

Platforms

- Ubuntu 24.04 platform upgrades, on page 1

Ubuntu 24.04 platform upgrades

Ubuntu platform upgrades

A platform upgrade is an operating system transition that

- addresses critical and high-security vulnerabilities through updated OS packages,
- extends the platform lifecycle by migrating to a newer Long Term Support (LTS) release, and
- maintains compatibility with existing network functions and infrastructure.

Transitioning to Ubuntu 24.04 LTS enables organizations to proactively address security concerns, avoid unsupported software risks after 2027, and minimize disruptions to ongoing network operations.

Before changing ISSM between CPS DRA 26.1 and 26.2, take and verify a Consul snapshot from the running 26.1 cluster (Consul 1.5.3). Take this snapshot before upgrading from 26.1, and keep it available for any later downgrade to 26.1. Consul 1.5.3 and Consul 1.18.0 are not backward-compatible at the Raft log level; a snapshot created on 26.2 cannot replace the required 26.1 snapshot.

Before taking the snapshot, verify that all seven Consul servers are HEALTHY by running `show consul-status`.

Save the snapshot by running `consul save-snapshot consul-1.5.3-backup`. Verify that it is listed by running `consul list-snapshots`.

When downgrading to 26.1, use the saved 26.1 snapshot and include the `consul-downgrade` option and saved snapshot name with the system downgrade CLI command.

Manage netplan network migration

Cisco Policy Suite is currently designed to use the ifstate networking service. While Netplan is a default in modern Ubuntu releases, it is not supported by the current DRA architecture. Ensure that Netplan remains disabled and that the ifstate service remains the primary networking service. Planning for a future migration to Netplan should be treated as a separate, subsequent effort.

Verify the ubuntu version

Confirm that the system is running Ubuntu 24.04 as expected.

Use this task after performing system updates or during initial validation to ensure the operating system matches your environment's deployment standard.

Follow these steps to verify the ubuntu version:

Procedure

Step 1 Access the system's command line interface.

Step 2 Run the following command: `#cat /etc/lsb-release` to display the OS version.

Step 3 Check that the output includes:

Example:

```
```
   DISTRIB_RELEASE=24.04
   ```
```

The system displays the current Ubuntu distribution details, confirming that Ubuntu 24.04 is installed.

Verify network service status

Confirm the operational status of the network service and ensure compliance with Cisco Policy Suite requirements.

Cisco Policy Suite relies on the ifstate networking service. It is essential to verify that the system is correctly configured and not using the unsupported Netplan service.

Follow these steps to verify the network service status:

Procedure

Step 1 Execute the **netplan info** command. If the system displays "command not found," the netplan service is correctly disabled.

Step 2 Check the status of the ifstate file by running the following command:

Example:

```
#cat /run/network/ifstate
```

Step 3 Review the output to confirm that the expected loopback and network interface configurations are listed.

You confirm that the network service is properly configured for the Cisco Policy Suite and operates using the supported ifstate service. The system is not using the unsupported Netplan service.

Restore consul cluster stability after downgrade

Restore stable operation to the Consul service following a downgrade, by purging potentially corrupted data and restarting all necessary services.

If a Consul cluster running version 1.18.0 does not stabilize after downgrading to version 1.5.3, a manual cleanup is required to remove incompatible Raft data and restore service health.

Before you begin

- The verified Consul snapshot from the running CPS DRA 26.1 cluster (Consul 1.5.3), taken before the upgrade. Record the snapshot name for downgrade recovery.
- Access to the Orchestrator CLI for all affected Consul nodes (consul-1 through consul-7).
- Administrative privileges to stop and start services.

Procedure

Step 1 Stop all supervisor-managed services on each Consul container (consul-1 through consul-7) by running `supervisorctl stop all`.

Step 2 Remove the incompatible Consul 1.18.0 Raft data from each container and restart it with an empty `/data` directory:

Example:

```
rm -rf /data/*
kill 1
```

Step 3 After each container restarts, start its services by running `supervisorctl start all`.

Step 4 Wait up to 5 minutes for at least five Consul containers to become HEALTHY and elect a leader.

Step 5 Restore the saved 26.1 snapshot by following the Consul snapshot restore procedure in the CPS Troubleshooting Guide. Do not use a snapshot created on 26.2.

Troubleshooting: Post-downgrade recovery

Resolving Diameter-endpoint and binding issues after Ubuntu 24.04 downgrade

Restore correct operation of diameter-endpoint and binding containers by correcting the resolver comment, restoring the HAProxy wrapper, repairing the generated server directive, and restarting services.

After downgrading, diameter-endpoint or binding containers may remain unhealthy because Ubuntu 24.04 adds a metadata comment to `/etc/resolv.conf`. The comment can corrupt the HAProxy server directive generated in `haproxy.cfg`. Apply these recovery steps inside each affected container.

Before you begin

- Access each affected diameter-endpoint or binding container from the orchestrator CLI.
- Administrative privileges to edit the container files and restart supervisor-managed services.

Procedure

Step 1 Remove the `# Overrides:` comment from `/etc/resolv.conf` by writing the filtered content back to the existing file:

Do not use `sed -i`; `/etc/resolv.conf` is a Docker bind mount and replacing it by rename can fail.

Example:

```
tmp=/tmp/resolv.conf.$$ && \
sed '/^# Overrides:/d' /etc/resolv.conf > "$tmp" && \
cat "$tmp" > /etc/resolv.conf && \
rm -f "$tmp"
```

a) Confirm that the comment is gone. Run `grep 'Overrides' /etc/resolv.conf`. The command should return no output.

Step 2 Check the sibling helper script without modifying it, and restore /usr/local/bin/run-haproxy as the wrapper for the HAProxy binary.

Example:

```
cat /usr/local/bin/run-haproxy-process
cat > /usr/local/bin/run-haproxy << 'EOF'
#!/bin/bash
exec /usr/local/sbin/haproxy -f /etc/haproxy/haproxy.cfg -W -db
EOF
chmod +x /usr/local/bin/run-haproxy
```

Step 3 Identify the malformed `lbvip02` entry in `/etc/haproxy/haproxy.cfg`:

The corrupt directive appears as two lines; the second line starts with `Overrides::11211`.

Example:

```
grep -n 'Overrides\|lbvip02' /etc/haproxy/haproxy.cfg
```

Step 4 Fix the corrupt server directive in `haproxy.cfg` with this Python command:

Example:

```
python3 -c "
content = open('/etc/haproxy/haproxy.cfg').read()
bad  = '         server lbvip02 lbvip02.weave.local\nOverrides::11211 check inter 10s resolvers dns
resolve-prefer ipv4'
good = '        server lbvip02 lbvip02.weave.local:11211 check inter 10s resolvers dns resolve-prefer
 ipv4'
open('/etc/haproxy/haproxy.cfg','w').write(content.replace(bad, good))
"
```

Step 5 Verify the fix by running `grep -n 'Overrides\|lbvip02' /etc/haproxy/haproxy.cfg`. The output must show one complete `lbvip02` directive ending in `:11211` and no `Overrides` entry.

Step 6 Validate the HAProxy configuration:

Example:

```
/usr/local/sbin/haproxy -f /etc/haproxy/haproxy.cfg -c
```

Step 7 Restart all supervisor-managed services:

Example:

```
supervisorctl restart all
```

Step 8 Confirm that the affected services are HEALTHY.

The affected diameter-endpoint and binding services are HEALTHY after the HAProxy configuration and supervisor services are restored.

Resolving HAProxy-Common Issues

Summary

After Consul cleanup and snapshot restoration, HAProxy-Common containers may remain in the STARTING state due to missing or stale crd-api/header keys. Follow these stages in order to restore the containers and verify service and Raft health.

The key components involved in the process are:

- Orchestrator service: Responsible for managing HAProxy and synchronizing crd-api/header keys.
- haproxy-common containers: Containers that may become stuck in the STARTING state after a Consul restore.
- Admin operator: Executes commands to restart services and verify container health.

Workflow

1. Restart the orchestrator on the host by running `docker restart orchestrator`. This triggers the HAProxySecuritySynchronization service to repopulate the `crd-api/header` key.
2. Wait approximately 30 seconds. Connect to the orchestrator container with `docker connect orchestrator`, run `supervisorctl status all`, and confirm that all programs are RUNNING. Then exit the container.
3. If any `haproxy-common` container remains in STARTING, identify it with `show docker service | tab | include haproxy-common` and restart each stuck container individually with `docker restart haproxy-common-s[ID]`.
4. Verify service health by running `show docker service | tab | exclude HEAL`. No output means all services are HEALTHY.
5. Verify Consul Raft by connecting to the orchestrator and running `consul operator raft list-peers`. Confirm that the output lists all seven servers, with one leader and six followers, and no `<nil>` addresses.

Result

HAProxy-Common containers are HEALTHY, all services pass the health check, and the Consul Raft peer list shows one leader and six followers.