
Ansible Tower Upgrade and Migration

Release Ansible Tower 3.4.0

Red Hat, Inc.

May 05, 2021

CONTENTS

1	Release Notes for Ansible Tower Version 3.4.0	2
1.1	Ansible Tower Version 3.4.0	2
2	Upgrading Ansible Tower	4
2.1	Upgrade Planning	4
2.2	Obtaining Ansible Tower	5
2.3	Setting up the Inventory File	5
2.4	The Setup Playbook	8
3	System Tracking Migration	10
4	Role-Based Access Controls	11
4.1	Enhanced and Simplified RBAC System	11
4.2	Specific Changes to Note	11
5	Job Template Changes	12
5.1	Scan Jobs	12
5.2	Prompt on Launch	12
5.3	Permissions/RBAC Notes	14
5.4	Surveys	14
6	Job Output View Changes	16
6.1	Details	16
6.2	Standard Out Pane	18
7	Using virtualenv with Ansible Tower	20
7.1	Preparing a new custom virtualenv	20
7.2	Assigning custom virtualenvs	21
8	Index	23
9	Copyright © 2019 Red Hat, Inc.	24
	Index	25

Thank you for your interest in Red Hat Ansible Tower. Ansible Tower 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 Ansible Tower 2.4.4 (or later) system to Ansible Tower 3.0 before you can upgrade to Ansible Tower 3.1.0.

The *Ansible Tower Upgrade and Migration Guide* discusses how to upgrade your Ansible Tower 2.4.4 (or later) system to the 3.0 version. The Ansible Tower 3.0 release introduced many updates and changes to Tower, including changes to the way upgrades run and a completely rewritten RBAC system. This guide covers these and other related changes that you should keep in mind as you plan to migrate your data and prepare for this upgrade.

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!

Ansible Tower Version 3.4.0; January 09, 2019; <https://access.redhat.com/>

RELEASE NOTES FOR ANSIBLE TOWER VERSION 3.4.0

1.1 Ansible Tower Version 3.4.0

- Added the ability for workflow templates to have workflows as nodes
- Added to workflow templates the ability to allow edges/links to have mix of always, success, or failure, therefore, removing edge conflicts to improve workflow behavior
- Added job slicing feature to job templates, where a job template will launch a workflow of a certain number of jobs, distributing inventory evenly amongst those jobs, and allowing work to execute in parallel among Tower instances
- Added optional and promptable inventory field to Workflow job templates (WFJTs)- the selected inventory will override the inventory of any Job Template node that prompts for an inventory, which includes inventory prompt-on-launch for standalone WFJTs, nested WFJT nodes, and WFJT schedules
- Added the ability to run Ansible Tower when Red Hat Enterprise Linux is configured in FIPS mode
- Added job event entries for `runner_on_start`, since it was introduced in Ansible core
- Deprecated several read-only API fields, pending removal in future release:
 - `ask_XXX_on_launch` fields in jobs API, duplicated with job template values
 - Inventory and groups `groups_with_active_failures`, not used by the Tower user interface
- Removed previously-deprecated read-only project field `scm_delete_on_next_update`
- Removed the restriction that workflows cannot be used in workflows
- Removed backwards-compatible support for deprecated legacy Auth Tokens. Starting in Ansible Tower 3.4, users who wish to authenticate with tokens must use the new OAuth 2.0 authentication mechanism introduced in Tower 3.3.
- Removed deprecated cluster management commands: `deprovision_node` → `deprovision_instance` and `register_instance` → `provision_instance`
- Removed deprecated `--name` option from `deprovision_instance`
- Fixed a deadlock scenario where new websocket connections would get 502's until daphne and runworker services were restarted
- Fixed the Workflow job templates to return a 400 response if a client attempts to launch without providing required survey variables using the API
- Improved the task runner by replacing celery with a custom Kombu-based dispatcher
- Improved validation survey specs on save, returning a 400 response in many new cases

- Improved the Workflow job nodes with missing `unified_job_template` to no longer unconditionally fail a workflow. Instead, the workflow job node is treated as a job that ran with status result of failure.
- Improved the Workflow job failure semantics. Before, if any leaf node failed then the workflow job was considered failed, else success. Now, if any node in the workflow fails without a failure handle path, then the workflow job is considered failed; else success.
- Improved workflows to now allow convergence nodes (nodes with multiple parents)
- Updated the environment variables used in job runs, which include the removal of `INVENTORY_HOSTVARS` and the ability for the inventory plugin and inventory unparsed behavior to be configurable
- Updated how project “sync” jobs (as opposed to project updates) work
- Updated the behavior of the Activity Stream associated with deleted inventory and displayed list of credentials used in job creation entry
- Updated the creation of schedules from the API to allow global schedules lists
- Updated GCE instances by moving from `GCE_PEM_FILE_PATH` to `GCE_CREDENTIALS_FILE_PATH`
- Updated Tower to allow use of organization roles with teams in the user interface
- Updated Configure Tower from the Settings menu to contain sub-menus and access to sub-tasks

UPGRADING ANSIBLE TOWER

Topics:

- *Upgrade Planning*
- *Obtaining Ansible Tower*
- *Setting up the Inventory File*
- *The Setup Playbook*

Note: You must upgrade your Ansible Tower 2.4.4 (or later) system to Ansible Tower 3.0 before you can upgrade beyond Ansible Tower 3.0. Refer to the [recommended upgrade path article](#) from your customer portal.

2.1 Upgrade Planning

This section covers changes that you should keep in mind as you attempt to upgrade your Ansible Tower Instance

- If you are not yet using a 2.4.x version of Ansible Tower, **do not** attempt to upgrade directly to Ansible Tower 3.0.x or 3.1.0. You must start with a system which has a version of Tower 2.4.x installed or the upgrade will fail.
- If you are not using using a 3.0.x version of Ansible Tower; you must upgrade to 3.0.x before you can upgrade beyond Ansible Tower 3.0.
- Ansible Tower 3.0 simplified installation and removed the need to run `./configure/` as part of the initial setup.
- The file `tower_setup_conf.yml` is no longer used. Instead, you should now edit the **inventory** file in the `/ansible-tower-setup-<tower_version>/` directory.
- Earlier versions of Tower used MongoDB when setting up an initial database; please note that Ansible Tower 3.0 has replaced the use of MongoDB with PostgreSQL.
- Clustered upgrades require special attention to instance and instance groups prior to starting the upgrade. Refer to the *Setting up the Inventory File* section.
- Changes in API authentication were made in Ansible Tower 3.3 to accommodate additional OAuth2 functionality. For more information, refer to [Token-Based Authentication](#) in the *Ansible Tower Administration Guide*.

2.2 Obtaining Ansible Tower

You may install standalone Tower or use the bundled installer:

- if you set up Tower on an environment with a direct Internet access, you can download the standalone Tower installer
- if you set up Tower on an environment without direct access to online repositories, or your environment enforces a proxy, you must use the bundled installer

Download and then extract the Ansible Tower installation/upgrade tool: <http://releases.ansible.com/ansible-tower/setup/>

```
root@localhost:~$ tar xvzf ansible-tower-setup-latest.tar.gz
root@localhost:~$ cd ansible-tower-setup-<tower_version>
```

To install or upgrade, start by editing the inventory file in the `ansible-tower-setup-<tower_version>` directory, replacing `<tower_version>` with the version number, such as 3.3.7 or 3.4.0 directory.

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 **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
rabbitmq_password='' ← create a new password here (alpha-numeric with no special characters)
```

Example Inventory file

- 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 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'

rabbitmq_port=5672
rabbitmq_vhost=tower
rabbitmq_username=tower
rabbitmq_password='password'
rabbitmq_cookie=rabbitmqcookie

# Needs to be true for fqdns and ip addresses
```

(continues on next page)

(continued from previous page)

```
rabbitmq_use_long_name=false
# Needs to remain false if you are using localhost
```

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'

rabbitmq_port=5672
rabbitmq_vhost=tower
rabbitmq_username=tower
rabbitmq_password=tower
rabbitmq_cookie=rabbitmqcookie

# Needs to be true for fqdns and ip addresses
rabbitmq_use_long_name=true
```

Example Inventory file for an external existing database

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

[database]

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

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

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

Example Inventory file for external database which needs installation

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

(continues on next page)

(continued from previous page)

```
[database]
database.example.com

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

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

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

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_password="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`

2.4 The Setup Playbook

Note: Ansible Tower 3.0 simplifies installation and removes the need to run `./configure/` as part of the installation setup. Users of older versions should follow the instructions available in the v.2.4.5 (or earlier) releases of the Tower Documentation available at: <http://docs.ansible.com/>

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
```

Note: Please note that a issue was discovered in Tower 3.0.0 and 3.0.1 that prevented proper system backups and restorations.

If you need to back up or restore your Tower v3.0.0 or v3.0.1 installation, use the v3.0.2 installer to do so.

SYSTEM TRACKING MIGRATION

System tracking feature was deprecated starting with Ansible Tower 3.2. However, you can collect facts by using the fact caching feature. Refer to [Fact Caching](#) for more detail. If you upgrade from a version earlier than Ansible Tower 3.0, you will find that your system tracking data (historical facts) has been migrated from MongoDB to PostgreSQL.

If you want to delete the old data in MongoDB, you can do so manually. First, connect to your mongo database using the mongo command line client, then run the following commands:

```
$ use system_tracking
$ db.runCommand( { dropDatabase: 1 } )
```

At this point, you can also remove the MongoDB packages.

ROLE-BASED ACCESS CONTROLS

Ansible Tower 3.0 has changed significantly around the way that the Role-Based Access Control (RBAC) system works. For the latest RBAC documentation, refer to the [Role-Based Access Controls](#) section in the Tower User Guide.

4.1 Enhanced and Simplified RBAC System

Based on user feedback, Ansible Tower both expands and simplifies its role-based access control. No longer is job template visibility configured via a combination of permissions on inventory, projects, and credentials. If you want to give any user or team permissions to use a job template, just assign permissions directly on the job template. Similarly, credentials are now full objects in Tower's RBAC system, and can be assigned to multiple users and/or teams for use.

A new 'Auditor' type has been introduced in Tower as well, who can see all aspects of the systems automation, but has no permission to run or change automation, for those that need a system-level auditor. (This may also be useful for a service account that scrapes automation information from Tower's API.)

4.2 Specific Changes to Note

There are a few changes you should keep in mind as you work with the RBAC system as redesigned for Ansible Tower:

- You no longer set the "team" or "user" for a credential. Instead, you use Tower's RBAC system to grant ownership, auditor, or usage roles.
- Deletion of job run data is now restricted to system and organization administrators.
- Projects no longer have multiple organizations. You *must* provide an organization when creating a new project through the API:

```
- projects/:id/organizations --> removed
```

- New Auditor type in Tower has been added which can see all aspects of the systems automation but does not have permission to run or change things.

JOB TEMPLATE CHANGES

Job templates have been updated in Tower to allow you more flexibility when creating and working with them.

5.1 Scan Jobs

If you maintained scan job templates in Ansible Tower 3.1.x and then upgrade to Ansible Tower 3.2, a new “Tower Fact Scan - Default” project is automatically created for you. This project contains the old scan playbook previously used in earlier versions of Ansible Tower.

5.2 Prompt on Launch

In prior versions of Ansible Tower, you could set “Prompt on Launch” against Extra Variables that you want to potentially pass through the job template. Starting with version 3.0, Ansible Tower allows you to prompt for an inventory selection, job type, and more.

Selecting “Prompt on Launch” means that even if a value is supplied at the time of the job template creation, the user launching the job will be prompted to supply new information or confirm what was entered in the job template originally.

The following job template settings allow for prompting at the time of launch:

- Job Type (run or check type jobs only, as scan jobs cannot be changed at the time of launch)
- Inventory
- Machine credential
- Limit
- Job Tags
- Extra variables

NEW JOB TEMPLATE
✕

DETAILS
COMPLETED JOBS
PERMISSIONS
NOTIFICATIONS

***NAME**

DESCRIPTION

***JOB TYPE**

Run

 Prompt on launch

***INVENTORY**

 Prompt on launch

***PROJECT**

***PLAYBOOK**

***MACHINE CREDENTIAL**

 Prompt on launch

CLOUD CREDENTIAL

NETWORK CREDENTIAL

FORKS

LIMIT

 Prompt on launch

***VERBOSITY**

JOB TAGS

 Prompt on launch

LABELS

EXTRA VARIABLES YAML JSON

1

 Prompt on launch

ADD SURVEY
CANCEL
SAVE

As you work with migrating your Tower 2.4.5 job templates to 3.x, please keep in mind the following:

- All “Prompt on Launch” fields are set to *False* by default after migrating to 3.x (new job templates created also have all “Prompt on Launch” fields set to *False* by default).
 - With one exception for those upgrading from 2.4.5 to 3.x: if a credential used in Tower 2.4.5 was null, the credential will be prompted for in 3.x.
- If you have Job Templates with a null credential, in the migration from 2.4.5 to 3.x, “ask_credential_on_launch” is set to *True*.
 - Note that there was no way to set a default credential in 2.4.5. However, in 3.x, you can set a default credential and select to prompt the user at launch time to confirm the default credential or change it to something new.
- All other “ask_xx_on_launch” prompts are set to *False*.
- Starting with Tower 3.x, if “ask_variables_on_launch” is set to *False*, extra variables passed at launch time (via UI or API) that are not part of an enabled survey are ignored.
- While there are no changes to how “ask_variables_on_launch” behaves, keep in mind that these variables combine with survey answers.

5.3 Permissions/RBAC Notes

Job template visibility is no longer configured via a combination of permissions on inventory, projects, and credentials. Admins who want to give any user or team permissions to use a job template can quickly assign permissions directly on the job template. Similarly, credentials are now full objects in Tower’s RBAC system, and can be assigned to multiple users and/or teams for use.

If a job template a user has been granted execution capabilities on does not specify an inventory or credential, the user will be prompted at run-time to select among the inventory and credentials in the organization they own or have been granted usage capabilities.

Users that are job template administrators can make changes to job templates; however, to make changes to the inventory, project, playbook, or credentials used in the job template, the user must also have the “Use” role for the project, inventory, and all credentials currently being used or being set.

5.4 Surveys

In prior versions of Ansible Tower, you had to select a checkbox to “Enable Survey” on the Job Template before a button appeared allowing you to “Create Survey”.

Enabling and creating surveys is much simpler in Ansible Tower.



At the bottom of each job template is a button () which opens a new dialog where you can enter your survey questions and responses.

NEW JOB TEMPLATE | SURVEY ON ✕

ADD SURVEY PROMPT

* PROMPT

DESCRIPTION

* ANSWER VARIABLE NAME ?

* ANSWER TYPE

MINIMUM LENGTH MAXIMUM LENGTH

DEFAULT ANSWER

REQUIRED

CANCEL ADD

PREVIEW

PLEASE ADD A SURVEY PROMPT ON THE LEFT.

CANCEL SAVE

Use the **ON/OFF** toggle button to quickly activate or deactivate this survey prompt.

Once you have entered the question information, click **Add** to add the survey prompt.

A stylized preview of the survey is presented, along with a **New Question** button. Click this button to add additional questions.

For any question, you can click on the **Edit** button to edit the question, the **Delete** button to delete the question, and click on the Up and Down arrow buttons to rearrange the order of the questions. Click **Save** to save the survey.

NEW JOB TEMPLATE | SURVEY ON

ADD SURVEY PROMPT

* PROMPT

DESCRIPTION

* ANSWER VARIABLE NAME ⓘ

* ANSWER TYPE

Choose an answer type ▼

REQUIRED

CANCEL ADD

PREVIEW

* WHICH GROUP(S) SHOULD INCLUDE THIS USER?

Enter groups, one per line.


⋮ ✎ 🗑

CANCEL SAVE

JOB OUTPUT VIEW CHANGES

With the update of the overall Tower user interface, it is worth noting the changes to how job results are displayed.

Job results for inventory syncs and SCM updates only show the Results and Standard Out of the job recently Run. Job results for playbook runs consist of Results, Standard Out, Details, and the Event Summary.





For more details regarding Job Results, refer to [Jobs](#) in the *Ansible Tower User Guide*.


6.1 Details

The **Details** area shows the basic status of the job (*Running*, *Pending*, *Successful*, or *Failed*), its start and end times, which template was used, how long the job run took, who launched it, and more. The buttons in the top right of the Details view allow you to relaunch or delete the job.

By clicking on these detail entries, where appropriate, you can view the corresponding job templates, projects, and other Tower objects.



DETAILS

STATUS	● Successful
STARTED	1/20/2017 1:27:48 PM
FINISHED	1/20/2017 1:27:54 PM
TEMPLATE	Demo Job Template
JOB TYPE	Run
LAUNCHED BY	admin
INVENTORY	Demo Inventory
PROJECT	● Demo Project
REVISION	347e44fea036c94d5f60e544de006453ee5c71ad
PLAYBOOK	hello_world.yml
MACHINE CREDENTIAL	Demo Credential
FORKS	0
VERBOSITY	0 (Normal)
EXTRA VARIABLES 	

1	
---	--


6.2 Standard Out Pane

The **Standard Out** pane shows the full results of running the Ansible playbook. This shows the same information you would see if you ran it through the Ansible command line, and can be useful for debugging. You can view the event summary, host status, and the host events. The icons at the top right corner of the Standard Out pane allow you to toggle the output as a main view () or to download the output ().

6.2.1 Events Summary


The events summary captures a tally of events that were run as part of this playbook:

- the number of plays
- the number of tasks
- the number of hosts
- the elapsed time to run the job template





6.2.2 Host Status Bar

The host status bar runs across the top of the **Standard Out** pane. Hover over a section of the host status bar and the number of hosts associated with that particular status displays.



6.2.3 Standard output view

The standard output view displays all the events that occur on a particular job. By default, all rows are expanded so that all the details are displayed. Use the collapse-all button () to switch to a view that only contains the headers for plays and tasks. Click the () button to view all lines of the standard output.

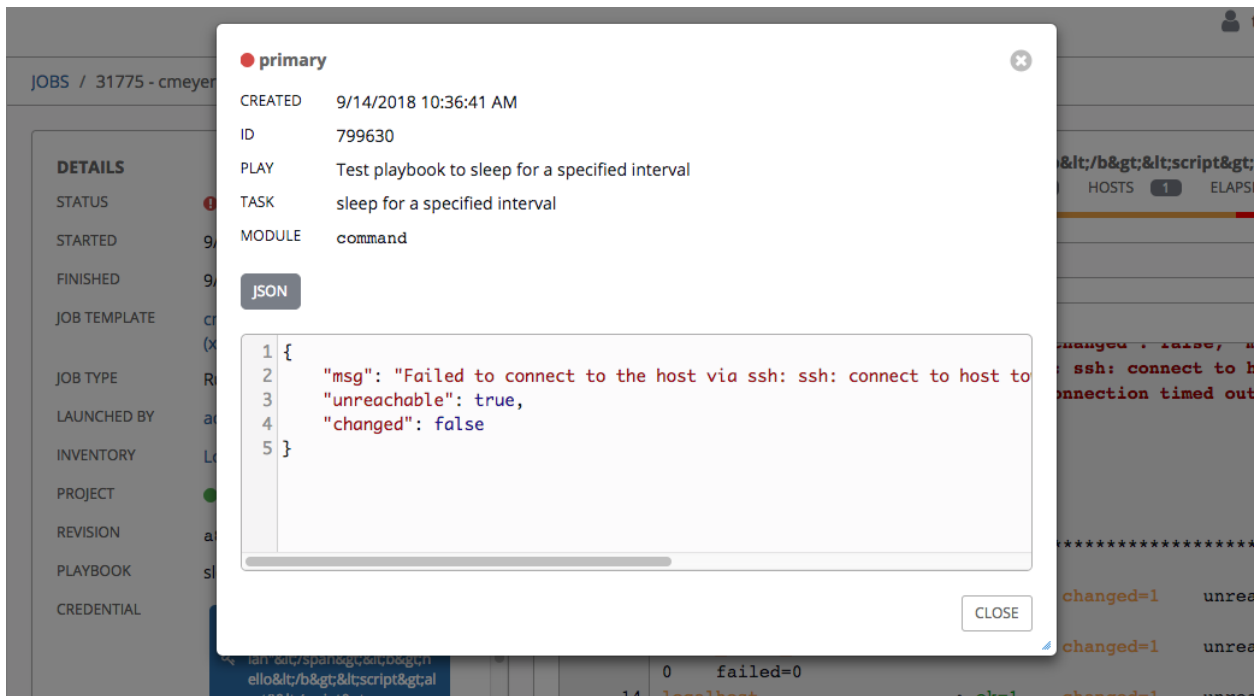
Alternatively, you can display all the details of a specific play or task by clicking on the arrow icons next to them. Click an arrow from sideways to downward to expand the lines associated with that play or task. Click the arrow back to the sideways position to collapse and hide the lines.



Click on a line of an event from the **Standard Out** pane and a **Host Events** dialog displays in a separate window. This window shows the host that was affected by that particular event.

6.2.4 Host Events

The **Host Events** dialog shows information about the host affected by the selected event and its associated play and task:



In the above example, this dialog box shows a single button for JSON, which is the only format given for the host event to view and display to you for this particular event. If Standard Out and Shell outputs are available, those buttons will display to allow you to see other output formats.

USING VIRTUALENV WITH ANSIBLE TOWER

Ansible Tower 3.0 and later uses *virtualenv*. *Virtualenv* creates isolated Python environments to avoid problems caused by conflicting dependencies and differing versions. *Virtualenv* works by simply creating a folder which contains all of the necessary executables and dependencies for a specific version of Python. Ansible Tower creates two *virtualenvs* during installation—one is used to run Tower, while the other is used to run Ansible. This allows Tower to run in a stable environment, while allowing you to add or update modules to your Ansible Python environment as necessary to run your playbooks. For more information on *virtualenv*, see the Python Guide to [Virtual Environments](#) and the *Python virtualenv* project itself.

By default, the *virtualenv* is located at `/var/lib/awx/venv/ansible` on the file system.

Tower also pre-installs a variety of third-party library/SDK support into this *virtualenv* for its integration points with a variety of cloud providers (such as EC2, OpenStack, Azure, etc.) Periodically, you may want to add additional SDK support into this *virtualenv*, which is described in further detail below.

Note: It is highly recommended that you run `umask 0022` before installing any packages to the virtual environment. Failure to properly configure permissions can result in Tower service failures. An example follows:

```
# source /var/lib/awx/venv/ansible/bin/activate
# umask 0022
# pip install --upgrade pywinrm
# deactivate
```

In addition to adding modules to the *virtualenv* that Tower uses to run Ansible, you can create new *virtualenvs* as described below.

7.1 Preparing a new custom *virtualenv*

Tower allows a different *virtualenv* to be specified and used on Job Template runs. To choose a custom *virtualenv*, first create one under `/var/lib/awx/venv/`:

```
$ sudo virtualenv /var/lib/awx/venv/my-custom-venv
```

Your newly created *virtualenv* needs a few base dependencies to properly run playbooks (Tower uses `memcached` as a placeholder for playbook artifacts):

```
$ sudo /var/lib/awx/venv/my-custom-venv/bin/pip install python-memcached psutil
```

From here, you can install additional Python dependencies that you care about, such as a per-*virtualenv* version of Ansible itself:

```
$ sudo /var/lib/awx/venv/my-custom-venv/bin/pip install -U "ansible == X.Y.Z"
```

Or you can add an additional third-party SDK that is not included with the base Tower installation:

```
$ sudo /var/lib/awx/venv/my-custom-venv/bin/pip install -U python-digitalocean
```

If you want to copy them, the libraries included in Tower's default virtualenv can be found using `pip freeze`:

```
$ sudo /var/lib/awx/venv/ansible/bin/pip freeze
```

In a clustered Tower installation, you need to ensure that the same custom virtualenv exists on every local file system at `/var/lib/awx/venv/`. Custom virtualenvs are supported on isolated instances. If you are using a custom virtual environment, it needs to also be copied or replicated on any isolated node you would be using, not just on the Tower node.

7.2 Assigning custom virtualenvs

Once you have created a custom virtualenv, you can assign it at the Organization, Project, or Job Template level:

```
PATCH https://awx-host.example.org/api/v2/organizations/N/  
PATCH https://awx-host.example.org/api/v2/projects/N/  
PATCH https://awx-host.example.org/api/v2/job_templates/N/  
  
Content-Type: application/json  
{  
    'custom_virtualenv': '/var/lib/awx/venv/my-custom-venv'  
}
```

An HTTP GET request to `/api/v2/config/` provides a list of detected installed virtualenvs:

```
{  
    "custom_virtualenvs": [  
        "/var/lib/awx/venv/my-custom-venv",  
        "/var/lib/awx/venv/my-other-custom-venv",  
    ],  
    ...  
}
```

You can also specify the virtual environment to assign to an Organization, Project, and Job Template from their respective edit screens in the Ansible Tower User Interface. Select the virtualenv from the **Ansible Environment** drop-down menu, as shown in the example below:

ORGANIZATIONS / Default

Default

DETAILS USERS NOTIFICATIONS

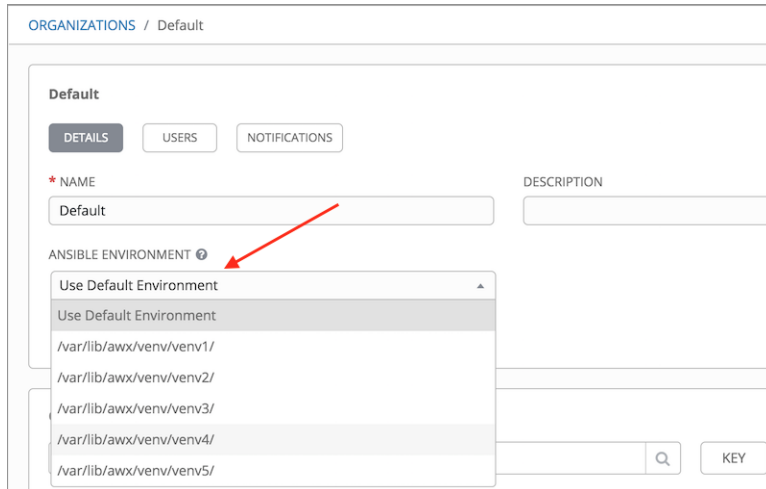
* NAME DESCRIPTION

Default

ANSIBLE ENVIRONMENT ⓘ

- Use Default Environment
- Use Default Environment
- /var/lib/awx/venv/venv1/
- /var/lib/awx/venv/venv2/
- /var/lib/awx/venv/venv3/
- /var/lib/awx/venv/venv4/
- /var/lib/awx/venv/venv5/

Q KEY



- genindex

COPYRIGHT © 2019 RED HAT, INC.

Ansible, Ansible Tower, 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.

A

Ansible, executing in a virtual environment, 20

I

installation script
 inventory file setup, 5
 playbook setup, 8
inventory file setup, 5

J

job output, sdout, 16
jobs
 event summary, 18
 events summary, 18
 host events, 19
 host status bar, 18
 host summary, 18
 job summary, 18

M

migration, 10
migration considerations, 10
MongoDB data removal, 10

P

permissions, 11, 12
playbook setup, 8
 installation script, 8
 setup.sh, 8
PostgreSQL data migration, 10
prompt on launch, 12

R

RBAC, 11, 12
release notes, v3.0, 2
roles, 11

S

scan jobs, 12
setup.sh
 playbook setup, 8

singleton roles, 11
surveys, 12
system tracking data, 10
system-wide roles, 11

U

upgrade, 4
upgrade considerations, 4

V

virtual environment, 20