Juniper Apstra Version 4.2.2 Release Notes

Fixed Apstra General Issues

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.

---

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.

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.

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 vtysh 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.

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"

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           -                             down           admin-down
off       25000          9100           Ethernet0

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           -                             down           admin-down
off       25000          9100           Eth1/1
```

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.

Workaround

Please clone device profile for EX4400-24MP-EM and then replace configuration of transformation for 1G over 0-23 ports

```
from
{"global": {"breakout": false, "fpc": 0, "pic": 0, "port": 0, "speed": "1G"},
"interface": {"speed": ""}}
to
{"global": {"breakout": false, "fpc": 0, "pic": 0, "port": 0, "speed": ""},
"interface": {"speed": "1G"}}
```

Please contact Apstra Support Team for more information

Apstra 4.1.1 upgrade to 4.2.1.1 or 4.2.2 disables Generate EVPN host route from ARP/IPv6 ND ARP feature in the fabric policy (AOS-53556)

When using `Generate EVPN host route from ARP/IPv6 ND ARP` feature in the virtual network policy, caution should be used when upgrading from Apstra 4.1.1 or 4.1.0 to Apstra 4.2.1.1 or 4.2.2.

Configuration for this feature is lost during upgrade.
4.1.1 -> 4.1.2 : OK (no issue observed)
4.1.1 -> 4.2.1.1 or 4.2.2 : Reset to default value as disable (losing configuration)

Workaround

Upgrade to Apstra 4.1.2 at first before moving to Apstra 4.2.1.1 or 4.2.2 for non-service impact. The other option with brief service impact is to enable this feature again in the fabric policy before making the operational mode in the upgraded Apstra controller go from maintenance mode to ready mode, commit the new change, and then change the operational mode to ready status, which will push two commits, one from upgraded changes and the other from enabling the feature for generating EVPN hosts. For further assistance, please contact the Juniper Apstra Support Team.

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 customerâ€™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 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.

Workaround

The user can add an external generic system to the other even port in the port group for the ACX7100-32C so Apstra will correctly add the "unused" configuration parameter on the odd interface. Please refer to https://apps.juniper.net/port-checker/acx7100-32c/ for more information.

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†’ Uncommited 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 expecetd behavior.

Apstra-CLI system turn-beacon-on Command with Juniper EX4400 (AOS-43034)

The Juniper EX4400 platform does not have a dedicated Locator/ID LED. The Junos `"show chassis beacon"` command will always return ON. All ports or connected ports will be lit in `"GREEN"` depending on the explicit beacon command. Also, it uses a 5-minute default timer, and CLI supports between 1 and 120 minutes. After a predefined time, the beacon status changes back to the default state in CLI. The switch port status is not changing based on Junos `"request chassis beacon"` command.

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.

Workaround

After ZTP processing, manually update banner in the device like the following example or contact Juniper Apstra Support team for further assistance.

```
sed -i s/"#*Banner.\*$"/"Banner \/etc\/ssh\/my_banner"/ /etc/ssh/sshd_config
cat >& /etc/ssh/my_banner << EOF
###########################################################################
##########################
# This device is for the exclusive use of XXXXX.
```

```
# All unauthorized access or configuration changes to this device are subject
to prosecution.
##############################################################################
###########################
EOF

service ssh restart
```

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.

Workaround

Use configlet to explicitly to add maximum-paths into user-defined VRF's BGP configuration

```
Example Configlet:
router bgp {{ bgpService.asn }}
  {% for vrf in security_zones | list | reject('in', ['default']) %}
   vrf {{ vrf }}
   address-family ipv4 unicast
   maximum-paths 64
  {% endfor %}
```

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

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 Link returns error with '<' not supported between instane 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.

Workaround

Please assign the proper hostname into the generic system, connected to leaf node via port channel port.

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.

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.

Workaround

After Apstra Device Agent is created via ZTP process, update the pristine configuration with the changes, which assign the mgmt interface with static IP address and default GW address or use custom script file that assigns static IP address and default GW address during ZTP process in case SONiC 4.1.2 and Apstra 4.2.1 or 4.2.2 is used. Please contact Juniper Apstra Support for more details.

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 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.

Workaround

Clone builtin Cisco 93108TC-FX3P device profile, modify model field value of Selector from 93108TC-FX3P to C?93108TC-FX3P, and the assign the new device profile into the device.

For the further support, please contact Apstra Support Team.

DeviceTelemeryAgent 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 executd 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.

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.

Workaround

To enable log rotation, follow these steps:

```
1. Create logrotate configuration file for each container (dhcpd, tftp,
status) in the /containers_data/logrotate/<container_name> directory.
2. Create logrotate configuration file for rsyslog.log for status container in
the ZTP VM (/containers_data/logrotate/status directory).
3. Add a script to the ZTP VM's cron.hourly directory that executes the
logrotate command inside each container via docker exec, ensuring the logs are
rotated according to the specified configuration.
4. Modify Docker Compose file (/etc/apstra_ztp/docker-compose.yml) to map
logrotate configuration file in the ZTP VM into file inside container.
5. Restart the containers to apply the changes
```

Please contact Juniper Apstra Support team for the further support.

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

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

Workaround

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

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

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

The primary reasons for this are:

1. Deprecation of Older Collectors: The EVPN IBA Telemetry probe and associated collectors were designed for earlier NOS versions. These collectors relied on commands and logic that are no longer supported by certain vendors (e.g., with EOS 4.25.3.1M, evpn_type3 collector failed to gather data due to a plugin error triggered by an unsupported command).
2. Python Dependency Transition: Apstra versions prior to 4.2.x were based on Python2, whereas 4.2.x and later versions have transitioned to Python3. This shift in underlying architecture breaks backward compatibility for older telemetry packages.
3. Multiple Breakpoints in Compatibility: Over several releases, the extensible telemetry framework has undergone significant updates, requiring users to update their custom packages to align with newer standards and dependencies.

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

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

Workaround

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

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

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

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.

Workaround

The user must use a route target where the second number is less than 65536 (e.g. 64512:65535).

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.

Workaround

Please use /api/blueprints/{blueprint_id}/experience/web/cabling-map in the REST API explorer for exporting cabling map instead of UI export action.

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.

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

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

Workaround

Please remove parenthesis or bracket characters from system's label

---

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

---

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.

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.

Workaround

Verify that the warnings indicated under Error are benign and expected.

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.

Workaround

Make at least one port-channel interface rendered in the device by assigning CT with virtual network or sub-interface to the port channel interface.

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.

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

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

Workaround

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

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)

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.

Workaround

Manual workaround: Click the user icon at the top right and select Switch tenants to change back to the Global tenant.

Permanent workaround: modify `/etc/opensearch-dashboards/opensearch_dashboards.yml` file to change this:

```
opensearch_security.multitenancy.tenants.preferred: [Private, Global]
```

To this:

```
opensearch_security.multitenancy.tenants.preferred: [Global, Private]
```

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.

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;
    }
}
```

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

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

Workaround

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

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.

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

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

Workaround

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

---

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

---

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 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
```

Workaround

To generate the Device Show Tech data, use the following command directly on SONiC devices:

```
sudo python3 /usr/bin/aos_show_tech --platform sonic
```

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.

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

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

Workaround

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

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

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

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

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

Workaround

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

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.

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.

Sustained High Disk Utilization Observed After SONiC NOS Upgrade (AOS-49288)

When upgrading to SONiC 4.1.2 via Apstra, users may notice sustained high disk utilization on their devices. The affected SONiC devices have significant space usage on the /host partition, ranging from 8 to 30 GB. This leaves little free space, potentially affecting system performance

and stability. The problem has been identified as FRR logging into the BGP overlay2 storage rather than the correct location at /var/log/frr/frr.log. The disk does not fill up immediately following the NOS upgrade; it takes approximately 2.5 to 3 months for pre-production and 1.5 to 2 weeks for production. The rate of disk space consumption is proportional to the number of connected servers and the frequency of log-triggering events.

```
admin@sonic:~$ df -h
Filesystem      Size  Used Avail Use% Mounted on
:
/dev/sda3        32G   30G     0 100% /host  <<<<<<<<<<<<<<<

root@sonic:~# du -sh /host/image-4.1.2-Enterprise_Standard/docker/*
:
28G  /host/image-4.1.2-Enterprise_Standard/docker/overlay2  <<<<<<<<<<<<<<<
```

Workaround

As a workaround, the SONiC environment can use the configlet to execute log rotation for containers. The configlet example below utilizes a 10-minute cron job for SONiC devices that are at least 4.1.0. Kindly utilize the below configlet example to create a configlet that utilizes a cron job every 10 minutes, then import it into the blueprint for every SONiC device. Please get in touch with the Juniper Apstra Support team for additional assistance.

Configlet for SONiC Device

```
Section: FILE
Template Text:
  {% if function.min_os_version('4.1.0') %}
  # Rotate FRR logs inside bgp container per 10 minutes
  */10 * * * * root docker exec bgp logrotate --verbose /etc/logrotate.d/frr >
/tmp/bgp-docker-logrotate.log 2>&1
  # if less aggressive rotation, such as hour job, uncomment the below line
for an hourly job and comment the above line for a 10 minute job.
  # 0 * * * * root docker exec bgp logrotate --verbose /etc/logrotate.d/frr >
/tmp/bgp-docker-logrotate.log 2>&1
  {% endif %}
Filename: /etc/cron.d/bgp-docker-logrotate
```

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

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.

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.

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

---

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

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

Workaround

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

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.

ZTP devices, which use python3, fails in getting ztp_py3.py file via tftp (AOS-47007)

In Apstra ZTP < 5.0.0, ZTP for Junos EVO devices would fail as the `'ztp_py3.py'` is not available over tftp to provision due to missing the right file permissions.

Workaround

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

Known 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"}
```

Workaround

For Apstra versions 4.2.1 and 4.2.2, user can apply given Apstra workaround configlet to all SONiC devices using the text below.

This configlet will migrate the certificate from the Apstra device agent to the new 4.1.0+ trust store and re-enable certificate authentication for the HTTPS RESTCONF server. The configlet should be assigned in the Apstra UI to all SONiC devices across all blueprints, applied under the system section.

Note: For devices that are not deployed or not assigned to a blueprint, it is strongly recommended to create and run the script manually:

```
#!/bin/sh -

exec > /tmp/fix_cert.out
exec 2>&1

date

grep ^build_version /etc/sonic/sonic_version.yml | grep -q 4.0. && exit 0


if test -f /etc/aos/aos_ca_cert.pem; then
    cp /etc/aos/aos_ca_cert.pem /home/admin/aos.crt
    sudo -u admin sonic-cli -c 'crypto ca-cert install home://aos.crt' -c
config -c 'crypto trust-store aos ca-cert aos' -c 'crypto security-profile
trust-store aos aos' -c 'ip rest security-profile aos' -c exit
fi

true
```

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.

Workaround

We recommend modifying the SSHD configuration according to the Ubuntu Security Team's

guidance on CVE-2024-6387 to mitigate the risk. Contact Juniper Apstra Support for assistance.

```
Set LoginGraceTime to 0 in /etc/ssh/sshd_config
sudo systemctl restart sshd
```

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.

Workaround

Please contact Juniper Apstra Support Team

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 rasied 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 (>=23.4R2-S4).

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.

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 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
```

Juniper ACX Incomplete Packet Exported to Sflow Collector (AOS-42800)

When operating Juniper ACX devices, you might encounter a situation where layer-3 packets are inaccurately identified as layer-2 packets. This can result in incomplete packets being exported to the flow collector.

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 Virutal 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.

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

In the Apsta 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 "Decvices / Managed Devices / Pristine Configuration" Apstra UI or using the Apstra-CLI "system pristine_config_append" command.

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

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

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

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

Workaround

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

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.

Workaround

Upgrade to Junos 22.2R3-S3.

SONiC log rotation may not work, causing /var/log to get filled up (AOS-44012)

Due to `/usr/sbin/` not being in the default PATH and an explicit PATH is not set in cronjob `/etc/cron.d/logrotate`, the regular recurring rotation that is done by `/usr/bin/log-rotate.sh` which is exercised in `/etc/cron.d/logrotate` does not work correctly.

Despite that, if `disk-log-rotate-daemon` is operational, it will also invoke /usr/bin/log-rotate.sh itself, this time with the correct PATH set. However, for this to happen, files `/var/log/rotate/disk/info` and `/var/log/rotate/disk/debug` have to be present in `/var/log`. If these files are removed for any reason by the user, then `/usr/bin/log-rotate.sh` will not be invoked by the `disk-log-rotate-daemon`.

Workaround

Please never manually delete files under `/var/log` in a SONiC device, especially the files mentioned in the description.

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 `vtysh` 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.

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.

Workaround

Avoid attaching hosts to EVPN VXLAN Stitching DCI Border leafs, or disable Symmetric IRB.

Known 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

Workaround

Manual workaround:

Click the user icon at the top right and select Switch tenants to change back to the Global tenant.

Permanent workaround:

```
modify /etc/opensearch-dashboards/opensearch_dashboards.yml file to change
following:

    from: opensearch_security.multitenancy.tenants.preferred: [Private,
Global]

    to: opensearch_security.multitenancy.tenants.preferred: [Global, Private]
```