

Juniper Apstra Version 5.0.0 Release Notes

New Features

Modular linecard device profiles now display which Device Profiles they are assigned to (RFE-2406)

Feature Category: Device Profiles

Modular linecard device profiles now display which Device Profiles they are assigned to.

Manage the Juniper MX304 using Apstra Freeform blueprint (RFE-2585)

Feature Category: Device Profiles

You can now manage the Juniper MX304 using Apstra Freeform blueprint. The MX304 is often used as a connection between data centers and other networks, such as public clouds or remote data centers. Equipped to support advanced networking protocols including BGP, OSPF, IS-IS, MPLS, segment routing, advanced traffic engineering, and network segmentation. With the MX304 and Apstra Freeform, you'll benefit from design flexibility, protocols, and a streamlined approach to data center management, all within a single solution.

Do not unassign a device when changing the Interface Map with a different Device Profile unless the model changes (RFE-2690)

Feature Category: Device Profiles

Apstra now provides the ability to change the Interface Map with a different Device Profile without unassigning the device in the active Blueprint assuming none of the values in the Device Profile "Selector" are changed (eg Model, OS Version regex, OS Family).

Device Profiles Extension to indicate the supported Reference Designs (Datacenter and/of Freeform) (RFE-2543)

Feature Category: Device Profiles

Device Profiles now include the list of supported Reference Design that they can be used in. You

now can use this field to know what device can be used what Reference Designs, example the Juniper MX304 is available for Freeform reference designs only and cannot be used in Datacenter Reference Designs. When authoring a new Device profile you can set the Ref Design Capabilities field to define its use.

Device Profile for Juniper EX4400-24X included with Apstra (RFE-2983)

Feature Category: Device Profiles

Device Profile for Juniper EX4400-24X included with Apstra.

Device Profile for Cisco Nexus 93600CD-GX (RFE-3210)

Feature Category: Device Profiles

New Device Profile for Cisco Nexus 93600CD-GX. Please note an NOS related issue where rollback fails when some breakout ports are used, documented in .

Added EVPN/VXLAN Leaf support for Dell-EMC Z9432F-ON (RFE-3063)

Feature Category: Device Profiles

You can now use the Dell-EMC Z9432F-ON device as a leaf in EVPN/VXLAN deployments.

Qualification for Arista EOS 4.28.7.1M (RFE-3215)

Feature Category: Device Operating Systems

Arista EOS 4.28.7.1M is now a qualified version.

New Apstra ZTP DHCP backend and improvements to UI workflow (RFE-2787)

Feature Category: Device Operating Systems

Replaced the Apstra ZTP backend from ISC DHCP to Kea DHCP bringing new capabilities and improved UI workflows for configuring DHCP and ZTP settings.

GA support of Juniper QFX5120 for Integrated DCI (VXLAN Stitching) (RFE-3177)

Feature Category: Device Operating Systems

EVPN VXLAN Stitching support for QFX5120 running 23.4R2 or greater is now fully supported (Apstra and platform).

Following incremental configuration is added, which requires a PFE restart: 'set forwarding-options evpn-vxlan vxlan-trans-vni-enable'

Allow HTTPS device OS image url (RFE-2967)

Feature Category: Device Operating Systems

Apstra now supports HTTPS URLs for registering os images with https type urls via the "provide image url" option.

Add the Apstra user who committed a Blueprint change to the Junos commit comments (RFE-2856)

Feature Category: Device Operating Systems

The Apstra user who committed a Blueprint change is now added to the Junos commit comments which can be checked by logging on the junos device and run "show system commit".

You are now able to add additional inter-switch mesh links to a collapsed fabric topology via the web UI (RFE-2875)

Feature Category: Design, Build, Operate

You are now able to add additional inter-switch mesh links to a collapsed fabric topology via the web UI by navigating into the blueprint ,click on the switch and add mesh link option is available.

Warn user when upgrading a device not in drain mode (RFE-2934)

Feature Category: Design, Build, Operate

When a user tries to do an NOS upgrade on a device that is a part of the blueprint and NOT in drain mode, the UI shows an additional warning.

View Loopbacks in Routing Zones (RFE-2142)

Feature Category: Design, Build, Operate

You can now view the loopbacks assigned to each device in each Routing Zone, which gives you a better understanding and improves visibility and management of routing configurations.

To access this feature, navigate to Staged/Active > Virtual > Routing Zones and click the VRF name.

Validate software package uploads (RFE-2500)

Feature Category: Design, Build, Operate

This feature improves your user experience by validating the file extension and checksum before uploading, preventing unnecessary wait time and reducing the risk of errors.

Checks compatibility: Before allowing the upload, it ensures that the software package is compatible with the switch's operating system. This prevents uploading packages that are not suitable for your specific device.

Validates checksum: Before performing the actual upload, it verifies the checksum of the package to guarantee that it is not corrupt or damaged. This ensures that the package is error-free and can be trusted.

User Defined Tags (RFE-1937)

Feature Category: Design, Build, Operate

Users can now use tags on physical and virtual constructs such as nodes, links, virtual networks and VRFs to help group and search based on additional user-defined metadata.

To add or remove tags on virtual networks, select virtual networks under the staged blueprint, select 1 or more virtual networks, then click on the new "tag" icon. In the Add/Remove Tags window type in the user defined tag you wish to apply to this virtual network. Then select the "Add/Remove Tags" button.

You can now use the tag as a filter and search criteria in Apstra.

Switch Port Constraint Validations (RFE-2872)

Feature Category: Design, Build, Operate

We've introduced two key enhancements to improve your experience when configuring switch port groups.

Warning: Potential Network Impact

When you try to change a non-master port within a switch port group, you'll now receive a warning. This alert ensures you're aware of potential network-impacting changes before they're applied, preventing unintentional effects on other ports in the group. This feature saves you from inadvertently changing the entire port group to an unintended speed.

Error-Proof Your Port Layouts

Our new feature also alerts you to potentially invalid port layouts, catching errors early on. This helps maintain consistency in new blueprints and prevents issues that could arise from invalid port layouts. This is especially crucial for platforms with complex port restrictions, where configuring one port can affect others.

Swagger API documentation in the Apstra ZTP UI (RFE-2777)

Feature Category: Design, Build, Operate

The Apstra ZTP server now has the REST APIs documented in Swagger API Specifications on the UI.

Support tag-driven filtering of leaf/leaf-pair during Virtual Network creation (RFE-2888)

Feature Category: Design, Build, Operate

You can now leverage system tags for filtering when you create Virtual Networks. This helps speed-up the definition of a Virtual network footprint in large-scale deployment. It nicely complements the tag-driven interface assignment at the Connectivity Template level.

Support for ESI-LAG on SONiC (RFE-2885)

Feature Category: Design, Build, Operate

You now can use ESI as a redundancy option for SONiC based blueprints with devices running

SONiC version 4.2.1. This capability is available at the Leaf layer. Blueprints must have the all racks running the same design option, mix of ESI and MLAG is not supported.

SNMP configlet for interface enrichment in Apstra Flow. (RFE-3232)

Feature Category: Design, Build, Operate

A sample SNMP configlet and property are provided in the product to give an example to enable SNMP for interface enrichment in Apstra Flow.

Search primitives in Connectivity Templates (RFE-2771)

Feature Category: Design, Build, Operate

Find primitives quickly in complex templates! We've added a new search function to the "Parameters" tab, allowing you to locate specific primitives quickly and save time.

Routing-Zone Aware Role-Based-Access (RFE-2887)

Feature Category: Design, Build, Operate

You now can express "Tenants" as a collection of Routing Zones and use them in the user's role creation through tenant-specific permissions. This lets you as an admin user define granular permissions enforcements restricting user operations to only the objects they are allowed to manipulate, according to the Tenant memberships.

A new "Tenant Permissions" role type is added, allowing to grant a user write permissions on a per tenant(s) for the following operations

- Manage Routing Zones.
 - Manage Virtual Networks.
 - Manage Virtual Networks Endpoints.
-

Re-order and add descriptions to routes in routing policy (RFE-2636)

Feature Category: Design, Build, Operate

You can now re-order the extra routes in your routing policy as needed. You can also add a

description to each route to provide context and clarity. The descriptions are present in Apstra and are not rendered to the local device configuration.

Provide option in ZTP UI to reset DHCP and ZTP settings (RFE-3000)

Feature Category: Design, Build, Operate

A new button is now available in the ZTP UI to load the default dhcp.conf and ztp.conf file settings.

Pending/Waiting Icon added for deployments (RFE-1591)

Feature Category: Design, Build, Operate

When a config deployment is pending, a "pending/waiting" yellow icon on the dashboard and active tabs will also now be displayed to indicate to users that the new config deployment is still in progress, and users should wait for its completion before attempting to deploy another config change.

Option to set Interface Operation State and Generic System Deploy mode when adding a new Generic System/External Generic System (RFE-2841)

Feature Category: Design, Build, Operate

When adding a new Generic System or External Generic System you now can set the following parameters:

- The "Deploy mode" of Generic System / External Generic System. This controls telemetry collection and validation of interface status.
 - The "Operation State" of each interface added. You can select Up / down.
-

New Export Virtual Networks to Connectivity Templates options (RFE-3095)

Feature Category: Design, Build, Operate

When you click "Export to Connectivity Templates", you now have more control over the output.

You can choose to create a single Connectivity Template (CT) that combines all your Virtual Networks (VNs) or generate one CT per VN.

New boot script for Apstra ZTP for initial VM configuration (RFE-2996)

Feature Category: Design, Build, Operate

A new `ztp_config` boot script is available on the Apstra ZTP VM to help during boot for changing the admin password, configuring network settings, and starting the ZTP service.

New API endpoints to Add/Delete a leaf or leaf-pair from an existing rack (RFE-3137)

Feature Category: Design, Build, Operate

The following 4 new API endpoints have been added

- Add new leaf to existing rack: `POST /api/blueprints/{blueprint_id}/rack-leaf/{rack_id}`
 - Delete leaf from rack: `DELETE /api/blueprints/{blueprint_id}/rack-leaf/{rack_id}/{leaf_id}`
 - Add new leaf-pair to existing rack: `POST /api/blueprints/{blueprint_id}/rack-leaf-pair/{rack_id}`
 - Delete leaf-pair from rack: `DELETE /api/blueprints/{blueprint_id}/rack-leaf-pair/{rack_id}/{rg_id}`
-

Moving one or multiple Virtual Networks between Routing Zones (RFE-2051)

Feature Category: Design, Build, Operate

You now can move a virtual network from its current Routing Zone (VRF) to a different one. This migration capability is available at the individual Virtual Network level or in the form of a batch operation by selecting multiple Virtual Networks. This operation can be disruptive from a traffic standpoint as it needs to delete and recreate the Virtual Network(s).

Modify interface descriptions via user interface (RFE-2865)

Feature Category: Design, Build, Operate

Apstra now provides a new button in the user interface to modify interface descriptions by navigating to staged> Blueprint>Physical>Interfaces and edit the interface description.

Logical design visualization (RFE-3163)

Feature Category: Design, Build, Operate

A new logical map visualization feature that lets you view all the Logical Devices, Interface Maps, Device Profiles, and Racks in your Template. The logical map provides an intuitive image that allows you to validate and verify the design elements match your expectations and reduce errors during the deployment phase.

In-Product Documentation (RFE-2648)

Feature Category: Design, Build, Operate

We're introducing context-sensitive guidance directly within the UI to streamline your learning and problem-solving processes. With In-App Guidance, you can access relevant documentation without leaving the UI. Hover over Tooltips for additional information, Deep Dives, and direct links to access relevant information.

Graphic of Logical Design Elements and Relationships (RFE-3164)

Feature Category: Design, Build, Operate

A new static graphic depicting the relationships of Logical Devices, Interface Maps, Racks, Templates, and Device profiles was added. This graphic makes it easy to understand and visualize the connections between different elements, allowing you to easily grasp the concepts and modeling developed by Juniper Apstra.

Enhance topology view for generic systems (RFE-2741)

Feature Category: Design, Build, Operate

When you add access switches to a topology, generic systems will now be rendered in their own column alongside the access switches. This change is designed to make your topology more readable and organized as it grows.

Device Serial Number is now included in the device context. (RFE-3224)

Feature Category: Design, Build, Operate

Devices Serial Numbers' are now included in the device context and can be used within Jinja configlet rendering.

Deployment or Config Deviation Anomalies Warning (RFE-2720)

Feature Category: Design, Build, Operate

Users will now be presented a warning notification if they attempt to commit changes from staged to active when a deployment or configuration irregularity is present. This is useful in scenarios when there is a deviation in the running configuration from the user's intent.

Automatically Collapse External Systems In View (RFE-2719)

Feature Category: Design, Build, Operate

We've made it easier to navigate and understand your blueprint by automatically collapsing external systems and removing external links from view. This means you can now focus on the essential components of your blueprint, without the clutter.

Apstra Time Voyager descriptions now support Markdown (RFE-3015)

Feature Category: Design, Build, Operate

Apstra now enables the Time Voyager description field to support Markdown providing more flexibility in descriptions for your save points with items like headers, categories, fonts, bullets, links, and more.

Add Descriptions to Virtual Networks (RFE-2869)

Feature Category: Design, Build, Operate

You can now add a free-text description to your virtual networks, providing additional context that goes beyond the network name. This description is not only visible in the Apstra UI but also reflected in the device configuration. This feature allows you to capture important details that can't

be conveyed through the network name alone, making it easier to understand and manage your virtual networks.

Add Descriptions to Routing Zones (RFE-2646)

Feature Category: Design, Build, Operate

This feature lets you add descriptions to Routing Zones. These descriptions are not only visible in the UI but are also applied to the device itself, providing additional context to your Virtual Routing and Forwarding (VRF) configuration.

Ability to Relocate Generic Systems (RFE-2855)

Feature Category: Design, Build, Operate

This feature allows you to relocate Generic Systems between internal and external, and vice versa. With this feature, you can now easily move Generic Systems without having to delete them from the blueprint and re-add them. This streamlined process saves you time and effort, making it easier to adapt to changing network requirements.

Support for multiple Values in Custom Collectors (RFE-2953)

Feature Category: Telemetry and Analytics

You now can define a custom telemetry service (Custom Collectors) with more than one value. This is particularly useful when the object you want to monitor has multiple value metrics associated with the same identity. As an example an interface has several counters for traffic statistics: TX Unicast, TX Broadcast, TX Multicast, and the same for RX. In such scenario a defined service will typically have one key, the interface name, and several values, one for each counter. This feature allows you to address such use-cases in Custom Collectors.

Support for multiple output Values in Custom IBA probes (RFE-2679)

Feature Category: Telemetry and Analytics

You can now reference custom telemetry collectors producing multiple output values in a user-defined IBA probe. This allows you to create custom IBA probes with richer context by extracting more telemetry data from devices with a single service.

New IBA Probe for monitoring the health of Virtual Networks and checking correct propagation and programming of EVPN Type2 routes (RFE-3069)

Feature Category: Telemetry and Analytics

You now have a new IBA probe pre-defined and auto-enabled for all Datacenter blueprints to monitor the health of Layer2 service for all Virtual Networks in the blueprint. The probe will leverage the intent context of the Virtual Networks provisioned along with the telemetry of the locally learned MAC address on every VTEP-enabled system to automatically derive expectations on the presence of those same MAC entries as remotely learned ones on other systems. With this probe you get a unique perspective on the health of your Virtual Network for every hosts they contain. By default the probe performs all the analytics provides the result but does not raise an anomaly. Use the predefined probe attribute to check the Anomaly Raising condition if you want the anomalies to be flagged.

Import/Export of IBA Dashboards (RFE-1070)

Feature Category: Telemetry and Analytics

You now can export and import any IBA dashboard. This allows you to easily move dashboards between from on instance to another and it includes the underlying probes with it. Note that as part of this change the "Widget" tab has been removed as it is no longer an independent entity, but rather a component of the Dashboards. This streamlines the process and reduces the number of tabs to interact with.

Import/Export capability for custom collectors (RFE-2959)

Feature Category: Telemetry and Analytics

You now have a UI workflow for easily exporting and importing telemetry service definitions including all the sub-components: Service schema and the service collectors. This will help you streamline the process of sharing custom collectors between instances.

Filter field of custom collectors now has Syntax highlighting for expressions (RFE-3144)

Feature Category: Telemetry and Analytics

Syntax highlighting is applied to the expressions used in custom collectors. This is displayed while editing.

Analytics report for device environmental data (RFE-3193)

Feature Category: Telemetry and Analytics

New analytics Environmental Data predefined report performs historical analytics on data from the `Device Environmental Checks` probe. With this report you will get a temperature health score of the devices in your blueprint, ranging from 1 (best) to 5 (worst) according to vendor threshold limits.

Adding examples of IBA processors (RFE-3174)

Feature Category: Telemetry and Analytics

When creating a new probe or editing an existing one, you now have examples for all the Analytics processors. Those are the processors receiving an input data and performing a given analytical function on it, as opposed to the source processors producing data. The purpose of providing examples for those processors in the probe creation menu is to better guide you in understanding the function and use of each one of these processors so you can better choose the right one for your probe's requirement.

Supported Upgrade Paths to Apstra 5.0.0 (RFE-3045)

Feature Category: Platform

This release supports upgrade paths from previous Apstra 4.2.X releases.

Users must use VM-VM for upgrades from Apstra 4.2.X releases. See the Apstra user guide for more information on Apstra upgrades.

Support Apstra Guest on VMware ESXi Version 8.0 (RFE-2942)

Feature Category: Platform

Ability to install Apstra on VMware ESXi version 8.0

Redesign of Developers Page (RFE-3227)

Feature Category: Platform

This is a revamp of the developer's page to make it look more modern, easier to read and visually appealing.

Include Apstra worker nodes in Show Tech log collection (RFE-1931)

Feature Category: Platform

The worker nodes as part of an Apstra cluster are now included in the Apstra Show Tech log collection to help troubleshoot issues across the entire cluster of your Apstra deployment.

Enhanced Account Management (RFE-3046)

Feature Category: Platform

Easily manage access and maintain compliance with your security standards using our new features:

Account Status Control: Enable or disable user accounts as needed, providing temporary access restrictions.

Password Expiry Management: Set global and per-user policies to manage password expiry, ensuring compliance with your corporate security policies.

Automatic Account Disablement: Disable accounts for users who don't change their password before it expires, allowing them to reset their password.

Apstra ZTP UI Log viewer for TFTP and DHCP services (RFE-2985)

Feature Category: Platform

Added a new Logs menu item in the ZTP UI which will provide users the ability to see the info, debug, and rsyslog log files related to TFTP and DHCP service directly in the UI to help troubleshoot issues faster.

Apstra Guest VM Support on Windows Server 2019 (RFE-2882)

Feature Category: Platform

Ability to install Apstra as a Guest VM on Windows Server 2019

AOS-SDK client library documentation (RFE-2657)

Feature Category: Platform

We are productizing AOS-SDK, a Python client library to interact with Apstra's RESTful APIs. Among the provided capabilities:

- Auto-completion and auto-documentation at the IDE level so you can easily navigate the structure of the client library and find the function you need to use. This documentation provides you with explanations about the use of a given class or method as well as high-level information on the expected arguments for the methods. Similar information is also provided on a static HTML page available from the "Developers Page".
 - Type-checking. With this, a function will have a strongly defined model used to perform input checks and identify missing elements or data type mismatches, in which case the execution is declined following a fail-early principle without placing any API call towards the server.
 - High-level generator functions for some of the most popular Day 2 workflows. These are Python methods consuming user intent arguments and outputting data structures ready to be consumed by some create operations. For example, you can use a "create_bound_to_payload" function and pass the result to your virtual_networks.create() function. With that, you operate at a higher level of abstraction than the low-level APIs and speed up your development journey.
-

Added ability to edit interface description (RFE-1618)

AOS now allows the user to edit interface descriptions with aos-cli.

Apstra Flow New Features

Display Flow license tier in the dashboard (AOSEXTRFE-15)

Feature Category: FLOW

The Apstra Flow dashboard displays the license tier to confirm which license level is used.

Improved error handling for hostname in yml config file (AOSEXTRFE-7)

Feature Category: FLOW

Improved error handling when incorrectly formatted hostname is provided and added ability to use IP address instead of hostname

New IFA dashboards (AOSEXTRFE-13)

Feature Category: FLOW

New dashboards to better visualize IFA flows.

RDMA header parsing (AOSEXTRFE-14)

Feature Category: FLOW

Parse RDMA headers for RoCEv2 traffic in AI data centers.

VM setup script for Apstra Flow (AOSEXTRFE-11)

Feature Category: FLOW

A VM setup script now exists at /usr/local/bin/startup for initial setup

Changed Features

Qualified switch operating systems with Apstra 5.0 (RFE-3099)

Feature Category: Device Operating Systems

The following updates have been made for switch operating systems qualified for the Apstra 5.0

release.

Juniper Networks:

Junos (All roles)

21.4R3-S7

22.2R3

22.4R3

23.4R2-S4

Junos Evolved for IP-Forwarder role (Spines in EVPN or any role in an IP-Fabric):

22.2R3-EVO

22.4R3-EVO

23.4R2-EVO

Junos Evolved for EVPN leaf roles:

22.2R3-EVO

22.4R3-EVO

23.4R2-EVO

Junos Interconnect Gateway Leaf:

22.4R3 (minimum)

23.4R2-S4

Junos Evolved Interconnect Gateway Leaf:

22.4R3-EVO (minimum)

23.4R2-EVO

Cisco Systems:

9.3(13)

10.2(6)

10.3(4a)

Arista Networks:

4.24.5M

4.28.7.1M

4.30.3M

Dell EMC & Edgecore:

Enterprise SONiC 4.1.2

Enterprise SONiC Edge Standard 4.1.2

Enterprise SONiC 4.2.1

Enterprise SONiC Edge Standard 4.2.1

VMware vCenter 8 support in the Apstra VMware integration (RFE-2743)

Feature Category: Design, Build, Operate

You can now use VMware vCenter version 8 for the Apstra VMware integration.

Support for VMware NSX 4.1.X (RFE-3097)

Feature Category: Design, Build, Operate

Incremented support of NSX to 4.1.X versions

Streamlined Resource Pool Creation (RFE-2862)

Feature Category: Design, Build, Operate

You can now create a resource pool directly from within a blueprint, saving you time and effort. After clicking the "Update assignments" button, click the new "Create resource pool" button, enter the details for the pool, then click "Create". You can then select the check box for the new pool to use for assigning resources.

Simplified LAG Creation (RFE-2838)

Feature Category: Design, Build, Operate

This streamlined process reduces the number of steps required to create a bonded interface (LAG) when adding a new internal or external Generic System. You can simply select the "Create LAG" option at the end of the wizard, and you're done!

Network Operating System (NOS) upgrade improvement (RFE-2685)

Feature Category: Design, Build, Operate

To provide you with a smoother and more reliable experience, all interfaces are automatically shut down during the upgrade process to prevent potential issues, such as traffic blackholing. For example, this improves the upgrade experience when static LAGs are in use.

Managed Device UI enhancement with active DP and initial DP (RFE-2684)

Feature Category: Design, Build, Operate

This UI change under managed device shows which device profile was auto-selected during the onboarding phase and the actual device profile currently in use. Managed Device UI enhancement with active DP and initial DP together tooltips

Interface Maps default selection is now Connected To (RFE-3170)

Feature Category: Design, Build, Operate

When creating Interface Maps, the new default behavior is to select all port roles in Connected To. There is also a new option to select All Port Roles. By default checking all the "Connected to roles" for the Logical Device, helps perform design and day-2 operations more efficiently and avoids issues related to missing port roles.

Improved experience creating Interface Maps (RFE-3169)

Feature Category: Design, Build, Operate

To provide a better user experience, we have rearranged how fields are presented in the Interface Map screen and added info messages to provide a better workflow. The "Name" block will also be auto-generated based on the Logical Device and Device Profile selected.

Improved device deploy/undeploy handling (RFE-2326)

Feature Category: Design, Build, Operate

We have improved the device deploy/undeploy handling to provide a more intuitive and efficient user experience. This feature adds an explicit "Not Set" value for device deploy mode, allowing easy unsetting without errors. Removing System IDs automatically sets deploy to "Not Set", enabling commit changes. The UI now displays a "Not Set" radio button, making it clear how to unset the mode. This change also fixes an issue with device deletion from the blueprint and normalizes the deploy mode spelling across product.

Filter Logical Devices by port capabilities (RFE-3032)

Feature Category: Design, Build, Operate

You can now easily find the Logical Devices you need by filtering them based on port

capabilities. This new feature allows you to quickly narrow down your search results to show only devices that match your specific requirements.

Enhanced Metadata Tag Search (RFE-2963)

Feature Category: Design, Build, Operate

When searching for metadata tags, you'll now see the device name alongside the interface name. This added context helps you quickly identify the exact device and interface associated with the metadata tag, making it easier to find what you need.

Enhanced DCI with ESI MAC MSB Alerts (RFE-3226)

Feature Category: Design, Build, Operate

When using DCI interconnect with both blueprints managed by our product, this feature proactively warns you when default ESI MSB values are used, encouraging you to change to non-default values.

Enhanced agent installation checks (RFE-2466)

Feature Category: Design, Build, Operate

An improvement to the agent installation process to ensure a smoother and more reliable device management experience. This feature detects and prevents potential issues by identifying specific pre-existing configuration statements that may conflict with the Apstra reference model. If any problematic elements are found, the agent installation process is aborted, and a detailed log is generated, clearly stating which specific configuration elements are causing the conflict.

Bulk Edit of Node Names (RFE-2431)

Feature Category: Design, Build, Operate

You can now bulk edit capabilities for switch names, leaf pair names, and rack names from a single screen.

Allowing the creation of Virtual Networks with an empty footprint (RFE-2921)

Feature Category: Design, Build, Operate

You now can create a Virtual Network with an empty footprint, i.e with no leaf/leaf-pair assigned to it at the creation time. This decoupling of the VN creation and the nodes assignment will allow you to more easily comply with existing workflows or be used by some API clients such as Terraform providers for Apstra.

Tag-driven anomaly configuration and predefined dashboard for Optical Transceiver Probe (RFE-2547)

Feature Category: Telemetry and Analytics

You now can use Interface tags to control the anomaly raising logic in the Optical XCVR probe through a set of 3 possible conditions (new fields) exposed in the probe menu: 1) "Do not raise anomaly" 2) "Alarm anomaly interface tags" and 3) "Warning and Alarm anomalies". You can choose none or any combination of those fields to granularly control the anomaly raising policy you wish for your blueprint.

You now can instantiate the predefined dashboard for Optical Transceivers which groups in the same view the interfaces exhibiting anomalous metrics along with the Optical transceivers meta-data: Vendor Name, Vendor Part Number, Vendor Serial Number, Media Type and Fiber Type.

Qualification of QSFP-DD800 optics on QFX5240-64QD (RFE-3190)

Feature Category: Telemetry and Analytics

QSFP-DD800 optics in QFX5240-64QD is now qualified and supported by the Optical transceivers IBA probe.

Predefined Dashboard for the EVPN Host Flapping probe (RFE-3050)

Feature Category: Telemetry and Analytics

You now have a predefined dashboard for the EVPN Host Flapping probe showing the list of flapping MACs along with the corresponding leafs as well as the count of historical anomalies for 30 days.

Migration from Native gRPC to gNMI (RFE-3039)

Feature Category: Telemetry and Analytics

As a result of this migration "Interface" telemetry service will use gNMI in Periodic mode for any device running Junos or Junos Evolved release $\geq 22.4R3$.

Enforcing unique labels for IBA probes (RFE-2546)

Feature Category: Telemetry and Analytics

Users must now create IBA probes with unique labels. Probes created with duplicate labels will not be accepted.

This change helps avoid confusions resulting from situations where the same probe is instantiated more than once with different scopes and different expectations. A example of such probe is the "Hot/Cold Interface" amongst few others. In those situations, the user now has to provide unique label identifiers to disambiguate the probes between them.

Enforcing unique labels for analytics dashboards (RFE-3199)

Feature Category: Telemetry and Analytics

Users must now create analytics dashboards with unique labels. Dashboards created with duplicate labels will not be accepted.

This change helps avoid confusions resulting from situations where the same Dashboard is created more than once with different scopes and different widgets.

Dynamic series Support for Analytics processors of type multi-input (RFE-3065)

Feature Category: Telemetry and Analytics

You now have support of Dynamic series in the Multi-Input processors including: "Ratio", "Comparision", "Set Comparision", "Substract", "Union" and "Logical Operator". This expands the set of possible custom probes you can create for example by mixing in the same pipeline Graph-driven data (Static series) with non Graph-driven data (Dynamic series).

Decoupling the device's OS Version from the RPC Schema version when defining a custom collector (RFE-3052)

Feature Category: Telemetry and Analytics

- When you create or edit a custom collector you now treat the device's OS version independently from the CLI schemas. With this change the CLI schema are decoupled from the Custom collector workflow to be treated as an optional aid rather than a primary key. This provides you with more flexibility in creating custom collector for any Junos version irrespective of the presence or not of a CLI schema for that version. Whenever existing, the CLI schema remains available for the user as an aid. But their absence no longer constitutes a blocker to use the feature.

Updating API Endpoint documentation to encourage using Connectivity Templates API endpoints for all Virtual Network operations (RFE-3116)

Feature Category: Platform

The documentation of the following API endpoints has been updated to more clearly describe their usage and the constraint on the Virtual Networks created with Connectivity Template API. The added documentation encourages to use the Connectivity Template instead:

- POST /api/blueprints/{blueprint_id}/virtual-networks/{virtual_network_id}/endpoints
 - PUT /api/blueprints/{blueprint_id}/virtual-networks/{virtual_network_id}/endpoints
 - DELETE /api/blueprints/{blueprint_id}/virtual-networks/{virtual_network_id}/endpoints/{endpoint_id}
-
- POST /api/blueprints/{blueprint_id}/virtual-networks
 - PATCH /api/blueprints/{blueprint_id}/virtual-networks/{virtual_network_id}
 - PATCH /api/blueprints/{blueprint_id}/virtual-networks-batch-patch
-
- POST /api/blueprints/{blueprint_id}/subinterfaces
 - PATCH /api/blueprints/{blueprint_id}/subinterfaces
 - DELETE /api/blueprints/{blueprint_id}/subinterfaces/{subinterface_id}
-
- POST /api/blueprints/{blueprint_id}/l3-dot1q-links
 - DELETE /api/blueprints/{blueprint_id}/l3-dot1q-links/{logical_link_id}

- POST /api/blueprints/{blueprint_id}/floating-ips
 - PATCH /api/blueprints/{blueprint_id}/floating-ips/{floating_ip_node_id}
 - DELETE /api/blueprints/{blueprint_id}/floating-ips/{floating_ip_node_id}
-
- POST /api/blueprints/{blueprint_id}/protocol-sessions
 - PUT /api/blueprints/{blueprint_id}/protocol-sessions/{protocol_session_id}
 - DELETE /api/blueprints/{blueprint_id}/protocol-sessions/{protocol_session_id}
-
- POST /api/blueprints/{blueprint_id}/static-routes
 - PATCH /api/blueprints/{blueprint_id}/static-routes/{static_route_id}
 - DELETE /api/blueprints/{blueprint_id}/static-routes/{static_route_id}

The following endpoints were also updated to clarify their use, since the previous update points to them

- PUT /api/blueprints/{blueprint_id}/obj-policy-import
- PATCH /api/blueprints/{blueprint_id}/obj-policy-batch-apply

Show Tech improvements for more logs on EOS and NXOS devices (RFE-2094)

Feature Category: Platform

Additional logs on EOS and NXOS devices are now included in the Apstra Show Tech support bundle.

For EOS:
show logging all

For NXOS:
show logging logfile

FIPS Mode Enhancements (RFE-3168)

Feature Category: Platform

You can now enable FIPS mode for the Apstra Zero-Touch Provisioning (ZTP) Server

streamlining the process. Additionally, Onbox agents now support FIPS mode and are automatically configured from Apstra during agent installation.

Documentation of platform-level API endpoints (RFE-3102)

Feature Category: Platform

Enhanced documentation of platform level API endpoints, applicable to both Datacenter and Freeform reference designs, to help developers during the creation of API workflows.

Adding guardrails on low-level graph API endpoints to protect from accidental use (RFE-2845)

Feature Category: Platform

The API endpoints performing low-level graph manipulations are now protected against accidental use with a new "allow_unsafe" request parameter to be set to "true" by the API client. Setting "allow_unsafe" to "true" is required for the Apstra server to accept and process the API call. The endpoint documentation has been updated.

The list of API endpoint is:

- POST /api/blueprints/{blueprint_id}/nodes
 - PATCH /api/blueprints/{blueprint_id}/nodes
 - PATCH /api/blueprints/{blueprint_id}/nodes/{node_id}
 - DELETE /api/blueprints/{blueprint_id}/nodes/{node_id}
 - POST /api/blueprints/{blueprint_id}/relationships
 - PATCH /api/blueprints/{blueprint_id}/relationships/{relationship_id}
 - DELETE /api/blueprints/{blueprint_id}/relationships/{relationship_id}
-

New User experience for the selection and IBA processors (RFE-2836)

You now have a more human-friendly GUI for the creation of IBA probes with a refactoredUI modal for the selection of IBA processors. This includes a categorisations of the processors by groups for easier selection and better understanding of their function. The processor's parameters and their descriptions have also been revisited for a more simple identification of their role.

Apstra Flow Changed Features

Adding sflow to the filter in threats dashboards (AOSEXTRFE-9)

Feature Category: FLOW

The threats dashboards now include sflow in the filter.

Additional health checks in the metrics endpoint (AOSEXTRFE-12)

Feature Category: FLOW

Additional health checks in the metrics endpoint for VM CPU and memory, collector version, Apstra connection status, and SNMP configuration status.

Enable SNMP V2 by default in config (AOSEXTRFE-10)

Feature Category: FLOW

SNMP V2 with public community string is now enabled by default in the Apstra Flow config to make it easier for users to see interface enrichment.

Include OpenSearch logs in the support bundle (AOSEXTRFE-8)

Feature Category: FLOW

The OpenSearch logs are now included in the support bundle to improve the troubleshooting experience.

Removed Features

Deprecation of Telemetry Only mode when creating system-agents (RFE-3084)

Feature Category: Device Operating Systems

- The "Operation Mode" parameter of System-Agent which used to have two options: Full Control (default) and Telemetry Only has been removed and the Telemetry Only mode has been deprecated. System-Agent both Onbox and Offbox now only support Full Control.

Removed L3 Edge Server Links option (RFE-2476)

Feature Category: Design, Build, Operate

The routing policy option "L3 Server Links" is no longer relevant or required.

Deprecating analytics widget of type "Anomaly Heatmap" (RFE-2205)

Feature Category: Telemetry and Analytics

The creation of analytics widgets is simplified by supporting the widgets of type "Stage" only and deprecating the ones of type "Anomaly Heatmap".

Tech Preview Features

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

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

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

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

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

Tech-Preview support of the Juniper QFX5130E-32CD (RFE-3253)

Feature Category: Device Profiles

You can now use QFX5130E-32CD as a Tech-Preview. This model has identical port layout as the QFX5130-32CD with lower scale.

Tech Preview - ACX7100 as an DCI gateway for Integrated DCI (RFE-3187)

Feature Category: Design, Build, Operate

Tech Preview support for ACX7100, running Junos 23.4R2, as DCI gateway for Integrated DCI use cases.

Support of Collapsed-Fabrics in SONiC blueprints (RFE-3247)

Feature Category: Design, Build, Operate

You can now deploy SONiC based blueprints based on Collapsed-Fabrics templates.

Support of Access-Switch pairs in SONiC blueprints (RFE-2392)

Feature Category: Design, Build, Operate

You can now deploy SONiC based blueprints with Access-Pairs.

History of blueprint anomalies (RFE-3214)

Feature Category: Telemetry and Analytics

You now can see the history of the blueprint anomalies in the Active tab of the blueprint. The Anomalies tab now shows a chart of the anomalies count over time, grouped per type of anomaly (Ex: BGP, Cabling, Route ..). You can zoom in and zoom out to look for a specific time interval in detail.

Selecting an individual anomaly you also can see the historical timeline of the anomaly to show the occurrences of that specific anomaly in the recent history. The default retention period is set to 30 days.

Exploratory Analytics Interface (RFE-2910)

Feature Category: Telemetry and Analytics

A new Exploratory Analytics interface is now available, providing capabilities for Exploratory Data Analysis (EDA), a common practice in data science that involves investigating large datasets to discover correlations, identify various patterns, and verify hypotheses. This interface offers a UI-driven approach to constructing Query-Based Analytics (QBA) queries, which is an extension of IBA primarily used for performing statistical analysis on historical data. It allows users to join data from multiple sources, such as different IBA probes. With this interface, you can select a metric from a given IBA probe, calculate statistical summaries like the Interquartile Range (IQR), and visualize distributions using Box-and-Whisker diagrams. A different use-case is one where you select two metrics to calculate linear regression and create scatter plots, helping to identify correlations between the two variables. These functionalities enable you to explore various data patterns and gain deep insights into your data, facilitating the future creation of new probes or reports, thereby empowering users to leverage data effectively and enhance decision-making and analytical capabilities.

Fixed Apstra General Issues

Alternate name of interface has the same name as the real interface name (AOS-46121)

In Apstra-configured SONiC devices, the Alternate Name of an interface is always the same as the real interface name (native mode) used in the SONiC OS's Linux kernel. Some customers prefer the standard interface name as an Alternate Name over the native interface name. Since version 5.0.0, AOS does not explicitly define the Alternate Name as the native interface name, so it is automatically filled in by SONiC, which follows the standard interface name for Alternate Name.

Example) Alternate Name as native interface name for "show interface status" command in the sonic-cli

```
sonic# show interface status
```

```
-----  
Name           Description           Oper           Reason  
AutoNeg  Speed           MTU           Alternate Name  
-----  
Ethernet0     -           9100           down           admin-down  
off           25000  
-----
```

Example) Alternate Name as standard interface name for "show interface status" command in the sonic-cli

```
sonic# show interface status
```

```
-----  
Name           Description           Oper           Reason  
AutoNeg  Speed           MTU           Alternate Name  
-----  
Ethernet0     -           9100           down           admin-down  
off           25000           Eth1/1  
-----
```

Resolution

For installations upgrading from Apstra versions earlier than 5.0.0, A full config apply after the alias fields have been left empty will restore the aliases' values to standard naming interface names. The full config apply will not be done automatically during upgrade, to avoid incurring downtime. Customers interested in this fix should schedule and execute a full config apply at a non-critical moment of their choosing.

Apply Config failed when 1G interface in the Juniper EX4400-24MP-EM is configured (AOS-45648)

When 1G interface is configured in the Juniper EX4400-24MP-EM, applying configuration fails because device profile for EX4400-24MP-EM has invalid setting for speed.

Apstra may incorrectly render Juniper ACX7100-32C Channelized Port Config (AOS-44243)

When configuring 10G/25G channelized port transformations on Juniper ACX7100-32C, Apstra

may incorrectly omit the "unused" configuration parameter on the odd interfaces within the port group if no generic systems are connected to channelized ports on the even interface.

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.

Resolution

SysDB was updated in AOS 4.2.2 and AOS 5.0 to fix virtualinfra agentid entity.

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.

Resolution

In the case of vCenter or NSXT related to the Virtual Infra Manager, skip the SSH connectivity check during the upgrade.

Banner not updated by ZTP's custom config in the SONiC device (AOS-45991)

After the SONiC device's banner is updated by a script file configured in the ZTP's custom config, the final stage of ZTP processing replaces the customized banner with another piece of information based on the success or failure of ZTP processing. As a result, unlike in the custom config of ZTP, the banner may not be updated properly.

Resolution

`ztp.py` has a `save_and_complete` function which calls the `_update_ssh_banner` function that updates the banner on the switch. If we have a banner written in `sonic_custom.sh` we should not call the `save_and_complete` function in `ztp.py`. This will help not overwriting the banner config on the switch that is being configured using `sonic_custom.sh` file.

BGP maximum-paths configuration not rendered in the user-defined VRF's BGP configuration for Cisco NXOS device (AOS-46667)

For Cisco NXOS devices, the BGP maximum-paths configuration is currently rendered in only default VRF's BGP configuration; the user-defined VRF's BGP configuration lacks this configuration.

Resolution

During upgrade to Apstra 5.0.0 or higher, maximum-paths is added to the rendered configuration for named VRFs under ipv4 and ipv6 unicast address families.

Deleting Link returns error with '<' not supported between instance of 'str' and 'NoneType' (AOS-47479)

When Deleting Link is executed by UI or by delete-switch-system-links API call for the port channel port, the backend recalculates port channel pool to reuse the port channel IDs. The recalculation uses sorting key comprising of generic system's hostname for comparison. If the hostname is null, comparison can fail with the exception because of incompatible type comparison.

Resolution

When hostname of generic system is null, comparison automatically converts it into empty string.

Dell SONiC devices after Apstra ZTP loses mgmt IPs if the ZTP server is not available (AOS-44712)

For Dell SONiC devices after Aptra ZTP, they would lose mgmt IPs if the ZTP server is not available because Apstra ZTP processing doesn't configure static IP address to mgmt interface.

Device Profile is not assigned when Cisco 93108TC-FX3P device is onboarded (AOS-45123)

When Cisco 93108TC-FX3P device is onboarded, it reports the hardware model differently depending on the version. The current built-in Device Profile for Cisco 93108TC-FX3P has the selector's model as 93108TC-FX3P. If the device reports the hardware model as C93108TC-FX3P (with prefix C), it can't be matched on the built-in device profile. Therefore, the device profile can't be associated with the onboarding device.

Disk Space Exhaustion Due to Unrotated Logs in Apstra ZTP VM Containers (AOS-44969)

Users may encounter disk space exhaustion in the ZTP VM because log rotation is not enabled for the '/var/log/messages' and 'syslog' files in the dhcpd, tftp, and status containers, as well as for the '/logs/rsyslog.log' file in the status container. Although the logrotate utility and a crontab file are available, log rotation is not activated by default. As a result, these logs can accumulate, quickly consuming disk space and potentially causing degraded system performance or service interruptions.

Resolution

Apstra 5.0.0 enables log rotation and optimization for the ZTP VM containers. Log rotation for the Docker container logs is now managed by the built-in mechanism controlled by the Docker daemon. Additionally, rsyslog logs are rotated using the logrotate utility, and a cron job has been scheduled to ensure that log rotation occurs regularly, preventing disk space exhaustion.

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.

Resolution

Add waiting for Transport Node update from NSXT and perform the tvlan to pnic relations.

Error Not Indicated in the ztp.json Configuration in the UI Editor (AOS-44598)

When the ztp.json configuration with an error is saved, a red dot appears for a second next to the error configuration, then suddenly disappears. The red dot eventually returns.

Export/Import Route Targets Under Routing Zone Introduce Additional Character in Config (AOS-44770)

Junos device config deployment fails for Junos BGP community configurations when the user defines import/export route targets for a routing zone where the second number is greater than

65535 (e.g. 64512:4200000000), as the Apstra rendered config appends an "L" string at the end of the assigned number.

Exporting and importing cabling map triggers validation errors for port-channel interfaces (AOS-45683)

When a cabling map is exported and then imported, UI generates validation errors for port-channel interfaces. The import process doesn't expect port-channel interfaces from the input data. However, the exporting process includes not only individual interfaces but also port-channel interfaces together as output data. Therefore, importing with input data, which has port-channel interfaces, triggers the validation error during the import process.

High CPU utilization on QFX10000 modular devices (AOS-23192)

On QFX10000 modular devices, high CPU utilisation can be observed when under Apstra management. This is caused by large volumes of telemetry data being collected at a high rate.

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.

Resolution

In 5.0.0, parenthesis or bracket characters are allowed for CSV importing validation.

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.

Interface description doesn't allow less-than character(<) or greater-than character(>)

(AOS-46743)

Interface descriptions with less-than characters (<) or greater-than characters (>) may cause deployment errors on certain NOS devices.

Resolution

Interface descriptions containing those characters is prohibited in 5.0.0. An upgrade to 5.0.0 converts those characters into another allowed character.

Interface descriptions are no longer allowed to contain the double quote character (AOS-45883)

Due to issues with configuration rendering across all supported platforms, the use of double quote characters in interface descriptions is no longer permitted on any managed device. In Apstra 5.0.0, user can edit interface descriptions through the UI, but earlier versions allowed editing raw graphs via the Apstra API. If a customer previously used a direct blueprint node API call to add an interface description with double quotes, these characters will be automatically converted to underscores during the 5.0.0 upgrade. Starting from Apstra 5.0.0, any attempt to include a double quote character in an interface description will be rejected by the Apstra API.

Resolution

Apstra 5.0.0 enforces the prohibition of double quote characters in interface descriptions to overcome configuration rendering issues. During the upgrade to 5.0.0, any existing double quotes in interface descriptions are converted to underscores. The Apstra API will also block any future attempts to use double quotes in interface descriptions.

Juniper QFX5240-64QD and QFX5240-64OD new breakout modes in Junos Evolved 23.4R2 and later (AOS-47020)

When an even port is channelized into a mode producing 8 logical interfaces (e.g 8x100G), Junos Evolved versions prior to 23.4R2 dictate that the immediate next odd port needs to be left unused. In the case of Apstra, channelizing the even port is all that is needed and Apstra itself will take care to set the next odd port to unused.

Starting from Junos Evolved 23.4R2 and beyond, more possible channelization options have been added, and also it is allowed to use the lower odd port even in the case where the even port above it is channelized into 8 logical interfaces. The limitation is that the total number of logical interfaces must not exceed 10. As an example, it is now legal to channelize the lower odd port to 2x400G, etc. The new possible breakout modes are available in the updated Device Profile for the QFX5240-64QD and QFX5240-64OD. Users must exercise caution not to use these new channelization options unless running Junos Evolved 23.4R2 and beyond.

Specifically, the additional channelization modes added with Junos Evolved 23.4R2 to the QFX5240-64QD are: 4x200G and 2x200G. Moreover 1x800G, 1x400G, 2x400G, 2x200G, 1x100G are now allowed for odd ports even in the case that the immediately preceding even port is channelized in 8-ply mode.

The additional channelization mode added with Junos Evolved 23.4R2 to the QFX5240-64OD is 4x200G. Moreover 1x800G, 2x400G are allowed for odd ports even in the case that the immediately preceding even port is channelized in 8-ply mode.

JUNOS device commit check feature using Apstra UI may incorrectly indicate an error when testing config (AOS-45715)

When using the commit check feature on the Uncommitted tab in Apstra UI, it may incorrectly indicate it experienced a red Error. RpcError(serverity: warning) when the JUNOS device issues a warning over the tested config while retrieving the config diff from the device.

Resolution

In Apstra 5.0.0, warnings emitted while retrieving the config diff of a commit check operation, no longer cause the operation to fail.

LAG Sustained Execution Failures Anomaly in the Device Telemetry Health Probe (AOS-46043)

LAG Telemetry Service should be enabled as long as the device has a port channel configuration. However, even if no port channel configuration is rendered in the device, Apstra assumes that at least one port channel interface exists in the leaf or access switch, enabling the LAG Telemetry service and causing the failure condition.

Resolution

LAG service is enabled in the leaf and access in the Apstra 5.0.0 only if there is intent for port channel configuration in the device (CT for virtual network or sub-interface is assigned to port channel, leaf switch with MLAG, leaf switch with ESI to access, access switch with ESI).

LAG telemetry collection may fail in the SONiC device when using static LAG (AOS-46129)

When a device has only one static LAG connection enabled, LAG telemetry collection for a leaf may fail to work properly in the SONiC device. If other non-static LAG connections are configured on the same device, the problem will not occur.

Resolution

This problem is corrected in Apstra 5.0.0.

New 4x10G transformation in Device Profiles for QFX5120-48Y, EX4650-48Y, QFX5100-24Q (AOS-46879)

In the device profile of the Juniper QFX5120-48Y and Juniper EX4650-48Y, a new explicit transformation 5 has been added for ports 48,49,50,51,52,53,54,55, which is the recommended way to achieve 4x10G channelization for those ports. The previous DP transformation 4 to achieve 4x10G channelization is also retained in the Device Profile for consistency and for any users who have already used that transformation in their blueprints.

Furthermore, transformation 3 has been added to all ports of the QFX5100-24Q device profile, allowing explicit configuration of the 4x10G channelization. When using 4x10G, explicit transformation is recommended.

Finally, customers encountering 4x10G channelization problems on models QFX5100-48T, QFX5200-32C, QFX5100-48S, QFX5100-96S are advised to contact Apstra Juniper Support for a customized DP that will be providing explicit 4x10G channelizations.

Resolution

New port breakout transformations were added to better support QFX5120-48Y, EX4650-48Y, QFX5100-24Q

New Device Profile based on new port layout for the QFX5240-64QD and QFX5240-64OD devices (AOS-47405)

It is noted that, starting with Junos Evolved 23.4R2-EVO, the port numbering layout of the QFX5240 models is different. Apstra 5.0.0 includes a new pair of device profiles that support the new layout and will only match any device running the new version or later. The previous version of Junos Evolved supporting the QFX5240 was the 22.2X100-EVO, and the previous device profile present in Apstra 4.2.1 will now be updated to match only against the 22.2X100-EVO.

Resolution

Since the physical layout of the ports changes between Junos Evolved versions, the transition involves a reboot of any QFX5240 devices. The customer is required to configure the new cabling.

Node in Stage tab shows uncommitted changes, but no changes in Uncommitted tab (AOS-47472)

The issue arises when certain actions, such as deleting and reverting links between nodes, cause nodes to incorrectly highlight as having uncommitted changes despite no actual differences. The root cause is a bug in the Cache Diff input plugin for system info, specifically related to the `uplinked_system_ids` attribute being an unsorted list which cause mismatches between staged and operational graphs.

Resolution

When comparing two sets of `uplinked_system_ids`, having them sorted simplifies the comparison process. This is especially useful for generating diffs, as it avoids false positives where the IDs are the same but in a different order. Issue will be fixed in 5.0.0

NOS Upgrade makes device isolated with missing management IP address in the Cisco NXOS device (AOS-48839)

During a NOS upgrade job for a Cisco NXOS device via Apstra, Apstra attempts to collect the new pristine configuration based on the target NOS version as soon as the device is rebooted. Even if the device is ready, it may not yet provide management IP address information when the pristine configuration is collected. Apstra pushes the rendered configuration with the newly collected pristine configuration, which lacks the management IP address and makes the device unreachable, resulting in service disruption.

Resolution

Makes pristine collection process to check whether management IP address is available or not with multiple attempts.

Not Able to See Apstra Flow Dashboards as Tenant Switched to Private (AOS-45813)

When logged into the Apstra Flow UI, occasionally, when the UI times out, the user is switched from the Global to Private tenant, rendering the dashboards inaccessible once they log in again.

PortChannel description not rendered in the SONiC device (AOS-46316)

The interface description, including the port channel, can be updated via an API. However, Apstra does not render a description of the SONiC device's port channel interface. As a result, the SONiC device contains no description for the port channel interface.

Resolution

This issue is fixed in Apstra 5.0.0.

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.

Resolution

Rack template creation via UI, Users can no longer set link to superspine speed value when superspine count is set to 0.

SONiC device Show Tech collection is failing due to a remote SSH command error (AOS-47972)

SONiC customers may encounter issues generating Device Show Tech due to a remote SSH command failure. When attempting to collect the device show tech data from the Apstra UI, users might see the following error. The error logs indicate that the SSH command to generate the show tech data fails with a return code of 124, which typically indicates a timeout.

```
2024-09-03 11:13:37,467 INFO:TASK: Generate device show tech
2024-09-03 11:13:37,468 INFO:command (timeout-350): service aos show_tech --
output-prefix=/tmp/911b3e70-show-tech
2024-09-03 11:19:27,475 ERROR:Failure reason: , return-code: 124
2024-09-03 11:19:27,475 ERROR:FAILED
2024-09-03 11:19:27,477 ERROR:Failed command: sudo service aos show_tech --
output-prefix=/tmp/911b3e70-show-tech
2024-09-03 11:19:27,477 ERROR:Remote ssh command failed
```

Resolution

The SONiC CLI did not complete due to paging, which caused a timeout. To resolve this, paging for SONiC CLI output has been disabled in Apstra 5.0.0.

SONiC device's error log for "Did not log record sleeping for 60s for vrf to be created" during ZTP processing (AOS-46485)

In the current ZTP implementation for SONiC devices, after the VRF is switched to management VRF during the ZTP process, the ZTP script running in the device sends status information to the ZTP server with a sleep time of 60 seconds. The change in VRF disconnects the device for a certain period and prevents it from sending logs to the ZTP server until reachability is restored. The failures in the delivery of logs would be recorded and displayed on the device as error messages. The error messages don't mean that the ZTP process failed. It can be ignored safely.

Resolution

The enhancement makes logs be updated to the ZTP VM, followed by sleep, before switching to VRF.

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.

Resolution

Apstra version 5.0.0 will be able to reach the helper over the VRF to which the DHCP/DHCPv6 client belongs, instead of the default VRF.

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 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.

vCenter Collector's Invalid Mac address Validation error (AOS-45831)

All PNICs in the dummy hypervisor, created by VMware mobility agent, have MAC addresses of none. A validation error is triggered by these invalid PNIC MAC addresses. Because data is not correctly gathered from the vCenter by the Virtual Infra Manager, Apstra's UI is unable to display the correct virtual machine visibility.

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

VirtualInfraGraphAgent Crash by removing transport zone on multiple transport nodes (AOS-45842)

VirtualInfraGraphAgent creates nodes and relationships based on data provided by Vcenter's collector. When the collector encounters validation errors due to invalid MAC addresses, it may produce incomplete transport vnet information that is separate from the hypervisor. The NSXT configuration change in the problem status, which includes removing the transport zone from multiple transport nodes, causes the dangling transport vnet to be released several times, causing the agent to crash.

Resolution

Deletion process for transport zone and transport nodes checks null checking to prevent double deletion issues.

VirtualInfraGraphAgent Crash from portgroup not managed by NSXT (AOS-45840)

The current Apstra implementation expects NSXT to exclusively control the ESXi hosts it is managing. When a portgroup is created in the ESXi hosts not by NSXT but by vCenter, restarting VirtualInfraGraphAgent can cause the cleanup process to reference invalid information in the portgroup. If NSXT with vCenters is registered in Apstra, NSXT should create and manage all portgroups.

Resolution

Before referencing, the VirtualInfraGraphAgent cleanup process in 5.0.0 checks for the presence of invalid information in the portgroup that vCenter created.

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.

Resolution

When tftp container init.sh starts it will now cp with -p to preserve ownership and permissions for 'ztp_py3.py'

ZTP UI Should Not Send Non Configured Params in API Payload (AOS-44596)

The UI should not configure the keys in the ztp.json configuration for system-agent-params if the user has not configured the values. The ZTP UI reports the payload with empty strings, which trips the schema validation. For the attached screenshot, I added Junos as the platform in the GUI for the configurator, and the UI generated a payload with the platform, agent_type, job_on_create, and profile. This is unexpected.

Fixed Apstra Security Issues

Apstra - SONiC 4.1.2 Misconfiguration of CONFIG_DB allows unauthenticated RESTCONF calls (AOS-46786)

SONiC devices configured by Apstra with SONiC release 4.1.0 or higher expose an unauthenticated RESTCONF HTTPS server due to unhandled changes in the configuration database schema between SONiC versions 4.0.0 and 4.1.0. This issue affects all Apstra releases prior to 5.0.0.

When a SONiC device running 4.1.0 or later is configured with "client_auth": "cert" but without a security profile, it will allow any remote HTTPS call to query or modify the device configuration without requiring authentication.

```
root@sonic:/etc/sonic# curl -k https://localhost/restconf/data/openconfig-system:system/state/hostname  
{"openconfig-system:hostname":"sonic"}
```

Resolution

The issue resolved in Apstra release 5.0.0 and mitigated during the upgrade process. The upgrade is designed to be compatible with the existing workaround configlet, which can be safely removed after the upgrade is completed.

OpenSSH Vulnerability CVE-2024-6387 in Apstra Server Version 4.2.x (AOS-47640)

Apstra server version 4.2.x, which is based on Jammy Ubuntu 22.04 includes OpenSSH version 8.9p1, is exposed to a security vulnerability identified as CVE-2024-6387.

Resolution

Apstra Server Version 4.2.2.1 and 5.0.0 will include a fixed OpenSSH package.

USN-6891-1(Python vulnerabilities) from Tenable scan report (AOS-48042)

Tenable scan reports USN-6891-1 (Python vulnerabilities) for both Apstra 4.2.x Controller and Apstra 4.2.x ZTP.

Resolution

Apstra 5.0.0 and 4.2.2.1 include fixed python packages for Ubuntu 22.04.

Fixed Third-Party Issues

Any blueprint commit results in the restart all BGP IPv4 and IPv6 peerings in any SONiC device (AOS-45866)

On any configuration push against a SONiC device, Apstra will always utilize the 'frr-reload.py' script that is accompanying the FRR routing daemon, to gracefully apply any configuration changes made to that daemon, if any. It has been observed that in SONiC versions 4.1.2, this always results in the restarting of all IPv4 and IPv6 peerings. This includes cases where

there is no change whatsoever to the FRR routing configuration.

Resolution

The culprit of the problem is in the way `frr-reload.py` interprets lines of style "address-family ipv4 unicast" against "address-family ipv4". Apstra has hitherto used the latter style when generating an `frr.conf` configuration file. From the point of view of FRR, these two styles have the exact same meaning and, in fact, typing the latter into `vtsh` results in the former getting inserted into the configuration. However, `frr-reload.py` erroneously compares the two styles verbatim and thinks sections of one style are different than sections of the other. This causes the "address-family ipv4 unicast" style sections in the running config to be removed every time `frr-reload.py` runs and replaced with "address-family ipv4" style sections found in the `frr.conf`, which are then changed by FRR to "address-family ipv4 unicast" style sections, and so on.

In Apstra 5.0.0 and 4.2.2, Apstra switches to rendering "address-family ipv4 unicast" explicitly to avoid this pitfall.

ARP entry may be missed in the MLAG pair during the Arista EOS upgrade from Arista 4.28.7.1M to EOS 4.30.3M (AOS-50665)

During the EOS upgrade from 4.28.7.1M to 4.30.3M in the Arista device MLAG pair, the device that would be upgraded first may exhibit incorrect ARP behaviour and miss receiving ARP entries from its MLAG peer. Stochastic traffic failures may also be observed between the two MLAG members.

Resolution

It should be noted that using an MLAG pair of EOS devices from different versions is not supported and will result in traffic loss. For more information, please refer to the vendor release notes. It is recommended to update both switches in an MLAG pair in quick succession.

BGP daemon may crash in the SONiC device when changing the LAG mode of a PortChannel (AOS-46058)

Due to a cleanup issue in FRR, the `bgpd` may crash when changing the LAG mode of a PortChannel interface. Operation is expected to return to normal after a few seconds.

Resolution

The issue also exists in community FRR. Please refer to vendor issue SONIC-90228 for details. This is expected to be addressed in SONiC 4.4.0.

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.

Resolution

The problem will be addressed in SONiC 4.1.0, which will be incorporating improvements making the VTEP IP change substantially easier.

ClusterHealthWriterAgent Error During Initial Deployment (AOS-44106)

On an initial deployment, the customer may see the following traceback error in `/var/log/aos/controller/ClusterHealthWriterAgent.err`:

```
IndexError: Aos::MetricLog::MetricLogWriter.newMetricLogWriter: error in
function call : Tac::RangeException("Error mounting event file:
/var/lib/aos/metricdb/cluster_health_info/container/utilization/meta-
1704279564839064-180-2024-01-03--10-59-24.839098.tel")
```

This issue is identified as an intermittent glitch while interacting with VM FileSystem within the "aos_controller_1" container. The ClusterHealthWriterAgent process will recover on its own after the restart without any action from the user.

Enabling of L2_NEXTHOP_GROUP in Trident-4 based SONiC devices (AOS-45548)

Special care must be taken for any SONiC leaf that is using the Trident-4 chipset (e.g. Dell Z9432F-ON, S5448F-ON) and belongs to an EVPN-VXLAN fabric on which ESI multihoming is used for the first time.

If an ESI multihoming pair is being added to that fabric for the first time and said fabric already has Trident-4 leaves that have already accepted EVPN routes, then all these leaves must be either rebooted or a full config apply must be executed against them.

Resolution

To properly manage next hops consisting of multiple VTEPs, the Trident-4 ASIC must be configured in "wide" mode. This mode reduces the ASIC's MAC/ARP capacity by half. If there are existing MAC/ARP entries created from EVPN routes, they must be replaced after the ASIC is switched to wide mode.

Packet Drops on Untagged Layer-2 Interfaces on EX4400, EX4650, and QFX5120 Platforms (AOS-42959)

Due to an outstanding bug in versions of Junos prior to 22.2R3-S3 on the EX4400, EX4650, and QFX5120 platforms, packets may be dropped on layer-2 interfaces configured with an untagged native VLAN.

QFX5120 devices must have their SFP28 interface speeds configured in quads (AOS-27254)

Juniper QFX5120 devices have their SFP28 ports configured in groups of 4 called 'quads'. Each interface within a quad must be set to the same speed. Quads are grouped as 0-3, 4-7, 8-11, etc. If there are interfaces with different speeds within the same quad, config deployment by Apstra will succeed but interfaces may not come up.

Please refer to the following article for more details:

<https://kb.juniper.net/KB37520>

Symmetric IRB on Enabled on Junos EVPN-VXLAN Stitching Fails to Forward Traffic (AOS-43921)

Combining Symmetric mode IRB with EVPN VXLAN Stitching is not recommended until an upcoming Junos release supports this feature. If Symmetric IRB is configured, local hosts attached to the EVPN-DCI border gateways will fail to generate the additional Type2 Mac:IP label corresponding to the L3 VNI, they will operate asymmetrically.

Apstra Config Rendering Changes

User-defined extra import or export host routes without CIDR prefixlength not included in configuration (AOS-47151)

User-defined host routes within extra import or export prefix-list entries in a routing policy were

not being included in device configuration if they did not have the /32 CIDR prefix length. For example, '3.3.3.3 le_mask=None ge_mask=None' was not rendered, while '3.3.3.3/32 le_mask=None ge_mask=None' was being rendered.

Resolution

Upon upgrade, prefix-lists will be re-rendered if they include non-CIDR qualified host entries and will appear during upgrade config preview. If these changes are not desired, delete the prefix-list entry from the blueprint before initiating the upgrade process. No configuration change will be performed in the pre-upgrade blueprint.

Fixed Apstra Flow Issues

Not Able to See Apstra Flow Dashboards as Tenant Switched to Private (AOSEXT-2, ESD-460)

When logged into the Apstra Flow UI, occasionally, when the UI times out, the user is switched from the Global to Private tenant, rendering the dashboards inaccessible once they log in again.

Affected Product Version

6.4.2

Resolution

Fixed in Product Version \geq 6.4.4.

Known Apstra General Issues

"On Device Configlet Preview" Might Emit Error if the Device Is Unassigned (AOS-43233)

Attempting to pull the configlet preview for specific blueprint device ("On Device Configlet Preview"), by clicking on the device label under the general configlet preview page might fail with a slightly misleading error, if the device is unassigned. Certainly trying to get a preview for a device which is not assigned is bound to cause an error, as a real preview for a device that doesn't exist isn't possible. However, the error emitted is slightly confusing.

[Juniper EX] Duplication of Interface Map(IM) on Adding New Access Switch to the Leaf (AOS-49715)

Users may encounter duplication of Interface Maps (IM) when adding a new access switch to the leaf in a collapsed fabric blueprint, specifically with Juniper EX series devices. This issue is due to a device profile configuration mismatch: the backend incorrectly generates an additional IM with a duplicate label because of inconsistencies in connector type information (RJ45 vs. rj45). Users can view these duplications in the Interface Maps section by navigating to Staged -> Catalog -> Interface Maps.

Workaround

To workaround this, delete the duplicate IMs from Staged -> Physical -> Catalog -> Interface Maps.

[SONiC] Golden Config Validation Error When Modifying FRR Log Level from "Informational" to "Notifications" (AOS-49660)

Apstra sets the FRR log level to "log syslog informational" by default in the SONiC device. When a customer attempts to change the log level to "log syslog notifications" using a configlet, the golden configuration validation fails due to a mismatch between the expected and running configurations. Apstra has made "log syslog informational" a prerequisite for golden config validation, using it as a verification key.

Workaround

It is recommended that users avoid changing "log syslog informational" to "log syslog notifications"

A large bump in the memory footprint for MetricQueryManagerAgent is seen when the 'Time Series' data option is selected on the Active->Anomalies tab (AOS-48770)

MetricQueryManagerAgent handles large historical data to serve the '/blueprints//anomalies-history' endpoint. Depending on the amount of data, the agent's memory footprint may increase significantly. The benchmark environment recorded a memory footprint of up to 2.2Gb for the agent. The memory footprint settles after the initial bump.

If the system administrator is concerned about the MetricQueryManagerAgent footprint's impact on the system's available memory, the following workaround is recommended.

Workaround

Restart MetricQueryManagerAgent and avoid using the 'Time Series' query.

Aggregated Interface Counters show aggregated traffic as 0 for the aggregated period when the AOS service restarts (AOS-45993)

There is a 2 minute interval for aggregating telemetry data. When the AOS goes down and comes back up, for the first 2-3 mins, data is reported as 0. Therefore, during this time, the aggregated telemetry data is 0 and we expect to see a flat line during this time. After this 2-3 mins window, first non-zero data is reported to AOS and that is when we can expect the AOS to populate the charts correctly.

AOS device agent installation fails on Junos and Junos EVO device with hostname as IP address (AOS-47453)

If the Junos and Junos EVO devices do not have a name-server configuration, AOS device agent installation fails when a hostname is used instead of an IP address. The missing name-server configuration in the device hostname prevents the management IP address from being resolved.

Workaround

Here are the two workarounds. Either of them can be used to resolve the problem.

Option 1: Configure the name server on the device. This allows the device to resolve the hostname and use the correct management IP address.

Option 2: When configuring system agents in AOS, use IP addresses rather than hostname.

AOS Show tech execution fails due to lack of space on the server (AOS-50912)

AOS show tech collection failed due to the increased size of the folder `/var/tmp/show_tech_tmp` and the AOS is transitioning into read-only operation mode.

Workaround

When show tech collection is failing due to lack of space in `/var/tmp/show_tech_tmp` folder, log in to the AOS server and delete the `/var/tmp/show_tech_tmp/controller_show_tech/` folder.

Apstra AnomalyGenerator crash with PrimaryKeyIndex violation newRow (AOS-52744)

When AnomalyGenerator tries writing to MetricDB and taking SysDB snapshots if vlanid is missing from the primary key it will cause the AnomalyGenerator to crash with the below trace information

```
Unique Index (PrimaryKeyIndex) violation newRow index matches that at row:
65536
python3.10: /Project/leblon/infra/TableTop.tin:206: void
Aos::DoubleLinkListHelper::addToList(Aos::RowIndexHelper&, U32): Assertion
`false && "!row"' failed.
Process 22875 died with signal 6 (SIGABRT) errno 0 code -6 (unknown)
```

Workaround

Add the following to disable all anomaly logging into the `/etc/aos/aos.conf` file and then restart the aos service.

```
[anomaly_metric_logging]
enable = 0
```

Apstra CLI device password change may fail due to task timeout (AOS-54971)

The scenario change-device-password CLI command securely updates device credentials by performing tasks like SSH checks, configlet staging, blueprint commits, and agent password updates. In the current Apstra CLI version, the system agent check has a 60 second timeout, while configlet staging is limited to just 20 seconds. If these operations take longer than expected, the command may fail with errors like:

```
Failure 1: Task Stage creation of Configlet for password change may fail with:
AssertionError: Timeout waiting for Wait that last task status is succeeded
```

```
Failure 2: Task Check System agent status may fail with:
409 Conflict: Agent is already running a job (check)
```

These failures occur when backend tasks exceed the current timeout settings, which is particularly noticeable in Apstra 4.2.x and later versions, where performance issues with configlet and configuration rendering are known. Additionally, longer durations in the check job can result from changes in the customer's environment.

Workaround

Manual intervention is required to cleanup and proceed as follows:

1. For System Agent Check Failure (409 Conflict):

- * Update the pristine configuration to reflect the new encrypted password for the user.
 - * Remove the temporary configlets named `change_pass_<...>_junos`.
 - * Manually commit the blueprint.
 - * Re-run agent check jobs to confirm the password update.
2. For Background Task Timeout (e.g., Configlet Staging Failure):
- * Revert the blueprint to the previous working version.

After completing the steps above, use the latest Apstra CLI image for the Apstra release (4.2.2, 5.0.X, 5.1.0, 6.0.0) and retry the scenario change-device-password command. The most recent Apstra CLI image can be obtained by contacting Apstra Support.

Apstra Onbox agent in Junos Evolved device with dual routing engines reports incorrect device facts after a switchover (AOS-48335)

When switching between routing engines in a Junos Evolved System (such as the PTX10008) with dual routing engines, the Apstra Onbox agent may report an incorrect management IP, management interface, or management MAC address.

This bug does not affect Offbox agent.

Workaround

It is recommended that the Apstra Onbox agent be restarted after a switchover using the command "request system application app aos restart node reX", where reX is the Master routing engine.

Apstra UI prevents setting "IP Links to Generic Systems MTU" to 9216 under fabric settings (AOS-53627)

Apstra customers 4.2.x and greater may encounter an issue where the MTU value for IP Links to Generic Systems cannot be set to 9216 via the UI. Although the UI states that only even values in the range 1280-9216 are accepted, the input of 9216 is incorrectly rejected, while 9214 is accepted.

Important Note:

1. This issue is applicable only to customers upgrading from Apstra 4.1.x to 4.2.x or later, where Fabric MTU remains disabled post-upgrade and customers who wish to continue without enabling the Granular MTU feature.
2. This issue is not applicable to customers with fresh 4.2.x deployments, where Fabric MTU is enabled by default, activating the Granular MTU feature.

Workaround

Although the UI blocks 9216, the backend API does accept this value. As a workaround, users can update the MTU value via the REST API:

1. Navigate to Platform â†’ Developers â†’ REST API Explorer.
2. Use the PATCH /api/blueprints/{blueprint_id}/fabric-settings endpoint with the following payload:

```
{
  "external_router_mtu": 9216
}
```
3. Verify the update by performing a GET on the same endpoint: GET /api/blueprints/{blueprint_id}/fabric-settings
4. Navigate to Blueprints â†’ Blueprint Name â†’ Uncommitted and check the diff
5. Commit the Blueprint

If further assistance is needed, please contact Apstra Support.

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. This is expected behavior.

Apstra ZTP version 5.0.x and lower may fail to start dhcp container (AOS-49934)

Apstra ZTP versions 5.0.x or below may fail to start the dhcp container service if a custom dhcp config is used where multiline /* */ or # bash comments are used. It may also fail to start where config file includes are used <? include "filehere" ?>. The dhcp init.sh container startup script fails to parse out comments and file includes so that socket_name path can be determined and created at startup.

Workaround

Recommend using only ZTP UI to configure DHCP configuration instead of manual changes into the DHCP configuration file.

Arista EOS device may leave VLANs with stale DHCPv6 relay entries In the event of a VRF change (AOS-50511)

In the case of a VLAN with DHCPv6 relay(s) configured in Arista EOS devices, changing the VRF may result in stale DHCPv6 relay commands remaining in the VLAN configuration. This is the result of an incorrect negation performed while the VRF was being changed.

Workaround

In the case of a small number of affected VLANs, manually deleting the DHCPv6 relays is probably the simplest solution. A full configuration apply against the device will also resolve the issue, but will cause traffic disruption.

BGP peers configured on IRB/Loopback interfaces may flap on Juniper EVO device during commit (AOS-59821)

An event involving BGP flaps from BGP peers configured on IRB/Loopback interfaces for Juniper EVO device has been reported during the commit with incremental changes (deleting an IRB/Loopback interface in the same VRF). The BGP flaps happen when the device's configuration is committed in override mode, which was Apstra's default setting prior to 6.1.0. However, using load update mode did not reveal the issue.

Workaround

If the EVO device is running as an off-box agent, please add key `load_mode` with value `update` into open options in the edit agent menu. Otherwise, recommend upgrading Apstra to 6.1.X, which uses load update as the default mode for commit.

Change Sonic MLAG rack to ESI rack is not supported in 5.0.0 (AOS-47476)

The change of the Sonic MLAG rack to the ESI rack is not supported in 5.0.0 as Day 2 operation. It is advised to rebuild a blueprint from the scratch using ESI rack.

Workaround

The customer is advised to rebuild the blueprint from the scratch using ESI rack.

Commit failures in the JUNOS and EVO devices are caused by newline characters in the virtual network description field (AOS-54198)

If the rendered device configuration for the VLAN description contains newline characters populated from the virtual network's description field, the commit operation fails because JUNOS and JUNOS-EVO devices do not support multi-line string for the VLAN description.

Workaround

Please use one-line formatted string in the description field of Virtual Network rather than multi-line string.

Configlet may not be applied in the SONiC device during the commit when system time moves back (AOS-46890)

When the SONiC device time is set back using the NTP configuration from configlet during the commit process, the device agent may reboot. When the agent reconnects, the device agent tries to reapply configuration changes that were interrupted by the previous commit. However, because the shell script file from configlet exists and matches the controller's information, it may be skipped rather than applied.

Workaround

Apply full push configuration into the SONiC device

Configlets using Jinja custom function.merge_vlans_to_list() fails with TypeError on datacenter interface device model allowed_vlans (AOS-52484)

When making use of the built-in jinja function on a custom configlet, `{{ function.merge_vlans_to_list(interface_model["allowed_vlans"]) }}` fails when used within a configlet. An error will be seen in rendered configuration previews "TypeError: unsupported operand type(s) for -: 'int' and 'str'"

Workaround

In the configlet, map the string values to integers, such as `{{ function.merge_vlans_to_list(interface_model["allowed_vlans"] | map("int")) }}`

Configuration anomalies in the SONiC device caused by the configlet when the device agent restarted after losing connection to the controller (AOS-50752)

While configlet is being applied to the SONiC device, if the device agent restarts after being disconnected from the controller, the agent executes any remaining changes and collects the running configuration as golden configuration to monitor for configuration anomalies. Because the process of applying configlet changes is still running independently of the agent, it introduces changes into the running configuration even when the golden configuration is collected by the agent. The following changes from the process cause configuration anomalies in the SONiC device.

Workaround

After reviewing the running configuration on the SONiC device, if all the changes from the configlet are correctly applied, the customer can safely accept changes to avoid further configuration anomalies.

CT(Connectivity Template) field may appear with empty value (AOS-56498)

While viewing/editing CT (Connectivity Template)s across blueprints, it's possible that the CT may incorrectly display empty field values.

Workaround

By clicking the browser refresh button, CT would display the correct data.

Dashboard shows 'Pending' Service Config for All Devices During Commits on Specific Device, with Delays in Larger Blueprints (AOS-51083)

Users experience confusion when committing changes for specific devices because the Dashboard shows 'Pending Service Config' for all devices, which can mislead them into thinking other devices are being updated as well. This is a known behavior in Apstra's current design. When a commit is made, all devices temporarily enter a 'Pending' state while the system determines which devices require changes. Even devices that don't need updates briefly show as pending, which can create the false impression that changes are being made. Additionally, as the number of devices in a blueprint increases, the delay becomes more noticeable because Apstra processes each device sequentially. This raises concerns about performance and efficiency when managing larger blueprints.

Workaround

There is no immediate workaround. The behavior is aligned with the current system design.

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.

Deleting Virtual Networks in CTs with Multiple VLANs - All Active Endpoints Unassigned (AOS-44623)

In version 4.2.0, Apstra introduces the capability for users to forcibly delete a Virtual Network, even if it has active endpoints. Apstra will initially display the interfaces to which the Virtual Network (VN) is currently allocated and prompt the user to confirm the deletion. It's important to note a limitation in the current design: if a user deletes a VN assigned in a CT where Multiple VLANs are present, all active endpoints will be unassigned.

Workaround

User should manually remove the specific VLAN from the CT before proceeding to delete it from the Staged > Virtual Networks section.

Deployment Performance degrades when draining or deploying (AOS-54006)

When the device is deployed or drained, Apstra showed a noticeably longer delay in finishing the operation than the Apstra 4.1.X release. The problem was linked to the significantly increased delay in the Jinja configuration rendering area following Apstra's migration from Python version 2 to version 3. Additionally, it affects the rendering configuration for the blueprint's configlet processing.

Workaround

Recommend upgrading to the Apstra 6.0.0 release, which addressed the issue. In case of 4.2.X customer, 2-step upgrade (4.2.X -> 5.0.1 -> 6.0.0) is required

Device NOS upgrade fails in Cisco modular devices (C9504, C9508) (AOS-49832)

The subsequent steps to gather pristine configuration based on the newly upgraded NOS and push full service configuration would fail, causing a service impact, even though the device NOS is upgraded during the NOS upgrade operation.

Workaround

After the switch upgrade fails, manually collect pristine and do a full configuration push

DeviceTelemetryAgent crash in the MAC Telemetry service for the JUNOS/EVO device (AOS-50058)

The JUNOS/EVO device uses GRPC for the MAC Telemetry service. During the GRPC processing, Apstra Controller uses device's credential information (username and password) to populate GRPC meta data. If the password includes non-printable ASCII characters, a validation error for invalid characters can lead DeviceTelemetryAgent to fail with a crash.

Workaround

Please use only printable ASCII characters for device's password to avoid validation error or use polling mechanism by disabling GRPC in the telemetry service

To disable GRPC service in the telemetry service, change `grpc_enabled = 0` in the `/etc/aos/aos.conf` file and then restart AOS service in the Apstra Controller.

```
[telemetry_global_config]
```

```
# Python multithreading enable/disable knob for telemetry collection
```

```
multithreading_config = 1
```

```
# Execution timeout for extensible telemetry collectors
```

```
command_timeout = 120
```

```
# Knob to enable/disable gRPC based service collectors
```

```
grpc_enabled = 0
```

```
# Space separate list of device models where gRPC based service collectors are
```

```
# disabled. The configuration is case insensitive. The device model can be
```

retrieved from Managed Devices page. Multiple models can be specified as:

ModelA ModelB ModelC

```
grpc_disabled_models = QFX5100-48T-6Q QFX5100-24Q-2P QFX5100-48S-6Q
```

DeviceTelemetryAgent.{pid}.log by gRPC trace logs filling up disk (AOS-51846)

DeviceTelemetryAgent.{pid}.log files in /var/log/aos/ in the offbox agents become large and can fill up the disk

Workaround

The following Python script can be added to run via crontab on an hourly basis, which will clean up older log files. This workaround needs to be applied to controller VM and worker VMs where offbox agents are running (nodes with offbox tags in the Platform/Apstra Cluster/Nodes).

```
# Copyright 2024-present, Apstra, Inc. All rights reserved.
#
# This source code is licensed under End User License Agreement found in the
# LICENSE file at http://apstra.com/eula

import configparser
import json
import os
import re
import shutil
import subprocess
import traceback

SystemIdPattern = re.compile(r'AOS_SYSTEM_ID=offbox, (.+), (.+)')

def update_aos_conf(task_id):
    aos_config = os.path.join(
        '/var/lib/aos/conf.d/task/offbox/',
        task_id,
        'aos.conf',
    )

    parser = configparser.ConfigParser()
    if os.path.isfile(aos_config):
        parser.read(aos_config)

    if not parser.has_section('logrotate'):
```

```

    parser.add_section('logrotate')

    if 'max_kept_backups' not in parser.options('logrotate'):
        parser.set('logrotate', 'max_kept_backups', '1')

        staging_file = aos_config + '.staging'
        with open(staging_file, 'w') as f:
            parser.write(f)

        shutil.move(staging_file, aos_config)
        return True

    return False

def refresh_logging_infra(container_id):
    subprocess.check_output([
        'docker', 'exec', container_id, 'pkill', '-HUP', 'DeviceKeeperAge'
    ])

def get_offbox_containers():
    containers = subprocess.check_output([
        'docker', 'ps', '-q', '--filter',
        'label=AOS_CLUSTER_APPLICATION=offbox',
    ]).decode()
    return containers.splitlines()

def get_task_ids():
    def extract_info(container_env):
        try:
            envs = json.loads(container_env)
        except ValueError:
            return None, None

        for env in envs:
            matched = SystemIdPattern.match(env)
            if matched:
                return matched.group(1), matched.group(2)

    return None, None

containers = get_offbox_containers()
if not containers:
    return

containers_env = subprocess.check_output([
    'docker', 'inspect', '--format', '{{json .Config.Env}}', *containers,
]).decode()
for line in containers_env.splitlines():
    task_id, container_id = extract_info(line)
    if not task_id:
        print('Failed to extract task id from container env:
        {}'.format(line))
        continue

```

```

        yield task_id, container_id

def main():
    for task_id, container_id in get_task_ids():
        try:
            if update_aos_conf(task_id):
                refresh_logging_infra(container_id)
        except:
            print('Failed to update aos.conf for container:
{}').format(container_id)
            traceback.print_exc()

main()

```

Even if the above workaround is applied, there is a chance of filling up partition. The below command can be executed with root permission to clean up logs quickly in the controller VM and worker VMs.

```

find /var/log/aos/task -name "*.log" -size +10M -print | grep
"DeviceTelemetry" | xargs -I {} sudo cp /dev/null {}

```

Disallow hyphens (-) in key and value names for telemetry service registry entries (AOS-50120)

Apstra introduced custom telemetry services in the 4.2.0 release. Users define a service schema to structure and store data, based on key and value from the CLI output. The UI doesn't allow hyphens in telemetry key and value names. However, the API allows them. If a telemetry service registry entry with a hyphen is created via the API, the upgrade to Apstra 5.x may fail with validation error.

```

File "/usr/local/lib/python3.10/dist-packages/lollipop/errors.py", line 182,
in raise_errors
    raise ValidationError(self.errors)
lollipop.errors.ValidationError: Invalid data: {'key': 'Unable to identify
"key" from schema'}

```

In Apstra 5.1.0, telemetry key/value names with hyphens are now disallowed.

Workaround

To resolve the issue, follow below steps:

1. Navigate to Analytics > Service Registry
2. Identify the service name that contains hyphens (-) in telemetry keys and telemetry values within the application_schema payload.
3. Edit the service entry:
 - Click the Edit action for the affected service
 - Replace all hyphens (-) with underscores (_)
 - Click the Update button to save changes
4. Retry the Apstra upgrade process by running `aos_import_state` again

Please contact Apstra Support for further assistance.

Event log page is rendering a blank page when clicking on the Apstra Web UI (AOS-50475)

Assigning a role with a space in the role name to a user and then removing it, while keeping the other role still assigned to the user, causes the "Event Log" page in the UI to display as a blank page.

Workaround

The hotpatch should be downloaded and run on the AOS server. Please contact the Juniper Apstra Support team for assistance with this process.

Expect traffic loss on NOS upgrade for breakout ports (AOS-46300)

While upgrading the Switch, all the front end ports are brought down till the intended configuration is pushed to the switch after upgrade. However if the switch had ports in non breakout mode when the switch was added to the fabric and subsequently it was converted to breakout ports such ports will be up briefly until the service configuration is pushed to the switch from AOS.

gRPC junos ephemeral database UI_EPHEMERAL_COMMIT filling partition (AOS-52519)

Apstra gRPC probing of JUNOS devices is causing large ephemeral database files which may fill the disk and cause issues accessing the device via SSH.

```
jtac-QFX5120-48Y-8C-r011 mgd[14126]: UI_EPHEMERAL_COMMIT: User 'root' has
requested commit on 'junos-analytics' ephemeral database
jtac-QFX5120-48Y-8C-r011 mgd[14126]: UI_EPHEMERAL_COMMIT_COMPLETED: commit
complete on 'junos-analytics' ephemeral database
```

Workaround

Add the following to JUNOS devices' pristine config or to the configlet to reduce the number of stored versions within the ephemeral database to reduce its size. Even if JUNOS 23.2R1 introduced the recycling feature of the ephemeral database, the issue can still happen in the Apstra Qualified JUNOS (not EVO) version, 23.4R2-S4.

```
set system configuration-database ephemeral purge-on-version 30
```

gRPC Periodic Response Timeouts in the Interface Telemetry Service (AOS-56175)

Because the JUNOS device might not send the full snapshot of interface-related data per reporting interval (default interval = 120 seconds) after initial synchronization, Device Telemetry Health detects continuous anomalies for gRPC Periodic Response Timeout in the interface telemetry service in a high-scale environment. The problem still exists even if the reporting interval is extended.

Workaround

It is advised to disable gRPC globally, which forces all gRPC-related telemetry services (interface, MAC) to switch to polling mode rather than gRPC, since the problem may occur at random on several devices.

To disable the gRPC service globally, change `grpc_enabled = 0` in the `/etc/aos/aos.conf` file and then restart AOS service in the Apstra Controller.

```
[telemetry_global_config]
# Python multithreading enable/disable knob for telemetry collection
multithreading_config = 1
# Execution timeout for extensible telemetry collectors
command_timeout = 120
# Knob to enable/disable gRPC based service collectors
grpc_enabled = 0
```

gRPC Sequence Number Overrun in the MAC Telemetry Service (AOS-56282)

Because the MAC telemetry service uses the gRPCOnChange mode, the device only sends

updates after initial synchronization. When the JUNOS device subscribes to the PATH (/network-instances/network-instance/mac-table/entries/entry) for MAC telemetry service, Apstra's gRPC client (Apstra) receives the first full data from two processes (l2ald, l2aldTM). During the initial synchronization, these processes use their own sequence number range (duplicate range), which makes the gRPC client think that the gRPC packets may be dropped internally. Granular sequence number handling will be introduced in 6.1.0 to address the existing sequence overrun issue.

Workaround

It is advised to disable gRPC globally, which forces all gRPC-related telemetry services (interface, MAC) to switch to polling mode rather than gRPC, since the problem may occur at random on several devices.

To disable the gRPC service globally, change `grpc_enabled = 0` in the `/etc/aos/aos.conf` file and then restart AOS service in the Apstra Controller.

```
[telemetry_global_config]
# Python multithreading enable/disable knob for telemetry collection
multithreading_config = 1
# Execution timeout for extensible telemetry collectors
command_timeout = 120
# Knob to enable/disable gRPC based service collectors
grpc_enabled = 0
```

gRPC telemetry services for mac and interface fails in the JUNOS device running 23.4R2-S3 (AOS-50216)

If the gRPC probe used to check the device's gRPC health status keeps failing, Apstra gRPC Telemetry services remain in a failed state on the JUNOS device running 23.4R2-S3. When the device is restarted, the issue becomes apparent. Even if a gRPC probe request is received, the device doesn't respond with the data. Because Apstra gRPC probe doesn't have a timeout mechanism for the gRPC probe request, Apstra gRPC continues to use the same TCP connection, which may have issues.

Workaround

Restarting offbox container via `"docker restart <offbox-container-name>"` in the controller or worker VM where offbox container is running

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.

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

Jinja configlet using interface shows no actual changes in the UI even if diff exists (AOS-48906)

When importing a configlet, even if there are uncommitted changes from other Jinja-based

configlets based on interface information, the logical diff output for the changes does not match the full node diff output. If the Jinja-based expression using interface uses an empty context, rendering will fail with empty information, resulting in an empty output with changes.

Workaround

For all Jinja-based configlets that use interfaces in the device context, wrap existing jinja statements with jinja if conditional statements to check interfaces are not empty for the configlet in the global catalog, and then remove/add the configlet to the blueprint to reflect the new changes.

```
==== original configlet using interface ====
{% for intf in interface.values()|selectattr('role', 'eq',
'l2edge')|rejectattr('part_of')|map(attribute='intfName')|list|sort %}
{% if loop.first %}
protocols {
    rstp {
{% endif %}
        replace: interface {{intf}};
{% if loop.last %}
    }
}
{% endif %}
{% endfor %}

==== corrected configlet using interface ====
{% if interface != '' %}
{% for intf in interface.values()|selectattr('role', 'eq',
'l2edge')|rejectattr('part_of')|map(attribute='intfName')|list|sort %}
{% if loop.first %}
protocols {
    rstp {
{% endif %}
        replace: interface {{intf}};
{% if loop.last %}
    }
}
{% endif %}
{% endfor %}
{% endif %}
```

Junos Device Profile and Linecard Profile "validations" field schema in interface settings is not fully enforced (AOS-50305)

Apstra 5.0.0 has added to some Juniper Device Profiles and Linecard Profiles a mechanism (using the validations field in the interface setting) to describe port group constraints and alert the customer of possible port breakout combinations that are prohibited. Because the schema

validation of these constraints is not correctly enforced, a customer may be able to create a custom device profile with incorrect or incomplete port group constraints, which could result in an aborted rendering of the Junos device configuration.

Workaround

Since the problem described could potentially involve a customer editing interface settings inside a Custom Device Profile in order to implement their own port group constraints, there is no workaround recommended. It is advised to avoid entering any invalid information in the validation field.

License information is not transferred during Apstra controller upgrade (AOS-52927)

The license information from the previous version of the Apstra controller is not transferred to the upgraded controller when it is upgraded. As a result, license information is missed in the upgraded Apstra controller.

Workaround

License files need to be copied from the old controller to the upgraded controller via the below command, and then restart the license container in the upgraded controller.

```
sudo scp -r admin@{old_controller_ip}:/var/lib/aos/.license /var/lib/aos/  
docker restart aos_license_1
```

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 is to restart the AOS services.

MAC anomaly not showing in the active > Anomalies tab (AOS-49936)

Even if it is shown in the Device> Managed Devices> Telemetry> Anomalies tab, the MAC anomaly was not showing in the active > Anomalies tab. It was identified that UI request to backend doesn't include anomaly type for MAC.

MAC query fails with type exception "expected TAC::String, got NoneType" (AOS-50112)

MAC querying might have crashed because it added all the systems to the index, although some systems might lack system_id if their deploy_mode is not deployed. Therefore, it led to KeyError as None (system_id) is not a valid value for the "TAC::String" field

Workaround

Please assign systems, where virtual network is associated, into deployed and commit blueprint

MetricDb migration fails silently for multi-step migration. Eg. 4.2.1 -> X -> Y, data is lost for release Y (AOS-54413)

Upgrade script in 4.2.2 and following releases introduced a regression that manifests itself in multi-step migration scenarios.

Step-1. Upgrade from 4.2.1 -> X (eg. 4.2.2) - causes the permission on '/var/lib/aos/metricdb/iba' folder and it's sub-directories to be 700.

Step-2. Upgrade from X (eg. 4.2.2) -> Y (eg. 5.0.1) - causes the silent failure in step that copies '/var/lib/aos/metricdb' folder and ALL its sub-directories to new VM.

The impact of this failure is that the 'Audit', 'IBA stage history' and 'Aos cluster health history' data is lost in the final upgraded AOS instance. The data from the previous release will be lost subsequently if there are further migration steps involved.

This issue affects all the releases starting upgrade from 4.2.2. If your upgrade source Apstra is at least 4.2.2, please apply the workaround suggested BEFORE performing the upgrade.

Workaround

Apply workaround fix(aos_54413_fix_metricdb_permissions.run:

<https://supportportal.juniper.net/sfc/servlet.shepherd/document/download/069Dp00000Gc0kCIAR>

) to the old version Apstra Controller Node * BEFORE * every upgrade, following the below steps.

1. Copy the bundle aos_54413_fix_metricdb_permissions.run to the source (old) Apstra controller node. The tool expects Apstra service to be running because it needs to get cluster node information from Sysdb.

2. Make it as executable and execute the bundle as sudo

```
admin@aos-server:~$ chmod 755 ./aos_54413_fix_metricdb_permissions.run
admin@aos-server:~$ sudo ./aos_54413_fix_metricdb_permissions.run
Verifying archive integrity... All good.
Uncompressing Fix for AOS-54413 for AOS >= 4.2.2 100%
AOS[2025-05-25_19:36:22]: Fixing controller node
AOS[2025-05-25_19:36:23]: Getting cluster node metadata
AOS[2025-05-25_19:36:24]: Fixing worker node: 10.28.75.6
Logs have been collected at:
/home/admin/aos_54413_fix_logs_20250525_193623.tar.gz
```

3. The absence of any errors means that the issue has been fixed. In case of errors during execution, please reach out Juniper Apstra Support Team.

If AOS instances upgraded without work-around and if old apstra VM is preserved, Contact Juniper Apstra support team to help with migrating MetricDB data.

Multiline banner motd or exec is not supported in the Cisco NXOS Device (AOS-40278)

A banner configured in the Cisco NXOS device must be single line. Multiline banner (motd or exec) is not supported.

Workaround

Configure single line banner (motd or exec)

NOS upgrade fails in the device running Arista EOS 4.30.3M (AOS-50386)

After installing a new image during the NOS upgrade process, a reboot command with the newly installed image is executed on the device running Arista EOS 4.30.3M. The device running 4.30.3M remains online for a significantly longer period of time without rebooting than the other release. As a result, the system agent in charge of the NOS upgrade job makes the assumption that the device has successfully rebooted with a new image. The NOS upgrade fails version comparison validation since the device still runs the old version without rebooting.

Workaround

For the device running Arista EOS 4.30.3M, manual NOS upgrade is advised rather than Apstra NOS upgrade via UI.

Onbox Device Agent Not Supported in Dual Routing Engine Junos-Evolved Devices (AOS-

43980)

It is not possible to install and use the Apstra onbox device system agent for Juniper dual routing engine devices running Junos-Evolved version 22.4R2.

Workaround

Please use the Apstra offbox device system agent.

Onbox or Offbox agents configured with FQDN as the management IP address cause deployment to be stuck (AOS-50888)

For deployment, the system agent attempts to map the agent's ID (UUID) to the system serial number using two step resolutions: system agent ID to management IP address mapping configured when the agent is created, and management IP address to system serial number mapping during device fact collection. When a system agent is configured as a FQDN rather than an IP address for the management IP address, the resolution from the agent's ID to the system serial is incorrect, causing deployment to stall until it is resolved.

Workaround

Recommend using IP address instead of FQDN for management IP address when device agent is created

PFE(Packet Forwarding Engine) in the QFX5120 platform restarts during NOS upgrade (AOS-57490)

When the Junos EVPN Next-hop and Interface count maximums parameter in the staged->Fabric settings->Fabric-policy is enabled, Apstra introduced modifying the default hardware settings for VXLAN routing's resource (next-hop and interfaces) for QFX5110, QFX5120, EX4650, and EX4400 devices in the rendered configuration () starting with version 4.2.0. Whenever configuration changes in VXLAN routing's resource, JUNOS triggers PFE automatic restarts to reflect new changes with service impact. The typical scenarios would be when the device becomes deployed, undeployed, or the device is in NOS upgrade. To prevent unnecessary PFE restarts in those scenarios, the configuration for VXLAN routing's resource needs to be included in the pristine configuration.

Workaround

If the Junos EVPN Next-hop and Interface count maximums parameter in the staged->Fabric settings->Fabric-policy is enabled, add the below configuration into the device's pristine configuration.

QFX5120 and EX4650 VXLAN routing's resource

```
forwarding-options {
  vxlan-routing {
    next-hop 45056;
    interface-num 8192;
    overlay-ecmp;
  }
}
```

QFX5110 VXLAN routing's resource

```
forwarding-options {
  vxlan-routing {
    next-hop 32768;
    interface-num 8192;
    overlay-ecmp;
  }
}
```

EX4400 VXLAN routing's resource (add overlay-ecmp if Junos EX-Series Overlay ECMP is also enabled)

```
forwarding-options {
  vxlan-routing {
    next-hop 16384;
    interface-num 6144;
    overlay-ecmp;
  }
}
```

PipelineAgent crash with invalid key and query expansion configuration in Extensible Service Collector (AOS-49084)

Any custom probe that utilizes the Extensible Service Collector processor must now have explicitly string types for its property key value types due to additional validation constraints. If this is not done, the UI will display a validation error pointing to the offending property key.

Workaround

For cases when the validation could not infer the value type of property keys, wrapping the property key value definition with str() should be enough to resolve the issue. I.e., if a particular property key was defined as "interface.label," then it should be changed to "str(interface.label)".

Pristine Config Update Fails with "System Already Parsed" Error for Junos Devices (AOS-52788)

Customers may encounter the following Server-side Validation Error in the Web UI when the pristine configuration contains multiple system stanzas which is not an expected behavior:

```
"Cannot parse config: system already parsed."
```

According to ScotchInventoryAgent logs, POST requests to update the pristine configuration failed with a 422 Unprocessable Entity error, indicating a validation issue:

```
2025-02-17 23:31:18,730 680:INFO:aos.scotch.libs.scotch_flask:request: POST
/api/systems/AN10555621/pristine-config HTTP/1.0 34074 bytes
2025-02-17 23:31:18,737 680:INFO:aos.scotch.libs.scotch_flask:response: 422 55
bytes 0.007347 seconds
```

Background of the issue:

1. In Apstra 4.2.x, gRPC was introduced to support Telemetry Streaming, and as a result, having two system blocks in the pristine configuration was expected in that release.
2. Starting from Apstra 5.0.0, enhancements were made to automatically merge multiple system stanzas in the pristine configuration during the NOS upgrade process.
3. If a customer chooses to remain on their current NOS version for an extended period, multiple system stanzas can exist in the pristine configuration without causing issues.

Workaround

If a customer chooses to remain on their current NOS version for an extended period and needs to forcefully update the pristine configuration, they should manually merge the system stanzas within the pristine configuration using the UI and then perform a Force Update.

For further assistance, please contact Juniper Apstra Support.

Remote MAC expectations are incorrectly suppressed on MLAG racks in SONiC when no VN endpoints are present (AOS-49085)

Every leaf (device that is a part of the VN network) device has its EVPN Type2 expectations validated by `Mac Monitor Probe`. By using EVPN type 2 advertisement ingestion, it indicates as missing the MAC addresses that should be on the devices but aren't.

It is observed in a SONiC environment that a device does not ingest the type 2 advertisements to the corresponding VLAN forwarding (MAC) table if it is not connected to a host for the given VN. In the event that the device is a part of the MLAG pair, the type 2 advertisements are ingested into the Mac Table even if it is not host-attached.

As of right now, the MacMonitor probe ignores any device from the previously mentioned type 2 validation that has no host attached. As a result, SONiC devices that are part of the MLAG but lack VN endpoints (hosts attached) are excluded from the type 2 missing MAC check. As a result, on such devices, the probe might not find any missing MACs.

However, in MLAG scenarios, no traffic is expected for the specified VLAN on the rack if a leaf device has no connected hosts, i.e., VN endpoints, so this issue has no functional impact.

Workaround

None

Rotation of `frr-reload.log` inside the `bgp` container for SONiC Device (AOS-49921)

Every time a config apply happens in a SONiC device managed by Apstra, the FRR daemon configuration is gracefully reloaded by the `frr-reload.py` script inside the `bgp` container. The output of that script is directed to the file `/var/log/frr/frr-reload.log` inside the same container. The size of that log file is not expected to ever become a concern, unless a customer performs many thousands of config apply operations with a rather large FRR configuration.

Workaround

If the rotation of the `/var/log/frr/frr-reload.log` inside the `bgp` container is desirable, Apstra 5.1.0 includes a predefined configlet that can activate an appropriate `logrotate` cronjob inside the `bgp` container in regular intervals. The customer can use and/or modify the configlet and use it in their blueprint(s). Please contact Juniper Apstra support if more help is required.

For versions of Apstra earlier than 5.1.0, the same configlet can be manually created and used by the customer. Please contact customer support for further details.

Route anomalies caused by incorrect expected nexthops for leaf loopback address in the 5 Stage Clos topology (AOS-49802)

The expected routes for the loopback address of the leaf node in the spine node are calculated by

using pod_label to determine whether the target leaf node and the current spine node are in the same pod. When the pod label is updated by UI, the ExpectationRenderer Agent may not update the new pod label information into all spine and leaf nodes, resulting in including the wrong nexthops into superspine nodes for leaf loopback address, even if the leaf node is directly connected from spine node.

Workaround

Please execute `sudo service aos restart` to restart ExpectationRenderer Agent

Some probe stages(EVPN Host Flapping per System stage) may not yield the right output based on the chosen aggregation method (AOS-47451)

In some cases, the frontend UI generates incorrect queries (without specifying per-metric aggregation) to retrieve time series data for stages. When this happens, the backend selects the default aggregation method based on the value type, which means that instead of no aggregation, average aggregation is used by default. As a result, the output does not match the aggregation method used.

SONiC BGP route collector may fail because of stale VRF entries in the FRR routing daemon (AOS-49833)

In some cases where a VRF has been created, used, and then deleted, the FRR bgpd daemon may still indicate the existence of that VRF. Example, the Vrf-PURPLE in this vtysh output:

```
leaf2# show vrf
vrf Vrf-PURPLE inactive
vrf Vrf-blue id 120 table 1001 (configured)
vrf Vrf-red id 122 table 1002 (configured)
vrf mgmt id 47 table 5000
```

The existence of Vrf-PURPLE confuses the Apstra BGP route collector, causing it to crash. In such a case, the BGP route telemetry will stop working.

Workaround

`service bgp restart` in the affected device has been observed to remove the stale entries and thus restores the operation of the Apstra BGP route telemetry. Please do note that doing such a restart will flap all BGP peerings and can momentarily affect traffic.

SONiC Device Profile for Dell S5448F-ON doesn't enforce device constraints for the total number of logical ports per pipeline (AOS-47664)

The Apstra Device Profile for the Dell S5448F-ON is not enforcing the device constraints for the total number of logical ports per pipeline. Customer should be aware of the constraints around pipelines in the specific model and should take care to not exceed 18 logical ports per pipeline.

For further information, please refer to the documentation of the vendor.

SONiC FRR restart or device reboot may cause configuration anomaly from rearrangement of FRR running configuration sections (AOS-49906)

Rebooting the device or restarting FRR in SONiC may cause the FRR running configuration sections (related with route-map) to be rearranged. The rearranging of sections will typically show a configuration deviation even if the running configuration is exactly the same as before.

Workaround

The anomaly can be eliminated by the user reviewing the deviation and accepting the changes. No further action is necessary.

Sonic ZTP cosmetic console python traceback error messages (AOS-49897)

The device console may display cosmetic Python traceback error messages indicating an unreachable network when Sonic ZTP is provisioning a device. In the event that the device modifies the VRF of its management interface and is unable to send logging messages to the ZTP server, this may occur.

Workaround

If the device completes the ZTP process, these messages in the console can safely be ignored.

Stalled poll timers caused by device reset can lead to the agent restart (AOS-49046)

This is a rare case in which a device reboots during a gRPC session, resulting in stale polling timers on the Apstra Agent side. When gRPC restarts, the stale timers are replaced by new timers, which trigger the handling timer for collection, resulting in an agent crash. After the agent restarts from the crash, the system functions normally without any further crashes.

System agent in the dual-re Junos EVO system may not work correctly when routing engine master switchover happens (AOS-43956)

if new system agent (onbox or offbox) for a Junos EVO system, which has dual routing engines, is not created with a master-only address, when routing engine master switchover occurs, the system agent and the device agent may not work correctly or introduce problems.

Example dual routing engine configuration:

```
re0:mgmt-0 {
  unit 0 {
    family inet {
      address 10.49.110.35/19;
      address 10.49.100.226/19 {
        master-only;
      }
    }
  }
}
re1:mgmt-0 {
  unit 0 {
    family inet {
      address 10.49.108.203/19;
      address 10.49.100.226/19 {
        master-only;
      }
    }
  }
}
```

In the above example, the common address for routing engines is 10.49.100.226. Installing against other management interface addresses will initially work, but will cause serious problems for the system agent if and when a routing engine master switch occurs.

Workaround

Please use the master-only address that is common in both routing engines' management interfaces when creating a system agent. For further support, please contact the Juniper Apstra Support Team.

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

The selected Device Profile for the device matching multiple Device Profiles may change after upgrading to Apstra 5.0.0 (AOS-46375)

If there are multiple device profiles for the same device model, upgrading to Apstra 5.0.0 may result in the selection of a different Device Profile due to new deterministic logic in the code. In rare cases, the Apstra device manager may identify the device as having a Device Profile that differs from the original Device Profile before the upgrade. If the device is used in the Blueprint, the change of the Device Profile causes it to enter an error state.

Workaround

The user can simply edit the assigned Device Profile from Managed Devices to use the correct Device Profile, which corresponds to what is used in the Blueprint for the same device.

The SONiC device with the Trident 4 chipset may experience traffic loss in the collapsed fabric topology (AOS-47356)

Apstra enables wide mode in any Trident 4 chipset for any SONiC devices (Z9432F-ON, S5448F-ON) that support ESI, as long as the CLOS topology uses ESI. If the topology is collapsed rather than CLOS, the Apstra reference design contains an error that prevents wide mode from being automatically enabled as needed. This can occasionally cause traffic drops for packets entering the fabric.

Workaround

It is possible to turn on wide mode manually. The below section must be added to the pristine configuration of any leaf device of the collapsed fabric. Then, a full configuration apply must be performed on the device. It should be noted that applying the full configuration is a traffic-affecting operation.

```
"SWITCH_RESOURCE_TABLE": {  
  "L2_NEXTHOP_GROUP": {  
    "l2-nexthop-group": "ENABLED"  
  }  
}
```

}

The UI page remains stuck in a loading state when navigating from the Stage tab to the Uncommitted tab (AOS-47430)

Upon navigating from the Stage tab to the Uncommitted tab, the page does not load and continues to display the loading spinner. The UI is continuously polling the API endpoint `/api/blueprints/tasks?mode=full`, which forces the backend to attach the complete task payload to every response. Since these payloads can contain large data information, the API responses become significantly heavier. Because the UI repeatedly polls this endpoint, it increases backend processing time and causes the page loading issue, particularly when tasks have large payloads.

UI doesn't allow IM (Interface Map) creation with mixed transformations for the same speed (AOS-47654)

When the Device Profile allows multiple transformations for the same speed and the IM (Interface Map) needs mixed transformation for the same speed across ports (for example, Port 1: 1X100G, Port 2: 4X100G, etc.), it is not possible to create an IM using the UI.

Workaround

Recommend using the same transformation for the same speed across ports in the IM (Interface Map) when IM is created in the UI

UI show tech collection for controller fails with "Show tech execution failed on target apstra_vm" (AOS-51964)

The current default timeout for collecting show tech via UI is 20 minutes (1200 seconds) in the `controller_timeout` value of the `show_tech` section of aos configuration file `/etc/aos/aos.conf`. The collection contains operations for exporting all binary log files from all running agents across containers into printable format files. Depending on the number of accumulated log files, the translation processing time may exceed the default timeout for collecting showtech, resulting in timeout failures for the UI showtech collection task.

Workaround

The timeout value for `controller_timeout` under the `show_tech` section of `/etc/aos/aos.conf` file should be increased to a higher value, such as 1 hour and 30 minutes, followed by restarting aos service (`sudo service aos restart`). Otherwise, a manual command (`aos_show_tech`) can be used to collect controller show tech.

UI showed server error: null while displaying racks (AOS-48937)

When displaying racks, Apstra generates statistical information for the rack by sorting through each leaf's position data in ESI or MLAG cases. When a rack is updated by inserting a generic system into one of the leaf nodes that make up ESI/MLAG, Apstra uses sorting criteria based on the label from the leaf nodes. This inconsistent sorting criteria leads to calculation statistics referring to non-existent keys, resulting in errors. This issue only arises when a generic system is connected to one leaf node of an ESI/MLAG pair via a single attachment for a specific speed and the other leaf node lacks an interface for the same speed.

Workaround

To avoid missing key issues, add a generic system at the same speed to the other leaf node in the same ESI/MLAG pair.

Unable to delete a Datacenter Blueprint entry from the Blueprint page when Apstra fails to create an unsupported blueprint (AOS-52300)

Users will not be able to create a Datacenter Blueprint using the EVPN reference design with IPv6 RFC:5549 as the Spine to Leaf Links Underlay Type, as this configuration is not currently supported. This limitation is tracked in RFE-1364 which is planned for release in 6.1.0. If the user attempts to create an unsupported blueprint using the steps below, the UI will throw a critical error message - EVPN is not allowed with spine-leaf IPv6 links addressing policy.

Steps to Reproduce:

1. Create a template with the overlay control protocol set to MP-EBGP EVPN.
2. Navigate to the Blueprint Creation page.
3. Select Datacenter as the reference design.
4. Set Spine-to-Leaf Links Underlay Type to IPv6 RFC:5549.
5. Attempt to create the blueprint.

However, the blueprint becomes locked and cannot be deleted, even with admin privileges. Deletion can only be performed via REST API Explorer or CLI. This issue was addressed in 6.0.0, ensuring that the delete function works correctly.

Workaround

Please follow the steps below to delete the failed and locked blueprint:

1. Navigate to Platform -> Developers -> REST API Explorer.
2. In the REST API Explorer, go to blueprint -> {blueprint_id} -> DELETE.
3. Select the locked blueprint ID and click Try It.
4. Expect a success response code to appear on the same page.

Please contact Apstra support for assistance if you encounter any issues during this process.

Unintended advertisement of all fabric VTEP loopbacks to external routers in default routing zone in VXLAN DCI environment (AOS-54864)

Integrated DCI feature(vxlan stitching) was introduced in Apstra 4.2.0. In versions 4.2.x and later, customers using this feature may encounter an issue where all VTEP loopback addresses from the fabric including those from non-border leaf devices are being advertised to external routers over BGP in the default routing zone.

This affects only VXLAN DCI Stitching deployments(Stitching requirement for VTEP loopbacks for only border leaf nodes vs OTT requirement for all VTEP loopbacks in the fabric). Even when customers configure routing policies to export only loopback of border leaf nodes, Apstra backend logic automatically includes all VTEP loopbacks. Due to the current design, Apstra does not differentiate between border and non-border leaf roles in this context, resulting in the unintended advertisement of all fabric loopbacks to external peers.

Engineering has confirmed this as a bug. The expected behavior is to advertise only the loopback addresses of border leaf switches to external routers in the default routing zone. There is no official workaround to modify this behavior through standard configuration. Engineering is actively working on a fix to address this issue in a future release. The only option is to use a custom configlet to override Apstra default export logic. Please reach out to Apstra Technical Support for assistance.

Update Link Speed to 10M in the Juniper EX4400 device not allowed in the UI (AOS-50182)

When 10 Mbps speed is selected via Update Link Speed, the user interface (UI) disables the update button so that it cannot be applied, even though the interface supports 10 Mbps.

Upgrade to 5.0.0 fails when JUNOS device's pristine config has system login message or announcement with # characters (AOS-49768)

The upgrade to 5.0.0 validates the device's pristine configuration. When the pristine configuration for the JUNOS device contains system login message or announcement with # characters, upgrade validation fails with an error, resulting in the upgrade failing. When performing a NOS upgrade or device agent upgrade in 5.0.0, the same validation error may occur.

Workaround

Please remove # characters in the system login message and announcement by updating pristine in the UI. For further support, please contact the Juniper Apstra Support Team.

Virtual Infra manager's information not cleared even if virtual infra manager is removed from Apstra Controller (AOS-53538)

When the virtual infra manager is removed from the Apstra controller, Apstra should have cleared any data related to the virtual infra manager. Because it's not cleared, when the same virtual infra manager is added back to Apstra later, old data is still used together with the new collected data from the virtual infra manager's collector. In some scenarios, when old, uncleaned data has an error condition, it can trigger continuous error even if newly collected data doesn't have an error condition.

Workaround

If the virtual infrastructure manager requires re-onboarding (removing and then adding back) from Apstra, the user must take the actions listed below.

1. Remove virtual infra manager from Apstra Controller (External Systems/Virtual Infra Managers).
 2. Restart the AOS service.
 3. Add the virtual infra manager back to to the Apstra
-

Virtual Network configuration changes do not properly reflect as changed after making a change in Apstra (AOS-40852)

After editing a VN (Virtual Network) configuration, if the same Virtual Network is open, none of the changes appear until the VN is opened again in the UI.

Workaround

After Saving the changes to the VN simply close the VN configuration and open it again. The second opening of the VN configuration page should reflect all changes that have taken place.

Virtual Network Endpoint View in the UI showing empty information (AOS-53783)

Since the UI misses polling of the node detail information to the Apstra backend, the Virtual Network Endpoints view of Generic System Node (Staged > Physical > Topology > Virtual Networks Endpoints) shows empty information.

Workaround

Refresh Web page in the browser to make the UI send requests explicitly to collect data

VirtualInfra telemetry service shows error message when a PNIC is unassigned from the VDS (Virtual Distributed Switch) (AOS-53537)

Apstra creates a relationship between the PNIC and the Link Discovery Policy (which determines which discovery protocol is used) configured in the VDS when a PNIC is assigned to a VDS (Virtual Distributed Switch). One PNIC may inadvertently become linked to two relationships without clearing out the previous relationship when a user moves a PNIC directly from one VDS to another VDS. An error message below appears when the PNIC becomes unassigned from VDS because there is more than one relationship between the PNIC and Link Discovery Policy that is invalid.

```
virtual_infra failed to collect data, plugin raised exception: {'item_iter':  
<aos.sdk.graph.graph.RelationshipIterator object at 0x7f60c03f9570>, 'items':  
[df57a2f0-969c-4dee-9831-a53526bd7d5a-[:policy]->4c8c4c31-df9a-4188-933f-  
6b5d67703a1f, df57a2f0-969c-4dee-9831-a53526bd7d5a-[:policy]->1787d662-4a41-  
4b8b-9b77-acac5508e771]}
```

Workaround

Instead of performing one direct migration action from one VDS to another, the problem can be avoided by two actions: unassigning the PNIC from the old VDS and then assigning it to the new VDS.

Procedures for fixing the errors as a workaround

1. Remove the virtual infra manager from not only the blueprint but also the External Systems/Virtual Infra Managers.
 2. Restart the AOS service.
 3. Add the virtual infra manager back to the External Systems/Virtual Infra Managers and then blueprint.
-

VNI column Hyperlink for MAC entry from the VLAN Type of the MAC Monitor Probe shows no associated virtual network in the Active Virtual Network (AOS-49081)

When a MAC entry is learned via Virtual Network, the VNI column in the MAC Address Table of MAC Monitor Probe includes a hyperlink to show details for the associated virtual network. In the case of a VLAN type Virtual Network, the VNI value is displayed as 0 rather than NA (Not

Available), and this is used to generate the incorrect filter for the Active->Virtual->Virtual Network Tab, which matches Virtual Network with VNI value with 0. As a result, the tab does not display any associated virtual networks.

When Generic System is added/deleted from leaf device in the rack, fail with an error not enough ports on leaf to connect (AOS-51116)

When a rack is built with leaf devices and generic systems, group labels for generic systems can have the same value as leaf or access switches' target_switch_label, contrary to the expectation that the group label should not be the same value as target_switch_label inside the rack. Any changes to the rack, such as adding a generic system or deleting an existing generic system, would fail due to the validation error caused by not meeting the above expectation.

Workaround

Change the group_label of the generic systems and the label in the rack_type_json to a different label that does not match the target_switch_label. For further assistance, please contact the Juniper Apstra Support team.

Worker nodes may remain in a failed configuration state after connectivity issues, requiring manual intervention for recovery (AOS-60167)

With current cluster design, worker nodes may fail to recover their configuration state after a temporary loss of SSH connectivity to the controller. This condition can be observed in the UI by navigating to Platform > Apstra Cluster > Nodes > Worker, where the following error may be displayed:

```
Configuration Error:  
ssh: connect to host 10.28.17.4 port 22: Connection refused
```

When the controller (ClusterManagerAgent) attempts to push configuration to worker nodes, it retries SSH connections up to three times. If all attempts fail, due to connection refused or no route to host, the node is marked with a FAILED Configuration State. Once this state is set, the system does not automatically retry configuration, even if SSH connectivity is later restored. Although worker nodes may resume sending keepalives and transition back to an active operational state, the configuration state remains in failed state indefinitely. Due to this the overall node state may continue to appear FAILED despite restored connectivity.

Workaround

Manually trigger a configuration synchronization using one of the following methods:

1. Navigate to Platform > Developers > REST API Explorer
2. Execute REST API: POST /api/cluster/worker/sync

[OR]

1. Restart AOS from the controller VM: `systemctl restart aos`

Known Third-Party Issues

ACX platform doesn't support export sflow over mgmt instance (AOS-58680)

When sFlow collector is configured with `mgmt_instance`, the configuration will be ignored in the ACX platform with a warning such as the below example.

```
sflow {
  polling-interval 10;
  sample-rate {
    ingress 10000;
    egress 10000;
  }
  source-ip 10.217.6.15;
  collector 10.217.0.165 {
    udp-port 6343;
    ##
    ## Warning: statement ignored: unsupported platform (ACX7024X)
    ##
  }
  routing-instance mgmt_junos;
```

Workaround

Please use a non-management instance for exporting sFlow until the ACX platform supports a management instance for sFlow export.

Anomalies are raised for interfaces on Juniper EX4400-48T devices running JUNOS 22.4R3 (AOS-56571)

Anomalies are raised due to mismatch in the operational status of interfaces due to interface status showing "unknown" on Juniper EX4400-48T devices running Junos 22.4R3.

Workaround

Restart the Apstra AOS service to collect the right interface status information. This issue is not observed in higher JUNOS versions. Recommend upgrading to an Apstra-qualified higher JUNOS version ($\geq 23.4R2-S4$).

Cisco N9K-C93600CD-GX rollback fails when using breakouts (AOS-47891)

The NXOS rollback feature on the N9K-C93600CD-GX device has significant limitations when the devices' interfaces are broken out.

Ports 1-24 in the model are organized into four-port groups: (1, 2, 3, 4), (5, 6, 7, 8), (9, 10, 11, 12), (13, 14, 15, 16), (17, 18, 19, 20), and (21, 22, 23, 24). When port 1 is broken out as 4x10G or 4x25G, port 3 is automatically broken out in the same mode, and vice versa. When any port in the quadruple is split into 2x50G, all four ports are automatically split in the same mode. Similarly, ports 26-28 are organized in pairs of two, i.e. (25, 26) and (27, 28). Both ports in the pair must operate in the same breakout mode.

In most cases where a breakout (or more than one) exists, rollback fails to generate a working rollback patch. The reason for this is that the breakouts cannot be reversed if the remaining broken-out interfaces in the same port group have not been shutdown first. For example, to negate the breakout of port 1, the broken-out interfaces of port 3 must be shutdown, and vice versa. It appears that the rollback logic shuts down the interfaces associated with the port whose breakout is being reverted (port 1 in the previous example), but fails to shut down other broken-out ports in the same port group (port 3).

Workaround

The safer way for the N9K-C93600CD-GX to be used with AOS is for the customer to avoid using breakouts altogether on the device.

No issue with rollback when ports 29-36 have been broken out has been observed. Breakouts on these ports can be rolled back

In the case that the last interface of a port-group is the only one used and broken out, would the nxos rollback feature (and rollback to pristine) be successful. However this is highly discouraged. In any other case the only way to reverting to pristine would be to manually shutdown all broken down interfaces before reverting to pristine (or using the rollback to a pristine config)

Cisco NXOS 10.3.4a qualified NOS upgrade path may require intermediary upgrade version (AOS-46319)

Previously, in the Apstra 4.1.2 release, Apstra qualified the 10.1.(2) version. If a device is

operating on this version and the upgrade version is 10.3.4a, the Zero Touch Provisioning (ZTP) would fail. This failure occurs because the image transition from 10.1.2 to 10.3.4 was not successful.

For instance, if the starting version is from 10.1.(2), the officially recommended Network Operating System (NOS) upgrade path is 10.1(2) -> 10.2(4)M -> 10.3(4a)M. Similarly, if the current release is 9.3(8) and the target release is 10.3(4a)M, the recommended path is 9.3(8) -> 9.3(13) -> 10.3(4a)M based on the official NOS upgrade paths provided by Cisco for the Nexus 9000 and 3000 series [https://www.cisco.com/c/dam/en/us/td/docs/dcn/tools/nexus-9k3k-issu-matrix/index.html#cur=9.3\(8\)&tar=10.3\(3\)F](https://www.cisco.com/c/dam/en/us/td/docs/dcn/tools/nexus-9k3k-issu-matrix/index.html#cur=9.3(8)&tar=10.3(3)F).

To upgrade to 10.3.4a, Apstra requires a minimum of 9.3(10) in the 9.x NOS version release train and 10.2(4) in the 10.2.x release train.

Workaround

Upgrade to a minimum version NXOS 9.3(10) in 9.x train or NXOS 10.2(4) in 10.2.x train before upgrading to 10.3(4)M

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

gRPC connection timeout consistently raised in the JUNOS and EVO device (AOS-50591)

The device telemetry health probe for the Junos and EVO devices running >= 23.4R3-S3 or >=22.3R2-S2 release shows abnormalities associated with the "gRPC connection reset" The Apstra telemetry collector is no longer compatible with gRPC due to a fix for XPATH in the specific NOS versions, which causes anomalies in the telemetry health probe.

Workaround

Polling can be used as an alternative to gRPC for telemetry service on devices running the specified NOS. To use polling globally, disable the gRPC service in the telemetry service, or make a REST API call (/api/systems//services/) to enable polling for a specific service on a

specific system.

To disable the gRPC service in the Telemetry Service, change `grpc_enabled = 0` in the `/etc/aos/aos.conf` file and restart the AOS service on the Apstra Controller.

```
[telemetry_global_config]
# Python multithreading enable/disable knob for telemetry collection
multithreading_config = 1
# Execution timeout for extensible telemetry collectors
command_timeout = 120
# Knob to enable/disable gRPC based service collectors
grpc_enabled = 0
```

gRPC server reset count anomalies in the JUNOS-EVO platform (AOS-53526)

gRPC server reset count anomalies are observed in the JUNOS-EVO platform when `gRPC Max Client connection limit error` occurs in the device due to the problem that gRPC stalled connections are not cleared. gRPC keepalive is not enabled by default on the JUNOS-EVO platform running 22.2R3 or 22.4R3, which is the cause of the problem. gRPC keepalive is enabled for 300 seconds in the `>=23.4R2-EVO` release to avoid a build-up of stalled gRPC connections.

Workaround

In JUNOS-EVO device running 22.2R3 or 22.4R3, apply the below configuration via `configlet` into the device to enable gRPC keepalive or upgrade the device to `>=23.4R2-EVO`. For further assistance, please contact the Juniper Apstra Support Team.

```
set system services extension-service request-response grpc grpc-keep-alive
300
```

gRPC timeout in Junos EVO device during loading high scale configuration (AOS-51023)

When the JUNOS-EVO device running `<=23.4R2-S3-EVO` is loaded with high-scale configuration, the gRPC subscription for Apstra's MAC telemetry service may fail with a timeout because the `l2ald` agent process is in hang status. While the device is in the status, the interface telemetry service may fail concurrently. The Device Telemetry Health probe would log anomalies resulting from continuous telemetry service failures.

Workaround

After the device is back to the stable normal status after internal service restart, the gRPC will be subscribed to the device correctly, and anomalies from the probe will be cleared. Polling can be used instead of GRPC for telemetry services to prevent issues from long delay by GRPC subscription.

Juniper EVO When Configured With DHCP Relay as Border Leaf Role, DHCP Packets May Be Discarded (AOS-43348)

When Juniper EVO device hosts DHCP servers in a border leaf role with DHCP relay configuration, DHCP may not work as intended due to an unresolved bug in Junos EVO which prevents DHCP packets from being processed correctly. Please refer to the following KB for dhcp relay limitations: https://supportportal.juniper.net/s/article/Juniper-Apstra-Support-for-Stateless-DHCP-Relay?language=en_US

Workaround

Using Apstra's configlet feature, create configlet to remove the rendered DHCP relay configurations and apply it to the Juniper EVO border leaf device.

Junos EVO Forwards DHCP for Virtual Network With DHCP Forwarding Disabled (AOS-43238)

Due to an outstanding bug in all available versions of Junos EVO, when two virtual networks are hosted on the same set of ESI leafs, one with DHCP enabled and one with DHCP disabled, the virtual network with DHCP disabled will also have DHCP requests forwarded.

Junos QFX 5K and 10K Device running 23.4R2-S1/S2/S3 fails in Interface telemetry service (AOS-50710)

Junos QFX5K and 10K devices running 23.4R2-S1/S2/S3 don't support the XPath `/interfaces/interface/subinterfaces/subinterface/state/admin-status` which makes Interface telemetry services fail

Workaround

Please change service type for interface telemetry service from gRPC periodic into polling via REST API call (`/api/systems/{system_id}/services/interface`).

JUNOS: EVPN MACs are limited to 200 IPs per MAC for bridge domain by default (AOS-57099)

A new default limit for the number of IP addresses per MAC per bridge domain in EVPN (mac-ip-limit) was added in Junos and Junos Evolved Releases 24.2R1, 23.4R2, and 23.2R2. 200 IPs per MAC is the default setting.

Clients who use EVPN fabrics, such as those with MAC-VRF deployments in fabrics managed by Apstra, might observe that MACs linked to more than 200 IP addresses cease to learn new IPs in the EVPN MAC-IP table. The Junos software release introduced this expected behavior.

The fabric can support more IPs per MAC while preserving per-bridge-domain enforcement by setting mac-ip-limit globally using an Apstra Configlet. No software fix is required.

Workaround

To adjust the limit in an Apstra-managed fabric, create a Configlet in Apstra with the below command, specifying the desired limit:

```
set protocols evpn mac-ip-limit <desired-value>
```

Import the Configlet into the blueprint and apply it to the relevant switches.

Note: Although the command is global in Junos, the limit is enforced per MAC per bridge domain, including inside MAC-VRFs.

L3 sub-interface doesn't work on both aggregate Ethernet and normal physical interface in the Juniper PTX Platform (AOS-48564)

L3 sub-interfaces don't work on aggregate Ethernet and normal physical interfaces with a 23.4R2 image on the Juniper PTX platform. This issue was not seen in the previous qualified NOS versions

Workaround

If a sub-interface needs to be configured, please use a qualified NOS version for 5.0.0, except 23.4R2 for Juniper PTX Platform.

L3 sub-interface triggers Deployment failure in the Juniper QFX52XX platforms running Junos-EVO (AOS-47861)

On Juniper QFX52xx platforms running the Junos-EVO image, the L3 sub-interface configurations applied via connectivity template (CT) are not rendered correctly over the aggregate Ethernet and regular physical interfaces. This results in deployment failures in affected blueprints.

This issue is tied to hardware constraints on Broadcom TH3, TH4, and TH5 ASIC platforms. The flexible-vlan-tagging feature required to support L3 sub-interfaces is not officially supported on the QFX52xx series (QFX5220, QFX5230, and QFX5240). While this command may have appeared in earlier Junos-EVO versions such as 22.4R2, its presence was unintentional and not validated for use. In Junos-EVO 23.4R2, the option is intentionally hidden, and any attempt to configure it will result in a commit error, reflecting Junipers enforcement of this limitation.

Apstra enforces this hardware limitation through validation preventing user getting into deployment error, explicitly blocking sub-interface configurations for affected platforms. Apstra returns the following error when validation fails:

```
"System {system} OS Family {os_family} with ASIC {asic} does not support subinterfaces due to hardware constraints. This error can be relaxed within validation policy in the event of a vendor-supplied OS update adding support."
```

will introduce changes in a future release to enable sub-interface configurations on Juniper TH3, TH4, and TH5 ASIC-based platforms (QFX5220, QFX5230, QFX5240). This support will not rely on the flexible-vlan-tagging feature, which is not officially supported on these platforms in any Junos-EVO release. Instead, Apstra uses vlan-tagging as per Junos Evo platform recommendation.

Workaround

Avoid using flexible-vlan-tagging on QFX52xx platforms in any release, as this feature is not officially supported. This issue is present on QFX5220, QFX5230, and QFX5240 models.

MAC Monitor Probe by GRPC collector-type reports the internal MAC addresses (AOS-47101)

In some Junos and Junos EVO releases, GRPC may report MAC addresses that are not visible in the "show ethernet-switching table" output. This issue has been fixed in newer versions. Those MAC addresses can be reported to the MAC Monitor Probe using the MAC Telemetry's GRPC collection. These MAC addresses are internal allocations used by Junos and Junos EVO for programming purposes only, and they have no effect on device functionality.

Workaround

Manual Reboot Required for "shared-tunnels" Configuration Following Junos Upgrade (AOS-45139)

In the Apstra 4.2 reference design change for MAC-VRF, the Junos "forwarding-options evpn-vxlan shared-tunnels" configuration is added via the Apstra rendered configuration. However, this command requires a device reboot to take effect with the Junos warning "Config: forwarding-options evpn-vxlan shared-tunnels has changed. A system reboot is mandatory". A user doing a Junos upgrade with Apstra may re-experience this issue after the device is upgraded.

Workaround

To avoid the need to a additional, manual reboot after a device Junos upgrade, the user can add the following configuration to the Apstra device system-agent pristine-configuration.

```
forwarding-options {
  evpn-vxlan {
    shared-tunnels;
  }
}
```

This can be done in the "Devices / Managed Devices / Pristine Configuration" Apstra UI or using the Apstra-CLI "system pristine_config_append" command.

MGD NETCONF issue causing loss of management access in the JUNOS QFX device >= 23.4R2 and <= 23.4R2-S3 (AOS-52766)

A netconf/mgd bug in QFX device running JUNOS >= 23.4R2 and <= 23.4R2-S3 could result in the loss of management access, including SSH, netconf, and gRPC, which Apstra uses to manage the device.

Workaround

If the device is running the exposed version, recommend upgrading to JUNOS >= 23.4R2-S4

Missing Local MAC address in the MAC Probe from Sonic MLAG pair (AOS-48979)

A member leaf in a Sonic MLAG leaf pair does not show the local MAC in its mac table, but

because its peer learns it as DYNAMIC, the probe generates expectations that this MAC will be present across the fabric.

Moving dual-connected (MLAG) to single connected Port-channel link makes member interfaces down with error state (AOS-48013)

When MLAG port-channel across dual nodes is unbundled into 2 x single connected port-channels, Apstra removes the "VPC" configuration on the interface port-channel. This change causes Cisco NXOS to trigger interfaces that remain down with the "IntFailErrDis" state as wrong behavior.

Workaround

Perform the below steps to recover

1. Remove the port-channel configuration
 2. Bounce the error-disabled interface.
 3. configure port-channel configuration again
-

Packet Drop for Juniper PTX Device running JUNOS EVO 23.4R2 into VTEP X.X.X.0/32 (AOS-49147)

PTX on EVO version 23.4R2 will drop the packets towards the vtep addresses x.x.x.0/32 (i.e 10.0.0.0), which has "0" in the last octet.

Workaround

Change the loopback address of all the vtep ip in the Blueprint which has "0" in the last octet (i.e 10.0.0.0) to non-zero value in the last octet (i.e 10.0.0.11)

Shutdown interface feature fails in rollback on Cisco NXOS device during NOS upgrade (AOS-48858)

Apstra 5.0 introduces a new system agent feature that shuts down device interfaces during the NOS (Network Operating System) upgrade process. This enhancement adds pristine configuration settings to disable the interfaces, relying on Cisco's rollback capability to restore the original configuration after the upgrade.

However, a known Cisco bug affects this process: if an interface is up (no shutdown) in the running configuration and has switchport enabled and interface shutdown from the feature in the pristine configuration, the rollback operation fails. As a result, this AOS extension exposes the existing Cisco rollback issue, which impacts the device's OS upgrade capability.

Workaround

This bug affects Cisco versions prior to 9.3.13. It was seen in both 9.3.11 and 9.3.8. The feature "Shutdown interface during NOS Upgrade" is compatible with 9.3.13, 10.1.2, and 10.2.5. For Cisco NXOS devices running the affected versions, the customer should enable "Skip Shutting Down Interface During Upgrade" in the system agent manager settings on the Managed Devices page.

SONiC cannot have BGP peerings using a vxlan-enabled Virtual Network endpoint (AOS-45640)

Because vxlan-enabled Virtual Networks in SONiC are not allowed to have a unique IP address and an anycast IP address simultaneously, it is not possible to establish BGP peerings where the local endpoint doing the connection is any such Virtual Network. This includes Virtual Networks used in ESI Multihomed LAGs, that must naturally always be vxlan-enabled.

Note, if a vxlan-enabled VLAN is not using an anycast address, it is perfectly possible to use said VLAN in a multihomed Layer-2 connection from both leafs in a redundancy group to an external device (such as a router), to establish multiple BGP peerings inside the same Ethernet Segment. In such a scenario and when using SONiC, the redundancy group can be MCLAG or ESI.

Static route for loopback address of external router not installed in the SONiC device (AOS-45557)

If a SONiC device removes and then re-adds an IP address, the device may fail to add a static route involving that address to the kernel routing table, even if the static route configuration exists. The `show ip route` output in `vttysh` in the SONiC device experiencing this issue may include the following output lines:

```
S>r 198.51.100.2/32 [1/0] via 192.168.0.9, Po1.4, weight 1, 01:59:25
B * 198.51.100.2/32 [20/0] via 10.0.0.3, Vlan201 onlink, weight 1, 01:59:25
```

Above, the "S" static route has been rejected by the kernel and was not installed.

Workaround

A full config apply will restore normal operation, in case such a problem occurs.

Traffic drop on tagged when L2 untagged and tagged VLAN are defined over the same port for Junos-Evo device (AOS-48905)

When both untagged and tagged VLANs are defined over the same port for L2 generic systems, traffic drop over tagged VLANs is observed on QFX 5230/5240 Junos-Evo platforms.