
Ansible Automation Platform Upgrade and Migration

Release Automation Controller 4.0.0

Red Hat, Inc.

Oct 24, 2021

CONTENTS

1	Release Notes for Automation Controller Version 4.0.0	2
1.1	Automation Controller Version 4.0	2
2	Upgrading to Ansible Automation Platform	4
2.1	Upgrade Planning	5
2.2	Obtaining the Installer	5
2.3	Setting up the Inventory File	5
2.4	Running the Setup Playbook	10
3	Upgrading to Execution Environments	12
3.1	Migrate legacy venvs to execution environments	12
4	Role-Based Access Controls	15
4.1	Organization field on Job Templates	15
5	Index	16
6	Copyright © Red Hat, Inc.	17
	Index	18

Thank you for your interest in Red Hat Ansible Automation Platform controller. automation controller is a commercial offering that helps teams manage complex multi-tier deployments by adding control, knowledge, and delegation to Ansible-powered environments.

Note: You must upgrade your automation controller to automation controller 3.8 before you can upgrade to automation controller 4.0.

We Need Feedback!

If you spot a typo in this documentation, or if you have thought of a way to make this manual better, we would love to hear from you! Please send an email to: docs@ansible.com

If you have a suggestion, try to be as specific as possible when describing it. If you have found an error, please include the manual's title, chapter number/section number, and some of the surrounding text so we can find it easily. We may not be able to respond to every message sent to us, but you can be sure that we will be reading them all!

Automation Controller Version 4.0.0; July 15, 2021; <https://access.redhat.com/>

RELEASE NOTES FOR AUTOMATION CONTROLLER VERSION 4.0.0

1.1 Automation Controller Version 4.0

Introduced

- Support for automation execution environments. All automation now runs in execution environments via containers, either directly via OpenShift, or locally via podman
- New PatternFly 4 based user-interface for increased performance, security, and consistency with other Ansible Automation Platform components

Added

- Added identity provider support for GitHub Enterprise
- Support for RHEL system crypto profiles to nginx configuration
- The ability to disable local system users and only pull users from configured identity providers
- Additional Prometheus metrics for tracking job event processing performance
- New `awx-manage` command for dumping host automation information
- Red Hat Insights as an inventory source
- Ability to set server-side password policies using Django's `AUTH_PASSWORD_VALIDATORS` setting
- Support for Centrify Vault as a credential lookup plugin
- Support for namespaces in Hashicorp Vault credential plugin

Updated

- OpenShift deployment to be done via an Operator instead of a playbook
- Python used by application to Python 3.8
- Nginx used to version 1.18
- PostgreSQL used to PostgreSQL 12, and moved to partitioned databases for performance
- The “container groups” feature to general availability from Tech Preview; now fully utilizes execution environments
- Insights remediation to use new Red Hat Insights inventory source rather than utilizing scan playbooks with arbitrary inventory
- Subscriptions display to count hosts automated on instead of hosts imported
- Inventory source, credential, and Ansible content collection to reference *controller* instead of *tower*

Deprecated

- None

Removed

- Support for deploying on CentOS (any version) and RHEL 7
- Support for Mercurial projects
- Support for custom inventory scripts stored in controller (use `awx-manage export_custom_scripts` to export them)
- Resource profiling code (`AWX_RESOURCE_PROFILING_*`)
- Support for custom Python virtual environments for execution. Use new `awx-manage` tools for assisting in migration
- Top-level `/api/v2/job_events/` API endpoint
- The ability to disable job isolation

UPGRADING TO ANSIBLE AUTOMATION PLATFORM

Automation Hub acts as a content provider for automation controller, which requires both an automation controller deployment and an Automation Hub deployment running alongside each other. The Ansible Automation Platform installer contains both of these. This section covers each component of the upgrading process:

- *Upgrade Planning*
- *Obtaining the Installer*
- *Setting up the Inventory File*
 - *Example Inventory files*
 - * *Example Standalone Automation Hub Inventory File*
 - * *Example Platform Inventory File*
 - * *Example Single Node Inventory File*
 - * *Example Multi Node Cluster Inventory File*
 - * *Example Inventory file for an external existing database*
 - * *Example Inventory file for external database which needs installation*
- *Running the Setup Playbook*

Note: All upgrades should be no more than two major versions behind what you are currently upgrading to. For example, in order to upgrade to automation controller 4.0, you must first be on version 3.8.x; i.e., there is no direct upgrade path from version 3.7.x or earlier. Refer to the [recommended upgrade path article](#) on the Red Hat customer portal.

In order to run automation controller 4.0, you must also have Ansible 2.10.

2.1 Upgrade Planning

This section covers changes that you should keep in mind as you attempt to upgrade your automation controller instance.

- Even if you already have a valid license from a previous version, you must still provide your credentials or a subscriptions manifest again upon upgrading to automation controller 3.8. See [Import a Subscription](#) in the *Automation Controller User Guide*.
- If you need to upgrade Red Hat Enterprise Linux and automation controller, you will need to do a backup and restore of your controller data (from Tower). Refer to [Backing Up and Restoring](#) in the *Automation Controller Administration Guide* for further detail.
- Clustered upgrades require special attention to instance and instance groups prior to starting the upgrade. See [Setting up the Inventory File](#) and [Clustering](#) for details.
- Prior versions of automation controller used the variable name `rabbitmq_host` during installation. If you are upgrading from a previous version of automation controller, and you previously specified `rabbitmq_host` in your inventory, simply rename `rabbitmq_host` to `routable_hostname` before upgrading. See [Clustering](#) for details.

2.2 Obtaining the Installer

Refer to [Choosing and obtaining a Red Hat Ansible Automation Platform installer on the Red Hat Customer Portal](#) for detail. Be sure to use your Red Hat customer login to access the full content.

2.3 Setting up the Inventory File

As you edit your inventory file, there are a few things you must keep in mind:

- The contents of the inventory file should be defined in `./inventory`, next to the `./setup.sh` installer playbook.
- For **installations and upgrades**: If you need to make use of external databases, you must ensure the database sections of your inventory file are properly setup. Edit this file and add your external database information before running the setup script.
- For **Ansible Automation Platform** or **Automation Hub**: Be sure to add an automation hub host in the `[automationhub]` group (Tower and Automation Hub cannot be installed on the same node).
- Tower will not configure replication or failover for the database that it uses, although Tower should work with any replication that you have.
- The database server should be on the same network or in the same data center as the Tower server for performance reasons.
- For **upgrading an existing cluster**: When upgrading a cluster, you may decide that you want to also reconfigure your cluster to omit existing instances or instance groups. Omitting the instance or the instance group from the inventory file will not be enough to remove them from the cluster. In addition to omitting instances or instance groups from the inventory file, you must also [deprovision instances or instance groups](#) before starting the upgrade. Otherwise, omitted instances or instance groups will continue to communicate with the cluster, which can cause issues with tower services during the upgrade.
- For **clustered installations**: If you are creating a clustered setup, you must replace `localhost` with the hostname or IP address of all instances. All nodes/instances must be able to reach any others using this hostname

or address. In other words, you cannot use the `localhost ansible_connection=local` on one of the nodes *AND* all of the nodes should use the same format for the host names.

Therefore, this will *not* work:

```
[tower]
localhost ansible_connection=local
hostA
hostB.example.com
172.27.0.4
```

Instead, use these formats:

```
[tower]
hostA
hostB
hostC
```

OR

```
hostA.example.com
hostB.example.com
hostC.example.com
```

OR

```
[tower]
172.27.0.2
172.27.0.3
172.27.0.4
```

- For **all standard installations**: When performing an installation, you must supply any necessary passwords in the inventory file.

Note: Changes made to the installation process now require that you fill out all of the password fields in the inventory file. If you need to know where to find the values for these they should be:

```
admin_password='' ← Tower local admin password
pg_password='' ← Found in /etc/tower/conf.d/postgres.py
```

Warning: Do not use special characters in `pg_password` as it may cause the setup to fail.

2.3.1 Example Inventory files

- For **provisioning new nodes**: When provisioning new nodes add the nodes to the inventory file with all current nodes, make sure all passwords are included in the inventory file.
- For **upgrading a single node**: When upgrading, be sure to compare your inventory file to the current release version. It is recommended that you keep the passwords in here even when performing an upgrade.

Example Standalone Automation Hub Inventory File

```
[automationhub]
automationhub.acme.org
[all:vars]
automationhub_admin_password='<password>'
automationhub_pg_host=''
automationhub_pg_port=''
automationhub_pg_database='automationhub'
automationhub_pg_username='automationhub'
automationhub_pg_password='<password>'
automationhub_pg_sslmode='prefer'
# The default install will deploy a TLS enabled Automation Hub.
# If for some reason this is not the behavior wanted one can
# disable TLS enabled deployment.
#
# automationhub_disable_https = False
# The default install will generate self-signed certificates for the Automation
# Hub service. If you are providing valid certificate via automationhub_ssl_cert
# and automationhub_ssl_key, one should toggle that value to True.
#
# automationhub_ssl_validate_certs = False
# SSL-related variables
# If set, this will install a custom CA certificate to the system trust store.
# custom_ca_cert=/path/to/ca.crt
# Certificate and key to install in Automation Hub node
# automationhub_ssl_cert=/path/to/automationhub.cert
# automationhub_ssl_key=/path/to/automationhub.key
```

Example Platform Inventory File

```
[tower]
tower.acme.org
[automationhub]
automationhub.acme.org
[database]
database-01.acme.org
[all:vars]
admin_password='<password>'
pg_host='database-01.acme.org'
pg_port='5432'
pg_database='awx'
pg_username='awx'
pg_password='<password>'
pg_sslmode='prefer' # set to 'verify-full' for client-side enforced SSL
# Automation Hub Configuration
#
automationhub_admin_password='<password>'
automationhub_pg_host='database-01.acme.org'
automationhub_pg_port='5432'
automationhub_pg_database='automationhub'
automationhub_pg_username='automationhub'
automationhub_pg_password='<password>'
automationhub_pg_sslmode='prefer'
# The default install will deploy a TLS enabled Automation Hub.
# If for some reason this is not the behavior wanted one can
```

(continues on next page)

(continued from previous page)

```
# disable TLS enabled deployment.
#
# automationhub_disable_https = False
# The default install will generate self-signed certificates for the Automation
# Hub service. If you are providing valid certificate via automationhub_ssl_cert
# and automationhub_ssl_key, one should toggle that value to True.
#
# automationhub_ssl_validate_certs = False
# Isolated Tower nodes automatically generate an RSA key for authentication;
# To disable this behavior, set this value to false
# isolated_key_generation=true
# SSL-related variables
# If set, this will install a custom CA certificate to the system trust store.
# custom_ca_cert=/path/to/ca.crt
# Certificate and key to install in nginx for the web UI and API
# web_server_ssl_cert=/path/to/tower.cert
# web_server_ssl_key=/path/to/tower.key
# Certificate and key to install in Automation Hub node
# automationhub_ssl_cert=/path/to/automationhub.cert
# automationhub_ssl_key=/path/to/automationhub.key
# Server-side SSL settings for PostgreSQL (when we are installing it).
# postgres_use_ssl=False
# postgres_ssl_cert=/path/to/pgsql.crt
# postgres_ssl_key=/path/to/pgsql.key
```

Example Single Node Inventory File

```
[tower]
localhost ansible_connection=local

[database]

[all:vars]
admin_password='password'

pg_host=''
pg_port=''

pg_database='awx'
pg_username='awx'
pg_password='password'
```

Warning: Do not use special characters in `pg_password` as it may cause the setup to fail.

Example Multi Node Cluster Inventory File

```
[tower]
clusternode1.example.com
clusternode2.example.com
clusternode3.example.com

[database]
dbnode.example.com

[all:vars]
ansible_become=true

admin_password='password'

pg_host='dbnode.example.com'
pg_port='5432'

pg_database='tower'
pg_username='tower'
pg_password='password'
```

Warning: Do not use special characters in `pg_password` as it may cause the setup to fail.

Example Inventory file for an external existing database

```
[tower]
node.example.com ansible_connection=local

[database]

[all:vars]
admin_password='password'
pg_password='password'

pg_host='database.example.com'
pg_port='5432'

pg_database='awx'
pg_username='awx'
```

Warning: Do not use special characters in `pg_password` as it may cause the setup to fail.

Example Inventory file for external database which needs installation

```
[tower]
node.example.com ansible_connection=local

[database]
database.example.com

[all:vars]
admin_password='password'
pg_password='password'

pg_host='database.example.com'
pg_port='5432'

pg_database='awx'
pg_username='awx'
```

Warning: Do not use special characters in `pg_password` as it may cause the setup to fail.

Once any necessary changes have been made, you are ready to run `./setup.sh`.

Note: Root access to the remote machines is required. With Ansible, this can be achieved in different ways:

- `ansible_user=root ansible_ssh_pass="your_password_here"` inventory host or group variables
- `ansible_user=root ansible_ssh_private_key_file="path_to_your_keyfile.pem"` inventory host or group variables
- `ANSIBLE_BECOME_METHOD='sudo' ANSIBLE_BECOME=True ./setup.sh`
- `ANSIBLE_SUDO=True ./setup.sh` (Only applies to Ansible 2.7)

The `DEFAULT_SUDO` Ansible configuration parameter was removed in Ansible 2.8, which causes the `ANSIBLE_SUDO=True ./setup.sh` method of privilege escalation to no longer work. For more information on become plugins, refer to [Understanding Privilege Escalation](#) and the [list of become plugins](#).

2.4 Running the Setup Playbook

The Tower setup playbook script uses the `inventory` file and is invoked as `./setup.sh` from the path where you unpacked the Tower installer tarball.

```
root@localhost:~$ ./setup.sh
```

The setup script takes the following arguments:

- `-h` – Show this help message and exit
- `-i INVENTORY_FILE` – Path to Ansible inventory file (default: `inventory`)
- `-e EXTRA_VARS` – Set additional Ansible variables as `key=value` or `YAML/JSON` (i.e. `-e bundle_install=false` forces an online installation)

- `-b` – Perform a database backup in lieu of installing
- `-r` – Perform a database restore in lieu of installing (a default restore path is used unless `EXTRA_VARS` are provided with a non-default path, as shown in the code example below)

```
./setup.sh -e 'restore_backup_file=/path/to/nondefault/location' -r
```

UPGRADING TO EXECUTION ENVIRONMENTS

If upgrading from older versions of automation controller to 4.0 or later, the controller has the ability to detect previous versions of virtual environments associated with Organizations, Inventory, and Job Templates; and inform you that you will need to migrate to the new execution environment model. A brand new installation of automation controller creates two virtualenvs during installation—one is used to run the controller itself, while the other is used to run Ansible. Like legacy virtual environments, execution environments allow the controller to run in a stable environment, while allowing you to add or update modules to your execution environment as necessary to run your playbooks. For more information, see [Execution Environments](#) in the *Automation Controller User Guide*.

3.1 Migrate legacy venvs to execution environments

You can have the exact same setup in an execution environment that you had in a prior custom virtual environment by migrating them to the new execution environment. Use the `awx-manage` commands in this section to:

- list of all the current custom virtual environments and their paths (`list_custom_venvs`)
- view the resources that rely a particular custom virtual environment (`custom_venv_associations`)
- export a particular custom virtual environment to a format that can be used to migrate to an execution environment (`export_custom_venv`)

1. Before you migrate, it is recommended that you view all the custom virtual environments you currently have running by using the `awx-manage list` command:

```
$ awx-manage list_custom_venvs
```

Below is an example output when running this command:

```
bash-4.4$ awx-manage list_custom_venvs
# Discovered Virtual Environments:
/var/lib/awx/venv/i_heart_ansible
/var/lib/awx/venv/testing
/var/lib/awx/venv/new_env_better_name

- To export the contents of a (deprecated) virtual environment, run the following command while supplying the path as an argument:
awx-manage export_custom_venv /path/to/venv

- To view the connections a (deprecated) virtual environment had in the database, run the following command while supplying the path as an argument:
awx-manage custom_venv_associations /path/to/venv

- Run these commands with `-q` to remove tool tips.
```

The above output shows three custom virtual environments and their paths. If you have a custom virtual environment that is not located within the default `/var/lib/awx/venv/` directory path, it will not be included here.

2. Use the `_associations` command to view what organizations, jobs, and inventory sources a custom virtual environment is associated with in order to determine which resources rely on them:

```
$ awx-manage custom_venv_associations /this/is/the/path/
```

Below is an example output when running this command:

```
bash-4.4$ awx-manage custom_venv_associations /var/lib/awx/venv/new_env_better_name
# Virtual Environments Associations:
inventory_sources:
- id: 15
  name: celery
job_templates:
- id: 9
  name: Demo Job Template @ 2:40:47 PM
- id: 13
  name: elephant
organizations:
- id: 3
  name: alternating_bongo_meow
- id: 1
  name: Default
projects: []

- To list all (now deprecated) custom virtual environments run:
awx-manage list_custom_venvs

- To export the contents of a (deprecated) virtual environment, run the following command while supplying the path as an argument:
awx-manage export_custom_venv /path/to/venv

- Run these commands with `-q` to remove tool tips.
```

3. Select a path for the virtual environment that you want to migrate and specify it in the `awx-manage export` command:

```
$ awx-manage export_custom_venv /this/is/the/path/
```

The resulting output is essentially the results of executing a `pip freeze` command. The example shows the contents of the selected custom virtual environment:

```
bash-4.4$ awx-manage export_custom_venv /var/lib/awx/venv/new_env_better_name
# Virtual environment contents:
ansible==2.9.0
cffi==1.14.5
cryptography==3.4.7
Jinja2==3.0.1
MarkupSafe==2.0.1
numpy==1.20.2
pandas==1.2.4
psutil==5.8.0
pycparser==2.20
python-dateutil==2.8.1
pytz==2021.1
PyYAML==5.4.1
six==1.16.0

- To list all (now deprecated) custom virtual environments run:
awx-manage list_custom_venvs

- To view the connections a (deprecated) virtual environment had in the database, run the following command while
supplying the path as an argument:
awx-manage custom_venv_associations /path/to/venv

- Run these commands with '-q' to remove tool tips.

bash-4.4$
```

Note: All of these commands can be run with a `-q` option, which removes the instructional content provided on each output.

Now that you have the output from this `pip freeze` data, you can paste it into a definition file that can be used to spin up your new execution environment using `ansible-builder`. Anyone (both normal users and admins) can use `ansible-builder` to create an execution environment. See [Building an Execution Environment](#) in the *Automation Controller User Guide* for further detail.

ROLE-BASED ACCESS CONTROLS

automation controller 3.7 contains minor updates to the Role-Based Access Control (RBAC) system. For the latest RBAC documentation, refer to the [Role-Based Access Controls](#) section in the Tower User Guide.

4.1 Organization field on Job Templates

Job templates in automation controller now include an organization field in the API. This is set on creation based on the organization of the project used by the Job Template, and cannot be changed. Because of this, a project's organization cannot be changed once it is in use by Job Templates.

This changes visibility and access to job templates. Previously, an admin of the organization that a job template's inventory belonged to would also be granted admin access to the job template. While existing permissions are preserved on an upgrade to automation controller 3.7, newly created jobs will only grant view access to the job template to the inventory admin in this scenario.

- genindex

COPYRIGHT © RED HAT, INC.

Ansible, Ansible Automation Platform, Red Hat, and Red Hat Enterprise Linux are trademarks of Red Hat, Inc., registered in the United States and other countries.

If you distribute this document, or a modified version of it, you must provide attribution to Red Hat, Inc. and provide a link to the original version.

Third Party Rights

Ubuntu and Canonical are registered trademarks of Canonical Ltd.

The CentOS Project is copyright protected. The CentOS Marks are trademarks of Red Hat, Inc. (“Red Hat”).

Microsoft, Windows, Windows Azure, and Internet Explore are trademarks of Microsoft, Inc.

VMware is a registered trademark or trademark of VMware, Inc.

Rackspace trademarks, service marks, logos and domain names are either common-law trademarks/service marks or registered trademarks/service marks of Rackspace US, Inc., or its subsidiaries, and are protected by trademark and other laws in the United States and other countries.

Amazon Web Services”, “AWS”, “Amazon EC2”, and “EC2”, are trademarks of Amazon Web Services, Inc. or its affiliates.

OpenStack™ and OpenStack logo are trademarks of OpenStack, LLC.

Chrome™ and Google Compute Engine™ service registered trademarks of Google Inc.

Safari® is a registered trademark of Apple, Inc.

Firefox® is a registered trademark of the Mozilla Foundation.

All other trademarks are the property of their respective owners.

INDEX

A

Ansible
 executing in a execution
 environment, 12

B

build
 execution environments, 12

E

executing in a execution environment
 Ansible, 12
execution environment, 12
execution environments
 build, 12

I

installation script
 inventory file setup, 5
 playbook setup, 10
inventory file setup, 5

M

migrate to execution environments
 virtual environments, 12

P

permissions, 15
playbook setup, 10
 installation script, 10
 setup.sh, 10

R

RBAC, 15
roles, 15

S

setup.sh
 playbook setup, 10
singleton roles, 15
system-wide roles, 15

U

upgrade, 4
upgrade considerations, 4

V

virtual environments
 migrate to execution environments,
 12