Oracle Linux Automation Manager 2.1

Installation Guide

F80482-05
November 2024

Oracle Linux Automation Manager 2.1 Installation Guide,

F80482-05

Contents

Preface

Oracle Linux Automation Manager 2.1: Installation Guide describes how to install Oracle Linux Automation Manager in single-host deployments or clustered deployments.

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
italic	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
`monospace`	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at https://www.oracle.com/corporate/accessibility/.

Access to Oracle Support for Accessibility

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit https://www.oracle.com/corporate/accessibility/learning-support.html#support-tab.

Diversity and Inclusion

Oracle is fully committed to diversity and inclusion. Oracle respects and values having a diverse workforce that increases thought leadership and innovation. As part of our initiative to build a more inclusive culture that positively impacts our employees, customers, and partners, we are working to remove insensitive terms from our products and documentation. We are also mindful of the necessity to maintain compatibility with our customers' existing technologies and the need to ensure continuity of service as Oracle's offerings and industry standards evolve. Because of these technical constraints, our effort to remove insensitive terms is ongoing and will take time and external cooperation.

1

Oracle Linux Automation Manager Requirements

This chapter describes the requirements for the systems to be used in an installation of Oracle Linux Automation Manager.

Oracle Linux Automation Manager Hardware Requirements

You can install Oracle Linux Automation Manager on a single machine or in a clustered setup in x86-64 Oracle Linux 8 hosts.

Oracle Linux Automation Manager does not require specific hardware; however, certain operations are memory intensive and require a certain amount of disk space and CPU. A minimum configuration is:

- 4 GB RAM
- 40 GB disk space (170 GB is recommended)
- Two core CPU

These requirements are for the minimum to run Oracle Linux Automation Manager. You must determine any additional hardware requirements and capacity based on your operational needs. For more information, see the upstream documentation.

2

Planning the Installation

This chapter provides information about planning your installation.

Oracle Linux Automation Manager Node Architecture

Oracle Linux Automation Manager supports a Service Mesh that provides a multi-service network that links nodes within a secure mesh. The Service Mesh can include up to 20 nodes that have the following node types:

- **Control Nodes**: These nodes provide management functions such as launching system jobs, inventory updates, and project synchronizations. Control nodes use ansible-runner which in turn uses Podman to run jobs within the **Control Plane Execution Environment** execution environments. The **Control Plane Execution Environment** execution environment references the `olam-ee` container image found on the Oracle Linux Container Registry. Control nodes do not run standard jobs.
- **Execution Nodes**: These nodes run standard jobs using ansible-runner which in turn uses Podman to run playbooks within **OLAM EE** execution environments. The **OLAM EE** execution environment references the `olam-ee` container image found on the Oracle Linux Container Registry. Execution nodes do not run management jobs. Execution nodes can also run custom execution environments that you can create using the Builder utility. For more information about custom execution environments, see Oracle Linux Automation Manager 2.1: Private Automation Hub Installation Guide. For more information about using Podman and the Oracle Linux Container Registry, see Oracle Linux: Podman User's Guide.
- **Hybrid Nodes**: Hybrid nodes combine the functions of both control nodes and execution nodes into one node. Hybrid nodes are supported in single host Oracle Linux Automation Manager deployments, but not in clustered multi-host deployments.
- **Hop Nodes**: You can use hop nodes to connect control nodes to execution nodes within a cluster. Hop nodes cannot run playbooks and do not appear in instance groups. However, they do appear as part of the service mesh.

Oracle Linux Automation Manager can manage a variety of different instance types, such as devices, servers, databases, network equipment, and so on. In general, Oracle Linux Automation Manager manages the following instance types:

- **Directly Managed Instances**: Directly managed instances (nodes) are any virtual, physical, or software instances that Oracle Linux Automation Manager manages over an ssh connection.
- **Indirectly Managed Instances**: Indirectly managed instances (nodes) include any identifiable instance not directly connected to Oracle Linux Automation Manager, but managed by a device that is directly connected to Oracle Linux Automation Manager.

For example, the following image illustrates all node and instance types and some of the ways that they can be related to one another.

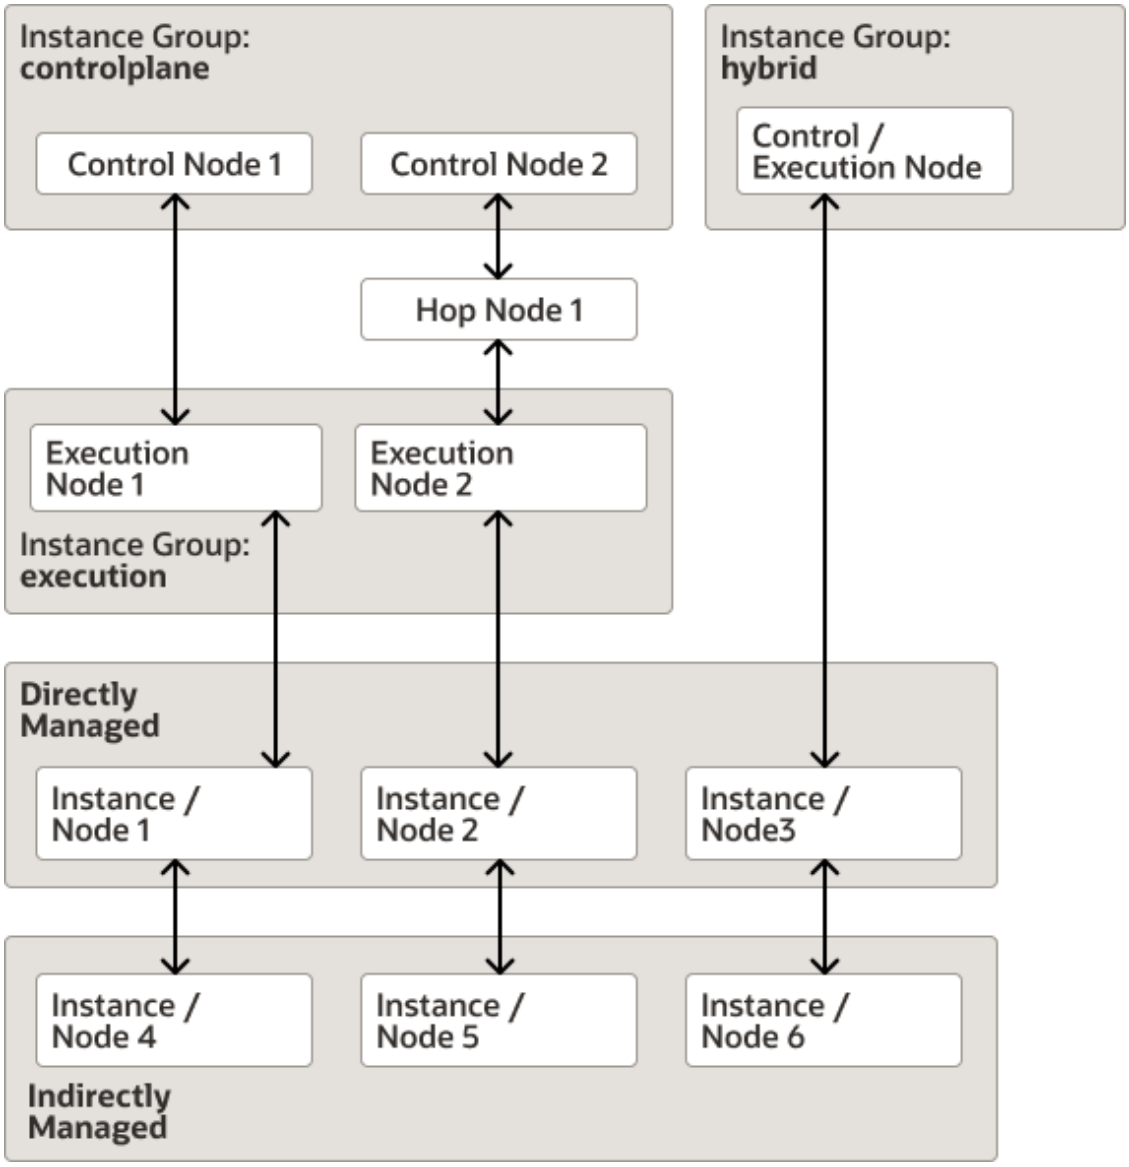

Installation Options

Oracle Linux Automation Manager provides three installation options:

- Standalone Installation: All components are on the same host, including the database.

Figure 2-1 Standalone Installation with Local Database

- Standalone Installation with Remote Database: All components are on the same host, with the exception of the database which is on a remote host.

Figure 2-2 Standalone Installation with Remote Database

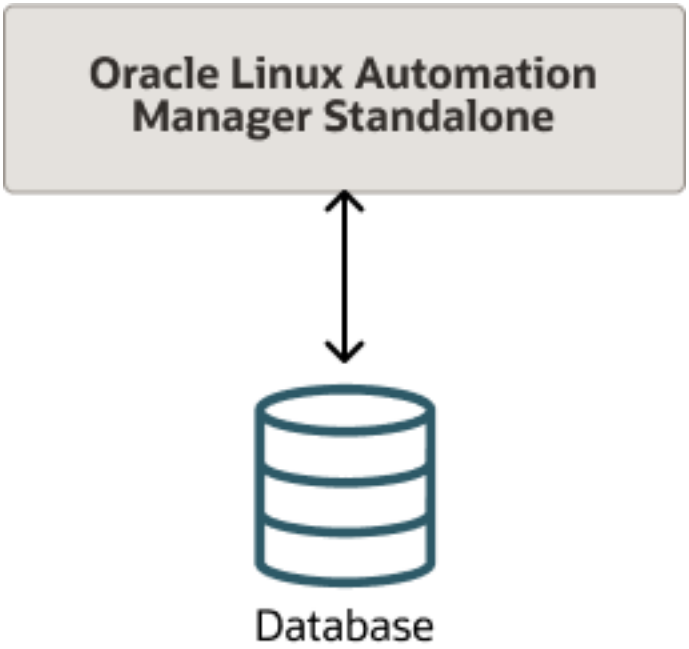

- Clustered Installation with Remote Database: Clustered installation can contain up to 20 nodes with one or more control node, one or more execution nodes, and one or more hop nodes all connected to one database. For example, the following shows a cluster with two control plane nodes and two execution plane nodes, each on separate hosts, and all of them connected to a remote database.

Figure 2-3 Clustered Installation with Remote Database

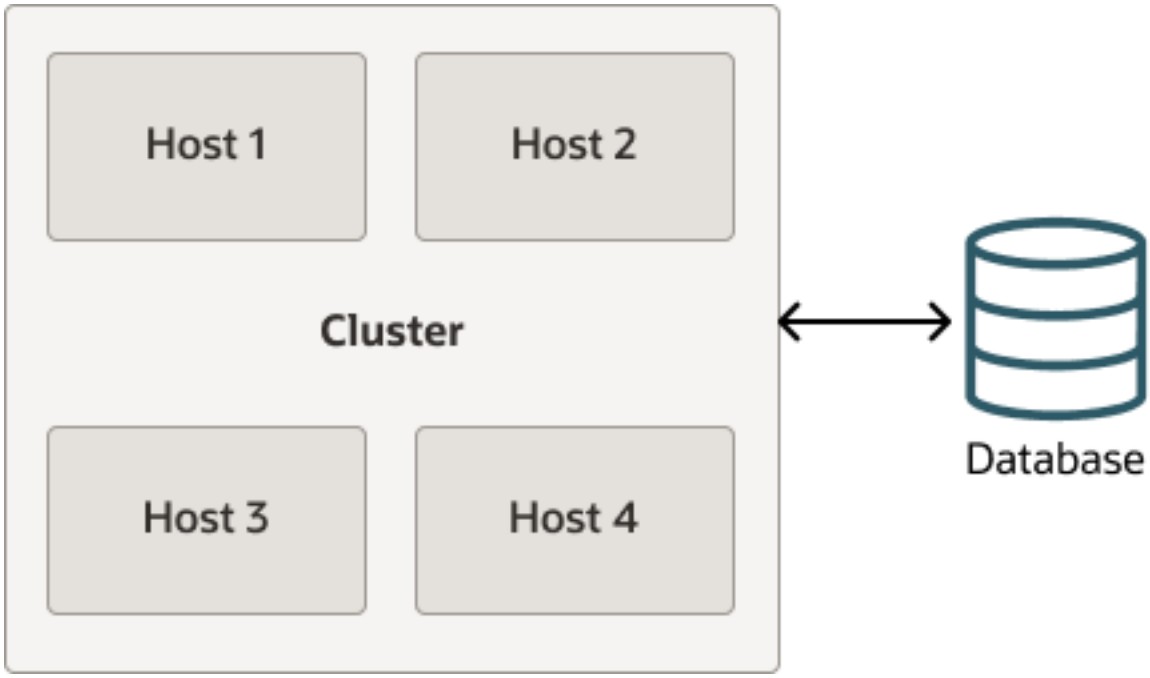

Service Mesh Topology Examples

There are a variety of ways that you can configure the Oracle Linux Automation Manager Service Mesh topology.

Example 1: Design the Service Mesh such that you have at least one backup control plane node and one backup execution plane node. For example, two control nodes and two execution nodes. Each execution plane node would have communication with both control plane nodes in case one of the control plane node were to fail. If the first execution plane node were to fail, the control plane node would switch to the second execution plane node.

Figure 2-4 Clustered Installation with Remote Database

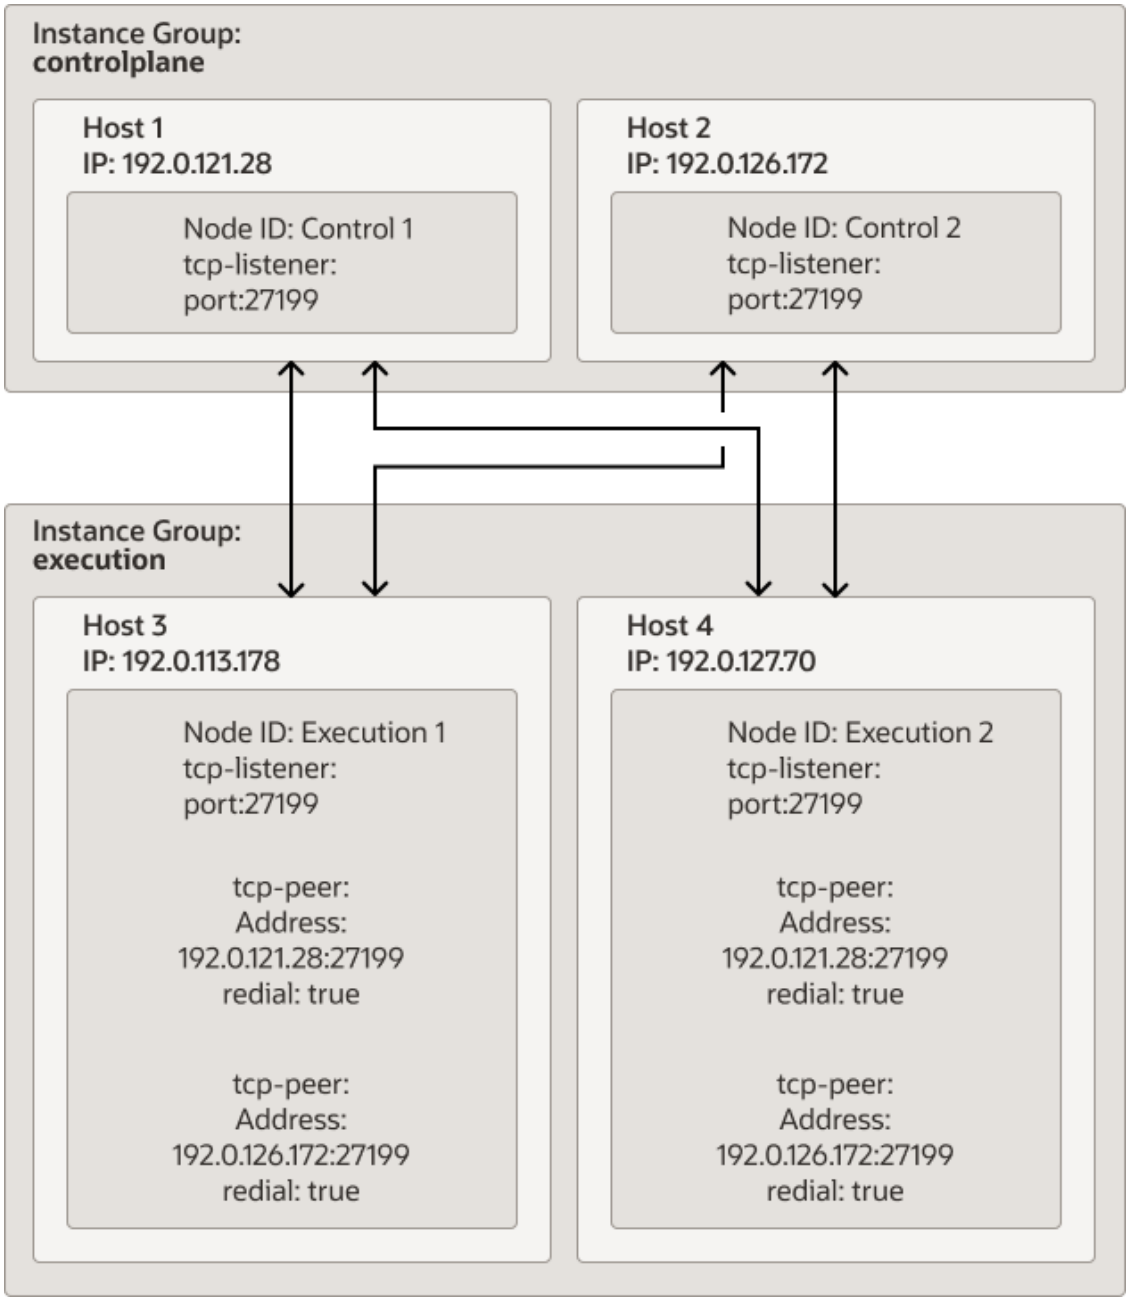

The high level steps to configure the Service Mesh are as follows:

1. Configure the `/etc/receptor/receptor.conf` file with the Node ID, tcp-listener, and tcp-peer addresses as required for each node. For more information about this task, see Configuring and Starting the Control Plane Service Mesh and Configuring and Starting the Execution Plane Service Mesh.
2. From a control plane node, log in as the awx user, and run the awx-manage command to do the following:

a. Provision each host's IP address or host name, and designate it as a control plane or execution plane node type. For example, the following commands provision two control plane and two execution plane nodes as illustrated in the figure above:

```
awx-manage provision_instance --hostname=192.0.121.28 --
node_type=control
awx-manage provision_instance --hostname=192.0.126.72 --
node_type=control
awx-manage provision_instance --hostname=192.0.113.178 --
node_type=execution
awx-manage provision_instance --hostname=192.0.127.70 --
node_type=execution
```

b. Register each node to either the `controlplane` or the `execution` instance group, based on the type of node you designated for each node. The awx-manage command refers to instance groups as queuenames. For example, the following commands create the controlplane and execution instance groups and associates the two control plane and two exection plane nodes to each instance group as illustrated in the figure above:

```
awx-manage register_queue --queuename=controlplane --
hostnames=192.0.121.28
awx-manage register_queue --queuename=controlplane --
hostnames=192.0.126.72
awx-manage register_queue --queuename=execution --
hostnames=192.0.113.178
awx-manage register_queue --queuename=execution --hostnames=192.0.127.70
```

c. Register the peer relationship between each node. Note that when you register a peer relationship between a source IP address to a target IP address, the peer relationship establishes bidirectional communication. For example, the following commands registers the host IP address of the execution nodes as the source and each tcp-peer connection are the targets, which are the control plane nodes:

```
awx-manage register_peers 192.0.113.178 --peers 192.0.121.28
awx-manage register_peers 192.0.113.178 --peers 192.0.126.172
awx-manage register_peers 192.0.127.70 --peers 192.0.121.28
awx-manage register_peers 192.0.127.70 --peers 192.0.126.172
```

The command must be run for each link you want to establish between nodes.

d. Register each instance group as the default queuename for either the control plane or the execution plane. This ensures that only control type jobs go to the control plane instance group and only Oracle Linux Automation Engine jobs go to execution plane instance group. To do this, you must edit the `/etc/tower/settings.py` file with the `DEFAULT_EXECUTION_QUEUE_NAME` and the `DEFAULT_CONTROL_PLANE_QUEUE_NAME` parameters.

For more information about these steps, see Configuring the Control, Execution, and Hop Nodes.

Example 2: Deploy as many control and execution plane nodes as you require such that you build in fail over in case any control or execution plane node fails. Ensure you don't exceed the 20 node limit for your cluster. Additional options you can consider are:

- In some cases you may have an execution node that cannot be directly connected to a control plane node. In such cases you can connect the execution node to another

execution node that is connected to the control node. This does introduce a risk such that if the intermediate execution node were to fail, then the connected execution node would become inaccessible to the control node.

- In some cases you may have an execution node that cannot be directly connected to a control plane node. In such cases you can connect the execution node to a hop node that is connected to the control node. This does introduce a risk such that if the intermediate hop node were to fail, then the connected execution node would become inaccessible to the control node.
- Establishing a peer relationship between control plane nodes. This ensures that control plane nodes are always directly accessible to one another. If no such relationship is established, then control plane nodes are aware of each other through connected execution plane nodes. For example, control A connects to control B through execution A which is connected to both.

Tuning Instances for Playbook Duration

Oracle Linux Automation Manager monitors jobs for status changes. For example, some job statuses are Running, Successful, Failed, Waiting, and so on. Normally the playbook being run triggers status changes as it makes progress in various ways. However, in some cases, the playbook will get stuck in the Running or Waiting state. When this happens, a reaper process automatically changes the state of the task from Running or Waiting to Failed. The default timer for when the reaper changes the status of a stuck job to the Failed state is 60 seconds.

If you have jobs that are designed to run longer than 60 seconds, then modify the **REAPER_TIMEOUT_SEC** parameter to the `/etc/tower/settings.py` file. Specify a time in seconds that is longer than the duration that your playbooks with the longest duration is expected to run. This avoids scenarios where the reaper mistakenly sets a long running playbook to the Failed state because the **REAPER_TIMEOUT_SEC** value has expired.

A possible scenario could occur if you run many short and long duration playbooks together with a reaper that has a long timeout value. If one or more of the short duration playbooks run for longer than expected, (for example, because of a network outage making it impossible for these playbooks to complete) the reaper continues to track the status of the stuck short duration playbooks until they either get unstuck and transition to the Successful state or until the reaper timeout value is reached. This scenario should cause no performance difficulties if only a few such failures were to occur. However, if hundreds of such failures were to occur at the same time, Oracle Linux Automation Manager would waste resources on tracking these stuck jobs and could degrade the performance of the host processing the jobs.

For more information about setting the **REAPER_TIMEOUT_SEC** parameter, see Setting up Hosts.

3
Preparing the Database and Hosts

The following chapter provides information about setting up the network firewalls, database, and hosts for your Oracle Linux Automation Manager installation. This chapter also discusses how to enable the repositories to install the Oracle Linux Automation Manager packages.

Setting Up the Network

This section contains information about the generic networking requirements for an Oracle Linux Automation Manager hosts, the database host and shows you an example of how to set up the network to enable the communication between the Oracle Linux Automation Manager host and the inventory hosts in an environment.

Setting Up the Firewall Rules

Oracle Linux 8 installs and enables `firewalld`, by default. Example commands to open the ports and to set up the firewall rules are provided below.

On the Oracle Linux Automation Manager hosts, run the following firewalld commands:

```
sudo firewall-cmd --add-port=27199/tcp --permanent
sudo firewall-cmd --add-service=http --permanent
sudo firewall-cmd --add-service=https --permanent
sudo firewall-cmd --reload
```

> **Note:**
>
> Port 27199 provides a TCP listener port for the Oracle Linux Automation Manager service mesh and must be open on each node in the mesh. The HTTP and HTTPS ports are for the Nginx server.

If you choose to install a remote database, open the following port on the host running the database:

```
sudo firewall-cmd --add-port=5432/tcp --permanent
sudo firewall-cmd --reload
```

Enabling Access to the Oracle Linux Automation Manager Packages

This section contains information on setting up the locations for the operating system on which you want to install the Oracle Linux Automation Manager software packages.

Enabling Channels with ULN

If you are registered to use ULN, use the ULN web interface to subscribe the system to the appropriate channels.

To subscribe to the ULN channels:

1. Log in to https://linux.oracle.com with your ULN user name and password.
2. On the Systems tab, click the link named for the system in the list of registered machines.
3. On the System Details page, click **Manage Subscriptions**.
4. On the System Summary page, select each required channel from the list of available channels and click the right arrow to move the channel to the list of subscribed channels. Subscribe the system to the following channels:
   - `ol8_x86_64_automation2`
   - `ol8_x86_64_addons`
   - `ol8_x86_64_baseos_latest`
   - `ol8_x86_64_UEKR6` or `ol8_x86_64_UEKR7`
   - `ol8_x86_64_appstream`
5. Click **Save Subscriptions**.

Enabling Repositories with the Oracle Linux Yum Server

If you are using the Oracle Linux yum server for system updates, enable the required yum repositories.

To enable the yum repositories:

1. Use the `dnf config-manager` tool to enable the `ol8_baseos_latest` repository.

   ```
   sudo dnf config-manager --enable ol8_baseos_latest
   ```

   > **Note:**
   >
   > This repository is typically enabled by default.

2. Install oraclelinux-automation-manager-release-el8:

   ```
   sudo dnf install oraclelinux-automation-manager-release-el8-2.1
   ```

3. Enable the following yum repositories including the Oracle Linux Automation Manager release 2 repository:
   - `ol8_addons`
   - `ol8_UEKR6` or `ol8_UEKR7`
   - `ol8_appstream`

   Use the `dnf config-manager` tool to enable the yum repositories and do one of the following:
   - If you are using ol8_UEK6, use the following command:

     ```
     sudo dnf config-manager --enable ol8_addons ol8_UEKR6 ol8_appstream
     ```

   - If you are using ol8_UEK7, use the following command:

     ```
     sudo dnf config-manager --enable ol8_addons ol8_UEKR7 ol8_appstream
     ```

Setting Up a Local or Remote Database

To setup a local or remote Postgresql database instance on Oracle Linux 8 for Oracle Linux Automation Manager single host or multi-host configurations, do the following:

1. Install and configure Oracle Linux 8 on a host.
2. If the database is remote, open the database port in the firewall as described in Setting Up the Firewall Rules.
3. Enable the `postgresql 12` or `postgresql 13` module stream.

   ```
   sudo dnf module reset postgresql
   sudo dnf module enable postgresql:12
   ```

   or

   ```
   sudo dnf module reset postgresql
   sudo dnf module enable postgresql:13
   ```

   > **Note:**
   >
   > For more information about the Postgresql 12 and 13 life cycle, see the appendix discussing the application life cycle for stream modules in Oracle Linux: Managing Software on Oracle Linux.

4. Install the database.

   ```
   sudo dnf install postgresql-server
   ```

5. Initialize the database:

   ```
   sudo postgresql-setup --initdb
   ```

6. In the `/var/lib/pgsql/data/postgresql.conf` file, switch the password storage mechanism from `md5` to `scram-sha-256`. For example, the following command makes the switch for you:

   ```
   sudo sed -i "s/#password_encryption.*/password_encryption = scram-
   sha-256/"  /var/lib/pgsql/data/postgresql.conf
   ```

7. Start the database using the following command that also ensures that the database restarts in case the host restarts:

   ```
   sudo systemctl enable --now postgresql
   ```

8. Ensure the database is running:

   ```
   sudo systemctl status postgresql
   ```

9. Create the database user accounts. For example:

```
sudo su - postgres -c "createuser -S -P awx"
```

10. Enter and confirm the password for the awx user.

```
Enter password for new role:
Enter it again:
```

11. Create the database.

```
sudo su - postgres -c "createdb -O awx awx"
```

12. As the root user, in the `/var/lib/pgsql/data/pg_hba.conf` file add the following line:

```
host  all  all 0.0.0.0/0 scram-sha-256
```

13. As the root user, in the /var/lib/pgsql/data/postgresql.conf file in the # CONNECTIONS AND AUTHENTICATION section, a line with the text `listen_addresses =` followed by the IP address or host name of your database in single quotes. For example:

```
listen_addresses = '<IP address or host name>'

#listen_addresses = 'localhost'         # what IP address(es) to listen on;
                                        # comma-separated list of
addresses;
                                        # defaults to 'localhost'; use '*'
for all
                                        # (change requires restart)
#port = 5432                            # (change requires restart)
```

In the previous example, *<IP address or hostname>* is the IP address or host name of the database.

14. Restart the database.

```
sudo systemctl restart postgresql
```

15. You are now ready to setup your hosts as described in Setting up Hosts.

Setting up Hosts

This section provides information for setting up one or more hosts intended to run Oracle Linux Automation Manager in any of the configurations listed in Installation Options.

To set up one or more hosts:

1. Install Oracle Linux Automation Manager.

```
sudo dnf install ol-automation-manager
```

2. If you are creating a cluster, choose the `/etc/tower/SECRET_KEY` from one node and replace the value of the `/etc/tower/SECRET_KEY` on all other nodes with the value from your chosen node. Ensure the file user and group ownership is `awx:awx` on all nodes. The

end result should be that all nodes have the same value in their `/etc/tower/SECRET_KEY` file.

3. Edit the `/etc/redis.conf` file to include the following lines:

```
unixsocket /var/run/redis/redis.sock
unixsocketperm 775
```

4. Edit the `/etc/tower/settings.py` file configure the `CLUSTER_HOST_ID` field:

```
CLUSTER_HOST_ID = "hostname or ip address"
```

In the previous example, *hostname or ip address* is the hostname or IP address of the system running Oracle Linux Automation Manager. If hostname is used, the host must be resolvable.

5. Replace the existing `DATABASES` fields with the following fields:

```
DATABASES = {
    'default': {
        'ATOMIC_REQUESTS': True,
        'ENGINE': 'awx.main.db.profiled_pg',
        'NAME': 'awx',
        'USER': 'awx',
        'PASSWORD': 'password',
        'HOST': 'database hostname or ip address',
        'PORT': '5432',
    }
}
```

In the previous example, *database hostname or ip address* is the hostname or IP address of the local or remote database. If hostname is used, the host must be resolvable. *password* is the password for your database, if you have configured one.

6. If you have playbooks designed to run longer than the default reaper timeout of 60 seconds, change the **REAPER_TIMEOUT_SEC** parameter to increase the timeout. For example,

```
REAPER_TIMEOUT_SEC=<longest_playbook_time>
```

In the previous example, <*longest_playbook_time*> is number of seconds that exceeds the duration of the longest playbook runtime.

7. Run the following commands on all hosts:

```
sudo su -l awx -s /bin/bash
podman system migrate
podman pull container-registry.oracle.com/oracle_linux_automation_manager/olam-
ee:latest
exit
```

> **Note:**
>
> After you finish installing Oracle Linux Automation Manager, you can configure whether you want your Execution Environments to always pull the latest `olam-ee` container image when running playbooks, or use some other option or custom image. For more information about these options, see Oracle Linux Automation Manager 2.1: User's Guide. For more information about Private Automation Hub, see Oracle Linux Automation Manager 2.1: Private Automation Hub User's Guide .

> **Note:**
>
> The previous command assumes that you are pulling the olam-ee image directly from the Oracle Container Registry. If you are using Private Automation Hub or have setup a custom container registry, you can pull the image from there instead. In addition, you can configure Oracle Linux Automation Manager to always pull from that container registry by replacing Oracle Container Registry path to your custom container registry path in the following fields in the `/etc/tower/settings.py` file:
>
> ```
> GLOBAL_JOB_EXECUTION_ENVIRONMENTS = [{'name': 'OLAM EE (latest)',
> 'image': 'container-registry.oracle.com/
> oracle_linux_automation_manager/olam-ee:latest'}]
> CONTROL_PLANE_EXECUTION_ENVIRONMENT = 'container-
> registry.oracle.com/oracle_linux_automation_manager/olam-ee:latest'
> ```

8. Run the following commands on one control host (in a clustered deployment) or on the single host (in single host deployment):

   ```
   sudo su -l awx -s /bin/bash
   awx-manage migrate
   awx-manage createsuperuser --username admin --email email
   ```

   In the previous example, *email* is the email address of the admin user.

9. Enter and repeat the password for the admin user.

   ```
   Password:
   Password (again):
   ```

10. Exit the awx user .

    ```
    exit
    ```

11. On all hosts, generate SSL certificates for NGINX:

> **Note:**
>
> The following instruction explains how to create a self-signed certificate for use by NGINX as part of Oracle Linux Automation Manager. It is recommended that on production systems you use CA signed certificates for this purpose. For more information on working with SSL certificates, see Oracle Linux: Managing Certificates and Public Key Infrastructure.

```
sudo openssl req -x509 -nodes -days 365 -newkey rsa:2048 -keyout /etc/tower/
tower.key -out /etc/tower/tower.crt
```

12. Remove any default configuration for NGINX. Edit `/etc/nginx/nginx.conf` to contain the following configuration:

```
user nginx;
worker_processes auto;
error_log /var/log/nginx/error.log;
pid /run/nginx.pid;

# Load dynamic modules. See /usr/share/doc/nginx/README.dynamic.
include /usr/share/nginx/modules/*.conf;

events {
    worker_connections 1024;
}

http {
    log_format  main  '$remote_addr - $remote_user [$time_local] "$request" '
                      '$status $body_bytes_sent "$http_referer" '
                      '"$http_user_agent" "$http_x_forwarded_for"';

    access_log  /var/log/nginx/access.log  main;

    sendfile            on;
    tcp_nopush          on;
    tcp_nodelay         on;
    keepalive_timeout   65;
    types_hash_max_size 2048;

    include             /etc/nginx/mime.types;
    default_type        application/octet-stream;

    # Load modular configuration files from the /etc/nginx/conf.d directory.
    # See http://nginx.org/en/docs/ngx_core_module.html#include
    # for more information.
    include /etc/nginx/conf.d/*.conf;
}
```

> **Note:**
>
> For advanced NGINX users, the Oracle Linux Automation Manager NGINX configuration file is located in `/etc/nginx/conf.d/ol-automation-manager-nginx.conf`. For example, you may use a different version of TLS or have different ciphers configured. If you have an existing customized NGINX setup, ensure that you also apply the `ol-automation-manager-nginx.conf` settings.

13. You are now ready to install Oracle Linux Automation Manager in a cluster or on a single host. For more information, see Installing Oracle Linux Automation Manager on a Single-Host Deployment and Installing Oracle Linux Automation Manager in a Clustered Deployment.

4

Installing Oracle Linux Automation Manager on a Single-Host Deployment

This chapter shows you how to set up a host and install the Oracle Linux Automation Manager software and includes an option for using a remote or local database.

Installing on a Single Host

This section provides instructions for installing the Oracle Linux Automation Manager on a single host where the database is local or on a remote host.

To set up the host:

1. On the Oracle Linux Automation Manager host, run the following commands:

   ```
   sudo su -l awx -s /bin/bash
   ```

2. Enter the following command:

   ```
   awx-manage provision_instance --hostname=<hostname or IP address> --node_type=hybrid
   ```

   In the previous example, *hostname or IP address* is the hostname or IP address of the system running Oracle Linux Automation Manager. If hostname is used, the host must be resolvable.

3. Run the following command to register the default execution environments, which are:
   - Control Plane Execution Environment
   - OLAM EE: (Latest)

   ```
   awx-manage register_default_execution_environments
   ```

4. Run the following command to create the default queue for standard jobs that run playbooks:

   ```
   awx-manage register_queue --queuename=default --hostnames=<hostname or IP address>
   ```

5. Run the following command to create the controlplane queue for Oracle Linux Automation Manager management type jobs.

   ```
   awx-manage register_queue --queuename=controlplane --hostnames=<hostname
   or IP address>
   ```

6. Exit the awx shell environment.

   ```
   exit
   ```

7. Remove any default configuration for Receptor. Edit `/etc/receptor/receptor.conf` to contain the following configuration:

   ```
   ---
   - node:
       id: <IP address>
   ```

```
- log-level: debug

- tcp-listener:
    port: port_number

- control-service:
    service: control
    filename: /var/run/receptor/receptor.sock

- work-command:
    worktype: local
    command: /var/lib/ol-automation-manager/venv/awx/bin/ansible-runner
    params: worker
    allowruntimeparams: true
#    verifysignature: true
```

In the previous example, *hostname or IP address* is the IP address of the host and *port_number* is the port number that this node is listening on. For example, you can call the host single1-192.0.121.30 where you provide a node name and the IP address of the node. And you could configure the tcp-listener list listen on port 27199.

8. Start the service:

```
sudo systemctl enable --now ol-automation-manager.service
```

9. Run the following command to preload data, such as:
   - Demo Project
   - Default Galaxy Credentials
   - Demo Organization
   - Demo Inventory
   - Demo Job template
   - And so on

```
sudo su -l awx -s /bin/bash
awx-manage create_preload_data
```

10. Exit the awx shell environment.

```
exit
```

11. The host is now ready. Using a browser, you can now log in as the admin user.

```
https://<hostname or IP address>
```

5

Installing Oracle Linux Automation Manager in a Clustered Deployment

This chapter discusses how to prepare hosts in an Oracle Linux Automation Manager multi-host deployment. When you prepare the hosts, you must install the Oracle Linux Automation Manager software packages and configure them as part of the Oracle Linux Automation Manager service mesh. Configure and start the Service Mesh nodes before configuring and starting the control plane and execution plane nodes.

Configuring and Starting the Control Plane Service Mesh

You configure each node in the control plane of a cluster by editing the `/etc/receptor/receptor.conf` file. This file contains the following elements:

- **node ID**: The node ID must be the IP address or host name of the host.
- **log-level**: Available options are: **Error**, **Warning**, **Info** and **Debug**. Log level options provide increasing verbosity, such that Error generates the least information and Debug generates the most.
- **tcp-listener port**: This is the port that the node listens for incoming tcp peer connections configured on other nodes. For example, if the node ID represents a control node that listens on port 27199, then all other nodes that want to establish a connection to this control node would have to specify port 27199 in the tcp-peer element they configure in their `/etc/receptor/receptor.conf` file.
- **control-service**: All nodes in a cluster run the control service which reports status and launches and monitors work.
- **work-command**: This element defines the type of work that can be done on a node. For control plane nodes, the work type is always Local. The command it runs is the Ansible Runner tool which provides an abstraction layer for running Ansible and Ansible playbook tasks and can be configured to send status and event data to other systems. For more information about Ansible Runner, see https://ansible-runner.readthedocs.io/en/stable/.

On each host intended for use as a control plane node, do the following:

1. Remove any default configuration for Receptor and edit `/etc/receptor/receptor.conf` to contain the following configuration control plane specific information:

```
---
- node:
    id: <IP address or host name>

- log-level: info

- tcp-listener:
    port: <port_number>

- control-service:
    service: control
    filename: /var/run/receptor/receptor.sock
```

```
- work-command:
    worktype: local
    command: /var/lib/ol-automation-manager/venv/awx/bin/ansible-runner
    params: worker
    allowruntimeparams: true
    verifysignature: false
```

In the previous example, *IP address or hostname* is the IP address or hostname of the node and *port_number* is the port number that this node is listening on. For example, you can use something like control1-192.0.121.28 where you provide a node name and the IP address of the node. And you could configure the tcp-listener list listen on port 27199. The worktype parameter must be `local` in control plane nodes.

2. Start the Oracle Linux Automation Manager mesh service.

```
sudo systemctl start receptor-awx
```

3. Verify the Service Mesh. For more information, see Viewing Service Mesh Status for a Cluster Node.

> **Note:**
>
> At this point in the process, the peer relationships between service mesh nodes have not been established yet. Status information only exists for the individual servers running the Service Mesh.

Configuring and Starting the Execution Plane Service Mesh

You configure each node in the execution plane of a cluster by editing the `/etc/receptor/receptor.conf` file. This file contains the following elements:

- **node ID**: The node ID must be the IP address or hostname of the host.
- **log-level**: Available options are: **Error**, **Warning**, **Info** and **Debug**. Log level options provide increasing verbosity, such that Error generates the least information and Debug generates the most.
- **tcp-listener port**: This is the port that the node listens for incoming tcp peer connections configured on other nodes. For example, if the node ID represents an execution node that listens on port 27199, then all other nodes that want to establish a connection to this execution node would have to specify port 27199 in the tcp-peer element they configure in their `/etc/receptor/receptor.conf` file.
- **tcp-peer port**: This element must include the hostname and port number of the host it is connecting with. For example, if this execution node needs to connect to more than one control plane node to provide redundancy, you would need to add tcp-peer elements for each control plane node that the execution node connects with. In the address field, enter the host name or IP address of the control plane node, followed by the port number. The redial element, if enabled, attempts to periodically reestablish a connection to the host if connectivity fails.
  You can also configure tcp-peer elements to include the hostnames and port numbers of other execution nodes or hop nodes based on your service mesh topology requirements.
- **control-service**: All nodes in a cluster run the control service which reports status and launches and monitors work.

- **work-command**: This element defines the type of work that can be done on a node. For execution plane nodes, the work type is always `ansible-runner`. The command it runs is the Ansible Runner tool which provides an abstraction layer for running Ansible and Ansible playbook tasks and can be configured to send status and event data to other systems. For more information about Ansible Runner, see https://ansible-runner.readthedocs.io/en/stable/.

On each host intended for use as an execution plane node, do the following:

1. Remove any default configuration for Receptor and edit `/etc/receptor/receptor.conf` to contain the following configuration execution plane specific information:

```
---
- node:
    id: <IP address or hostname>

- log-level: debug

- tcp-listener:
    port: <port_number>

- tcp-peer:
    address: <hostname or IP address>:<target_port_number>
    redial: true

- tcp-peer:
    address: <hostname or IP address>:<target_port_number>
    redial: true

- control-service:
    service: control
    filename: /var/run/receptor/receptor.sock

- work-command:
    worktype: ansible-runner
    command: /var/lib/ol-automation-manager/venv/awx/bin/ansible-runner
    params: worker
    allowruntimeparams: true
    verifysignature: false
```

   In the previous example,

   - *IP address or hostname* is the IP address or hostname of the node.
   - *port_number* is the port number that this node is listening on.
   - *target_port* is the port number of the peer node that you are configuring this node to connect with.
   - *hostname or IP address* is the hostname or IP address of the execution, control, or hop node being connected with.
   - The worktype parameter must be `ansible-runner` in execution plane nodes.

   If the execution environment is associated with more than one control, execution, or hop node, enter additional `- tcp-peer:` nodes for instances that the execution host is associated with.

2. Start the Oracle Linux Automation Manager mesh service.

```
sudo systemctl start receptor-awx
```

3. Verify the Service Mesh. For more information, see Viewing Service Mesh Status for a Cluster Node.

> **Note:**
>
> At this point in the process, the peer relationships between service mesh nodes have not been established yet. Status information only exists for the individual servers running the Service Mesh.

Configuring and Starting the Hop Nodes

You configure each hop node in the cluster by editing the `/etc/receptor/receptor.conf` file. This file contains the following elements:

- **node ID**: The node ID must be the IP address or hostname of the host.
- **log-level**: Available options are: **Error**, **Warning**, **Info** and **Debug**. Log level options provide increasing verbosity, such that Error generates the least information and Debug generates the most.
- **tcp-listener port**: This is the port that the node listens for incoming tcp peer connections configured on other nodes. For example, if the node ID represents an execution node that listens on port 27199, then all other nodes that want to establish a connection to this execution node would have to specify port 27199 in the tcp-peer element they configure in their `/etc/receptor/receptor.conf` file.
- **tcp-peer port**: This element must include the hostname and port number of the host it is connecting with. For example, you might configure your hop node to connect to a control node as the intermediate node between the control node and an execution node. In the address field, enter the host name or IP address of the control plane node, followed by the port number. The redial element, if enabled, attempts to periodically reestablish a connection to the host if connectivity fails.
- **control-service**: All nodes in a cluster run the control service which reports status and launches and monitors work.
- **work-command**: This element defines the type of work that can be done on a node. Hop nodes do not run playbooks. However, you must configure the default fields. The work type for hop nodes is always `ansible-runner`.

On each host intended for use as a hop node, do the following:

1. Remove any default configuration for Receptor and edit `/etc/receptor/receptor.conf` to contain the following configuration with hop node specific information:

```
---
- node:
    id: <node IP address or hostname>

- log-level: debug

- tcp-listener:
    port: <port_number>

- tcp-peer:
    address: <control hostname or IP address>:<target_port_number>
    redial: true

- tcp-peer:
```

```
    address: <control hostname or IP address>:<target_port_number>
    redial: true

- control-service:
    service: control
    filename: /var/run/receptor/receptor.sock

- work-command:
    worktype: local
    command: /var/lib/ol-automation-manager/venv/awx/bin/ansible-runner
    params: worker
    allowruntimeparams: true
    verifysignature: false
```

In the previous example,

- *node IP address or hostname* is the IP address or hostname of the node.
- *port_number* is the port number that this node is listening on.
- *target_port* is the port number of the peer node that you are configuring this node to connect with.
- *control hostname or IP address* is the hostname or IP address of the control nodes that the hop node is connecting with.

If the hop node is associated to more than one control node, enter additional `- tcp-peer:` nodes for each instance that the hop node is associated with.

2. Start the Oracle Linux Automation Manager mesh service.

```
sudo systemctl start receptor-awx
```

3. Verify the Service Mesh. For more information, see Viewing Service Mesh Status for a Cluster Node.

> **Note:**
>
> At this point in the process, the peer relationships between service mesh nodes have not been established yet. Status information only exists for the individual servers running the Service Mesh.

Configuring the Control, Execution, and Hop Nodes

To configure the control plane, execution plane, and hop nodes, on one control plane host do the following steps which applies to all Oracle Linux Automation Manager instances:

1. Run the following commands:

```
sudo su -l awx -s /bin/bash
```

2. Do the following:
   - Repeat the following command for each host you want to designate as control node type, changing the IP address or host name each time you run the command:

     ```
     awx-manage provision_instance --hostname=<control hostname or IP address> --
     node_type=control
     ```

In the previous example, *control hostname or IP address* is the hostname or IP address of the system running Oracle Linux Automation Manager. Your choice of host name or IP address must match with the host name or IP addressed used when you configured the `/etc/receptor/receptor.conf` file node ID (see Configuring and Starting the Control Plane Service Mesh). If hostname is used, the host must be resolvable.

- Repeat the following command for each host you want to designate as execution node type, changing the IP address or host name each time you run the command:

  ```
  awx-manage provision_instance --hostname=<execution hostname or IP address> --
  node_type=execution
  ```

  In the previous example, *execution hostname or IP address* is the hostname or IP address of the system running Oracle Linux Automation Manager. Your choice of host name or IP address must match with the host name or IP addressed used when you configured the `/etc/receptor/receptor.conf` file node ID (see Configuring and Starting the Execution Plane Service Mesh). If hostname is used, the host must be resolvable.

- Repeat the following command for each host you want to designate as the hop node type, changing the IP address or host name each time you run the command:

  ```
  awx-manage provision_instance --hostname=<hop hostname or IP address> --
  node_type=hop
  ```

  In the previous example, *hop hostname or IP address* is the hostname or IP address of the system running Oracle Linux Automation Manager. Your choice of host name or IP address must match with the host name or IP addressed used when you configured the `/etc/receptor/receptor.conf` file node ID (see Configuring and Starting the Hop Nodes). If hostname is used, the host must be resolvable.

3. Run the following command to register the default execution environments, which are:
   - Control Plane Execution Environment
   - OLAM EE: (Latest)

   ```
   awx-manage register_default_execution_environments
   ```

4. Run the following command to create the controlplane instance groups (specified as a queue in the command) and associate it to a control plane host. Repeat the command with the same queue name for each control plane host in your cluster:

   ```
   awx-manage register_queue --queuename=controlplane --hostnames=<control hostname or
   IP address>
   ```

5. Run the following command to create instance groups and associate it to an execution plane host. Repeat the command with the same queue name for each execution plane host in your cluster:

   ```
   awx-manage register_queue --queuename=execution --hostnames=<execution hostname or
   IP address>
   ```

6. Run the `awx-manage list_instances` command to ensure each host you registered are available under the correct instance group. For example, the following shows the IP addresses of two control plane and three execution plane nodes running under the controlplane and execution instance groups. The nodes are currently not running, and therefore do not show available capacity or heartbeat information.

   ```
   awx-manage list_instances
   [controlplane capacity=0]
   ```

```
    192.0.119.192 capacity=0 node_type=control version=?
    192.0.124.44 capacity=0 node_type=control version=?

[execution capacity=0]
    192.0.114.137 capacity=0 node_type=execution version=ansible-runner-???
    192.0.117.98 capacity=0 node_type=execution version=ansible-runner-???
    192.0.125.241 capacity=0 node_type=execution version=ansible-runner-???
```

> **Note:**
>
> Hop nodes do not appear in this list because they are not associated to any instance group.

7. Run the following command to register the Oracle Linux Automation Manager service mesh peer relationship between each node in the cluster:

   ```
   awx-manage register_peers <execution or hop hostname or IP address> --
   peers <execution, hop, or control hostname or IP address>
   ```

   This command must be run for each pair of nodes to requiring a peer relationship. For example, the peer relationships being established in the example described in Service Mesh Topology Examples shows the command being run twice for each execution node so that each execution node is connected to a different control node. This ensures that each execution node always has a backup control node if one of the control nodes were to fail.

   Additional topologies are possible, such as those where an isolated execution node must peer to a hop node, and the hop node must peer to a control node. In this case the command must be run one time to peer the execution node with the hop node, and again to peer the hop node with the control node.

8. Exit the awx shell environment.

   ```
   exit
   ```

9. For each control and execution plane host, edit the `/etc/tower/settings.py` file with the following:

   ```
   DEFAULT_EXECUTION_QUEUE_NAME = 'execution'
   DEFAULT_CONTROL_PLANE_QUEUE_NAME = 'controlplane'
   ```

Starting the Control, Execution, and Hop Nodes

To start the control, execution, and hop nodes, do the following:

1. Start the service on each node:

   ```
   sudo systemctl enable --now ol-automation-manager.service
   ```

2. On one control plane node, run the following command to preload data, such as:
   - Demo Project
   - Default Galaxy Credentials
   - Demo Organization
   - Demo Inventory

- Demo Job template
- And so on

```
sudo su -l awx -s /bin/bash
awx-manage create_preload_data
```

> **Note:**
>
> This command only needs to be run one time because the preloaded data persists in the database that all cluster nodes connect with.

3. Run the `awx-manage list_instances` command to ensure that the control and execution plane nodes are now running and show available capacity and display heartbeat information. For example, the following shows all control and execution plane instances running, with available capacity, and active heartbeat information.

```
awx-manage list_instances
[controlplane capacity=270]
    192.0.119.192 capacity=135 node_type=control version=19.5.1
heartbeat="2022-09-22 14:38:29"
    192.0.124.44 capacity=135 node_type=control version=19.5.1
heartbeat="2022-09-22 14:39:09"

[execution capacity=405]
    192.0.114.137 capacity=135 node_type=execution version=19.5.1
heartbeat="2022-09-22 14:40:07"
    192.0.117.98 capacity=135 node_type=execution version=19.5.1
heartbeat="2022-09-22 14:40:35"
    192.0.125.241 capacity=135 node_type=execution version=19.5.1
heartbeat="2022-09-22 14:40:55"
```

> **Note:**
>
> Hop nodes do not appear in this list because they are not associated to any instance group.

4. Exit the awx shell environment.

```
exit
```

Configuring TLS Verification and Signed Work Requests

Oracle recommends that you secure your Service Mesh communication within your cluster with TLS verification and signed work requests sent between cluster nodes. TLS verification ensures secure communication in the Service Mesh network and signed work requests ensure secure job execution.

The following procedure enables TLS for an existing Oracle Linux Automation Manager cluster. Complete the following tasks before doing this procedure:

- Setting up Hosts
- Configuring and Starting the Control Plane Service Mesh

- Configuring and Starting the Execution Plane Service Mesh
- Configuring and Starting the Hop Nodes
- Configuring the Control, Execution, and Hop Nodes
- Starting the Control, Execution, and Hop Nodes

To configure TLS verification and signed work requests, do the following:

1. On each host in the cluster (each execution, hop, and control plane nodes), to enable signed work requests, add the following text to the `/etc/tower/settings.py` file.

   ```
   RECEPTOR_NO_SIG = False
   ```

2. From one of your control nodes, in the `/etc/tower` folder, do the following:

   - If you are using IP addresses for the `node_id` field, run the following commands to create the certs folder and generate TLS certificates:

     ```
     sudo mkdir -p certs
     sudo receptor --cert-init commonname="test CA" bits=2048 outcert=certs/
     ca.crt outkey=certs/ca.key
     node=<node_id>; sudo receptor --cert-makereq bits=2048
     commonname="$node test cert"  ipaddress=$node nodeid=$node
     outreq=certs/$node.csr outkey=certs/$node.key
     node=<node_id>; sudo receptor --cert-signreq req=certs/$node.csr
     cacert=certs/ca.crt cakey=certs/ca.key outcert=certs/$node.crt
     ```

     In the previous example, *node_id* is the IP address of the node you are creating keys for that you set in the `/etc/receptor/receptor.conf` file for the execution, hop, or control plane nodes.

   - If you are using a host name for the `node_id` field, run the following commands to create the certs folder and generate TLS certificates:

     ```
     sudo mkdir -p certs
     sudo receptor --cert-init commonname="test CA" bits=2048 outcert=certs/
     ca.crt outkey=certs/ca.key
     node=<node_id>; sudo receptor --cert-makereq bits=2048
     commonname="$node test cert"  dnsname=$node nodeid=$node
     outreq=certs/$node.csr outkey=certs/$node.key
     node=<node_id>; sudo receptor --cert-signreq req=certs/$node.csr
     cacert=certs/ca.crt cakey=certs/ca.key outcert=certs/$node.crt
     ```

     In the previous example, *node_id* is the host name of the node you are creating the keys for that you set in the `/etc/receptor/receptor.conf` file for the execution, hop, or control plane nodes.

3. After the second command, type `yes` to confirm that you want to sign the certificate. For example, the following generates certificates for a cluster with two hosts:

   ```
   node=192.0.250.40; sudo receptor --cert-makereq bits=2048
   commonname="$node test cert"  ipaddress=192.0.250.40 nodeid=$node
   outreq=certs/$node.csr outkey=certs/$node.key
   node=192.0.250.40; sudo receptor --cert-signreq req=certs/$node.csr
   cacert=certs/ca.crt cakey=certs/ca.key outcert=certs/$node.crt
   Requested certificate:
   ```

```
  Subject: CN=192.0.250.40 test cert
  Encryption Algorithm: RSA (2048 bits)
  Signature Algorithm: SHA256-RSA
  Names:
    IP Address: 192.0.250.40
    Receptor Node ID: 192.0.250.40
Sign certificate (yes/no)? yes

node=192.0.251.206; sudo receptor --cert-makereq bits=2048
commonname="$node test cert" ipaddress=192.0.251.206 nodeid=$node
outreq=certs/$node.csr outkey=certs/$node.key
node=192.0.251.206; sudo receptor --cert-signreq req=certs/$node.csr
cacert=certs/ca.crt cakey=certs/ca.key outcert=certs/$node.crt
Requested certificate:
  Subject: CN=192.0.251.206 test cert
  Encryption Algorithm: RSA (2048 bits)
  Signature Algorithm: SHA256-RSA
  Names:
    IP Address: 192.0.251.206
    Receptor Node ID: 192.0.251.206
Sign certificate (yes/no)? yes
```

4. From the `/etc/tower/certs` folder, run the following commands to generate certificates for work request signing and verification:

```
sudo openssl genrsa -out signworkprivate.pem 2048
sudo openssl rsa -in signworkprivate.pem -pubout -out signworkpublic.pem
```

5. From the cd `/etc/tower/` folder, run the following command to change the certs folder ownership and all files within the folder:

```
sudo chown -R awx:awx certs
```

6. Check that you have all the files you need in the `/etc/tower/certs` folder. For example, the following shows the generated key information for a four node cluster.

```
ls -al
total 68
drwxr-xr-x. 2 awx awx 4096 Sep 12 18:26 .
drwxr-xr-x. 4 awx awx  132 Sep 12 16:49 ..
-rw-------. 1 awx awx 1180 Sep 12 18:19 192.0.113.178.crt
-rw-------. 1 awx awx 1001 Sep 12 18:19 192.0.113.178.csr
-rw-------. 1 awx awx 1679 Sep 12 18:19 192.0.113.178.key
-rw-------. 1 awx awx 1176 Sep 12 18:20 192.0.121.28.crt
-rw-------. 1 awx awx 1001 Sep 12 18:20 192.0.121.28.csr
-rw-------. 1 awx awx 1675 Sep 12 18:20 192.0.121.28.key
-rw-------. 1 awx awx 1180 Sep 12 18:20 192.0.126.172.crt
-rw-------. 1 awx awx 1001 Sep 12 18:19 192.0.126.172.csr
-rw-------. 1 awx awx 1679 Sep 12 18:19 192.0.126.172.key
-rw-------. 1 awx awx 1176 Sep 12 18:19 192.0.127.70.crt
-rw-------. 1 awx awx 1001 Sep 12 18:19 192.0.127.70.csr
-rw-------. 1 awx awx 1675 Sep 12 18:19 192.0.127.70.key
-rw-------. 1 awx awx 1107 Sep 12 16:54 ca.crt
-rw-------. 1 awx awx 1679 Sep 12 16:54 ca.key
```

```
-rw-------. 1 awx awx 1675 Sep 12 18:26 signworkprivate.pem
-rw-r--r--. 1 awx awx  451 Sep 12 18:26 signworkpublic.pem
```

7. On each node in the cluster, in the `/etc/tower` folder, create a certs folder and change the ownership and group of the certs folder to `awx:awx`:

```
sudo mkdir -p certs
sudo chown -R awx:awx certs
```

8. Copy over the ca.crt, node specific .crt, csr, and key files, and the signworkprivate.pem, and signworkpublic.pem files to each node in the cluster.
9. For each control plane node, add the following lines to the `/etc/receptor/receptor.conf` file:

```
---
- node:
    id: <IP address or host name>

- log-level: debug

# Add the tls: control that specifies the tls-server name for the listener
- tcp-listener:
    port: 27199
    tls: controller

# Add the TLS server configuration
- tls-server:
    name: controller
    cert: /etc/tower/certs/<IP address or host name>.crt
    key: /etc/tower/certs/<IP address or host name.key
    requireclientcert: true
    clientcas: /etc/tower/certs/ca.crt

- control-service:
    service: control
    filename: /var/run/receptor/receptor.sock

# Add the work-signing and work-verification elements
- work-signing:
    privatekey: /etc/tower/certs/signworkprivate.pem
    tokenexpiration: 30m

- work-verification:
    publickey: /etc/tower/certs/signworkpublic.pem

# Set verifysignature to true.
- work-command:
    worktype: local
    command: /var/lib/ol-automation-manager/venv/awx/bin/ansible-runner
    params: worker
    allowruntimeparams: true
    verifysignature: true
```

In the previous example, *IP address or host name* is the host name or IP address of the control plane host. If host name is used, the host must be resolvable.

10. For each execution plane node, add the following lines to the `/etc/receptor/receptor.conf` file:

```
---
- node:
    id: <execution IP address or host name>

- log-level: debug

- tcp-listener:
    port: 27199

# Add tls: client that specifies the tls-client name.
- tcp-peer:
    address: <hostname or IP address>:27199
    redial: true
    tls: client

- tcp-peer:
    address: <hostname or IP address>:27199
    redial: true
    tls: client

# Add the tls-client element.
- tls-client:
    name: client
    rootcas: /etc/tower/certs/ca.crt
    insecureskipverify: false
    cert: /etc/tower/certs/<execution IP address or host name>.crt
    key: /etc/tower/certs/<execution IP address or host name.key

- control-service:
    service: control
    filename: /var/run/receptor/receptor.sock

# Add the work-verification element.
- work-verification:
    publickey: /etc/tower/certs/signworkpublic.pem

# Set verifysignature to true.
- work-command:
    worktype: ansible-runner
    command: /var/lib/ol-automation-manager/venv/awx/bin/ansible-runner
    params: worker
    allowruntimeparams: true
    verifysignature: true
```

In the previous example,

- *execution IP address or host name* is the IP address or host name of the node
- *hostname or IP address* is the host name or IP address of the execution, control, or hop node you are peering with.

11. (If required) For each hop node, add the following lines to the `/etc/receptor/receptor.conf` file:

```
---
- node:
    id: <node IP address or hostname>

- log-level: debug

# Add the tls: control that specifies the tls-server name for the listener
- tcp-listener:
    port: 27199
    tls: controller

# Add tls: client that specifies the tls-client name.
- tcp-peer:
    address: <control hostname or IP address>:27199
    redial: true
    tls: client

# Add the tls-client element.
- tls-client:
    name: client
    rootcas: /etc/tower/certs/ca.crt
    insecureskipverify: false
    cert: /etc/tower/certs/<node IP address or hostname>.crt
    key: /etc/tower/certs/<node IP address or hostname>.key

- work-verification:
    publickey: /etc/tower/certs/signworkpublic.pem

# Add the work-signing and work-verification elements
- work-signing:
    privatekey: /etc/tower/certs/signworkprivate.pem
    tokenexpiration: 30m

# Add the TLS server configuration
- tls-server:
    name: controller
    cert: /etc/tower/certs/<node IP address or hostname>.crt
    key: /etc/tower/certs/<node IP address or hostname>.key
    requireclientcert: true
    clientcas: /etc/tower/certs/ca.crt

- control-service:
    service: control
    filename: /var/run/receptor/receptor.sock


# Set verifysignature to true.
- work-command:
    worktype: local
    command: /var/lib/ol-automation-manager/venv/awx/bin/ansible-runner
    params: worker
```

```
allowruntimeparams: true
verifysignature: true
```

In the previous example,

- *node IP address or host name* is the IP address or host name of the node
- *control hostname or IP address* is the host name or IP address of the control plane host you are peering with.

12. On each node, restart the Service Mesh and Oracle Linux Automation Manager.

```
sudo systemctl restart receptor-awx
sudo systemctl restart ol-automation-manager
```

13. Verify the Service Mesh. See Viewing the Service Mesh for more information.

6

Adding or Removing Nodes to an Existing Cluster

This chapter provides instructions for adding or removing nodes to and from an existing cluster.

Adding a New Control Plane Node to a Cluster

To add new control node to a cluster, do the following:

1. Prepare the new hosts, as described in Setting Up the Network and Enabling Access to the Oracle Linux Automation Manager Packages.
2. Configure the host, following the instructions in Setting up Hosts. Do not run the `awx-manage migrate` or `awx-manage createsuperuser`. These only need to be run when initially creating the cluster.
3. Set up the service mesh for the control plane node, by following the instructions in Configuring and Starting the Control Plane Service Mesh.
4. Set up the service mesh for the execution plane nodes you want to connect to your new control plane node, by following the instructions in Configuring and Starting the Execution Plane Service Mesh.
5. Set up the hop nodes you want to connect to your new control plane node, by following the instructions in Configuring and Starting the Hop Nodes .
6. Provision the node as the control node type, register the node to an appropriate instance group (called a queuename in the command), and establish the peer relationships between the execution, hop, and the control nodes as described in Configuring the Control, Execution, and Hop Nodes.
7. Start the control plane node as described in Starting the Control, Execution, and Hop Nodes. Do not run the command to create preloaded data.
8. If required, apply TLS verification and signed work requests as described in Configuring TLS Verification and Signed Work Requests.

Adding a New Execution Plane Node to a Cluster

To add a new execution node to a cluster, do the following:

1. Prepare the new hosts, as described in Setting Up the Network and Enabling Access to the Oracle Linux Automation Manager Packages.
2. Configure the host, following the instructions in Setting up Hosts. Do not run the `awx-manage migrate` or `awx-manage createsuperuser`. These only need to be run when initially creating the cluster.
3. Set up the service mesh for the execution plane node, by following the instructions in Configuring and Starting the Execution Plane Service Mesh.
4. Provision the node as the execution node type, register the node to an appropriate instance group (called a queuename in the command), and establish the peer relationships

between the execution node and the control plane nodes or between the execution node and the hop nodes as described in Configuring the Control, Execution, and Hop Nodes.

5. Start the execution plane node as described in Starting the Control, Execution, and Hop Nodes. Do not run the command to create preloaded data.
6. If required, apply TLS verification and signed work requests as described in Configuring TLS Verification and Signed Work Requests.

Adding a New Hop Node to a Cluster

To add new hop node to a cluster, do the following:

1. Prepare the new hosts, as described in Setting Up the Network and Enabling Access to the Oracle Linux Automation Manager Packages.
2. Configure the host, following the instructions in Setting up Hosts. Do not run the `awx-manage migrate` or `awx-manage createsuperuser`. These only need to be run when initially creating the cluster.
3. Set up the hop nodes you want to connect to your control plane nodes, by following the instructions in Configuring and Starting the Hop Nodes .
4. Set up the execution nodes you want to connect to your new hop node, by following the instructions in Configuring and Starting the Execution Plane Service Mesh.
5. Provision the node as the hop node type, and for any new execution nodes, register the execution node to the `execution` instance group (called a queuename in the command), and establish the peer relationships between the execution, hop, and the control nodes as described in Configuring the Control, Execution, and Hop Nodes.
6. Start the hop node and execution nodes as described in Starting the Control, Execution, and Hop Nodes. Do not run the command to create preloaded data.
7. If required, apply TLS verification and signed work requests as described in Configuring TLS Verification and Signed Work Requests.

Removing a Node from a Cluster

To remove a node from a cluster, do the following:

1. Log on the node you want to remove.
2. Stop Oracle Linux Automation Manager on the node.

   ```
   sudo systemctl stop ol-automation-manager.service
   ```

3. Stop the service mesh.

   ```
   sudo systemctl stop receptor-awx
   ```

4. Delete the `/etc/tower/SECRET_KEY` file.
5. Open the `/etc/tower/settings.py` file and remove the database password from DATABASES node or remove any configuration that provides a password for your database, if you are using alternative approaches.

6. From any control plane node, verify that the node you want to remove no longer shows capacity or heartbeat information. For example, the following shows the node with IP address 192.0.124.44 has zero capacity and no heartbeat information.

```
sudo su -l awx -s /bin/bash
awx-manage list_instances
[controlplane capacity=126]
        192.0.119.192 capacity=126 node_type=control version=19.5.1
heartbeat="2022-10-20 06:55:44"
        192.0.124.44 capacity=0 node_type=control version=19.5.1

[execution capacity=126]
        192.0.114.137 capacity=126 node_type=execution version=19.5.1
heartbeat="2022-10-20 06:56:20"
```

7. Deprovision the instance from the cluster.

```
awx-manage deprovision_instance --hostname=<IP address or host name>
```

   In the previous example, <*IP address or host name*> is the host you want to remove from the cluster.

8. Check the status of the remaining control and execution plane nodes to verify that the deprovisioned instance no longer appears. For example, the deprovisioned node with IP address 192.0.124.44 from the previous example no longer appears:

```
awx-manage list_instances
[controlplane capacity=126]
        192.0.119.192 capacity=126 node_type=control version=19.5.1
heartbeat="2022-10-20 06:55:44"

[execution capacity=126]
        192.0.114.137 capacity=126 node_type=execution version=19.5.1
heartbeat="2022-10-20 06:56:20"
```

9. Exit the awx shell environment.

```
exit
```

10. If required, remove any `tcp-peer` nodes pointing to the deprovisioning node in the `/etc/receptor/receptor.conf` files of the remaining cluster nodes, the restart the nodes.

```
sudo systemctl restart receptor-awx
```

7

Viewing the Service Mesh

This chapter describes methods to view service mesh information.

Viewing Service Mesh Status for a Cluster Node

This section provides instructions for obtaining Service Mesh status information about a node in an Oracle Linux Automation Manager cluster, such as:

- **Node ID**: The node ID must be the IP address of the host.
- **System Information**: Such as CPU count and system memory.
- **Connections**: A list of IP address or host names that the node is connected with and the number of hops required to reach them, listed as the cost. Cost is defined by each customer.
- **Known Node** and **Known Node Connections**: A list of all known nodes in the cluster and the further connections known to each node listed.
- **Route**: This parameter list the route by which a node connects to another node. If the node is the same, then node is directly connected. If the nodes are different, then there is one or more hop or execution plane nodes between the nodes.
- **Node Service**: The Control Service run on every node in the cluster. It reports node status and monitors work being performed on the node.
- **Node Work Types**: The work types are Local for control plane nodes and ansible-runner for execution plane nodes.

To view service mesh status, from any host in the cluster, do the following:

1. Run the following command to obtain status information about the service mesh:

   ```
   sudo receptorctl  --socket /var/run/receptor/receptor.sock status
   ```

   For example, the following command shows the status for a four host cluster where peer relationships have been established:

   ```
   sudo receptorctl  --socket /var/run/receptor/receptor.sock status
   Node ID: 192.0.121.28
   Version: +g
   System CPU Count: 4
   System Memory MiB: 15583

   Connection         Cost
   192.0.113.178    1
   192.0.127.70     1

   Known Node         Known Connections
   192.0.113.178    {'192.0.121.28': 1, '192.0.126.172': 1}
   192.0.121.28     {'192.0.113.178': 1, '192.0.127.70': 1}
   192.0.126.172    {'192.0.113.178': 1, '192.0.127.70': 1}
   ```

```
192.0.127.70     {'192.0.121.28': 1, '192.0.126.172': 1}

Route              Via
192.0.113.178    192.0.113.178
192.0.126.172    192.0.113.178
192.0.127.70     192.0.127.70

Node               Service   Type       Last Seen             Tags
192.0.113.178    control   Stream     2022-09-02 18:06:14   {'type':
'Control Service'}
192.0.121.28     control   Stream     2022-09-02 18:06:33   {'type':
'Control Service'}
192.0.126.172    control   Stream     2022-09-02 18:06:20   {'type':
'Control Service'}
192.0.127.70     control   Stream     2022-09-02 18:06:25   {'type':
'Control Service'}

Node               Work Types
192.0.113.178    ansible-runner
192.0.121.28     local
192.0.126.172    local
192.0.127.70     ansible-runner
```

Viewing Service Mesh Cluster Status

Using the api/v2/mesh_visualizer/ API call, you can view status information about each node in your service mesh cluster and details about available links setup between each node from the perspective of Oracle Linux Automation Manager.

To get cluster node and link details, do the following:

1. Log in to your Oracle Linux Automation Manager server with a user account.

   ```
   https://<hostname or IP address>/api/login/
   ```

   > **Note:**
   >
   > In the previous example, <*hostname or ip address*> is the hostname or IP address of the system running Oracle Linux Automation Manager . If hostname is used, the host must be resolvable.

2. In the response area, click one of the /api/v2 links. This performs a get request that lists all available resources.
3. Click the `/api/v2/mesh_visualizer/` link.
   The Mesh Visualizer get response appears. For example:

   ```
   HTTP 200 OK
   Allow: GET, HEAD, OPTIONS
   Content-Type: application/json
   Vary: Accept
   X-API-Node: 192.0.121.28
   X-API-Product-Name: AWX
   X-API-Product-Version: 19.5.1
   ```

```
X-API-Time: 0.019s

{
    "nodes": [
        {
            "id": 1,
            "hostname": "192.0.121.28",
            "node_type": "control",
            "node_state": "healthy"
        },
        {
            "id": 2,
            "hostname": "192.0.127.70",
            "node_type": "execution",
            "node_state": "healthy"
        },
        {
            "id": 3,
            "hostname": "192.0.126.172",
            "node_type": "hop",
            "node_state": "healthy"
        }
    ],
    "links": [
        {
            "source": "192.0.127.70",
            "target": "192.0.121.28"
        },
        {
            "source": "192.0.126.172",
            "target": "192.0.121.28"
        },
        {
            "source": "192.0.127.70",
            "target": "192.0.126.172"
        }
    ]
}
```

8

Installing Oracle Linux Automation Manager CLI

You can install the Oracle Linux Automation Manager CLI on the same machine you installed the Oracle Linux Automation Manager server or on another Oracle Linux 8 machine. For information about installing the Oracle Linux Automation Manager CLI, see Oracle Linux Automation Manager 2.1: CLI and API Reference Guide.

9
Upgrading Oracle Linux Automation Manager

The following chapter provides instructions for upgrading Oracle Linux Automation Manager.

Upgrading a Release 1.0.X to a Release 2.0 Single Host Deployment

To upgrade a single host instance of Oracle Linux Automation Manager release 1.0.x to a single host instance of Oracle Linux Automation Manager Release 2.0, do the following:

1. Log in to a terminal for the Oracle Linux Automation Manager Release 1.0.x version you want to upgrade.
2. Back up the `/etc/tower/SECRET_KEY` file to a secure location. For example, you can copy the file to your home directory:

   ```
   sudo cp /etc/tower/SECRET_KEY ~
   ```

3. Stop Oracle Linux Automation Manager.

   ```
   sudo systemctl stop ol-automation-manager
   ```

4. Log in to the user account that controls the database.

   ```
   sudo su - postgres
   ```

5. Export the database using the following command that creates a script file containing all the necessary SQL commands and input data to restore the databases. For example, this command creates the `olam1.dump` file in your database home directory.

   ```
   pg_dumpall > olamv1.dump
   ```

6. Exit the user account that controls the database.

   ```
   exit
   ```

7. Stop the database server.

   ```
   sudo systemctl stop postgresql
   ```

8. Remove (and optionally backup) existing database data directory. For example, the following command removes and creates a backup file in the home directory.

   ```
   sudo mv /var/lib/pgsql/data/ ~/data.old
   ```

9. Remove the current version of the database.

   ```
   sudo dnf remove postgresql
   ```

10. Enable the `postgresql 12` or `postgresql 13` module stream.

```
sudo dnf module reset postgresql
sudo dnf module enable postgresql:12
```

or

```
sudo dnf module reset postgresql
sudo dnf module enable postgresql:13
```

> **Note:**
>
> For more information about the Postgresql 12 and 13 life cycle, see the appendix discussing the application life cycle for stream modules in Oracle Linux: Managing Software on Oracle Linux.

11. Enable the Oracle Linux Automation Manager Yum repos for Release 2 as described in Enabling Access to the Oracle Linux Automation Manager Packages.
12. Update Oracle Linux Automation Manager.

```
sudo dnf update ol-automation-manager
```

> **Caution:**
>
> If you have installed the `ansible` package from another repository (for example, EPEL) the installation and upgrade process overwrites this package with the `ansible-core` package.

> **Note:**
>
> The following message that is observed during the upgrade process can safely be ignored as it does not indicate any failure:
>
> ```
> ValueError: File context for /var/run/tower(/.*)? already defined
> ```

13. Install the database.

```
sudo dnf install postgresql-server
```

14. After the update completes, set up the database.

```
sudo postgresql-setup --initdb
sudo systemctl start postgresql
sudo su - postgres
psql -d postgres -f olamv1.dump
exit
```

15. Run the following command to see if the database is available:

```
sudo su - postgres -c "psql -l |grep awx"
```

The output should resemble something like this:

```
awx       | awx      | UTF8     | en_US.UTF-8 | en_US.UTF-8 |
```

16. Replace `/etc/tower/settings.py` with `/etc/tower/settings.py.rpmnew`. For example:

```
sudo mv /etc/tower/settings.py /etc/tower/settingsold.py
sudo mv /etc/tower/settings.py.rpmnew /etc/tower/settings.py
```

17. In the `/etc/tower/settings.py` file, set the `CLUSTER_HOST_ID` as follows:

```
CLUSTER_HOST_ID = "hostname or ip address"
```

In the previous example, *hostname or ip address* is the hostname or IP address of the system running Oracle Linux Automation Manager. If hostname is used, the host must be resolvable.

18. In the `/etc/nginx/nginx.conf`, remove all existing configuration and replace it with the following text:

```
user nginx;
worker_processes auto;
error_log /var/log/nginx/error.log;
pid /run/nginx.pid;

# Load dynamic modules. See /usr/share/doc/nginx/README.dynamic.
include /usr/share/nginx/modules/*.conf;

events {
 worker_connections 1024;
}

http {
 log_format  main  '$remote_addr - $remote_user [$time_local] "$request" '
                   '$status $body_bytes_sent "$http_referer" '
                   '"$http_user_agent" "$http_x_forwarded_for"';

 access_log  /var/log/nginx/access.log  main;

 sendfile            on;
 tcp_nopush          on;
 tcp_nodelay         on;
 keepalive_timeout   65;
 types_hash_max_size 2048;

 include             /etc/nginx/mime.types;
 default_type        application/octet-stream;

 # Load modular configuration files from the /etc/nginx/conf.d directory.
 # See http://nginx.org/en/docs/ngx_core_module.html#include
 # for more information.
```

```
 include /etc/nginx/conf.d/*.conf;
}
```

19. In the `/etc/nginx/conf.d/ol-automation-manager-nginx.conf`, remove all existing configuration and replace it with the following text:

```
upstream uwsgi {
    server unix:/var/run/tower/uwsgi.sock;
}

upstream daphne {
    server unix:/var/run/tower/daphne.sock;
}

server {
    listen 443 default_server ssl;
    listen 127.0.0.1:80 default_server;
    listen [::]:443 default_server ssl;
    listen [::1]:80 default_server;

    # If you have a domain name, this is where to add it
    server_name _;
    keepalive_timeout 65;

    ssl_certificate /etc/tower/tower.crt;
    ssl_certificate_key /etc/tower/tower.key;

    ssl_session_timeout 1d;
    ssl_session_cache shared:SSL:50m;
    ssl_session_tickets off;

    # intermediate configuration. tweak to your needs.
    ssl_protocols TLSv1.2;
    ssl_ciphers 'ECDHE-ECDSA-AES256-GCM-SHA384:ECDHE-RSA-AES256-GCM-
SHA384:ECDHE-ECDSA-CHACHA20-POLY1305:ECDHE-RSA-CHACHA20-POLY1305:ECDHE-
ECDSA-AES128-GCM-SHA256:ECDHE-RSA-AES128-GCM-SHA256:ECDHE-ECDSA-AES256-
SHA384:ECDHE-RSA-AES256-SHA384:ECDHE-ECDSA-AES128-SHA256:ECDHE-RSA-AES128-
SHA256';
    ssl_prefer_server_ciphers on;

    # HSTS (ngx_http_headers_module is required) (15768000 seconds = 6
months)
    add_header Strict-Transport-Security max-age=15768000;
    # add_header Content-Security-Policy "default-src 'self'; connect-src
'self' ws: wss:; style-src 'self' 'unsafe-inline'; script-src 'self'
'unsafe-inline' *.pendo.io; img-src 'self' *.pendo.io data:; report-uri /
csp-violation/";
    # add_header X-Content-Security-Policy "default-src 'self'; connect-
src 'self' ws: wss:; style-src 'self' 'unsafe-inline'; script-src 'self'
'unsafe-inline' *.pendo.io; img-src 'self' *.pendo.io data:; report-uri /
csp-violation/";

    location /favicon.ico { alias /var/lib/awx/venv/awx/lib/python3.8/site-
packages/awx/ui/build/static/media/favicon.ico; }

    location /static/ {
```

```
        alias /var/lib/awx/venv/awx/lib/python3.8/site-packages/awx/ui/
build/static/;
    }

    location /websocket {
        # Pass request to the upstream alias
        proxy_pass http://daphne;
        # Require http version 1.1 to allow for upgrade requests
        proxy_http_version 1.1;
        # We want proxy_buffering off for proxying to websockets.
        proxy_buffering off;
        # http://en.wikipedia.org/wiki/X-Forwarded-For
        proxy_set_header X-Forwarded-For $proxy_add_x_forwarded_for;
        # enable this if you use HTTPS:
        proxy_set_header X-Forwarded-Proto https;
        # pass the Host: header from the client for the sake of redirects
        proxy_set_header Host $http_host;
        # We've set the Host header, so we don't need Nginx to muddle
        # about with redirects
        proxy_redirect off;
        # Depending on the request value, set the Upgrade and
        # connection headers
        proxy_set_header Upgrade $http_upgrade;
        # proxy_set_header Connection $connection_upgrade;
        proxy_set_header Connection upgrade;
    }

    location / {
        # Add trailing / if missing
        rewrite ^(.*[^/])$ $1/ permanent;
        uwsgi_read_timeout 120s;
        uwsgi_pass uwsgi;
        include /etc/nginx/uwsgi_params;
    }
}
```

20. Remove any default configuration for Receptor. Edit `/etc/receptor/receptor.conf` to contain the following configuration:

```
---
- node:
    id: <hostname or ip address>

- log-level: debug

- tcp-listener:
    port: 27199

#- work-signing:
#    privatekey: /etc/receptor/work_private_key.pem
#    tokenexpiration: 1m

#- work-verification:
#    publickey: /etc/receptor/work_public_key.pem

#- tcp-peer:
#    address: 100.100.253.53:27199
#    redial: true
```

```
#- tls-server:
#    name: mutual-tls
#    cert: /etc/receptor/certs/awx.crt
#    key: /etc/receptor/certs/awx.key
#    requireclientcert: true
#    clientcas: /etc/receptor/certs/ca.crt

- control-service:
    service: control
    filename: /var/run/receptor/receptor.sock

- work-command:
    worktype: local
    command: /var/lib/ol-automation-manager/venv/awx/bin/ansible-runner
    params: worker
    allowruntimeparams: true
#    verifysignature: true
```

In the previous example, *hostname or ip address* is the hostname or IP address of the system running Oracle Linux Automation Manager. If hostname is used, the host must be resolvable.

21. Prepare an Oracle Linux Automation Manager deployment as the awx user. Do the following:

    a. Run the following commands:

    ```
    sudo su -l awx -s /bin/bash
    podman system migrate
    podman pull container-registry.oracle.com/oracle_linux_automation_manager/olam-ee:latest
    awx-manage makemigrations --merge
    awx-manage migrate
    awx-manage register_default_execution_environments
    ```

    > **Note:**
    >
    > After you finish upgrading Oracle Linux Automation Manager, you can configure whether you want your Execution Environments to always pull the latest `olam-ee` container image when running playbooks, or use some other option. For more information about these options, see Oracle Linux Automation Manager 2.1: User's Guide.

    b. Exit the awx shell environment.

    ```
    exit
    ```

22. Restore the /etc/tower/SECRET_KEY file. For example:

    ```
    sudo cp ~/SECRET_KEY /etc/tower/SECRET_KEY
    ```

23. In the `/etc/tower/settings.py` file, add the following lines:

    ```
    DEFAULT_EXECUTION_QUEUE_NAME = 'tower'
    DEFAULT_CONTROL_PLANE_QUEUE_NAME = 'tower'
    ```

24. Restart NGINX.

    ```
    sudo systemctl restart nginx
    ```

25. Start Oracle Linux Automation Manager.

```
sudo systemctl start ol-automation-manager
```

Upgrading Release 2.0 to Release 2.1

This upgrade is necessary for Oracle Linux Automation Manager to make use of Private Automation Hub execution environment container images and collections.

1. On all Oracle Linux Automation Manager 2.0 nodes, log in to a terminal.
2. Run an update:

```
sudo dnf clean all
sudo dnf update oraclelinux-automation-manager-release-el8
sudo dnf update ol-automation-manager ol-automation-manager-cli uwsgi
```

3. Pull the latest olam-ee image using.

```
sudo su -l awx -s /bin/bash
podman pull container-registry.oracle.com/oracle_linux_automation_manager/
olam-ee:latest
exit
```

4. Add the following parameter to `/etc/tower/settings.py` file:

```
# OLAM Reaper Job Status Tracking
REAPER_TIMEOUT_SEC = 60
```

   The `REAPER_TIMEOUT_SEC` parameter specifies the time in seconds befor a job in the Running or Waiting state is considered stuck and transitioned by the reaper into the Failed state. You can modify this parameter in cases where you have playbooks that run longer than 60 seconds. For more information about this parameter, see Oracle Linux Automation Manager 2.1: Installation Guide.

5. Restart the following services on all nodes:

```
sudo systemctl restart ol-automation-manager
sudo systemctl restart nginx
sudo systemctl restart receptor-awx
```

6. If required, you can change the default execution environment location for new custom execution environment instances to your private automation hub instance. In addition, you might download `olam-ee` from `container-registry.oracle.com` to host it on Private Automation Hub to use for control plane instances. To set these default settings, do the following:
   a. Edit the `/etc/tower/settings.py` file as follows:

```
GLOBAL_JOB_EXECUTION_ENVIRONMENTS = [{'name':
'<customer_execution_environment>)', 'image':
'<private_automation_hub_hostname_or_ip_address>/
<customer_execution_environment>:latest'}]
CONTROL_PLANE_EXECUTION_ENVIRONMENT =
'<private_automation_hub_hostname_or_ip_address>/olam-ee:latest'
```

b. Register the new default environments using the following commands:

```
sudo su -l awx -s /bin/bash
awx-manage register_default_execution_environments
exit
```

c. Restart Oracle Linux Automation Manager on all nodes.

```
sudo systemctl restart ol-automation-manager
```

Migrating a Single Instance Deployment to a Clustered Deployment

To migrate a single host instance deployment of Oracle Linux Automation Manager to a clustered deployment, do the following:

1. If you need to upgrade your single instance host to release 2.0, complete the upgrade procedures in Upgrading a Release 1.0.X to a Release 2.0 Single Host Deployment.
2. Verify that the upgraded instance is working.
3. In a terminal, stop Oracle Linux Automation Manager.

   ```
   sudo systemctl stop ol-automation-manager
   ```

4. Create a database dump file.

   ```
   sudo su - postgres
   pg_dumpall > olamv2upg.dump
   ```

5. Open the firewall port on the remote database as described in Setting Up the Firewall Rules.
6. Complete the procedures for setting up a remote database in Setting Up a Local or Remote Database with the following exceptions:
   a. Before starting the procedure, copy over the dump file to the remote database. For example, using scp.
   b. After starting the database in step 7, import the dump file:

      ```
      sudo su - postgres
      psql -d postgres -f /dirwithbackup/olamv2upg.dump
      exit
      ```

   c. Skip steps 8 through 10 for creating the database user account and creating the database because these are already part of the dump file.
   d. Continue the procedure at step 11.
7. On the remote database, reapply the password to the database user account:

   ```
   sudo -u postgres psql
   \password awx
   ```

8. Enter and confirm the password for the awx user.

```
Enter new password for user "awx":
Enter it again:
exit
```

9. Restart the database.

```
sudo systemctl restart postgresql
```

10. Return to the upgraded instance, and in the `/etc/tower/settings.py` file, replace the existing `DATABASES` fields with the following fields:

```
DATABASES = {
    'default': {
        'ATOMIC_REQUESTS': True,
        'ENGINE': 'awx.main.db.profiled_pg',
        'NAME': 'awx',
        'USER': 'awx',
        'PASSWORD': 'password',
        'HOST': 'database hostname or ip address',
        'PORT': '5432',
    }
}
```

In the previous example, *database hostname or ip address* is the hostname or IP address of the remote database. If hostname is used, the host must be resolvable. *password* is the password for your remote database, if you have configured one.

11. Stop the local database.

```
sudo systemctl stop postgresql
```

12. Open the ports used for the Service Mesh.

```
sudo firewall-cmd --add-port=27199/tcp --permanent
sudo firewall-cmd --reload
```

13. Start Oracle Linux Automation Manager.

```
sudo systemctl start ol-automation-manager
```

14. Remove the local database, because it is no longer needed.

```
sudo dnf remove postgresql
```

15. Run the following command:

```
sudo su -l awx -s /bin/bash
```

16. Remove the `tower` instance group (queue name) because this is not used in Oracle Linux Automation Manager release 2.

```
awx-manage remove_from_queue --queuename tower --hostname <hostname or IP
address>
```

In the previous example, *hostname or IP address* is the hostname or IP address of the system running Oracle Linux Automation Manager.

17. Run the following commands.

```
awx-manage provision_instance --hostname=<hostname or IP address> --
node_type=control
awx-manage register_queue --queuename=controlplane --hostnames=<hostname
or IP address>
exit
```

In the previous example, *hostname or IP address* is the hostname or IP address of the system running Oracle Linux Automation Manager. Your choice of host name or IP address must match with the host name or IP addressed used when you configured the `/etc/receptor/receptor.conf` file node ID (see Configuring and Starting the Control Plane Service Mesh). If hostname is used, the host must be resolvable.

18. In the `/etc/tower/settings.py` file, replace the following lines.

```
DEFAULT_EXECUTION_QUEUE_NAME = 'tower'
DEFAULT_CONTROL_PLANE_QUEUE_NAME = 'tower'
```

with these lines.

```
DEFAULT_EXECUTION_QUEUE_NAME = 'execution'
DEFAULT_CONTROL_PLANE_QUEUE_NAME = 'controlplane'
```

19. Restart Oracle Linux Automation Manager.

```
sudo systemctl restart  ol-automation-manager.service
```

20. The original upgraded node is now converted into a control node. You must now add one more execute node for the upgraded cluster to be fully functional. For all other members of the cluster, follow the procedures described in Preparing the Database and Hosts, with the exception of setting up a remote database, because this is already completed. Then, follow the procedures for installing and configuring all other hosts as part of the cluster, as described in Installing Oracle Linux Automation Manager in a Clustered Deployment.

Migrating Playbooks to Oracle Linux Automation Engine Release 2.0

Test your Oracle Linux Automation Engine release 1.0.x playbooks to verify whether they function properly with Oracle Linux Automation Manager release 2.0. You may need to update your playbooks because the upstream projects have made changes such as, the number of modules, some modules have become collections, and some modules have been consolidated into other modules or collections.