

Juniper Apstra Version 4.1.0 Release Notes

New Features

Support for EX4300-48MP with 40/100G uplink modules (RFE-2060)

Feature Category: Device Profiles

Added support for 40/100G uplink module for the Juniper EX4300-48MP device.

Support for Dell N3248T switch (RFE-2393)

Feature Category: Device Profiles

Support for Dell N3248T has been added in Apstra 4.1.0.

Support for Arista 7280QRA-C36S switch (RFE-2427)

Feature Category: Device Profiles

Support for Arista 7280QRA-C36S has been added.

Support for Arista 7280CR3-32P4 switch (RFE-2442)

Feature Category: Device Profiles

Support for Arista 7280CR3-32P4 has been added.

Support Arista 7280CR3K-32D4 platform (RFE-2288)

Feature Category: Device Profiles

Added support for Arista switch 7280CR3K-32D4.

Support Arista 7280CR3-96 platform (RFE-2278)

Feature Category: Device Profiles

Added support for Arista switch 7280CR3-96.

New workflow for modular Device Profiles (RFE-1856)

Feature Category: Device Profiles

You can now autonomously create Device Profiles for modular switches, by selecting a chassis model and one or multiple line card models.

To support this feature and the associated workflow, 4.1.0 introduces two new concepts: Chassis Profile and Line Card Profile, both of which are visible in two new dedicated sub-tabs under Device Profile menu.

With that, you can create a new Device Profile of type "Modular" as a composition of a selected Chassis Profile and different Line Card Profiles to be inserted in specific slot numbers. The list of built-in Chassis and Line Card Profiles shipped in 4.1.0 are:

- Chassis: Juniper QFX10008, Juniper QFX10016, Arista DCS7504, Arista DCS7508
- Line cards: Juniper QFX10000-30C, Juniper QFX10000-30C-M, Juniper QFX10000-36Q, Juniper QFX10000-60S-6Q, Arista DCS-7500R-36Q, Arista DCS-7500-36CQ.

Several examples Modular Device Profile configurations are also shipped in 4.1.0. You can also filter the list of Device Profiles by their type "Monolithic" or "Modular".

Future Apstra releases will introduce new supported Chassis and/or Line Card Profiles, similar to the addition of new monolithic Device Profiles.

New Interface Maps for the Juniper QFX10002, QFX5220 and QFX5130 devices (RFE-2483)

Feature Category: Device Profiles

The new interface maps are:

- QFX10002 IM's for 36x40 LD
- QFX10002 IM's for 72x40 LD
- QFX5220-32CD for 32x400G LD
- QFX5130-32CD for 32x400G LD

New default Logical Devices, Interface Maps and Device Profiles for generic systems (RFE-2034)

Feature Category: Device Profiles

Added the default LD, IM, DP for 2x25, 2x100, 36x40 generic systems.

Support for new Junos and Junos Evolved versions (RFE-2378)

Feature Category: Device Operating Systems

New supported versions:

- 21.2R3, 21.2R3-EVO
 - 20.4R3-S2, 20.4R3-S2-EVO
-

Jinja syntax highlighting and validation in configlets (RFE-2091)

Feature Category: Device Operating Systems

You can benefit from automatic syntax highlighting when editing or viewing a configlet. This helps to author and read complex Configlets, particularly when multiple Property Set variables are involved or when Jinja control structures, such as loops and conditionals, are used.

The UI also performs syntax validation and raises a validation error, preventing you from saving the configlet, if the Jinja syntax is incorrect.

Enterprise SONiC 3.4.1 support (RFE-2340)

Feature Category: Device Operating Systems

Enterprise SONiC 3.4.1 is now supported for all roles.

VMware VLAN config mismatch detection with trunk logical switches (RFE-2118)

Feature Category: Design, Build, Operate

When enabling vSphere or NSX-T virtual infra and deploying the probe for "Hypervisor & Fabric VLAN Config Mismatch" probe, the probe now also detects VLAN mismatch for hypervisor trunk switches, where the VLAN is configured inside the VM and not on the host virtual switch/port-group level.

UI-based validation of VN and VRF name length (RFE-2130)

Feature Category: Design, Build, Operate

As a user sets a name for a new VRF or a new VN, the UI immediately validates the name length and provides the user with an error message if it exceeds the max number of characters allowed. Previously this was validated only after the VRF/VN gets created which was not user-friendly.

Support vCenter integration with vCenter 7.0 (RFE-1779)

Feature Category: Design, Build, Operate

The Apstra integration with vCenter has been qualified with vCenter 7.0.

Support of System tags in the application scope of a configlet (RFE-2090)

Feature Category: Design, Build, Operate

You can apply configlets based on node tags (such as for all GS tagged as "firewalls") in addition to be applied on link tags. This allows for more user-friendly definition of the configlet application scope, where instead of listing all interfaces, you can list a tag.

Support of incremental configuration view on Blueprint with Juniper devices (RFE-1632)

Feature Category: Design, Build, Operate

For Junos-based blueprints, when uncommitted changes are staged, the user can view the incremental configuration when selecting a system. The Apstra built-in Reference Design

configuration is displayed in diff mode with added and removed lines. The configlet-driven configurations will list the previous list of configlets with their respective configurations followed by the new list of configlets with their respective configurations.

Support of Import/Export button in the configlet page (RFE-1460)

Feature Category: Design, Build, Operate

You can import and export configlet definition files for ease of sharing between different Apstra instances.

Support for Junos Dynamic BGP peering in Default Routing Instance (RFE-2177)

Feature Category: Design, Build, Operate

This new feature allows you to dynamically establish BGP peering to a group of remote neighbors within a specified range of IPv4 or IPv6 addresses rather than by specifying individual remote or local neighbor IP addresses.

Support for IPv6 Virtual Networks in 5-Stage EVPN Blueprints. (RFE-1548)

Feature Category: Design, Build, Operate

Previous to Apstra 4.1.0, 5-Stage EVPN Blueprints only support IPv4 virtual networks in the EVPN overlay. You can now deploy Virtual Networks for IPv6 applications in 5-Stage EVPN Blueprints.

Support creating port-channels from S/SS to external systems (RFE-2107)

Feature Category: Design, Build, Operate

Prior to 4.1.0, when attaching an external FW/router to a spine/super-spine, users were forced to build individual IP links instead of a single bond interface for all the links (there could easily be 8/16 links sometimes based on the required N/S traffic bandwidth). This 4.1.0 feature allows users to design and build a LAG for all these links, using a single Connectivity Template with a single application point to manage.

Search for Virtual Networks based on VLAN ID (RFE-2264)

Feature Category: Design, Build, Operate

Users can now search for Virtual Networks based on VLAN ID(s).

Routing Zone Constraints (RFE-1546)

Feature Category: Design, Build, Operate

Routing Zone Constraints allow you to constrain what server-facing interfaces connect to what Routing Zones. Day-2 operators would be prevented from connecting a server to the wrong network, and assure that a given server never gets added to the wrong network. The constraint can be defined in various ways such as a list of allowed VRFs, a list of excluded VRFs, a maximum number of VRFs allowed, etc. Once the constraint is defined, you can enforce the constraint on server-facing interfaces using connectivity templates of the type "Routing Zone Constraint". Finally, you can use Apstra granular permission model to allow a group of users to connect servers to networks ("Manage virtual network endpoints" is enabled) but forbid them to change the interface constraints ("Make any change to staging blueprints" is disabled).

Reserve VLAN ID per Virtual Network (RFE-2161)

Feature Category: Design, Build, Operate

When creating a Virtual Network with a specific VLAN ID, the user now has the option to also reserve the VLAN ID so that it can only be used for that Virtual Network and no others. The option is exposed when a VLAN ID is specified and the "Reserve across blueprint" box is checked.

Read-only mode for Apstra Server (RFE-2117)

Feature Category: Design, Build, Operate

A third operation mode for Apstra Server is introduced called "read-only" that can be enabled/disabled by Apstra admin users to temporarily block all users from performing design and blueprint changes, for example in the context of doing a platform backup or upgrade.

RBAC: new permission to restrict users with blueprint creation rights to delete a blueprint

(RFE-1967)

Feature Category: Design, Build, Operate

A new permission is introduced in the user permissions model to explicitly block users from deleting blueprints even though they have the permission to create them.

Rack capacity exposed for 3-stage designs (RFE-1456)

Feature Category: Design, Build, Operate

The rack capacity is now also exposed for 3-staged designs under Staged/Active>Physical>Racks, which exposes the number of deployed and remaining Generic Systems available per Rack.

Option to preserve Connectivity Templates when forming and breaking LAGs (RFE-1719)

Feature Category: Design, Build, Operate

It is a very common operation to have to break a LAG towards a server into individual links, for example to re-bootstrap the server, then to have to reform the LAG again from the individual links, all keeping the same VLAN allocation. This is now possible via the new "Form LAG" and "Break LAG" tasks that are exposed in the topology view when selecting generic system-facing interfaces. The option to preserve connectivity templates on any of the links is presented to the user.

Option to exchange only EVPN Route Type RT-5 prefixes (interface-less model) for DCI (RFE-1875)

Feature Category: Design, Build, Operate

When extending layer-2 networks between data center fabrics you now have the option to exchange only EVPN Route Type RT-5 prefixes (interface-less model). This is useful when there is no need to exchange all host routes between data center locations. This results in smaller requirements for the routing information base (RIB), also known as the routing table, and the forwarding information base (FIB), also known as the forwarding table, on DCI equipment.

New monitoring dashboard to aggregate the health of multiple blueprints (RFE-1915)

Feature Category: Design, Build, Operate

A new monitoring dashboard is introduced on top of the blueprints list to see in one glance the overall health of all the blueprints, such as the anomaly status, deployment status, and uncommitted changes. The new tiles are interactive and can also be used to filter the blueprints based on health criteria (for example, show only the blueprints with anomalies).

Mouse-over elements in the topology view to show additional information (RFE-2230)

Feature Category: Design, Build, Operate

Additional node and interface information is shown when mousing over these elements in the topology view (staged and active). For example, the configured port-channel ID range is exposed on generic systems and the applied Connectivity Templates are exposed on interfaces.

Managed Devices and System Agents menus merged (RFE-2175)

Feature Category: Design, Build, Operate

The Managed Devices and System Agents menus are consolidated into a single view for ease of operations on managed devices and removing confusion on which menu does what.

Expose disk utilization and disk space warning when uploading new NOS images (RFE-2064)

Feature Category: Design, Build, Operate

In the menu where new NOS images are uploaded, we now expose the current usage of the Apstra Server partition where all the NOS images are stored, which allows users to see how much space is left for uploading new NOS images.

Additionally, when trying to register a new OS Image, users will be shown a warning if the partition has under 5GB of free space.

Easy access to "Device Model" from the UI (RFE-2092)

Feature Category: Design, Build, Operate

You can easily access the "Device Model" of a system in the blueprint to find out what variable you can leverage in your Configlets. The Device Model is a nested dictionary of variables, accessible through the "Device Context" hyperlink when selecting a system.

The query tab provides dynamic search capabilities to quickly search through keys or values and identify the variables of interest.

Note that the syntax is case-sensitive. For example, a search of the key word "bgp" will provide information on the BGP configuration of the switch as well as the BGP sessions (protocol_sessions), while a search on the key word "BGP" will provide the list of BGP route-maps such as "BGP-AOS-Policy". It is important to note that the use of this variables as built-in property-sets inside a configlet must also respect the case-sensitive attribute of the Device model.

Ease day-2 operations at the nodes and links level (RFE-2229)

Feature Category: Design, Build, Operate

Specific Day-2 tasks are exposed at the node and interface level in the topology view to make Day-2 operations such as adding new Generic Systems and adding/modifying links as easy as possible.

Day-2 operation: Update of the spine/ss layer (RFE-1924)

Feature Category: Design, Build, Operate

New day-2 operations were introduced to update the spine or super-spine layer:

1. Allow increasing the number/count of spine-superspine links for a specific pod.
2. Allow changing the speed of spine-superspine links for a specific pod.
3. Allow changing the number of SS per SS plane.
4. Update the LD/IM/DP of the spine or super-spine layer (for all the spines, or all the super-spines).

These are all POD-level operations. For per node LD/IM/DP modifications, AOS CLI should be used.

Custom Logical Device port numbering (RFE-1879)

Feature Category: Design, Build, Operate

The concept of "Display ID" field for every port is added to Device Profiles. This is a new field the Apstra UI uses to customize the display of the port numbering for a Logical Device (to address that some devices start port numbering at 0, others at 1).

Collect additional logs as part of Show Tech (RFE-1987)

Feature Category: Design, Build, Operate

Improve information collected as part of Show Tech for faster troubleshooting and better customer support.

Cluster Health summary to include IBA and offbox agents health (RFE-2198)

Feature Category: Design, Build, Operate

The health of IBA and offbox agents is now also continuously exposed to GUI users as part of the bottom left Apstra Cluster Health summary widget. If any of the agents are erroring, the corresponding bullet will turn red to capture the logged-in users' attention.

Bulk edit for System Agents (RFE-1554)

Feature Category: Design, Build, Operate

Bulk edit for onbox/offbox agents and agent profile settings is supported. For example, allow bulk update/reset of credentials in case the credentials used by the agent are incorrect or forgotten and fail device authentication.

Automatically save and persist user preferences in Web UI (RFE-2146)

Feature Category: Design, Build, Operate

When a user configures their favorite views (such as displayed table columns, show/hide links, or default topology view), the preferences are automatically saved in a user profile and persist after logout/login.

Apstra Show Tech and API version to include hotfix version (RFE-1438)

Feature Category: Design, Build, Operate

If a hotfix is applied to an Apstra Server, the hotfix version is tracked using a 4th digit (ex: 4.0.0.1) and exposed in the Apstra show tech, API and UI.

Apstra Server automatic self-integrity check (RFE-2254)

Feature Category: Design, Build, Operate

Checksum verification of all AOS container images during an upgrade and at Apstra VM and services startup is added.

Advanced search for Connectivity Templates (RFE-1955)

Feature Category: Design, Build, Operate

Connectivity Templates can now be searched based on the type and any CT parameters. The new advanced search is available under the Connectivity Templates tab.

New IBA probe to monitor EVPN Host flaps (RFE-2052)

Feature Category: Telemetry and Analytics

You can leverage this probe to detect human errors such as a switch attached to two leafs but not through an ESI-LAG.

The probe will report the list of MAC addresses per VLAN that are flapping, i.e appearing on multiple leafs.

The probe will raise an anomaly for any MAC being alternately learned from local and remote switches more often than it is allowed by constraints configured in the probe definition. Metric Logging is enabled on this probe with a default retention of 7 days.

The probe is supported on Juniper and Arista devices.

IBA probe to monitor Optical transceiver statistics (RFE-2234)

Feature Category: Telemetry and Analytics

4.1.0 introduces a new IBA probe to monitor Optical statistics. The provide collects the following metrics:

- Tx Power Level (dBm)
- Rx Power Level (dBm)
- Tx Bias (mA)

- Temperature (C)
- Voltage (V)

The first three metrics are lane specific, while the last two are global to the transceiver.

For example, a 40Gbps QSFP+ transceiver runs over 4 lanes of 10Gbps, even if the interface speed is 40G. In this case the probe will show Tx Power, Rx Power and Tx Bias for each one of the 4 lanes, while it shows Temperature and Voltage once for the entire transceiver.

The probe raises an anomaly If one metric exceeds the Warn or Alarm thresholds in High or Low end. Metric Logging is enabled on this probe with a default retention of 30 days.

The probe is supported on Juniper and Arista devices.

IBA probe to monitor BGP session flaps (RFE-2045)

Feature Category: Telemetry and Analytics

New IBA probe to monitor BGP flaps. The probe monitors all BGP sessions (fabric session and external sessions) on all address families (v4, v6 and overlay) and reports the following metrics:

- FSM state: possible values being Idle, Openconfirm, Opensent, Active, In_transition, Connect and Estabilshed
- Flap count: Number of flaps counted since the beginning.
- Flap count increment: Number of new flap counts since the last collection period.

The probe raises an anomaly for any session exceeding a user-defined Threshold flap value over a user-defined Anomaly Time Window. Metric Logging is enabled on this probe with a default retention of 30 days.

The probe is supported on Juniper and Arista devices.

Changed Features

Remove "Routing Policies" option from Template creation view (RFE-1425)

Feature Category: Design, Build, Operate

The option called "Routing Policy (import)" is removed from the template creation due to its

redundancy with Connectivity Templates. Since Apstra release 4.0, Routing policies (import/export) are configured per BGP peering session using Connectivity Templates.

Exposing the Connectivity Template information in the Protocol Sessions view (RFE-2173)

Feature Category: Design, Build, Operate

You now can see what Connectivity Template has been used to create a BGP Session by looking at the Connectivity Template column under "Protocol Sessions" view. You can click on the Hyperlink to get directly to the Template detail.

Support of Access Switches in "Leafs Hosting Critical Services: Utilization, Trending, Alerting" Probe (RFE-2257)

Feature Category: Telemetry and Analytics

The following modifications have been applied to the Probe definition to appropriately support Access switch based topologies in this IBA probe :

- The graph queries used in this probe have been updated to include all fabric links from Leaf switches as well as those from Access Switches, ie "Access_to_Leaf" links.
 - The stages are augmented with a new column labeled "Remote System Role" to indicate if the link terminates on a Generic System or on an Access Switch.
-

Support of Access Switches in "Critical Services: Utilization, Trending, Alerting" Probe (RFE-2256)

Feature Category: Telemetry and Analytics

The following modifications have been applied to this IBA Probe definition to appropriately support Access Switch based topologies:

- The first stage is split into two stages: one for "Leaf_to_Generic" links and another one for "Leaf_to_Access_to_Generic" links. The two stages are then grouped together to merge both results.

- The stages are augmented with a new column labeled "Remote System Role" to indicate if the link terminates on a Generic System or on an Access Switch.

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.

Support of ESI for Access Switches - Phase 1 (RFE-2270)

Feature Category: Design, Build, Operate

You can now support network topologies with pairs of Access Switches, which allow attachment of dual-homed Servers/Generic Systems to the Access layer using ESI-LAG.

This is supported on 3-Stage, 5-Stage, and Collapsed Fabric Blueprints. Day2 topology changes are available through Add/Edit/Remove Racks.

Requirements for the switch model acting as Access Switch are:

- EVPN-VxLAN with VTEP support is required on the Access Switches.
- L2 VxLAN only is required, L3 VxLAN (RIOT) is not required, and will continue to be available only at the leaf layer.

Fixed Apstra General Issues

5-Stage Spine-Superspine Link IP Networks Not Exported (AOS-18570)

In a 5-Stage AOS Blueprint, the Spine-Superspine link IP networks will not be exported to External Routers by default and there is not an explicit option to export these.

Adding and Deleting System Agents May Cause SystemAgentManager Crash (AOS-21671)

If a user creates a range of more than (20) System Agents and then immediately deletes them, the AOS Controlled SystemAgentManager process may repeatedly crash due to database corruption. The user will not be able to perform any System Agent functions until the issue with the database is resolved.

AOS does not restart its processes following memory issues (AOS-19797)

AOS has internal mechanisms to ensure its various processes are running correctly. These mechanisms may fail if the server has experienced Memory issues (insufficient RAM). Users may experience erratic AOS behaviour as AOS may no longer ensuring all required processes are up and running following such an event.

AOS is unable to render additional import/export Route Targets for VNs or SZs in Junos

blueprint (AOS-18392)

Where the remote side has already been configured, or the users expect AOS to use or import a different route-target. For example DCI where the distant end exists, and AOS is expected to match the established RT or import another RT. We can't change the RT as the import/export rendering from AOS is missing. This is because AOS can't add import and export RTs to the relevant policies to import these RTs in addition to autogenerated RTs used in the fabric.

AOS upgrade may fail after restoring from backup (AOS-16734)

An AOS "in-place" upgrade may fail on an AOS server that was restored from a backup after a previous upgrade attempt. This is due to the presence of the /var/lib/aos/db/blueprint_backups directory.

Apstra Cisco Device "Show Tech" Incomplete (AOS-28992)

If a user uses a username besides "admin" to install an Apstra device system agent on a Cisco NX-OS device, when running a "Show Tech" from the Apstra UI or from the device CLI using `aos_show_tech`, device "show" commands will be incomplete with "sudo: unknown user: admin" errors.

Apstra Device OS Upgrades for Arista EOS May Fail (AOS-28603)

Users using Apstra to upgrade Arista NOS between earlier EOS versions such as 4.21.7M and newer versions such as 4.25.3.1M may experience upgrade failures with "RemoteSshCommandException: Remote ssh command failed" errors.

Apstra Device Upgrade Fails on Juniper QFX10002-60C (AOS-28661)

Due to "vmhost" differences in the QFX10002-60C platform, Apstra device (NOS) upgrade for the Juniper QFX10002-60 is not supported.

Apstra ZTP Fails to Upgrade Junos on QFX10002-60C (AOS-29421)

Due to missing support for the "request vmhost software add" command, Apstra ZTP currently cannot be used to upgrade Junos on the Juniper QFX10002-60C platform.

Apstra ZTP Not Matching Cisco Nexus 3000 Model (AOS-28680)

Due to a Cisco Nexus 3000 platform change, Apstra ZTP 4.0 versions will not properly match the "N3K" model number (e.g. N3K-C36180YC-R) in the ztp.json file.

Arista Interface Remains Down After Moved to Layer3 Peer Port-channel (AOS-25357)

An Arista EOS interface will remain down after being removed from a port-channel and added to a Layer3 peer port-channel in the same Apstra commit operation.

Arista Spine Does Not Advertise the Default Route to Superspine (AOS-29991)

In Arista fabrics managed by Apstra, the spine does not advertise the default route learned from a leaf to the superspine when the route is received from an external router from the leaf and the superspine at the same time.

Attaching vxlan VN to an nxos leaf in non-evpn blueprint leads to incorrect config (AOS-27785)

In static VXLAN (non-EVPN) blueprints, for a NXOS leaf that has a single L3-enabled VXLAN bound to it, and that VXLAN is not bound to any other rack, then the configuration fails to apply. The statement fabric forwarding mode anycast-gateway under the SVI interface vlan cannot be applied because the global statement feature fabric forwarding is missing.

Cannot assign Connectivity templates after the external generic system property change (AOS-29272)

After a connectivity template is assigned to a generic system, changing any of the following property values (hostname, name or loopback) of that external Generic system will not throw any errors immediately but will result in "node already exists" error next time while assigning any connectivity template to any endpoint.

Dash "-" Cannot Be Used for a Property Name in a Property-Set (AOS-28123)

If an Apstra user uses a dash "-" for a property name there will be a build error due to missing properties.

Default Route Verification Anomalies for Dynamic BGP Peers to External Generic System (AOS-25937)

If a user configures dynamic BGP peers to Generic Systems using Connectivity Templates, Apstra will, by default, check for a default BGP route and raise anomalies.

Device Profile for Cisco 9348GC_FXP has incorrect column_id's (AOS-21464)

Device Profile for Cisco 9348GC_FXP has incorrect column_ids for 1/10/25Gbps transformation ports (49,50,51,52). User cannot create an Interface Map with 25Gbps ports.

Disabling IPv4 Virtual Gateway in Junos ESI Rack (AOS-25311)

Apstra will allow the user to disable the IPv4 Virtual Gateway for a Virtual Network assigned to a Junos ESI rack. This should not be allowed and doing so will break IP connectivity.

EOS 4.22 Upgrades May Cause EVPN To Fail (AOS-20482)

If a user is using AOS to upgrade Arista EOS devices to a new version of EOS 4.22, the absence of the `service routing protocols model multi-agent` EOS configuration from the AOS Device Pristine Configuration will cause EVPN on the device to fail after the upgrade is complete.

IBA LAG Imbalance Probe May Show Intermittent Junos 0bps Traffic Rate (AOS-27393)

Due to Junos hardware interface collection issues, interface traffic rate may show 0bps intermittently in LAG Imbalance Probe.

In Sonic, Changing LAG mode from LACP Active to Static or vice-versa does not take effect on device (AOS-29469)

In Sonic, Changing LAG mode from LACP Active to Static or vice-versa does not take effect on device without any deployment errors on the device.

Junos "set" or "delete" based Junos Configlet (AOS-28172)

While creating set or delete based interface configlets for Junos, user should use the correct radio button, and ensure not to use the "set" or "delete" keyword in the configlet.

Junos EVO Images Cannot Be Uploaded to Apstra (AOS-29671)

Apstra does support the NOS upgrade of Juniper Junos EVO devices, however, the user cannot upload Junos EVO ISO files to "OS Images" in Apstra. The user will get an "Invalid file extension .iso for JUNOS image" error.

Junos Unnumbered (link-local) peering only available in default VRF, and must be point-to-point interface (AOS-24976)

Junos 21.1R1 only supports peer-auto-discovery in the default VRF, it is not available in a routing-instance (vrf). Additionally, Juniper only supports point-to-point interfaces for unnumbered peering. If any extra IPv6 router advertisements are generated on the link, Juniper may pick a random ipv6 neighbor on the link to establish BGP with. Ensure that no shared broadcast domains are used for point-to-point unnumbered BGP peers. Other NOS vendors only peer with ipv6 neighbors which generate the RA-Bit, but Junos does not look at the RA bit for peer discovery: The first discovered peer is considered the interface BGP peer, whether it speaks bgp/sends ra or not.

NOS upgrade for SONiC may fail with inability to collect pristine (AOS-28677)

NOS upgrade for SONiC devices may fail after the actual NOS upgrade has been carried out with a pristine config collection problem printed in the log.

Resolution

When there is a NOS upgrade, the following sequence is being carried out:

1. AOS agent is stopped.
2. Pristine config file is copied to /etc/sonic/config_db.json but no config reload -y is issued because this is a NOS upgrade and there will be a reboot anyway.
3. AOS agent is started again.
4. New sonic image is installed.
5. Device is rebooted
6. Pristine is collected from the new version

7. AOS agent is installed on the new version

Between 3 and 4 and depending on the timing, the AOS agent may do a config save and restore the service config, which will be picked up by the sonic_installer in step 4. This will cause step 6 to fail, as the config used on the new version is not the pristine config, but the AOS service config.

NXOS Upgrades Failing With Rollback to Pristine Error (AOS-29499)

If an Apstra user who has upgraded from an earlier version with managed Cisco NX-OS devices uses a number sign character "#" as a delimiter for a "banner" configuration, they may get "Rollback to pristine" error when upgrading Cisco NX-OS.

Offbox Device Agent Reports a Liveness Anomaly After Changing AOS Worker Node Tag (AOS-31678)

When performing Apstra cluster operations with tags when many offbox Device Agents move between Apstra worker nodes, DeviceKeeperAgent may stop functioning and report device liveness anomalies.

Resolution

This issue was due to an issue with an internal database component which has been resolved in Apstra 4.1.0.

Optical Transceiver Probe does not show interfaces (AOS-29945)

On Arista EOS devices running Onbox Agents, when using a 40GBASE-AR4 transceiver with 10G breakouts, the interfaces will not be present in the Optical Transceiver Probe. This issue is caused by a limitation of the transceiver, which does not show txPower levels correctly. Apstra does not correctly handle this which results in the interfaces not being listed at all.

Performance degradation when using Active / Query / Mac page and filtering by VLANs (L2 VNIs) (AOS-28016)

Performance benchmarks show that making filtering only by VXLANs (L2 VNIs) on the Active / Query / Mac page shows that it's slower in `4.0.1` than in `3.3.0.2` in around 5.5 times (122 ms vs 22 ms) on a topology with 30 racks with MLAG-pairs and 2500 VXLANs each hosted on 5 racks with 289 MACs per leaf and 58 MACs per VXLAN.

Port speed change on a breakout port results in deployment error on devices running Sonic (AOS-29505)

For devices running Sonic, port speed change on a breakout port results in a deployment error("error occurred processing config") after configuration is committed in Apstra.

Prefixes of interconnection links of a Router with Blueprint Devices, may get advertised back to the Router (AOS-24717)

When a an external Router is connected to a Blueprint topology using separate connections to multiple devices (example: external Router connected to leaf1 and leaf2 switches of a topology), the prefixes of those interconnections may travel through the EVPN fabric and get advertised back to the Router via other links.

Example: Given a BGP peering between the Router and Leaf1 and if the interconnection link prefix of the Router to Leaf1 is 10.0.0.0/24, with 10.0.0.2 being the address of leaf1 in that subnet, then the subnet route 10.0.0.0/24 and the more specific route 10.0.0.2/32 may get advertised all the way through the EVPN fabric to Leaf2, which can advertise them back to the Router, via its peering with Leaf2. This can have the adverse effect of the Router learning a more specific route to Leaf1 address 10.0.0.2 of Leaf1 on the wrong link and via Leaf2.

Revert to Pristine Config may get stuck in in_progress state for SONiC devices (AOS-29258)

When a SONiC device is reverted to its Pristine Configuration, the agent is erroneously not restarted after the restoration finishes. As a result, the device gets stuck in a NOCOMMS state and the Revert to Pristine task is not able to successfully conclude. The act of Reverting to Pristine the SONiC works correctly, however, the agent isn't started after the Reverting concludes.

SONiC Apstra Device OS Upgrade Fails (AOS-28250)

Using Apstra 4.0.x to upgrade SONiC Device OS fails with "Conflicting pre-AOS configuration found in device" error.

Telemetry Tab in L2 server devices is misrepresented (AOS-10412)

Is the user is configuring AOS Agents on L2 Server Devices in "Full Control" mode, extra telemetry data will be seen which has no impact on the server.

The Device Profile for the 5120-32C model contains an error causing incorrect rendering of 'channel-speed' configuration (AOS-28699)

When using interface et-0/0/31 on the 5120-32C Device Profile, incorrect 'channel-speed' configuration is rendered ("speed 40g").

Resolution

Device Profile for QFX5120-32C fixed.

Unable to Change Leaf Interface Descriptions After Upgrade to Apstra 4.0.1 (AOS-28617)

During an upgrade from Apstra 3.2 or 3.3 to Apstra 4.0.1, the Apstra rack upgrade plugin will copy the current interface descriptions to static entries in the Apstra configuration graph database. Due to this, if a user wants to change an interface description by editing the generic system node label, the Apstra configuration builder will not update any interface description when a generic system node label is changed.

Unexpected error messages while trying to apply or to unapply dependant CTs (AOS-26742)

Customers might face unexpected error messages while trying to apply or to unapply multiple dependant Connectivity Template in a single batch.

Case 1: 'IP Link' and 'BGP Peering (Generic System)' CTs created and applied. BGP Peering is applied to the ip_link node created by the first CT. If user open 'Application Endpoints' modal window and try to unapply both CTs simultaneously the unexpected error message will be shown.

Case 2: 'IP Link' CT is created and applied. 'BGP Peering (Generic System)' CT is created. If user open 'Application Endpoints' modal window and try to unapply IP Link and apply BGP Peering to the node which is about to be deleted then unexpected error message will be shown.

The similar cases might be observed for other CTs with dependant input and output nodes.

Upgrade Failure When Apstra Server SSH Password Contains Pipe Character (AOS-28811)

Apstra VM-VM upgrade using aos_import_state does not work if the original Apstra SSH server password contains a Pipe "|" character.

User Is Not Able to Auto-Allocate Default IPv4 for Inter-Rack Virtual Network During Update (AOS-8899)

IPv4/6 subnet address field for Virtual Network could be specified in the following ways:

- explicitly (f.e. 11.1.0.0/24 or fc01:a05:fab:feed::/48),
- with netmask (f.e. /24 or /64),
- do not specify (i.e. value will be allocated from assigned IP pool).

There is a corner case when the subnet is set with netmask (f.e. /28), which belongs to the assigned IP pool (f.e. 10.0.0.0/24). In such a case in an attempt to remove the previously allocated value with custom netmask (/28) nothing will happen (i.e. new value with netmask /24 will not be allocated, because /28 already matches IP pool 10.0.0.0/24).

Worker node in a cluster can end up in Failed state after reboot (AOS-28547)

When rebooting a Worker node VM in an Apstra cluster, the node can end up in a Failed state under 'Platform > Apstra cluster'. This is caused by certain internal files not being removed on the worker node when starting the Apstra service. This issue is fixed in Apstra 4.1.0 and later.

Fixed Third-Party Issues

Arista EOS IPv6 Neighbor Advertisements (AOS-25130)

It has been observed that EOS 4.24.5M may ignore IPv6 neighbor advertisements from directly connected hosts, if the neighbor solicitation was not initiated by the EOS device itself. Consequently, the EOS device may not have the directly connected host in its IPv6 neighbor list and thus may not advertise the appropriate type 2 route to the rest of the devices in the fabric.

Changing vlan id on a subinterface in SONiC 3.3.0 may result in traffic loss (AOS-26117)

Changing the vlan id of a subinterface in SONiC 3.3.0 or 3.4.0 or 3.4.1 may result in a traffic loss, due to a SONiC NOS bug.

When creating a subinterface, AOS itself will always choose the assigned vlan id as the subinterface id of the subinterface, which essentially means that the subinterface of Ethernet0

using `vlan_id=15` will be always `Ethernet0.15`. Using a different `vlan id` means the creation of another subinterface. So, when using the UI or the high level API of AOS can never result in a subinterface changing `vlan id` and the bug can never be triggered.

However, if a customer uses the raw graph patch API to alter the name of the subinterface and break the rule that `subinterface_id=vlan_id`, then a change where the `vlan id` of a subinterface is altered might be brought about. Specifically, consider the example where `Vlan 4` is assigned to `Ethernet0.1` and must be changed to something else, whereas `Vlan4` itself must be assigned to `Ethernet0.2`. To accomplish that, one must delete `Vlan 4` from `Ethernet0.1`, and then assign it to `Ethernet0.2`. However, due to the specific bug, SONiC may fail to delete `Vlan 4` from `Ethernet0.1`, which will cascade and make the assignment of `Vlan 4` to `Ethernet0.2` also fail.

Cisco Nexus C36180YC Deployment Errors (AOS-28472)

Due to issues with this Cisco Nexus device, Apstra, by default, will not be able to install its device System Agent or successfully deploy configuration on the device.

Dynamic BGP Peering Connectivity Template for Junos fails with CommitError (AOS-25262)

When Dynamic BGP Peering Connectivity Template is applied for Junos device, deployment fails with the following error : "Apply config failed: CommitError(edit_path: [edit routing-instances blue protocols], bad_element: bgp, message: error: Error in group XXXX: peer AS number must be configured for an external peer error: configuration check-out failed)"

Resolution

This issue is resolved in Junos release 21.2R1+

Junos 21.2R2-S2-EVO LLDP neighbors might not show during the initial device deployment (AOS-29951)

During the initial deployment of a device running Junos 21.2R2-S2-EVO, the LLDP neighbors might not show up.

Junos EVO Device Channelized Interface May Not Come Up (AOS-28538)

For Juniper devices running Junos EVO, there is a potential issue where when a channelized port is configured, an individual interface link may not come up.

QFX10002-60C LAG Hashing Not Working for Layer2 Forwarding (AOS-30085)

Juniper QFX10002-60C devices with ESI LAGs may, by default, not hash layer2 traffic with single source/destination IP/MAC endpoints but with different TCP source/destination ports.

SONiC Routes Incorrectly to Reloaded Leaf (AOS-20886)

Due to a SONiC bug in SONiC-OS-3.3.0-Enterprise_Advanced, if a SONiC spine is connected to an MLAG pair of leafs, and one leaf is reloaded, SONiC will route incorrectly to the leaf with the higher metric once the BGP session is reestablished.

Unexpected Rx Discard Count on All Links (AOS-20787)

An AOS user using Broadcom SONiC, may experience incorrect AOS anomalies due to unexpected Rx discard counts from SONiC on all links which receive control-plane packet from peers. This only affects telemetry. No data plane traffic is dropped.

Known Apstra General Issues

"management_ip" field not visible in rendered config preview (AOS-24760)

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

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

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

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

Workaround

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

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

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

Workaround

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

EX4400:

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

QFX5110:

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

QFX5120:

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

If the Apstra device system agent installation and deployment in the Apstra blueprint is already done, the user can use Apstra Configlets to add the configuration. Note, "set" Configlets are only supported in Apstra 4.0.2 and later. For Apstra 4.0.1 and earlier, the user will need to use a Configlet with a "hierarchical" Junos configuration.

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

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

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

All AOS deployments running a specific version have the same set of secret keys. This is

potentially a security flaw as a user having access to an AOS VM of a version can get access to secret keys installed in a different VM as they are all the same.

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.

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

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

Workaround

N/A

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

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

Workaround

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

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

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

Workaround

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

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

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

Workaround

For customers with disk space of partition less than 50% , this issue will not cause any problems. If the disk space usage is higher, customers can increase the disk space of partition /dev/mapper/aos--server--vg-root on Source VM before the upgrade.

Apstra Sysdb Crash (AOS-34904)

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

Apstra ZTP Failure During Junos Upgrade with Console Special Characters (AOS-43732)

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

Workaround

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

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

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

Workaround

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

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

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

Workaround

Contact Juniper Support for an updated Device Profile.

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

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

Workaround

Apstra Deployment Agent timeouts can be increased. Please contact Juniper Support for assistance.

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

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

Workaround

Apstra Deployment Agent timeouts can be increased. Please contact Juniper Support for assistance.

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

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

Workaround

None.

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

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

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

BGP Flapping probe is not supported on Cisco and SONiC devices. If it is instantiated on a blueprint with Cisco or SONiC devices the probe will list all BGP sessions but the "Flap count" and "Flap count increment" will permanently show the value of "zero" which is misleading .

Workaround

Not instantiating this probe on Blueprints which includes devices others than Juniper and Arista.

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

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

Workaround

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

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

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

Workaround

Please contact JTAC Apstra support.

Cannot Edit System in Managed Devices (AOS-30144)

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

Workaround

Editing the Managed Device from the popup in the devices list, or selecting its checkbox in the list and clicking the edit button above doesn't exhibit the problem.

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

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

Workaround

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

Changing System Admin Status Fails With Validation Errors (AOS-32806)

When changing the system admin status via the managed system system page the update with fail with `Server-side Validation Errors
" not found`.

Workaround

The user can edit the system from the managed devices pages uses the 3 dots for the system you want to change.

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

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

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

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

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

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

Workaround

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

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

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

Workaround

Contact Juniper Support

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

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

Workaround

If edits are required on predefined payloads, please first clone and then proceed to make required changes to this cloned payload.

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

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

Workaround

Please contact Juniper Support for an updated Apstra Device Profile for the Arista DCS-7280CR3-96.

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

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

FFE Operations in POD Based Blueprints Involve the Formation of Different rack_type With the Same ID (AOS-31854)

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

Workaround

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

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

Workaround

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

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

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

Workaround

Manually set the Generic Systems' Deploy Mode (Staged > Physical > Nodes > Set Deploy Mode) before committing any change. If no changes are needed, you can click 'Revert' on the Uncommitted tab to revert all changes. If Generic Systems did inadvertently get set to "Not assigned", this will not impact network operations as configurations on the switch side are not dependent on this setting.

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

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

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

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

Workaround

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

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

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

Workaround

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

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

Due to an issue with the Apstra EVPN reference design, problems with inter-vrf routing via an external router can occur if there is a single spine (either by design, failure, undeploy, or drain) and the route is originated on a non-border-leaf. The one spine will drop the route due to as-path loop.

Workaround

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

Interface TAG Not Propagated to Untagged Subinterface (AOS-28087)

If an Apstra 4.0.0 or 4.0.1 user is using tag data in configlets, interface tag information is not available for untagged subinterfaces.

Workaround

If the user needs to use tag data in configlets, they must use tagged subinterfaces.

IPv4 AFI Is Enabled With IPv4 and IPv6 Addressing Types (AOS-30636)

BGP Generic System Connectivity Templates with unnumbered BGP sessions can be used to achieve RFC5549 (advertising IPv4 prefixes with IPv6 next-hops) use-case by enabling IPv4 AFI (IPv6 AFI disabled) with IPv6 addressing type. In this scenario, Apstra will configure a new IPv6 BGP peer inside the IPv4 address family. If the user also enables IPv4 addressing type along with IPv6 addressing type, Apstra will not configure IPv6 BGP peer under IPv4 address family. In the current design, Apstra prefers IPv4 BGP peer over IPv6 BGP peer inside the IPv4 address family. This will result in only IPv4 peers (advertising IPv4 prefixes with IPv4 next-hop). Though network intent is achieved, the user could expect IPv6 peer since IPv6 addressing type is enabled.

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

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

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

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

Workaround

The customer can enable the IPv4 SAFI in the Connectivity Template to remedy that problem.

Juniper QFX 5110 can not forward VXLAN packets out Layer3 interfaces for external connectivity (AOS-19030)

Due to an ASIC limitation on the QFX 5110, documented at <https://www.juniper.net/documentation/us/en/software/junos/evpn-vxlan/topics/concept/evpn-vxlan-qfx5110-l2-vxlan-l3-logical.html> this platform should not use L3 connectivity points within EVPN security zones.

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

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

Workaround

Contact Juniper Support for an updated Device Profile.

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

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

Workaround

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

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

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

Junos EVPN_IMPORT policy-statement config rendering change (AOS-33852)

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

Workaround

Correct this behavior via configlet or upgrade to Apstra 4.1.2.

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

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

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

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

Workaround

The user can use an Apstra Configlet to apply the necessary configuration for "protocols bgp group l3clos-l-evpn family evpn signaling prefix-limit maximum (value)" or contact Juniper Support for a hotfix patch.

Missing Upgrade Plugin for node_to_node_if_counter Processor (AOS-40850)

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

Workaround

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

MLAG telemetry expectations are not correct when one MLAG leaf is drained (AOS-12820)

When draining one MLAG leaf in Maintenance Mode, MLAG telemetry expectations are not correct. No anomalies should be seen even though AOS will shutdown the server facing links and port-channels. However, anomalies are seen because of incorrect expectations.

Workaround

None needed. These anomalies do not indicate any operational problems.

No port-group support for the Dell S5248F-ON and S5296F-ON models (AOS-29040)

Certain port speed configurations (e.g. 1G mode) for the Dell S5248F-ON and S5296F-ON models require the usage of the port-group feature in SONiC, which is not supported by the Juniper Apstra device profiles.

Workaround

The user can configure an Apstra Configlet to use SONiC-CLI to add a port-group and change and interface speed from 10G to 1G. The user will need to deploy the port in Apstra as 10G ports using the default SONiC Device Profile. An example SONiC System Configlet to change Ethernet0:

Template Text: `sudo -u admin sonic-cli -c 'config' -c 'port-group 1 speed 10000' -c 'interface Ethernet0' -c 'speed 1000' < /dev/console`

Negation Template Text: `sudo -u admin sonic-cli -c 'config' -c 'interface Ethernet0' -c 'speed 10000' < /dev/console`

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

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

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

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

Workaround

Contact Juniper Support for a workaround.

PUT to chassis profile would update chassis_info of already created modular DP but not top level fields of hw/sw cap and selector (AOS-34143)

Scope of impact of this issue is limited to those modular device profiles, the chassis profiles of

which were edited after the create of the modular DP. In such a scenario, the changes to fields of software capabilities, hardware capabilities and selector would reflect in the chassis_info section of the modular DPs, but not the top level sections by the same name.

Workaround

Upgrade to 4.1.2 fixes this issue.

ref_count of protocol_session from dynamic BGP peering was not correctly set (carried with null) when upgrading from 4.0.0 to 4.x.x (AOS-33904)

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

Workaround

Manually correct using curl or UI after finding out protocol_session node from QE.

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

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

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

3. For the the collected IDs of protocol_session nodes from step2, Choose Step 4 using UI or Step 5 using curl in the terminal

4. UI: Platform > Developers, Click REST API Documentation button, select PATCH
`/api/blueprints/{blueprint_id}/nodes/{node_id}`

a. Update node.

b. fill out blueprint ID and node ID (from Step 2), body as {

`"ref_count":1`

}

Execute it. it should be done for all protocol_session nodes collected from step 3.

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

a. from Step 1

b. from Step 2.

6. Execute Graph Query to check out whether protocol session has ref_count value as 1.

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

The result should show reference count value as 1.

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

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

Rendered Config for Static LAG Changes (AOS-38375)

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

Rendered Configuration Error for Juniper_EX4400-48T_EM-4Y and Juniper_EX4400-48T (AOS-35037)

Rendered configuration error is seen for Juniper_EX4400-48T_EM-4Y and Juniper_EX4400-48T when the 100Mbps transformation is used.

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

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

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

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

Workaround

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

```
match(
node('ep_group', name='corrupted_group')
.in_('ep_affected_by')
.node('ep_application_instance')
.out_('ep_top_level')
.node('ep_endpoint_policy', name='ct', policy_type_name='batch')
.having(
node('ep_group', name='corrupted_group')
.in_('ep_affected_by')
.node('ep_application_instance')
```

```

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

```

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

3. Assign each CT to proper application endpoints.

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

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

Workaround

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

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

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

DeploymentProxyAgent.

Workaround

Upgrade to 4.1.1 if possible otherwise hotfix squashfs image must be used on the device to work around the issue.

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

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

Workaround

The user can delete the files in the `/var/log/aos` directory to immediately restore the device system agent. Please contact Juniper Support for procedures to modify the logrotate configuration for SONiC devices for a permanent workaround.

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

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

Workaround

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

Alternatively, a full config apply has been observed to put the DHCPv6 and DHCPv4 relay in the correct VRF as well. Do note however, that any subsequent incremental manipulation of the DHCP helper configuration will negate the correct VRF and reset it to default, necessitating another full config apply.

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

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

Due to the caching logic used by the collector for the Apstra VxLAN Flood List Validation IBA Probe, the probe may fail under certain conditions where the collection of VLAN to VNI mapping failed but the last update was not reset. This will cause the collector to incorrectly believe that the cache is empty and therefore it did not post any data.

User Cannot Change Prefix Type for Dynamic BGP Peer Connectivity Template (AOS-30419)

When a dynamic BGP peer template is created in a Connectivity Template (CT) with no subnet and assigned to SVI it auto derives the neighbor from the SVI address. In Apstra 4.1.0 when you try to change this CT to use a “user-defined” subnet, it errors with “It is prohibited to change prefix type from user-defined to neighbor based and vice versa”

Workaround

To work around this issue, the user can delete the Dynamic BGP peer from the CT and create another one in the CT with desired parameters and then update, or the user can unassign the CT, and change the Dynamic BGP CT params to whatever you like and then assign it back.

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

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

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

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

Workaround

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

Known Apstra Security Issues

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

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

Workaround

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

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

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

Known Third-Party Issues

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

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

Workaround

Do not use BGP passwords in SONiC BGP peerings.

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

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

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

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

Workaround

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

Change the slice value from 96 to 300 in both the swsswait.sh and wait_for_replay_done.sh in their respective containers:-

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

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

Follow same steps to update slice value in wait_for_replay_done.sh file

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

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

Apstra Using Cisco NX-OS Depreciated "soft-reconfiguration inbound always"

Configuration Causing Packet Loss (AOS-35513)

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

Workaround

Contact Cisco TAC for a fix release.

Cisco C9348GC NXOS Rollback Failures (AOS-34250)

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

Workaround

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

Cisco NX-OS EVPN Route Installs Incorrect/Random Next-hop (AOS-30463)

A Cisco NX-OS bug (CSCvz75734) may affect Apstra users using NX-OS versions 9.3(6), 9.3(7), 9.3(8) where evpn route imported into vrf with bogus next-hop on a VTEP causing traffic to black-hole.

Workaround

Contact Juniper Support for a workaround.

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

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

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

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

Workaround

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

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

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

Workaround

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

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

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

Workaround

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

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

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

Workaround

The user can resolve the telemetry issue by executing the following command from the Junos cli

```
restart l2-learning gracefully
```

Junos Link-local Session Does Not Come Up After the Leaf Reboots (AOS-29127)

Due to a Junos issue, if a user creates a BGP link-local session over SVI and assigns it to a leaf to server port-channel, the session comes up, but if the leaf device is rebooted, the BGP session will not come up.

Workaround

The user can do any configuration commit operation or restart rpd.

Junos Underlay Sessions Flap When Device Put in Drain Mode (AOS-29885)

Due to a Junos issue, if the deploy mode of a leaf device is set to drain, the underlay BGP session will flap.

Junos Upgrade of QFX10002-36Q to 21.2R3-S2 Leaves Device in Failed State (AOS-32598)

Using Apstra to upgrade Junos on QFX10002-36Q to 21.2R3-S2 leaves the device in a failed state.

Workaround

The user will need to manually remove the device from the Apstra blueprint and manually upgrade the device.