system to an access-switch pair.

---

### **Generic System task to update Port Channel ID range (RFE-2395)**

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

You're now able to update the Port Channel ID range used by a Generic System when selecting the Generic System in the Topology view.

---

### **Documentation of custom Jinja filters and functions (RFE-2557)**

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

An in-product help page is now available for all Jinja functions supported. This includes all custom Jinja filters and functions added in the product and available to use in Configlets and/or Config Template along with examples on how to use them. This is grouped in the following categories:

- Jinja Common Filters: Available for both Reference Designs.
- Jinja Common Functions: Available for both Reference Designs.
- Freeform Jinja functions: Available for Freeform only.

---

### **Create Routing Zones and Virtual Networks at bulk (RFE-2231)**

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

In 4.1.2, you're now able to create and update hundreds of Routing Zones (RZ) and Virtual Networks (VN) in one go via the new file export/import functions available in the GUI. Export the file template for RZ or VN, populate the information using your favorite table editor tool, then simply re-import the file via the GUI to preview your desired changes. The file syntax and data will be validated before creating the bulk create/update task requested.

---

### **Bulk unassign Connectivity Templates from the Table View (RFE-2486)**

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

You are now able to bulk unassign Connectivity Templates from the recently added "Application Endpoints">"Table View". This is useful for example for removing multiple Virtual Networks from multiple interfaces in one go.

---

### **Bulk assign multiple Virtual Networks to multiple racks (RFE-1873)**

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

You can use this new GUI modal to manage the allocation of your VNs to your racks more simply, for example, you can assign multiple Virtual Networks to a newly added rack, or, assign multiple new Virtual Networks to a set of existing racks in one go.

---

### **Blueprint resource allocations for Freeform (RFE-2539)**

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

You can use Apstra's resources allocation capability on Freeform blueprints to facilitate allocation of resources at scale and automatically on any size of Blueprint. Resources represent some value of type: IPv4, IPv6, ASN, VLAN, VNI or Integer which is either manually allocated or automatically allocated (from Resource Allocation Group, Local Pool). You can use Resource Generators to automate the management of Resources in case of a large number of repetitive patterns. The Generators will rely on Graph queries to define the scope of application of a resource and will be kept in synch as the Blueprint evolves. i.e this allows the resources allocation schema defined to automatically react to topology changes, by providing or releases resources on the fly.

---

### **Added support for "local-as" option in Connectivity Template (RFE-2368)**

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

This feature simplifies the work and impact of migration or merger activities. This allows you to preserve the global switch autonomous system number while allowing you to use a "local-AS" number to establish a peer relationship. The local-as statement can be used to simulate the AS number already in place of the network fabric.

---

### **Add spines as a day-2 operation (RFE-1807)**

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

You're now able to increase the number of spines in a deployed 3-stage or 5-stage blueprint.

---

### **Add RBAC functionality to Freeform Blueprints (RFE-2545)**

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

Add support of RBAC for Freeform blueprints with Freeform specific roles.

---

### **Intent-Based Analytics probe to monitor switch environmental health for Juniper devices (RFE-2283)**

**Feature Category:** Telemetry and Analytics

Apstra now supports an IBA probe for monitoring environmental health metrics on Juniper devices including:

- Power Supply count and status
  - Power Supply Fan status
  - Switch Fan count and status
  - Power Supply temperature
  - Switch Temperature
  - Airflow direction
- 

### **Intent-Based Analytics probe to monitor health of telemetry services (RFE-2306)**

**Feature Category:** Telemetry and Analytics

Apstra now supports an IBA probe for monitoring the health of telemetry services and be alerted when any telemetry service on any devices is starting to perform poorly. The probe reports on

- Execution time and any degradation over time
- Sustained execution failures
- Sustained execution timeouts



- Sustained execution overruns.

---

## **Supported upgrade paths to 4.1.2 (RFE-2488)**

### **Feature Category:** Platform

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

---

## **Subnet-based whitelisting for UI access (RFE-2290)**

### **Feature Category:** Platform

As an Apstra Admin, you can now allow specific IP subnets or hosts to access the Apstra GUI.

---

## **Platform resource shortage warnings (RFE-2604)**

### **Feature Category:** Platform

Apstra Server platform shortage in disk space and memory will now generate warnings and errors, and be clearly indicated in the GUI for all users. Apstra Server will enter into Read-only Mode if the resources are critically short and users will not be fabric changes until addressed by Administrators.

---

## **Force delete a Device Agent (RFE-2365)**

### **Feature Category:** Platform

In 4.1.2, you will be presented with a new option to force delete an agent that has lost communication with the remote device.

---

## **Increase Numer of Super Spines and Uplinks (RFE-2501)**

### **Feature Category:** Apstra CLI

Using Apstra CLI, you can now update the number of superspines for a given plane, one plane at a time. Additionally, you can decrease the number of superspine uplinks from a given plane to a certain pod. This is for example useful for seamlessly increasing the number of superspines while reusing existing spine uplinks.

For more information, see the Apstra CLI help menu for the updated commands: "blueprint --blueprint superspine-config update --help", and "blueprint --blueprint spine-config update --help".

---

## **CLI command for Junos config check (RFE-2607)**

### **Feature Category:** Apstra CLI

Apstra-CLI now provides a command to executes a config-check on a Juniper based system, allowing you to verify that a candidate configuration is acceptable by the switch before committing the blueprint. This is useful in the context of Freeform blueprints as well as Datacenter ones where configlets are used.

- CLI command for Datacenter Reference design blueprints:
- CLI command for Freeform blueprints:

## **Changed Features**

### **Apstra provide device, chassis, and linecard profiles are immutable (RFE-1455)**

#### **Feature Category:** Device Profiles

The Apstra provided device profiles, chassis profiles, and linecard profiles are no longer able to be modified. This is to prevent user modifications from being in conflict and potentially wiped out during upgrades when applying updated product device profiles, chassis profiles, and linecard profiles. Users will need to create a copy of the desired device/chassis/linecard profile to make any edits.

---

### **Support of advanced Data Structures in Property-Sets (Lists and/or Dictionaries) (RFE-2388)**

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

Datacenter Reference design Blueprints now supports advanced Property-Sets allowing the definition of lists, dictionaries or any nested combination of.

---

### **Prevent User-Created Resources From Having Duplicate Labels (RFE-665)**

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

The Apstra UI now prevents users from creating resources (templates, resource pools, logical device maps, etc) with a duplicate label or display name improving the product usability.

---

### **Freeform full screen canvas optimizations and improvements (RFE-2688)**

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

Users can now view, edit, and add nodes plus links while in full screen mode for the Freeform canvas editor.

---

### **Enhanced search inside Config Template (RFE-2675)**

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

Searching inside Config Template in Freeform now supports scrolling between different results where previously it was only showing the the number of results.

---

### **Enhance uncommitted changes view for Freeform Blueprints (RFE-2549)**

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

Uncommitted changes in Freeform blueprints now hilights added and removed lines for more clarity.

---

### **Apstra event log now includes syslog config changes and user modifications (RFE-2599)**

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

The Apstra event log now contains information about syslog config changes and any user management modifications. Go to Platform and Event Log to see logs of new events.

---

### **Add support of YAML in Property-Sets (RFE-2571)**

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

Property-Sets now accepts YAML format in addition to JSON format. Capabilities are the same irrespective of the selected format. The addition of YAML supports facilitates migration of customer's existing key/value data structures from existing tools over to Apstra's Property-Sets. This is available for both Datacenter and Freeform blueprints.

---

### **Ability to use .(dot) in the node label in UI (RFE-2664)**

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

Allowed you to use " . " dot syntax for node names in the UI.

---

### **1G support on SONiC devices (RFE-2383)**

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

You can now support 1G interface speeds on SONiC devices. Some devices such as S5248, S5296 and N3248 requires the speed to be changed in groups of 4 instead of on a port-basis, which is now supported.

---

### **Intent-based Analytics probe for Optical transceivers now supports Juniper 400G interfaces (RFE-2266)**

**Feature Category:** Telemetry and Analytics

Apstra now supports monitoring of 400G interfaces for Juniper devices through the Optical transceivers Intent-based Analytics probe.

---

### **Enhancements to Optical transceivers IBA probe (RFE-1446)**

**Feature Category:** Telemetry and Analytics

Monitoring of optical metrics anomalies (Warning and Alarm) now distinguishes the High and Low variant of each anomaly. The thresholds are also represented on line graphs when selecting Times Series data so you can monitor the evolution of the monitored metric with respect to its four vendor defined thresholds: High\_Warn, Low\_Warn, High\_Alarm, Low\_Alarm. The probe also reports information on the transceiver such as: Vendor name, Vendor model, Part Number and Serial Number.

---

**Enhancement to interface names section for the IBA probes with "Specific Interfaces" input (RFE-2235)**

**Feature Category:** Telemetry and Analytics

When instantiating "Hot/Cold Interface Counters (Specific Interfaces)" or "Interface Flapping (Specific Interfaces)" you get the list of active interfaces for any added device as well as the remotely connected system. You no longer have to independently search for the interface names based on the connected device.

---

**Add support of "OR" operation in "Query Tag" Filter for IBA probes (RFE-2224)**

**Feature Category:** Telemetry and Analytics

You can now use the "Query Tag" Filter in IBA probe to express and "And" or an "Or" condition for Tags based selection. This increases the granularity in the probe's customisation.

---

**Increased default password complexity rules for Web and CLI users (RFE-2291)**

**Feature Category:** Platform

Default complexity for passwords has been increased in 4.1.2 but remains customizable by Apstra admins. This change won't affect current passwords and be enforced only for new deployments/new passwords.

---

**Support NSX-T Manager version 3.2 with DVS mode (RFE-2403)**

With Apstra 4.1.2, you will be able to create Virtual Infra Managers for NSX-T Manager version 3.2.x using DVS mode, as well as adding multiple Virtual Infra Managers per blueprint. This is useful in case you have multiple NSX-T Managers or multiple vCenter Servers hosted in the same

fabric blueprint.

The workflow to add new Virtual Infra of type "NSX-T" will now ask you to also provide the vCenter compute managers information (address and credentials), see the User Guide for more information. Note that NSX-T version 3.0 with NVDS mode is no longer supported in Apstra 4.1.2.

## 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.*

### **Dual Routing Engines support on Juniper modular devices (RFE-2362)**

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

Apstra now supports in Tech Preview Juniper Modular chassis with two management cards (Routing Engines). This includes updated support in ZTP for IP Address assignment and the corresponding pristine configuration.

Limitations:

- To benefit from Graceful-Restart, you need to enable it under "routing-options" for both default and non-default VRFs. Use a configlet to achieve that. Refer to
- NOS upgrades process remains unchanged. Both Routing Engine cards are upgraded simultaneously, NSSU support is not implemented.
- Graceful additions/removal of line cards are not implemented. Systems undeploy/redeploy will be required.

## **Fixed 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.

### **Resolution**

This issue is resolved in Apstra 4.1.2

---

### **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.

---

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

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

---

### **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.

---

### **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.

---

### **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

---

### **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.

---

### **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.

---

### **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.

---

## **Device pristine config becomes polluted after NOS upgrade failure (AOS-32484)**

When a NOS upgrade fails while installing, it is possible to re-install the device agent before the device has been rebooted on the new NOS version. This can cause a polluted state after the device is rebooted to the new NOS version.

---

## **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.

### **Resolution**

EVPN VxLAN type-5 probe fixed in Apstra 4.1.2 resolved probes expectations in the routing tables of ESI connected leafs.

---

## **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.

---

### **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.

---

### **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".`

---

### **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.

---

### **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.

---

### **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.

---

### **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.

---

### **IPv6 ECMP Multipath added to Juniper configuration (AOS-33192)**

A missing 'multipath' command is added to the IPv6 RIB underneath routing-options for all VRFs in 4.1.2 which permits ECMP for IPv6 routes. This will allow the router to load IPv6 traffic through multiple equidistant interfaces.

---

### **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".

---

### **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.

#### **Resolution**

On systems running Junos, the EVPN\_IMPORT policy-statement used for custom import & export-targets on routing zones has been modified to replace the 'accept' with 'then next term' statement when matching communities. This change permits the entire EVPN\_IMPORT policy to

be processed and resolves an issue where multiple routing-instances could not import routes from others on some EVPN import/export topologies.

---

### **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.

---

### **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.

---

### **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.

---

### **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)".

#### **Resolution**

The default `commit_timeout` value for Junos is changed from 60 seconds to 120 seconds in Apstra 4.1.2.

---

#### **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.

---

#### **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.

---

#### **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.

---

#### **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.

---

#### **Upgrade NOS Job to 4.27.4M Failed in Arista Device (AOS-38926)**

When Apstra NOS Upgrade to 4.27.4M is executed for Arista Device via Apstra UI, the upgrade job can end in a "Failed" state because the service configuration push operation happens too early before the rebooted device with new image finishes programming hardware resources and becomes fully operational when the device's pristine configuration has `hardware tcam \ system profile vxlan-routing`.

---

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

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

---

### **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.

---

### **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.

---

### **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.

---

### **When doing a Full Config Apply on SONiC, FRR restart may delay, which may cause a config deviation anomaly (AOS-27385)**

When doing a Full Config Apply on SONiC, the AOS agent issues a general `config reload -y` in the device, which will restart all containers, including the `bgp` container that contains the FRR

routing daemon. Restarting FRR may, in certain configurations, actually delay by up to a minute. By the time FRR restart finishes, the AOS Device Deployment Agent will already have collected the device configuration, which will contain only a minimal FRR configuration, e.g.:

```
----- vtysh -c "show run" -----
expectedRunningConfig:
log syslog informational
!
router bgp 64512
!
```

After a minute, the correct and full FRR configuration will appear, which will cause a config deviation anomaly for the device, since the comparison between the current full and the previous incomplete configuration will obviously come out negative.

## Resolution

In AOS 4.1.0, the verification procedure for the device configuration will be made to account for this eventuality and reject the partial FRR configuration, causing a retry until the full FRR configuration appears.

## Fixed Third-Party Issues

### **ACX7100 on Junos 22.2R2-EVO Switches Incorrectly Modify Destination IPv6 Address for Inner VxLAN Routed Packets After EVPN Trombone Routing (AOS-34809)**

If a VTEP sends an EVPN Type5 routed packet to a remote leaf, and that remote leaf then uses an EVPN Type2 route for the destination and has to re-encapsulate to another remote vTEP, the ACX7100 may rewrite the inner IPv6 address of the destination IPv6 address incorrectly within the VXLAN packet. This flow would be seen for something like leaf3 -> leaf1 > leaf2, in which the packet is rewritten from leaf1 > leaf2. For example, route a05:fab:4::/64 is known on leaf3 via two VTEPs on leaf1 & leaf2 which are both hosting the destination VN. The destination IPv6 address a05:fab:4::aaaa:aaaa:aaaa:aaaa is known on leaf2. If leaf3 -> a05:fab:4::aaaa:aaaa:aaaa:aaaa selects leaf1 for ECMP, the packet arrives in on leaf1 which then does a Type2 lookup to the destination a05:fab:4::aaaa:aaaa:aaaa:aaaa. When leaf1 puts the packet on the wire back to leaf2, it will rewrite the last part of the IPv6 address to a05:fab:4::aaaa:aaaa:aaaa:0 - causing traffic failure. See Junos bug PR1695877 for more information.

---

### **Apstra Using Cisco NX-OS Deprecated "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.

---

### **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.

---

### **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.

---

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

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

---

### **For EX4300 platform Fan speed is shown as 0 (AOS-34373)**

EX4300 does not support 'show chassis fan' CLI. Fan speed information is not available in any other CLI. Because of this in environment probe fan speed for this platform is shown as 0 RPM

---

### **Junos 21.2R3.8 False Positive VxLAN Floodlist Anomaly (AOS-30712)**

Apstra users running or upgrading to Junos 21.2R3.8 may experience false-positive anomalies from the Apstra IBA VxLAN Floodlist probe. Due to a Junos bug, incorrect telemetry is received in Apstra from Junos causing the anomaly. The Junos forwarding plane programming is correct and the user should not see network traffic impact.

---



### **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.

---

### **SONiC Hostcfdg may not respond to CONFIG\_DB key events (AOS-32963)**

The hostcfdg daemon in SONiC Cyrus up to 4.0.2 may not respond to CONFIG\_DB key changes. For example, issuing the command "config vrf del mgmt" may not delete the mgmt vrf. This may cause Apstra agent jobs (install, uninstall, NOS upgrade) to fail unexpectedly.

#### **Resolution**

This defect has been corrected in SONiC 4.0.3. Please refer to vendor bug SONIC-64655.

---

### **SONiC Supports Layer3 Subinterfaces Only in Trident 3 and Trident 4 Devices (AOS-36354)**

Device subinterface config is supported only on TD3x7(BCM56870), TD3x5(BCM56770) and TD4 (BCM56880) devices for SONiC.

---

### **When a Junos device is unassigned from the blueprint, the hostname is not deleted (AOS-27871)**

When a Junos device is unassigned from the blueprint, the hostname remains and you will need to manually modify it from the device CLI.

## **Known Apstra General Issues**

### **"Device Telemetry Health" Probe May Not Raise Anomalies (AOS-35948)**

If the "optical\_xcvr" service fails on Cisco NXOS devices, the "Device Telemetry Health" probe will not raise anomalies for these sustained telemetry failures.

---

### **"Virtual gateway needs to be enabled" Build Errors After Upgrading to Apstra 4.1.2 (AOS-35094)**

After upgrading to Apstra 4.1.2, when using the Junos platform for ESI racks, the user must enable virtual gateway on a layer3 virtual network. If the user has any VNs without the virtual gateway enabled, they will have the following build error:

```
Virtual gateway needs to be enabled on junos platforms for the VXLAN L3  
Virtual network (name) on ESI rack
```

#### **Workaround**

The user must enable the virtual gateway to resolve the build error and commit after upgrading to Apstra 4.1.2.

---

### **'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.

---

### **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 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 recovers.

## 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 when Virtual Infra Manager is removed and then added (AOS-45546)

When using AOS <= 4.2.1.x and removing & adding a Virtual Infra Manager (vcenter / nsxt), the Apstra backend database SysDB will have an agentId mismatch within entities related to Virtual Infra (Vcenter/ nsxt). When these entities are present in the Sysdb database, Apstra SysDB will crash repeatedly, rendering the Apstra GUI inaccessible.

## Workaround

Please contact JTAC for support to clean sysdb with provided patch  
esc\_339\_generic\_all\_versions.run

---

## 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 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 Failure 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.

---

### **Arista EOS BGP Unnumbered Not Enabled in Non-default VRFs (AOS-25041)**

On Arista EOS devices, BGP Unnumbered sessions are not enabled in non-default VRFs

---

### **Arista 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 Graceful Restart Not Enabled by Apstra (AOS-35171)

Apstra only configures graceful-restart under protocols bgp. Apstra does not configure graceful-restart under routing-options for the default routing-instance. Due to a Junos bug, graceful-restart does not work with non-default routing instances.

## Workaround

Configure Junos system set configlet to enable graceful-restart under routing-options in default routing-instance.

```
set routing-options graceful-restart
```

---

## BGP Session Flap IBA Probe With Dual Routing Engine Devices (AOS-35924)

The Apstra BGP Session Flap IBA Probe does not show BGP flaps for a node when BGP restarts gracefully with dual routing engines, which results in BGP flaps for adjacent peers. This is expected as per the current implementation. The BGP session goes down and comes up, but the forwarding state is kept unchanged Junos doesn't regard it as a flap.

---

## Blueprint Differences May Remain After Upgrading (AOS-36653)

After upgrading to Apstra 4.1.2, raw blueprint differences in the blueprint "Uncommitted" tab may remain after reverting an "add pod" operation after upgrading to Apstra 4.1.2. This is a cosmetic error and can be ignored. The issue will go away after the next blueprint commit.

---

## Blueprint Node Post API With Invalid JSON String Causes BlueprintCacheProducerAgent Crash (AOS-37977)

When blueprint Node POST API is called to create logical device in a blueprint with an invalid JSON string value, Apstra didn't validate the JSON string and responded with success. The invalid JSON string, saved into the Apstra graph database, later triggered BlueprintCacheProducerAgent to crash continuously because of JSON parsing error.

## **Workaround**

Node Patch API call is made to fix invalid JSON string saved into Apstra graph database. The customer must verify JSON string embedded inside the payload before calling POST/PATCH API call.

---

### **BlueprintDiffProducerAgent Crash When System Hostname Is Empty and Importing Configlet via Hostname (AOS-37837)**

When configlet is imported with hostname based condition, configlet assumes that all systems should have hostname information. When hostname is set to null via clicking clear button in the UI while editing hostname field in the system, BlueprintDiffProducerAgent will crash due to missing hostname.

## **Workaround**

Set the hostname via UI or using API via PATCH /api/blueprints/{blueprint\_id}/nodes/{node\_id} with data {"hostname":""}

---

### **BuilderAgent Crash by changing label of global anycast vtep interface node via REST API (AOS-40762)**

BuilderAgent uses only one automatically generated label for global anycast vtep interface node per blueprint. If the label is changed or modified to different value by API call, it can make more than one entries exist and then cause BuilderAgent to crash from violation of one unique entry.

## **Workaround**

Delete global anycast vtep interface node with non-default value via Blueprint Nodes Delete API call.

---

### **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.

---

### **Bulk Operation for VXLAN virtual networks without "Reserve across blueprint" configuration makes manually assigned VLAN ID changed (AOS-39517)**

For the VXLAN virtual networks which don't have a "Reserve across blueprint" configuration, Bulk operation to add/delete those virtual networks to/from multiple nodes makes manually assigned VLAN ID for each virtual network to be changed to a new VLAN ID from VLAN Pool.

#### **Workaround**

Use export and import virtual network feature at first for bulk operations. Export virtual network, assign the right VLAN ID to reserved\_vlan\_id column for each virtual network row in the exported CSV file, and import the updated CSV file back. After importing, use assign virtual network function icon for bulk operations from UI.

---

### **Cabling Anomalies From Juniper Dual Routing-Engine Switchover (AOS-35771)**

When a Juniper Dual Routing-Engine has a routing-engine switchover, the LLDP process will restart, causing temporary cabling false anomalies in Apstra.

---

### **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'
```

---

### **Clone Operations on EX4400 Device Profiles Removes Connector Type (AOS-36848)**

If an attempt is made to clone the device profile of a device of the Juniper EX4400 family, the Apstra UI will not be able to recognize the connector types in the original device profile and will leave the connector type empty for all the affected ports. The clone cannot be created unless the connector types for every affected port have an acceptable value.

### **Workaround**

The user can manually fill in the missing connector type in each port to an acceptable value, which can be cumbersome. Instead of the original RJ-45, the user can select rj45 and instead of lgbasex the user can select SFP.

Alternatively, the user can contact Juniper Support for instructions on how to circumvent the Apstra UI and use the Apstra API directly to carry out the clone operation.

---

### **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'].

---

### **Continuous MLAG Anomalies With Cisco NXOS Offbox Agent (AOS-38836)**

When servers, connected to VPC switches via MLAG interfaces, are deployed, Apstra raised MLAG anomalies continuously for MLAG interface state as unknown even if it's up state. The issue is related with parsing error of Cisco VPC command output. It can happen only if Cisco device agent is used with offbox agent and servers connected to MLAG interfaces are deployed mode to track link status.

---

### **Deleting Routing zone fails with "Protocol endpoint for protocol session is orphaned" error message (AOS-43808)**



After a CT (connectivity template) with dynamic BGP peering and BGP Prefix Dynamic Neighbor information is assigned to the SVI interface for a system, if the system is removed from the virtual network later, the CT becomes unassigned status, which allows the user to delete the CT. After the CT is removed later, protocol\_session becomes orphaned from the associated CT. it can lead to failure in deleting the routing zone.

### **Workaround**

Deleting protocol\_session via Blueprint Node Delete API or before modifying the virtual network for pruning system, update CT's assignment at first.

---

### **Device Environmental Checks IBA Probe Not Working With Juniper EX3400-48T (AOS-38598)**

Apstra users will see sustained execution failure anomalies from telemetry health for Juniper EX3400-48T devices after instantiating the Environmental Checks IBA Probe.

### **Workaround**

Contact Juniper Support.

---

### **device-profile-clone POST API Missing Validation (AOS-33606)**

In the device-profile-clone API POST payload, accept a new field for original\_device\_profile\_id. Validate that the device profiles in the original interface maps listed in "clone\_interface\_maps" match this original\_device\_profile\_id.

This is to make it airtight that the interface maps being cloned are using the same device profile as the one that is being cloned.

Currently, the Apstra UI is not vulnerable to this potential mismatch. However, with direct API POST, this discrepancy can potentially exist and go unnoticed, causing unexpected catastrophic side effects such as an Arista device profile getting swapped out for Juniper in the cloned IM.

---

### **DeviceTelemetryAgent Process' continuous crash in the offbox agent (AOS-37597)**

It is presumed that the aos\_cache file contains data if it is there when the DeviceTelemetryAgent process starts. There is a situation where the file is left empty due to an unexpected agent process crash. It makes the agent process crash repeatedly with an exception when it restarts and attempts to read data from the empty file.

## Workaround

Please follow the below steps in the Apstra Controller.

```
1. verify empty cache files
   for file in $(find /var/lib/aos/tmp/task/offbox/ | egrep
'aos_cache.*.json') ; do ls -la "$file" ; done
2. If 0 byte size cache files exist run the command below to remove them.
   for file in $(find /var/lib/aos/tmp/task/offbox/ | egrep
'aos_cache.*.json') ; do echo "$file" ; rm "$file" ; done
3. reset TAC Backoff timer for Offbox agents
   for offbox in $(docker ps | egrep aos-offbox | awk '{print $1}' ) ; do
docker exec $offbox tacspawnerctl resetBackoff --spawnerAddress
tbl://aos_spawner_sock ; done
```

---

## Duplicate Entries shown in the virtual infra inventory (AOS-47478)

The transport VLAN node would remain uncleared in the Graph Database while the Virtual Infra Manager for NSX-T manager with unreachable vCenter or without vCenter is created and added to the blueprint, and then removed from the blueprint and deleted. The same transport VLAN node would be re-created, and duplicate entries would exist if the identical virtual Infra manager for NSX-T was created and added back to Blueprint again.

## Workaround

restart AOS service via `sudo service aos restart` and then stalled duplicate entries will be automatically cleared

---

## EVPN IBA Telemetry probe is deprecated and incompatible with the latest NOS versions (AOS-52013)

The EVPN IBA Telemetry probe was originally implemented in older Apstra versions (3.x) but has been deprecated in later versions due to advancements in telemetry collection frameworks. Customers upgrading from older versions (3.x) to Apstra 4.x and then to later might encounter issues with the EVPN IBA Telemetry probe as it may no longer function as expected.

The primary reasons for this are:

```
1. Deprecation of Older Collectors: The EVPN IBA Telemetry probe and
associated collectors were designed for earlier NOS versions. These collectors
```

relied on commands and logic that are no longer supported by certain vendors (e.g., with EOS 4.25.3.1M, evpn\_type3 collector failed to gather data due to a plugin error triggered by an unsupported command).

2. Python Dependency Transition: Apstra versions prior to 4.2.x were based on Python2, whereas 4.2.x and later versions have transitioned to Python3. This shift in underlying architecture breaks backward compatibility for older telemetry packages.

3. Multiple Breakpoints in Compatibility: Over several releases, the extensible telemetry framework has undergone significant updates, requiring users to update their custom packages to align with newer standards and dependencies.

Since the EVPN IBA Telemetry probe is deprecated, customers experiencing issues with it are recommended to consider using standard predefined probes, such as EVPN Type-3/Type-5 Route Validation, while being mindful of their limitations.

Limitation: For the predefined Type-3 and Type-5 probes to work, all leaf nodes in the topology must be of the same platform. Mixed-vendor environments are not supported by these probes.

## Workaround

If the topology consists of all leaf nodes in the blueprint, users can perform the following actions to mitigate issues with the deprecated EVPN IBA Telemetry probe:

1. Perform Apstra upgrade, post upgrade the system enters maintenance mode
2. Delete EVPN IBA Telemetry Probe from all blueprints
3. Remove any related custom telemetry packages added to the agent profile
4. Change Maintenance mode to Normal to resume regular operation
5. If onbox, initiate onbox install for all agents
6. Post upgrade, start using standard predefined probes such as EVPN Type-3/Type-5 Route Validation.

The EVPN IBA Telemetry Probe has been deprecated. Customers experiencing issues with this probe are recommended to replace it with the standard predefined probes, such as the EVPN Type-3/Type-5 Route Validation probes. For these probes to work properly, all leaf nodes in the topology must be from the same platform, as mixed-vendor environments are not supported. To address the challenges of migrating from the EVPN IBA Telemetry probe to the Type-3/Type-5 probes in mixed-vendor leaf node environments, an enhancement request ( : [IBA][Probe] Expand EVPN probes to support mixed-vendor leaf nodes) has been submitted as a long-term solution. To prioritise , please contact sales/PLM.

---

## EX4400 models in Apstra-unsupported VCP mode have different interface naming for PIC 1 (AOS-43264)

It is noted that if an EX4400 device is set to VCP mode, the interface numbering of PIC 1 will be different from what a any existing builtin Apstra Device Profile for an EX4400 device has.

Specifically, et-0/1/1 does not exist in VCP mode, but instead becomes et-0/1/2.

Please note that Apstra version 4.2.1 does not support any Juniper EX4400 model in VCP mode.

---

### **Freeform 'Interfaces' Dictionary Element 'neighbor\_interfaces' Schema Changed From 4.1.1 -> 4.1.2 (AOS-34249)**

The freeform 'interfaces' device schema for Jinja configuration rendering has been altered from 4.1.1 -> 4.1.2.

#### **Workaround**

Apstra is not able to upgrade user-defined jinja automatically, so this must be done manually.

- 1) Execute AOS Upgrade using the command line argument --ignore-config-validation-results
  - 2) Post-upgrade, the blueprint will have config rendering anomalies, and deployment may fail â€œ it will not push any configuration to devices due to a jinja rendering error.
  - 3) Modify the user-defined jinja to accommodate the new modeling semantic. Underneath the 'interfaces' dict key, 'neighbor\_interfaces' has been changed from a dictionary to a list of dictionaries representing multiple remote interfaces to support N+1 LAG.
  - 4) Deploy the blueprint
- 

### **GBP Policies on Juniper Now Include Annotations Containing the Policy Descriptions (AOS-33532)**

Configuration for Junos GBP (Access-list) policies in 4.1.2 now include annotations for individual firewall-filter rules, indicated via /\* \*/.

Note that annotations will not appear in incremental config diff previews.

---

### **High interface hold timer value rendered in the Collapsed fabric reference design may affect PXEboot (AOS-42437)**

Customers may observe servers on a collapsed fabric failing to PXEboot where interface is rendered with a large hold time for up event as part of the collapsed fabric reference design

#### **Workaround**

Use a configlet to reduce the interface hold-timer.

---

## **IBA Hypervisor MTU Threshold Check Anomaly Telemetry Error (AOS-37312)**

The Apstra user may get IBA Hypervisor MTU Threshold Check related anomalies, with the pnic related traceback "KeyError" messages in DeviceTelemetryAgent.

### **Workaround**

To work around this issue, the user will need to modify the Hypervisor MTU Threshold Check IBA probe, Hypervisor low MTU anomalies stage to have stage output with "Pnic Name".

---

## **Importing CSV file for virtual network fails with error "Invalid CSV header order" (AOS-47014)**

Importing virtual networks from a CSV file fails if the `bound_system_` column header, where a virtual network is bound, contains parenthesis or bracket characters. These characters may originate from the system label and are not permitted by CSV header validation.

### **Workaround**

Please remove parenthesis or bracket characters from system's label

---

## **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.

---

## **Incorrect Incremental Device Configuration When Replacing Device Profile With a New Line Card (AOS-42128)**

When the user replaces an Interface Map with an updated modular chassis Device Profile with removed line card configuration, Apstra renders the correct, complete configuration. However, the incremental device configuration may be incorrect, leaving the interface configuration for ports on the removed line card.

---

## **Incorrect Interface Naming for Juniper EX4400-24MP EX4400-48MP Device Profiles (AOS-41912)**

Apstra device profiles for Juniper EX4400-24MP EX4400-48MP device will incorrectly name device "mge" interfaces (e.g. mge-0/0/0) as "ge" interfaces (e.g. ge-0/0/0).

### **Workaround**

The user can manually clone the device profile and correct the interface names in the new device profile. Contact Juniper Support for assistance.

---

## **incorrect routing policy applied when assigning/unassigning endpoints in a CT with multiple BGP peerings and distinct routing policies (AOS-51549)**

When a Connectivity Template (CT) includes a single Virtual Network (VN) primitive, multiple BGP peering primitives, and distinct routing policies, incremental configuration changes can occur during endpoint assignment and unassignment. These operations may unexpectedly swap or alter import/export routing policies, potentially disrupting routing configurations and causing traffic interruptions on commit.

In CTs with multiple BGP peering primitives, all BGP primitives are managed under a Batch policy. Apstra does not guarantee the execution order within a Batch Policy, especially during unassign and assign operations, where resources are allocated from a pool, and assignment order is unpredictable. This can lead to the unexpected swapping of routing policies.

It's important to note that this issue occurs only when the CT contains multiple BGP peering primitives with distinct routing policies. CTs with a single VN primitive and a single BGP peering primitive (and associated routing policies) do not experience this behavior.

### **Workaround**

To mitigate this issue, the following workaround is recommended:

1. Delete the distinct routing policies from the Virtual Network primitive of the affected Connectivity Template (CT).
2. Create separate Connectivity Templates (CTs) for each routing policy, ensuring that each CT corresponds to one routing policy.
3. Assign each CT to the appropriate protocol endpoints.

This workaround will help avoid the unexpected swapping of routing policies during endpoint assignment and reassignment. Please contact Apstra Support Team for more information

---

### **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.

---

### **Jinja Custom Filters Unavailable for Apstra Configlets (AOS-37781)**

Jinja custom filters listed in the Jinja documentation are unavailable for Apstra Configlets. Making use of these will raise blueprint property set validation errors.

---

### **Jinja Extensions jinja2.ext.do and jinja2.ext.loopcontrols Unavailable for Apstra Configlets (AOS-37559)**

Jinja extensions jinja2.ext.do and jinja2.ext.loopcontrols are not available for Apstra Configlets, raising template syntax errors when used.

---

### **Jinja Functions Unavailable for Apstra Configlets (AOS-34661)**

Jinja functions listed in the Jinja documentation are unavailable for Apstra Configlets. Making use of these will raise blueprint property set validation errors.

---

### **Juniper PTX10000 Additional Configurations Required for Tunnel Termination (AOS-34677)**

In Apstra set `forwarding-options tunnel-termination` needed to be added by Junos Configlet for fabrics containing leaf PTX10000 devices to advertise type 5 EVPN routes.

---

### **Juniper QFX10008 may fail to report rpm fan telemetry (AOS-40731)**

Juniper QFX10008 may fail to report fan telemetry causing telemetry issues within Apstra.

#### **Workaround**

Upgrade to Apstra 4.2.0 so that fan rpm telemetry failure is handled correctly.

---

### **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 EVO Leafs Flooding DHCP Messages (AOS-33697)**

Due to a bug in Junos 22.2-EVO, DHCP discover messages are incorrectly flooded across EVPN fabric.

---

### **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.

---

**Logical diff section continues to display link changes, even if configlet based on tag is removed (AOS-49983)**



Despite the configlet being applied to some nodes based on tags, there were some link changes in the logical diff section. The logical diff tab continued to display the changes even after the configlet was reverted, and there was nothing to commit in the uncommitted tab section.

### **Workaround**

The current diff plugin is not handling the system tag relationship , so it is not able to compute the difference. The workaround to restart the AOS services.

---

### **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.

---

### **No Validation Error While Creating Virtual Infra Manager (AOS-35959)**

When the user creates a Virtual Infra manager by either providing an incorrect IP address or an incorrect username/password for vCenter or NSX-T, Apstra will not raise any validation error. Just the state shows disconnected.

### **Workaround**

The user will need to correct the IP or credentials if the Virtual Infra manager is disconnected.

---

### **Node ID for Virtual Infra Manager Device Is Empty in Liveness Anomalies (AOS-36792)**

Node ID for virtual infra manager device is empty in liveness anomalies. The system ID and IP or hostname will be listed in the virtual infra section.

---

### **NSX-T 3.2 Integration Does Not Handle LAG Mode (AOS-38557)**

The Apstra NSX-T collector may crash because it gets a "none" value when the VDS uplink mode is configured as LAG mode.

### **Workaround**

The Apstra NSX-T collector will need an update to handle this scenario. Contact Juniper Support.

---

### **OffboxAgentManager crashes continuously after restoring backup (AOS-37265)**

After restoring backup in the new AOS vm, OffboxAgentManager crashes are seen. This issue is related with system Key difference between old AOS server where backup (taken by show tech) was executed, and new AOS server where restore is executed. This issue can happen for not only OffBoxAgentManager but also OnBoxAgentManager as long as Agent Profile's credential is used.

### **Workaround**

Please contact JTAC Apstra support.

---

### **Overlapping VNI Values Causing Error Messages to Not Show (AOS-42824)**

After a VNI is manually assigned to security zone, if the VNI value belongs to the dynamic VNI pool for Virtual network, Virtual network may be assigned with the same VNI value. When this duplicate VNI error condition occurs, UI doesn't show error messages correctly.

### **Workaround**

Fix duplicate VNI issue by assigning unique static VNI value, which shouldn't come from dynamic VNI pool where virtual network uses, to Security Zone.

---

### **Platform ACL Does Not Allow Loopback and Docker Networks (AOS-44009)**

When the Platform ACL feature is enabled, and the default rule (0.0.0.0/0) is set to deny, the Apstra UI and system agents cannot make necessary REST API calls to the Apstra controller.

### **Workaround**

The user must allow access from loopback (127.0.0.0/8) and docker (172.17.0.0/16) networks.

---

### **Predefined External Route Probe Is Not Displaying External Routes From Different BGP**

## Peers (AOS-36501)

Due to an incorrect built-in graph query, the predefined external route IBA probe is not displaying external routes from different BGP peers toward external routers.

### Workaround

The user can update the graph query of the predefined external route IBA probe to correct.

```
node('system', name='system', role='leaf', system_id=not_none(),
deploy_mode=is_in(['deploy', 'drain']))
.out('hosted_interfaces')
.node('interface', if_type='ethernet')
.out('composed_of')
.node('interface', if_type='subinterface', name='system_if')
.in_('layered_over')
.node('protocol_endpoint')
.in_('instantiates')
.node('protocol_session')
.out('instantiates')
.node('protocol_endpoint')
.out('layered_over')
.node(type=is_in(['interface', 'ip_endpoint']), name='neighbor_if')
.ensure_different('system_if', 'neighbor_if')
.where(lambda neighbor_if: bool(neighbor_if.ipv4_addr or
neighbor_if.ipv6_addr))
```

---

## Previous Apstra Version Agent Logs Lost After Apstra VM to VM Upgrade (AOS-36365)

After VM to VM upgrade, the agent page job section shows the Job history from the previous Apstra VM with a view icon to download/view the file. Actually, agent job files are not copied in VM to VM upgrade, so the files do not exist in the current AOS VM. The Apstra UI is reporting the error, but the UI is not good.

---

## QFX10008 Junos Upgrade Fails When PristineConfig Has Chassis GRES Enabled (AOS-36054)

If the user is using Apstra to upgrade Junos on a Juniper QFX10008 and has "chassis redundancy graceful-switchover" (GRES) enabled, the Junos upgrade will fail.

warning: This operation is not permitted with GRES enabled. Please disable GRES and try again.

### Workaround

The user must disable GRES before installing the Apstra device system agent, so this configuration is not in the Apstra device pristine configuration.

---

### **Rack-based template is not shown in the selection during creating blueprint (AOS-47522)**

When a rack-based template is created, two fields related to the 5-Staged Clos architecture are Links per Superspine Count and Link to Superspine Speed. The Link to Superspine Speed can be set to a non-null value even if Superspine Count is set to 0. The template is recognized as a component template to create a pod-based template rather than an independent rack-based template if Link to Superspine Speed is set to a non-null value. Therefore, it is not shown in the pop-down field for template as a selectable choice during blueprint creation.

#### **Workaround**

Edit the rack template and remove the Link to superspine speed value

---

### **Reverting to Pristine Configuration Can Make the Device Appear to Have a Config Deviation (AOS-32950)**

If a deployed device is unassigned from its blueprint or even if the blueprint is altogether deleted and subsequently the device is reverted to its Pristine Configuration, then the device will appear as if its configuration has deviated from the expected one. This behaviour is cosmetic and has no functional impact.

#### **Workaround**

To make the deviation disappear, it is suggested to press accept changes in the config telemetry tab of the affected device(s).

---

### **Route Anomalies Triggered in Juniper Junos-EVO Device When PIM Is Enabled Into Interface by Configlet (AOS-39289)**

Route anomalies are reported in Juniper Junos-EVO device when PIM protocol is enabled via configlet function, even if all route entries correctly exists in the device. This issue is triggered by collector's parsing error during collecting routes in the Juniper Junos-EVO device.

#### **Workaround**

Disable route telemetry service for Juniper Junos-EVO devices via adding serial number of EVO

device into the field route\_disable\_devices in the builtin\_telemetry\_disable section of aos.conf file, followed by restarting Apstra service. Contact Juniper Support.

---

### **Route-targets with leading zeroes cause upgrade failures to Apstra 4.2.1.x due to schema validation changes (AOS-42738)**

Leading zeroes in route-targets, such as 0555:0555, were accepted in device configurations. However, the route-maps rendered from these configurations failed to react to the generated routes because the leading zeroes were stripped during route generation (e.g., 0555:0555 became 555:555). This mismatch caused the rendered route-maps to be misaligned with the generated routes, resulting in forwarding failures, telemetry collection errors, route installation issues, and routing policy mismatches.

To address this inconsistency, Apstra 4.2.1 introduces stricter schema validation that disallows EVPN RD type:value fields containing leading zeroes. While this change ensures consistent route handling, it may cause upgrade failures if existing configurations include route-targets with leading zeroes.

#### **Workaround**

Before upgrading to Apstra 4.2:

1. Identify route-targets with leading zeroes in all blueprints (Edit Routing Zone)
  2. Update affected Route Target Policies by removing leading zeroes (e.g., change 0555:0555 to 555:555).
  3. Review and commit the changes in all blueprint.
  4. Proceed with the VM-to-VM upgrade.
- 

### **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
```

---

## **Running aos\_show\_tech on SONiC Device in Telemetry Only Mode May Make Device Unresponsive (AOS-36855)**

Running the aos\_show\_tech script on a SONiC device running in Telemetry Only mode generates a very large checkpoint\_config.tar and maxes out the available space in /tmp, leading to an unstable system.

### **Workaround**

Apstra recommends not running SONiC devices in Telemetry Only mode. If this issue occurs, the system will need to be rebooted to recover.

---

## **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.

---

### SONiC Device ARP and Neighbor Suppression Enabled for L2 only VxLAN (AOS-42546)

For SONiC devices, even if layer2-only VxLAN is configured, ARP and neighbor suppression functions will be enabled, making ARP packets sent to the CPU, potentially causing traffic to be dropped.

---

### 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 Displayed MTU for Member of PortChannel May Differ From Real Configured MTU (AOS-44229)**

In Apstra 4.1.2, due to an error in the config\_db.json rendering, a member Ethernet interface of a PortChannel that has MTU different than 9100 (the default MTU for SONiC interfaces) can display an MTU of 9100, instead of the inherited MTU from the parent PortChannel. The real MTU of the Ethernet member is same as the PortChannel's (as can be seen via ifconfig), but displaying the MTU may show 9100 instead.

This bug can happen on an Apstra 4.1.2 controller and can also be carried over to an Apstra 4.2.1 or 4.2.2 controller via upgrade.

### **Workaround**

If the customer wants to display the correct MTU, they can initiate a full config apply in the Apstra 4.2.1 or 4.2.2 controller. Apstra 4.2.1 or 4.2.2 will render the correct MTUs for both PortChannel and its members. To avoid a full config apply, please ask support for a fixer script that can restore the correct MTUs in the SONiC configuration database without doing any real change.

---

### **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.

---

### **SONiC port speed changes for device models Dell S5248F-ON and S5296FON after upgrades to AOS 4.1.2 may fail (AOS-33739)**

AOS 4.1.2 introduces support for port-groups in SONiC running on device models Dell S5248F-ON and S5296FON. For the feature to function correctly, the first commit after upgrading to 4.1.2 must not be a speed affecting change.

#### **Workaround**

After upgrading to AOS 4.1.2 and only for devices that are Dell S5248F-ON or S5296FON, please create a trivial sonic system configlet, e.g. one executing the benign command "true" and apply. This will ensure that the first commit to be done to the device after the upgrade is not a speed change. After this is done, the new device model will contain the necessary information to effect speed changes, making them safe to use.

---

### **sustained error anomaly for Optical Xcvr collector in the Telemetry Health Probe raised for skipped interface (AOS-41017)**

If the interface does not support all of the required information for optical transceivers, the Optical Xcvr collector is unable to parse the information, so it attempts to skip the interface and logs the reason in the error. The mechanism increases the collector's error count, which then causes anomalies in the Device Telemetry Health probe. The publishing error mechanism was originally designed to detect real collector issues. Skipping interfaces that may not provide the necessary information should not result in collector errors.

---

### **The Range Check processor stage displays no data when the minimum anomalous value is set to 0.1 (AOS-49137)**

Floating-point precision discrepancies can cause problems in the integration between IBA and metricdb when configuring IBA probes. To be more precise, the live data that was obtained from IBA is queried using metricdb using a trie-based matcher. However, minor variations in floating-point values (such as 0.1 being read as 0.10000000149) could cause metricdb to fail to match the desired keys. This can cause probes to miss crucial data when querying specific values.

#### **Workaround**

None

---

### **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 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.

---

### **Virtual Network Validation Error 'Virtual gateway IP allowed only if IPv4 subnet specified' when IPv4 subnet as netmask and Virtual G/W address as static IP address (AOS-43352)**

When IPv4 subnet information is configured with netmask information in the Virtual network, Apstra assumes that Virtual G/W address should be dynamically provisioned from dynamic IP pool. If static Virtual Gateway IP address is configured together with netmask in the subnet field, it would trigger validation error not to use static Virtual Gateway IP address

#### **Workaround**

Assign static IPv4 block into IPv4 subnet field with static Virtual Gateway IP address or clear Virtual Gateway IP address in the Virtual Network.

---

### **VirtualInfraGraphAgent May Restart When VMware NSX-T Collection and Configuration Coincide (AOS-38903)**

If an Apstra user is using VMware NSX-T, a possible race condition may occur, which will cause the Apstra VirtualInfraGraphAgent process to restart, possibly causing temporary system anomalies. This occurs when a telemetry collection is in progress, and the user makes a

configuration change resulting in a possible inconsistent state. The Apstra VirtualInfraGraphAgent process will recover on its own, and virtual-infra functionalities will not be affected.

---

### **VMware NSX-T Collection Failures May Cause Issues Showing Latest Virtual Infrastructure Details (AOS-38899)**

Under certain, rare conditions, if there are VMware NSX-T telemetry collection failures and, during this time, a logical-switch-ID changes, Apstra may correctly receive updates from vCenter but will have missed updates from NSX-T. This will cause Apstra to be unable to correlate details for logical-switch-ID. This situation will cause Apstra not to correctly show the latest virtual infrastructure details. Apstra will show no errors and will not recover from the situation alone.

#### **Workaround**

To resolve this issue, the user can restart the Apstra virtual\_infra agent. For example ...

```
root@aos-server:/home/admin# ps -ef | grep -i virtual_infra
root 24017 21621 2 01:32 ? 00:00:05 /usr/bin/python /usr/bin/aos_agent --
class=aos.reference_design.extension.virtual_infra.agent.virtual_infra_graph_a
gent.VirtualInfraGraphAgent --name=VirtualInfraGraphAgent
root 24045 23967 0 01:35 pts/0 00:00:00 grep --color=auto -i 24017
root@aos-server:/home/admin#
root@aos-server:/home/admin#
root@aos-server:/home/admin# kill -9 24017
```

---

### **When connectivity template with custom routing policy is applied with BGP session, rendered config may create BGP session with VRF's routing policy instead of custom routing policy (AOS-43501)**

When connectivity template with custom routing policy with BGP session is unassigned, Apstra doesn't clean up old BGP protocol endpoint information correctly, used to determine the effective routing policy for each BGP session. These stalled information makes the following assigning CT operation to use VRF's policy instead of custom routing policy.

#### **Workaround**

Restart AOS service RenderConfigDiffAgent process. Contact Juniper Support.

---

**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.

---

### Worker node encounters continuous NodeFileMonitorAgent crash (AOS-46491)

The NodeFileMonitorAgent of the worker node encounters a SIGTERM signal, leading to continuous restarts approximately every 1-2 minutes. NodeFileMonitorAgent on startup scans /var/lib/aos/metricdb directory to create a file registry in SysDB. In Apstra 4.1.2, the implementation uses Python to scan the file system and instantiate file registry entities. It takes too long to process a large number of files, resulting in missing Keepalive heartbeat messages.

### Workaround

As a temporary workaround, to address the recurring NodeFileMonitorAgent restart issue, users are advised to remove files within the /var/lib/aos/metricdb/iba directory on the affected worker node as mentioned below:

STEP\_1: Login to Failed Worker Node

```
ssh admin@
```

```
sudo su
```

STEP\_2: Check Total Samples File Count from /var/lib/aos/metricdb/iba

```
total_file_count=$(find /var/lib/aos/metricdb/iba type f -name "samples *.tel" | wc -l)
```

```
echo $total_file_count
```

STEP\_3: Check Old File Count, start with deleting 15 days old files, reduce as needed

```
old_file_count=$(find /var/lib/aos/metricdb/iba ctime +15 -name samples *.tel | xargs -d '\n' ls -l | wc -l)
```

```
echo $old_file_count
```

STEP\_4: Delete 15 days old files, reduce as needed

```
find /var/lib/aos/metricdb/iba ctime +15 -name samples *.tel | xargs -d '\n' rm -rf
```

STEP\_5: Check Total Samples File Count from /var/lib/aos/metricdb/iba

```
total_file_count=$(find /var/lib/aos/metricdb/iba type f -name "samples *.tel" | wc -l)
```

```
echo $total_file_count
```

If total\_file\_count > 100000 (100k) files, then repeat STEP\_2 to STEP\_5 again.

Repeat until the total\_file\_count falls below 100000 (100k)

---

### **ZTP devices, which use python3, fails in getting ztp\_py3.py file via tftp (AOS-47007)**

In Apstra ZTP < 5.0.0, ZTP for Junos EVO devices would fail as the 'ztp\_py3.py' is not available over tftp to provision due to missing the right file permissions.

#### **Workaround**

```
chmod +r /containers_data/tftp/ztp_py3.py
```

---

### **ZTP Fails When Apstra ZTP VM Is Located in the Different Subnet From Device and NOS Upgrade Is Performed During ZTP Processing (AOS-40946)**

After NOS is upgraded via ZTP processing, during the 2nd phase ZTP processing, JUNOS device may not install default route entry populated into the default routing table even if router option is received by DHCP server. Apstra ZTP assumes that default route is installed in the routing table when DHCP router option is bound into management interface. Because default route entry is missing, Apstra ZTP process in the device can't retrieve ztp.json file from TFTP server which led into failing ZTP processing.

#### **Workaround**

The currently available manual workaround is to initiate zeroize device again after NOS is upgraded into the intended version via ZTP. Contact Juniper Support for an Apstra ZTP 4.1.2 patch.

---

### **ZTP for EVO Device Succeeds Even if Default Gateway Information Is Missing by DHCP (AOS-40147)**

Apstra ZTP process requires devices to learn default gateway information by DHCP. The default gateway information is used to configure the default static route entry in the management routing instance. ZTP processing for Juniper EVO device succeeds without failing, even if default gateway information is missing. After ZTP processing, Juniper EVO device doesn't have a default static route in the management routing instance and becomes isolated from the other networks.

#### **Workaround**

The first step is to correct DHCP configuration in the Apstra ZTP VM by adding DHCP option routers with the right default gateway information, followed by restart dhcpd (docker restart

dhcpcd). The second step is removing devices from Apstra Server and restarting the ZTP process for device (zeroize).

## Known Apstra Security Issues

### Apstra VM SSHd Terrapin Vulnerability (AOS-44494)

Default sshd ciphers and mac exchanges included as part of Apstra base OS are vulnerable to terrapin attack for chacha20-poly1305@openssh.com and etm mac exchange hmac-sha2-512-etm@openssh.com,hmac-sha2-256-etm@openssh.com,umac-128-etm@openssh.com

#### Workaround

To mitigate the Apstra Base OS SSHd Terrapin vulnerability, it is recommended that the most affected cipher and mac exchange be removed from the /etc/ssh/sshd\_config file, or upgrade to 4.2.1.1

```
Edit /etc/ssh/sshd_config:
sudo nano /etc/ssh/sshd_config
```

REPLACE:

```
Ciphers chacha20-poly1305@openssh.com,aes256-gcm@openssh.com,aes128-
gcm@openssh.com,aes256-ctr,aes192-ctr,aes128-ctr
MACs hmac-sha2-512-etm@openssh.com,hmac-sha2-256-etm@openssh.com,umac-128-
etm@openssh.com,hmac-sha2-512,hmac-sha2-256,umac-128@openssh.com
```

WITH:

```
Ciphers aes256-gcm@openssh.com,aes128-gcm@openssh.com,aes256-ctr,aes192-
ctr,aes128-ctr
MACs hmac-sha2-512,hmac-sha2-256,umac-128@openssh.com
```

Restart SSHd:

```
sudo systemctl restart sshd
```

---

### SSH Terrapin Vulnerability Workaround Using SSH aes128-gcm or aes256-gcm Ciphers Is Not Supported by Apstra Paramiko SSH Client (AOS-44336)

Apstra uses the Paramiko SSH client library to access Junos devices. The Apstra version of Paramiko does not yet support the SSH ciphers aes128-gcm@openssh.com and aes256-gcm@openssh.com. Access from Apstra to the Junos device will not function properly if it is set up to use just these SSH ciphers.

## Workaround

All Apstra versions can work with Junos devices via SSH with aes256-ctr cipher and hmac-sha2-256 or hmac-sha2-512 for hmacs, minimizing the impact of the SSH Terrapin vulnerability

Please use an Apstra system set configlet for Junos devices to configure SSH ciphers and MAC encryption, or upgrade to 4.2.1.1.

```
set system services ssh ciphers aes256-ctr
set system services ssh macs [ hmac-sha2-256 hmac-sha2-512 ]
```

## 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 NOS Upgrade fails in the Cisco NXOS device with rollback command failure (AOS-46278)

Apstra renders the configuration with `spanning-tree vlan <vlan-id> priority 8192` for each vlan defined in the Cisco NXOS 9.3(8), 9.3(10) device. For Apstra rendered spanning-tree priority configurations, the Cisco NXOS device generates configurations using a combined format that includes range statements. When a rollback command is executed for a NOS upgrade, full push configuration, or un-deployment, it may fail if the reverse command from rollback exceeds 512 bytes.

## Apstra Rendered Configuration

```
spanning-tree vlan 100 priority 8192
spanning-tree vlan 101 priority 8192
spanning-tree vlan 102 priority 8192
```

```
...  
spanning-tree vlan 2999 priority 8192
```

## Cisco NXOS Device Configuration

```
spanning-tree vlan 100-105,201-209,211-232,501-506,509,511-512,517,520,525-  
526,528,530,533,535,999,1030,1032-1036,1040,1043-1044,1046,1050,1053-  
1054,1056,1099,1101-1102,1104,1113-1114,1116,1123-1124,1126,1173-  
1174,1176,1203-1204,1206,1208,1211,1213-1214,1223-1224,1226,1228,1233-  
1236,1253-1254,1263-1264,1283-1284,1286,1303-1304,1306,1900,1991,  
2001-2002,2030-2031,2033,2040,2043-2044,2046,2050,2053-2054,2056,2093-  
2094,2099,2113-2114,2116,2123-2124,2126,2133-2134,2153-2154,2156,2173-  
2174,2176,2203-2204,2206,2223-22  
24,2226,2233-2236,2253-2254,2263-2264,2283-2284,2286,2303-2304,2306,2991,2999  
priority 8192
```

## Workaround

1. Login to the device and then execute the below commands (it may introduce service impact because STP priority changes back to default value)

```
switch# configure  
switch(config)# no spanning-tree vlan 1-3967 priority  
switch(config)# end
```

2. then do NOS upgrade, undeploy device, or full push configuration for the device in the Apstra UI

---

## Arista EOS Config Deployment Fails Due to Device TCAM Configuration (AOS-34631)

Arista EOS device may fail configuration deployment if the following team change is pushed by Apstra with the standard EOS "System" configlet.

```
hardware tcam  
system profile vxlan-routing
```

## Workaround

It is recommended that the following config be added to EOS devices before adding it to an Apstra deployment so that it is part of its pristine config.



```
hardware tcam
system profile vxlan-routing
```

---

### **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 Authentication With Prefix Peering Does Not Work on 22.2R2.12-EVO (AOS-37492)**

When creating Dynamic BGP Peering with Apstra Connectivity Templates with IPv4 and IPv6 AFI and IPv4 and IPv6 modes enabled, BGP sessions with password authentication with prefix peering do not work on Juniper Junos 22.2R2.12-EVO .

#### **Workaround**

Upgrade to Junos 22.4R1.11-EVO.

---

### **Changing interface speed and adding same interface to an aggregated ethernet interface in the same commit may fail in Junos Evolved (AOS-40063)**

If the member physical interfaces that are going to comprise an ae are initially at different speeds, naturally we have to change the speeds to the desired ae speed and then create the ae from the constituent interfaces.

If the changing of the speed and the creation of the ae is in the same commit, the commit may fail with "Interface aeX with child links of mixed speed but link-speed mixed is not configured". This has been observed in Junos Evolved.

## **Workaround**

Please change the speed of all interfaces to be the same in a separate initial commit, then create the ae interface in a subsequent commit.

As a general note, it is highly recommended to test such configuration changes on Junos by using the commit-check feature before proceeding to the actual commit. The commit-check feature will detect problems such as this, without the need to attempt a blueprint commit operation that might fail for some Junos devices.

---

## **Changing VTEP Address as a Day-2 Operation in SONiC May Fail (AOS-34891)**

Changing the IP address of the VTEP of a SONiC device as part of a day-2 operation, may fail and the device may be left in a failed deployment state. The danger is more acute in cases of Blueprints with a large number of VRFs and VNs.

If the IP address of the VTEP is being set for the first time, e.g. when the blueprint is deployed for the first time, this release note does not apply. Only cases where the VTEP has already been set to an initial IP address and VNIs have been created using it.

## **Workaround**

It is suggested to apply full config after attempting to alter the VTEP IP address and getting a failure.

---

## **Default Route Nexthop Programmed as Drop in SONiC (AOS-35085)**

In SONiC versions up to 3.5.4 and 4.0.2, under rare timing conditions, the deletion and re-addition of a subinterface belonging to a non-default VRF may result in missing hardware routes to next hops behind that interface. This has been observed in a case where the default route towards an external router was missing from the hardware programming. Please refer to vendor bug SONIC-67559.

Unfortunately, it is relatively difficult to detect the occurrence of this bug, as the route is missing from the hardware, but present in the routing daemon and the kernel of the device.

## **Workaround**

A full config apply is necessary to resolve this problem.

---

## **Deploying Quickly After Draining the Fabric Node May Result Into EVPN Session Flap Due**

### **to Slow Processing of the Route Policies (AOS-37596)**

If a user deploys back leaf/spine/sspine/access(in case of ESI access switches) quickly after draining the node, there is a possibility of an EVPN session flap between the node and evpn bgp peer. It can be due to high CPU usage when performing drain and deploy operations. Slower platforms like Juniper QFX5100 are particularly vulnerable to this issue (it took around a min to apply drain route policy under normal conditions).

---

### **DHCPv4 Offer Dropped Client Between Junos and Junos-EVO (AOS-33701)**

Due to an existing issue in Junos 22.2R2-EVO, DHCPv4 offer packets are dropped when a virtual network is extended by Junos and Junos-EVO rack devices.

---

### **Firewall function in the Junos device may not work correctly when Security policy rule with tcp-established used (AOS-45677)**

When a rule with the tcp-established option exists in the Security Policy, even if the Apstra correctly renders the device configuration into firewall function, a Junos device running less than 22.2 version may fail to function properly because the entries are incorrectly programmed in the hardware.

#### **Workaround**

Upgrade the Junos version to at least qualified NOS version 22.2R3

---

### **IPv6 Communication Is Not Established Between VM/Host on L2 VXLAN Use Case (AOS-38715)**

IPv6 communication between vm/host on L2 vxlan use case is not established. The affected SONiC versions are 4.0.1 and 4.0.2.  
SONiC 4.0.3 or 4.0.5 are not affected.

#### **Workaround**

Remove neigh-suppress config when using issue-impacted SONiC versions

---

### **IPv6-Enabled Interfaces Can Have Traffic-Affecting Side Effects in SONiC Cyrus 4.0.2 (AOS-34342)**

Having an interface (e.g. a VN) that is IPv6-enabled in a device in the blueprint, can cause serious traffic loss in SONiC 4.0.2 devices in the network. Traffic affected is not limited to IPv6-enabled interfaces. Under specific circumstances, even external devices such as routers outside of the AOS Blueprint could trigger the same problem if they are IPv6-enabled.

Full description and more accurate scope of the bug can be found in vendor bugs SONIC-66452 and SONIC-66504.

## Workaround

If there is a strict need for an IPv6-enabled interface in an installation using SONiC Cyrus 4.0.2, please use the attached configlet, consult with support for a workaround or refer to vendor with bug ids SONIC-66452 and SONIC-66504. Important note, the configlet must be applied to all SONiC Cyrus 4.0.2 devices.

```
{% if os_version and os_version.startswith('4.0.2-') %}
ip prefix-list APIPA seq 5 permit 169.254.0.0/16 le 32
!
route-map DENY_APIPA deny 10
match ip address prefix-list APIPA
!
route-map DENY_APIPA permit 20
!
{% for sz_info in security_zones.itervalues() | sort(attribute='vrf_name') %}
{% if sz_info['sz_type'] == 'evpn' %}
router bgp bgpService['asn'] vrf 'Vrf-' + sz_info['vrf_name']
address-family ipv4 unicast
table-map DENY_APIPA
{% elif sz_info['vrf_name'] == 'default'%}
router bgp bgpService['asn']
address-family ipv4 unicast
table-map DENY_APIPA
{% endif %}
{% endfor %}
{% endif %}
```

---

## 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

---

### **Junos BGP Graceful Restart Not Working for Non-default Routing Instances (AOS-35849)**

Currently, for dual routing engine Junos devices, Apstra enables BGP graceful restart. Due to a current issue in Junos 22.2R2, graceful restart is not enabled for non-default routing instances.

---

### **Junos EVO on ACX7100 Failed Arp Resolution/Installation for VLANs Enabled Within MAC-VRF Instances (AOS-33571)**

Currently, when using Junos EVO on the ACX7100 platform If a VLAN which is not VXLAN-enabled is assigned within the MAC VRF routing instance, ARP learning fails. VXLAN enabled and non-VXLAN enabled VLANS, cannot be assigned to the same physical interface.

---

### **Loss of IPv6 traffic on the Juniper\_EX4400-48T platform (AOS-38006)**

Under certain circumstances, IPv6 traffic might not work on certain vxlan-enabled vlans of a Juniper\_EX4400-48T device.

#### **Workaround**

The only known workaround so far is rebooting the switch. Contact Juniper Support.

---

### **NXOS BGP Crashes When Removing and Reapplying Dynamic BGP Connectivity Template (AOS-44239)**

NX-OS 9.3(11) BGP crashes when removing and reapplying dynamic BGP CT as follows,

- Unassign CT from VLAN interfaces
- Edit CT and specify IPv4 subnet for BGP Prefix Dynamic Neighbors
- Reassign CT to VLAN interfaces
- Edit VN, verify that Secondary IP Allocation mode has changed from 'forced' to 'enabled', and remove IPv4 addresses from leafs
- Commit config

```
BGP-3-ASSERT: bgp- [27133] ../routing-sw/routing/bgp/bgp_peer.c:2004:
Assertion `*prev_peer' failed.
SYSMGR-2-SERVICE_CRASHED: Service "bgp" (PID 27133) hasn't caught signal 11
(core will be saved).
```

## **Workaround**

Manually shutdown the BGP peers before changing the dynamic BGP connectivity template to avoid the BGP crash.

---

## **QFX5120 Dropping Packets Destined to VRF IPv6 Loopback (AOS-33770)**

Juniper QFX5120 running Junos 22.1 and earlier versions contain a bug that will drop packets destined to the IPv6 VRF loopback interface.

## **Workaround**

Any Junos 22.1 version is not supported by Apstra. It is recommended to use Junos 22.2R2.

---

## **SONiC SAG Vlan may be down even though Vlan interface is set to oper up (AOS-35706)**

Under very rare circumstances, a Vlan interface may be set to oper up in SONiC, but in reality the underlying interface (visible via the command ``ip link show dev VlanXYZ``) may be down.

SAG currently listens to APPL\_DB events for VLAN readiness but internally also checks if STATE\_DB is updated with VLAN ready. If STATE\_DB is not ready, SAG will remain inactive. If intfmgd is scheduled between APPL\_DB and STATE\_DB updates by vlanmgrd, we will hit the issue where SAG will be inactive on this VLAN.

It is somewhat difficult to detect this issue. It is expected that if the bug happens on a SONiC device, there will be type-3 extensible probe anomalies present which cannot be accounted for without noticing (via `ifconfig` or `ip link show dev ...`) that the underlying vlan interface is not UP.

Please refer to vendor case SONiC-68179.

## **Workaround**

A workaround would be to remove and reconfigure SAG ip addresses on the VLAN. Also, setting the kernel vlan interface up via `ifconfig` or `ip link` is also sufficient to remedy the situation.